

NORTHWESTERN UNIVERSITY

A Search for Scalar Ultralight Dark Matter Using Optical Cavities

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Department of Physics & Astronomy

By

Andra-Maria Ionescu

EVANSTON, ILLINOIS

August 2025

© Copyright by Andra-Maria Ionescu 2025

All Rights Reserved

ABSTRACT

A Search for Scalar Ultralight Dark Matter Using Optical Cavities

Andra-Maria Ionescu

Scalar field ultralight dark matter (ULDM) couples to ordinary matter to induce time oscillations in fundamental constants such as the fine structure constant α and the electron mass, m_e , which can be detected as oscillating strains in solids. We describe and demonstrate a new ULDM detection method which searches for resulting differential shifts in the resonant frequencies of two Fabry-Pérot cavities in the spectrum of the beat note between two laser beams resonant with them. Two possible implementation schemes are considered.

The first scheme involves a comparison of two cavities with rigid spacers of different lengths. Its realization and first results constitute the main focus of this work. A new apparatus was built, featuring high-finesse cryogenic cavities, designed for optimal stability to vibrations. A custom 4 K cryostat reduces the thermal noise, such as that associated with thermal fluctuations of the mirror surfaces, by almost one order of magnitude compared to a room-temperature setup. Multiple stages of horizontal and vertical vibration isolation decouple the cavities from laboratory vibrations. A room-temperature ultra-stable optical cavity filters the laser frequency noise. The first results from this newly-built apparatus improve by one to two orders of magnitude upon previous bounds from direct detection experiments at frequencies around 5 kHz and between 20 kHz and 90 kHz.

A second implementation scheme, where one of the cavities has its mirrors attached to a rigid spacer, while the other is made of two suspended mirrors, could provide an avenue for improving the sensitivity in future generations. Preliminary research on the construction of the suspended cavity is presented toward the end of this work.

ACKNOWLEDGEMENTS

There are so many people without the help of whom my journey through grad school would not have been possible.

First, I would like to thank my advisor, Professor Gerald Gabrielse, for taking a chance on me as a first-year graduate student with little research experience. His guidance, support and patience over the years allowed me to grow as a scientist and gain independence and confidence in my abilities. I am also grateful to the other co-PIs on the project, Professors Andrew Geraci and Tim Kovachy, for their advice and insights while working on this experiment.

I would also like to thank the other members of my committee, Professors Michael Schmitt and Tim Kovachy, for their support during my graduate studies, and for their insightful questions and suggestions regarding this thesis.

I am greatly indebted to Tejas Deshpande, with whom I closely collaborated on this project. I learned a lot from working with him, and without his dedication to the experiment these results would not have been possible. I also thank the other graduate students who worked on this project, Nick Miller and Aaron Wang, and a few junior students with whom I had the opportunity to work closely: Mark Nguyen and Prady Parshi.

I am deeply grateful to Mason Marshall, who patiently mentored me during my first years as a graduate student, and showed me the ropes in the lab while working together on the proton magnetic moment experiment. He and Kathryn Marable taught me many essential practical skills, and working with them was a true pleasure. I also acknowledge Xing Fan and Tharon Morrison, who helped me hugely in my transition from Penning traps to lasers and optics.

During my time in the Gabrielse group, I have been privileged to cross paths with numerous other graduate students, postdocs and research faculty, whose help, both academically and through their friendship, and insightful conversations have been invaluable: Daniel Ang, Yao Chen, Collin

Diver, Sam Fayer, Bingjie Hao, Wayne Huang, Nathan Jones, Gunn Khatri, Boran Kuzhan, Daniel Lascar, Siyuan Liu, Nathaniel McDonough, Cole Meisenhelder, John Mitchell, Thomas Myers, Edouard Nottet, Elise Novitski, Geev Nahal, Cris Panda, Lily Soucy, Benedict Sukra, Eric Tardiff, Maya Watts, Melissa Wessels, Daniel Zambrano, Junfeng Zhou, as well as many undergraduate students. Special thanks go to Benedict for constantly pushing me to work on my thesis every time we ran into each other in the lab.

I received a lot of help along the way from the staff of the machine shop: Stan Cotreau and Steve Sansone at Harvard University, and Salomon Rodriguez, Thang Bui, John Bussan, Bob Golenia, and Dinka Spirovska at Northwestern University. I am also thankful to the Northwestern Department of Physics & Astronomy staff, especially to Laura Nevins, whose assistance on many logistical issues has been tremendously valuable.

Many outstanding instructors contributed to my education over the years. In particular, my deepest appreciation goes to my high school physics teacher, Mrs. Corina Dobrescu, who nurtured my curiosity and passion for physics and inspired me to study it in college and beyond.

Last but not least, I am profoundly grateful to my parents, Irina and Radu Ionescu, and to my sister, Roxana Ionescu, for their unwavering and unconditional love and support, and for the patience with which they waited for me to finish my grad school journey. This thesis is dedicated to them.

TABLE OF CONTENTS

Acknowledgments	3
List of Figures	12
List of Tables	21
Chapter 1: Introduction	22
1.1 Dark matter overview	22
1.2 Ultralight scalar field dark matter	26
1.2.1 Scalar field Lagrangian	27
1.2.2 Oscillation of fundamental constants	28
1.2.3 Effect on the Bohr radius	30
1.2.4 Stochastic properties of the bosonic field	31
1.2.5 Relaxion	34
1.3 Current bounds on bosonic scalar field DM	35
1.4 Overview of work presented	39
Chapter 2: Optical Detection of a Dark Matter-Induced Strain	42
2.1 Fabry-Pérot cavities as strain detectors	42

TABLE OF CONTENTS	7
2.1.1 Optical cavities fundamentals	42
2.1.2 Pound-Drever-Hall locking	46
2.1.3 Dynamic response	47
2.2 Detector response	54
2.2.1 Mechanical transfer function	54
2.2.2 Beat note frequency modulation and differential transfer function	58
2.3 Noise Sources	60
2.3.1 Thermal Noise	61
2.3.2 Shot Noise	64
2.3.3 Vibration Noise	65
2.3.4 Noise estimates for the current apparatus and future prospects	67
Chapter 3: The Cryogenic Cavities	71
3.1 General design principles and fabrication	71
3.2 Geometry and support position optimization	74
3.2.1 Vibration sensitivity of the optical cavities	74
3.2.2 Finite element modeling	77
3.2.3 Effect of machining tolerances and axes misalignment	82
3.3 Suspension design and construction	84
3.3.1 Mounting	85
3.3.2 Four-wire suspension	86
3.3.3 Eddy current damping	87

TABLE OF CONTENTS	8
3.4 Cooling down the sapphire cavities	89
3.4.1 Heat extraction and thermal links	89
3.4.2 Temperature sensors and wiring	91
3.4.3 Cool down results	93
3.5 Optical properties	93
3.5.1 Cavity parameters	93
3.5.2 Crystalline coatings birefringence	94
Chapter 4: Optical Layout and Detection	97
4.1 Overview of the experimental setup	97
4.2 Laser system	99
4.3 Two-cavity setup	101
4.4 Integration of the room temperature reference cavity	104
4.4.1 ULE Cavity	105
4.4.2 Mode alignment	106
4.4.3 Stabilization scheme	107
4.4.4 Filtering scheme	110
4.5 Detection	111
4.5.1 Beat note optics	111
4.5.2 Delay line	113
Chapter 5: Ultra-low vibration cryostat	119

5.1	Cryostat	119
5.2	Vibration Isolation	125
5.2.1	Inverted Pendulum	126
5.2.2	GAS Filter	128
5.2.3	Vibration isolation stack	130
5.2.4	GAS filter performance	132
Chapter 6: Mechanical Resonances of the Sapphire Cavities		135
6.1	Room temperature measurements	137
6.1.1	Detection via electro-mechanical coupling	138
6.1.2	Optical detection	139
6.1.3	Results	141
6.2	In-situ measurements	144
6.2.1	Response to ambient vibrations	145
6.2.2	Response to external drive	145
6.2.3	Corroboration with finite element analysis	149
6.2.4	Quality factor estimations	151
Chapter 7: Improved direct bounds on ultralight dark matter		153
7.1	Data acquisition	153
7.2	Strain spectrum	154
7.2.1	Discrete Fourier transform review	154

TABLE OF CONTENTS	10
7.2.2 Conversion to strain	155
7.2.3 Noise floor properties	157
7.3 Investigation of strain spectrum peaks	158
7.4 Methodology for placing constraints on ultralight dark matter	159
7.4.1 Parameterization in terms of a DM signal	160
7.4.2 Bayesian analysis	161
7.4.3 Implementation details	165
7.4.4 Injection of a simulated dark matter signal	165
7.5 Results and discussion	167
Chapter 8: Toward the realization of a non-rigid optical cavity	171
8.1 Vibration isolation	171
8.2 Pendulum oscillation modes and transfer function	173
8.3 Development and characterization of a suspended mirror	174
8.3.1 Pendulum prototype	174
8.3.2 Measurements of the pendulum motion	176
8.3.3 Cavity with one suspended mirror	178
8.4 Active position control	178
8.4.1 Actuation requirements	179
8.4.2 Optical Sensors and Electro-Magnetic actuators (OSEMs)	180
8.4.3 OSEM characterization	183
8.4.4 First steps towards active control	184

TABLE OF CONTENTS	11
8.5 Challenges associated with the non-rigid cavity	186
Chapter 9: Conclusions	188
References	190
Appendix A: Open loop measurements of cavity transfer functions	199
Appendix B: Delay line frequency discriminator	203
Appendix C: List of parameters used for noise projections	207

LIST OF FIGURES

1.1	Diagram representing some of the main dark matter candidates, spanning more than 80 orders of magnitude in mass	25
1.2	Dark matter distribution as a function of the dark matter velocity (left) and angular frequency (right), displaying the characteristic asymmetrical profile	33
1.3	Simulation of the amplitude of the dark matter field. Shaded region illustrates the average field amplitude, given by Eq. (1.6) (corresponding to the limit where the total observation time $T_{tot} \gg \tau_c$). Inset shows the field in the $T_{tot} \ll \tau_c$ limit, where the field has a sinusoidal behavior	34
1.4	Limits on the possible density of a relaxion halo around the Earth (red) and Sun (blue) taken from reference [33], expressed as a ratio to the local dark matter density in the SHM $\rho_{DM} = 0.4 \text{ GeV}/c^2$	35
1.5	Current (solid lines) and projected (dashed lines) bounds on d_{m_e} . Details in text	36
1.6	Diagram of the experiment for the two versions of the experiment involving either (a) two rigid cavities or (b) a rigid and a non-rigid cavity	40
2.1	Diagram of a Fabry-Pérot cavity illustrating the incident, reflected, transmitted and intra-cavity (circulating) fields	43
2.2	Left: Standard configuration for the implementation of PDH locking. Right: Example of PDH error signal as a function of the frequency offset from one of the cavity resonances f_{res} , normalized by the FSR, for a finesse of 1000	46
2.3	Reflected, transmitted and intra-cavity fields	51
2.4	Plots of the cavity dynamic responses. Note that the frequency modulation response also applies to amplitude modulation	54

2.5 Mass-spring toy model for the cavity deformation 55

2.6 Qualitative behavior of the magnitude and phase of the cavity transfer function and overall detector response. Here, the optical poles are much lower in frequency than the mechanical ones and $\eta = 0.01$ 60

2.7 Spectrum of vibration noise measured on the optical table. The displacement noise falls approximately as f^{-2} above around 100 Hz and as f^{-4} above around 1000 Hz 66

2.8 (a) Noise analysis for current experiment parameters. Depicted are thermal noise contributions from the mirror coatings and substrates and cavity spacers, the photodetector shot noise and the predicted noise from vibrations. (b) Limits on d_{me} set by each of the considered noise sources, for a four-day measurement time. Bounds set by equivalence principle violation tests ("EP") and a theoretical limit ("naturalness") are included for comparison 69

2.9 Projected limits on d_{me} corresponding to fundamental and vibration noise, presented alongside current limits from equivalence principle violation tests and the naturalness bound for (a) two rigid cavities and (b) a rigid and a non-rigid cavity 70

3.1 (a) CAD model of the 15 cm cavity and 7.5 cm cavity (front view); (b) cavity side view, highlighting the support holes and the mirror surfaces; the mirror curvature is exaggerated for clarity; (c) support holes position for the 15 cm cavity; (d) support holes position for the 7.5 cm cavity 73

3.2 Horizontal deformation of the cavity under a vertical force. Arrows illustrate the direction and magnitude of the displacement in a qualitative sense (not drawn to scale) 75

3.3 Horizontal deformation of the cavity under a horizontal force parallel with the optical axis. Points placed symmetrically with respect to the center are displaced by similar amounts. Arrows illustrate the direction and magnitude of the displacement in a qualitative sense (not drawn to scale) 77

3.4 Comparison of the harmonic analysis of the cavity deformation with a prediction inferred from the static approximation 79

3.5 Implementation of the supports in the finite element simulation; green shaded regions show the areas over which boundary conditions are applied in the case of (a) point supports and (b) disk supports 80

LIST OF FIGURES

14

3.6 Transfer functions corresponding to the length change and mirror tilt of the short cavity as a function of the position of the supports 81

3.7 Transfer functions corresponding to the length change and mirror tilt of the long cavity as a function of the position of the supports 82

3.8 Transfer functions corresponding to the length change and mirror tilt of the long cavity as a function of the vertical position of the cut plane and of the hole depth . . 84

3.9 Left: horizontal short cavity on its support during assembly. Right: horizontal sapphire cavities side by side inside the cryostat 85

3.10 Rigid cavity support and suspension. Top inset: support pin and spring. Bottom inset: section through eddy current damper 86

3.11 Rigid cavity wire assembly 87

3.12 Diagram illustrating the principle behind eddy current damping 88

3.13 Left: thermal links diagram. Right: estimates of the time required to cool the sapphire cavities as a function of the number of thermal links (per cavity) between the suspension plate and the cradle 91

3.14 Pictures of the temperature sensors and 4-wire measurement diagram 92

3.15 Left: position of temperature sensors on the cavities and cradles. Right: cooling curves, plotted together with the estimated thermal conductivity of the heat links . . 94

3.16 Transfer functions of the sapphire cavities, demonstrating their behavior as low-pass filters with corner frequencies at 4.7 kHz (the 15 cm cavity) and 7.4 kHz (the 7.5 cm cavity) 95

3.17 Measured birefringence splitting for the 15 cm cavity (left) and 7.5 cm cavity (right) 95

4.1 Picture of the cryostat and of the optical setup 98

- 4.2 Comparison of the frequency noise in the beat note between the two cavity frequencies measured between pick-offs from the incident light and transmitted light. The dashed curve shows a theoretical prediction of the beat note in transmission obtained by applying a digital low-pass filter to the incident data. For simplicity, the pole of the filter was chosen at the 15 cm cavity pole (4.7 kHz). The prediction and measurement show good agreement up to 30 kHz. Despite the extra noise above 30 kHz (potentially caused by optics downstream of the pick-offs, or by the differential filtering by the two cavities, which have different poles), the transmission measurement scheme showcases a notable advantage 100
- 4.3 Pictures of the cold cavities optical setup (spread between two breadboards). The first breadboard (left) is placed on the optical table and contains the EOM, the AOM double-pass setups and the mode-matching telescopes. The second breadboard (right) is placed on the cryostat stand and contains optics for picking up light for the incident beat note, steering the beam into the cavities and the detection of the cavity reflections 103
- 4.4 Pictures of the ULE cavity and its enclosure 105
- 4.5 Diagram of the alignment of the resonant modes of the three cavities. Left: Alignment of the modes over multiple free spectral ranges. The green box marks the periodicity of the mode structure. Thicker arrows are used to represent the modes the laser was in resonance with when the wavelength was measured (the measured wavelengths are annotated). The light blue box highlights the modes selected for the experiment, with two options for the 15 cm cavity. Right: closer view of the two triplets of modes selected as options and the offsets between them 107
- 4.6 Left: diagram of the ULE cavity optical setup in its final configuration. Right: Picture of the ULE cavity inside its vacuum enclosure and the optical setup 108
- 4.7 Evolution of the optical setup. (a) Without the ULE cavity; (b) With the ULE cavity used only for stabilization; (c) with the ULE cavity filtering all the laser light 112
- 4.8 Plots of the beat note frequency noise without the ULE cavity, with the stabilization scheme and with the filtering scheme 113
- 4.9 Beat note detection setup in transmission (left) and using the pick-up from the incident beams (right) 114
- 4.10 Schematic of the evolution of the delay line frequency discriminator setup during each of the three main stages of the experiment 117

4.11	Diagram of the optical setup and detection chain	118
5.1	CAD drawing of a section through the cryogenic apparatus, illustrating the main components of the system (the cold chamber with the optical cavities, the vibration isolation stack and the pulse tube refrigerator)	120
5.2	Schematic (left) and picture (right) of the cryostat	122
5.3	Pictures of thermal links inside and outside the cold chamber	123
5.4	Left: anchoring of temperature sensors and heaters leads at 40 K and 4 K. Also visible are the thermal links cooling the 4 K plate. Right: Custom cryogenic connector at the 4 K plate	124
5.5	Cryogenic-compatible window mounts for the radiation shields	125
5.6	Inverted pendulum working principle diagram and pictures of a test inverted pendulum outside of the cryostat. The close-up picture shows the thin flexure at the top. A similar structure is found at the bottom of the leg (inside the inverted pendulum tower)	127
5.7	GAS Filter	130
5.8	CAD drawing of the room temperature vibration isolation stack, showcasing the seven stages of vibration isolation (left), together with a photo of the vibration isolation system (top right) and a simplified diagram (bottom right)	131
5.9	Example of transfer function of a test GAS filter	133
5.10	Left: magic wand diagram. For simplicity, only one blade and one wand are represented. Right: transfer functions before and after the integration and tuning of the magic wand	134
6.1	Catalog of the lowest resonant modes of the (a) 15 cm and (b) 7.5 cm sapphire cavities, calculated using ANSYS. The longitudinal modes are highlighted in red	136
6.2	Room temperature setup for resonance measurements. Left: piezo glued to the side of a spare spacer. Right: spacer with piezo seated in its cradle; a fiber interferometer detects the motion of a mirror glued to one end of the spacer	137

6.3	(a) Mechanical transfer function and impedance frequency response of a piezo-electric element. (b) Equivalent circuit diagram, highlighting the electrical and mechanical components. (c) Equivalent circuit diagram, highlighting the series and parallel resonant circuits	140
6.4	Frequency dependence of the impedance magnitude and phase for a piezo hanging from a wire or attached to the 15 cm spacer. Dashed lines mark the expected values of the spacer resonant frequencies (red for the longitudinal mode, black for other modes). The inset shows a close-up of the resonance around 35 kHz	140
6.5	Optical detection scheme	141
6.6	Resonances of the 15 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode) . . .	142
6.7	Resonances of the 7.5 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode) . . .	143
6.8	Noise peaks in the transmission spectra of the two cavities. Theoretical frequencies of the mechanical resonances of the short cavity (blue) and long cavity (orange) are represented by vertical lines (for the main ones, the mode shapes are also shown) .	146
6.9	Beat note noise peaks around the longitudinal resonant frequencies of the two cavities. The theoretical values of the longitudinal frequencies are shown as solid vertical lines. Other resonances in the plotted frequency range are marked by dashed lines	146
6.10	Lock-in detection diagram	148
6.11	Candidates for the short cavity longitudinal mode identified through lock-in measurements	148
6.12	Candidates for the long cavity longitudinal mode identified through lock-in measurements. Scans with the drive off are also presented for reference	149
6.13	Resonances of the 7.5 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode) . . .	150

- 6.14 Resonances of the 7.5 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode) . . . 151
- 7.1 (a) Strain noise in the 0-90 kHz frequency range (average over 10^5 FFTs); b) Strain noise in a narrow window around 50 kHz of a width comparable to the dark matter characteristic linewidth at this frequency; the orange curve shows the dark matter lineshape; the dashed line marks the average value of the noise within the pictured window 156
- 7.2 Statistical distribution of the strain noise DFT over a window of width $\sim 11/\tau_c$ around 50 kHz. The real and imaginary parts (panels (a) and (b)) fit to a Gaussian distribution. The magnitude (panel (c)) is Rayleigh distributed. The red lines represent fits to the histograms. 157
- 7.3 (a) Downsampled amplitude spectral density (ASD) of the noise floor, alongside the noise mean value and the 95% confidence level threshold for dark matter detection. The dashed rectangle highlights the peak that is magnified in the adjacent plot; (b) Example of the profile of a noise peak at 73.75 kHz 159
- 7.4 Detector response (black), conservatively estimated as the minimum of the transfer functions obtained using lower (blue) and upper (orange) bounds of the mechanical resonances of the cavities. Shaded regions mark the uncertainty intervals for each cavity's resonance 161
- 7.5 (a) Plot of the frequency noise PSD with an added artificial dark matter-like signal (the peak marked with red) corresponding to $d_{m_e} = 10^6$. Top right inset: plot of the simulated dark matter signal, displaying the expected asymmetric lineshape; (b) Recovered dark matter signal after passing the injected data through the analysis algorithm, with 95% and 99.5% confidence level (CL); (c) Scan over a range of d_{m_e} values (computed at the 95% CL), with the dashed line marking the expected values, to which the extracted result is consistently close 166
- 7.6 New bounds on d_{m_e} in the context of a) the Standard Halo Model and b) a relaxation star bound to Earth. Also shown are limits from two other direct ULDM detectors, cavity-fiber [49] and Cs-cavity [38], as well as indirect limits from EP violating tests [18] and a theoretically motivated target [97] 168
- 7.7 Comparison of bounds on d_{m_e} in the SHM over a wider range with other direct and indirect observations and theoretical limits. Details in text 169

8.1	Predicted vibration noise for a single-stage and a double-stage pendulum, compared to the expected total sum of thermal and shot noise	172
8.2	a) Toy model for the double pendulum; b) Transfer function for a horizontal displacement; the dashed line indicates the expected slope for a f^{-4} roll-off; c) Pendulum resonant modes and the approximate frequencies where they can be expected to occur	174
8.3	Left: diagram of a test mass and a recoil mass with a damping mechanism. Center: CAD model of the double-stage pendulum. Right: view of the test mass and recoil mass with the front cover removed	176
8.4	Schematic representation of the optical lever, showing the conversion of longitudinal and angular displacements of the mirror to displacements of the beam at the photodetector	177
8.5	Picture of the setup used for measurements of the pendulum motion and example of a measured trace	177
8.6	Setup of the low-finesse optical cavity with a suspended mirror, as well as a picture of light coupled in the TEM ₀₀ mode, detected by an InGaAs camera, and an example of the transmission signal detected by a photodetector	179
8.7	Diagram of an OSEM showing the main components	181
8.8	Pictures of OSEM components during assembly: (a) LED glued to its holder; (b) photodiode glued to its holder; (c) coil during winding process; (d) photodiode viewed from the LED side; (e) flag with magnet inside; (f) fully assembled OSEM .	182
8.9	Dependence of the PD voltage on the flag position. Inset: setup for the OSEM sensor tests	183
8.10	Left: picture of the OSEM on the pendulum during the actuation tests (the top OSEM, which was used for the test, is highlighted in green). Right: plot of the measured flag position against the coil current	184
8.11	Output of the OSEM photodiode and of the optical lever quadrant detector as the PID loop for stabilizing the pendulum is turned off and on	185
8.12	Comparison of the amplitude spectral density of the flag displacement with the pendulum active control on and off	186

- A.1 Measurement scheme for the open loop transfer function. The response of the optical cavity, together with the optics and electronics involved in the Pound-Drever-Hall locking setup is measured for one cavity while the laser is locked to the other cavity. The red box marks the device under test (DUT) 200
- A.2 Bode plot of the measured open-loop transfer function of the 7.5 cm cavity, optical setup, AOM, photodetector 201
- A.3 Measured transfer functions for the 7.5 cm (left) and 15 cm (right) cavities. The magnitudes were fitted to a single pole with corner frequencies at 7.4 kHz for the 7.5 cm cavity and 4.7 kHz for the 15 cm cavity 202
- B.1 Schematic of the delay line frequency discriminator 203
- B.2 Left: magnitudes of the beat note carrier and of the sideband from an applied 60 kHz dither before demodulation. Right: 60 kHz sideband after passing through the delay line frequency discriminator 206

LIST OF TABLES

3.1	Geometrical parameters of the two cryogenic cavities	83
3.2	Length and tilt transfer functions for a 0.1 mm longitudinal or transverse displacement of the cavity support holes	84
3.3	Optical properties of the cryogenic cavities	96
4.1	Summary of ULE cavity characteristics	106
6.1	Summary of driven response tests for the longitudinal mode candidates	149
7.1	Data acquisition parameters	154
8.1	Actuation requirements	180
C.1	Experimental parameters used in the noise projections in Fig. 2.8 and Fig. 2.9 . . .	208

CHAPTER 1

INTRODUCTION

The overwhelming success of the Standard Model is evidenced by the accuracy with which its prediction had been tested in both low-energy table-top experiments and at high energy in particle accelerators, but numerous questions still remain unsolved. A major grievance is that the Standard Model, as it stands now, does not account for a majority of the mass observed in the Universe.

Studies of gravitational interaction on galactic and cosmological scales suggest that about 85% of the total matter content of the Universe and around a quarter of its critical energy density is composed of non-baryonic, non-luminous matter known as "dark matter". However, despite decades of experimental efforts, no direct detection of dark matter has ever occurred, and its composition and non-gravitational interactions with Standard Model particles so far remain elusive. Understanding the nature and properties of dark matter represents one of the salient current challenges in fundamental physics.

1.1 Dark matter overview

The most widely accepted cosmological model to date is the Λ CDM (Lambda Cold Dark Matter) model. Within this paradigm, the Universe is flat, expands with an acceleration characterized by the cosmological constant Λ , and consists of 4.9% baryonic matter, 68.5% dark energy and 26.5% dark matter (DM) [1]. Of the properties of the latter, very little is understood. Its interactions with electro-magnetic forces, if at all existent, are very weak (hence "dark"). It is also color neutral, although it is possible that it is subject to weak interactions. It has a long lifetime, on a scale similar with the age of the Universe, since dark matter which determined the distribution of the cosmic microwave background must still exist today. In order to be consistent with

the formation of large-scale structures in the Universe, it also needs to be non-relativistic (cold). Additionally, it is expected to have very weak self-interactions, as inferred from studies of galaxy cluster mergers [2].

The existence of dark matter is concurrently supported by a variety of independent astronomical observations at different scales. We enumerate some of the most prominent ones here.

- 1. Velocity dispersion of galaxies:** The first hints of the presence of undiscovered non-luminous matter came from Zwicky's study of galaxy clusters [3]. Based on observed velocity dispersion of galaxies in the Coma cluster and employing the Virial theorem, he inferred values for the mass of the galaxies that were much greater than those of the visible objects. While the disagreement was later ameliorated when additional mass from a halo of hot gas was discovered, a significant gap - about a factor of 6 - between the predicted and measured mass remained, a finding that was later confirmed in other galaxy clusters, as well.
- 2. Galactic rotation curves:** Additional evidence for dark matter on galactic scales has emerged from the rotational velocities of stars in spiral galaxies. In equilibrium, balance between gravitational and centrifugal forces imply a $1/\sqrt{r}$ dependency on radius of the velocity of stars situated at the periphery of the galaxy. This expectation was, however, contradicted by Doppler shift spectral line measurements, which revealed the rotation curves to be flat far beyond the region where most of the mass is concentrated ([4, 5]). These observations are consistent with the presence of a halo of non-baryonic matter of a radius much larger than that of the visible matter.
- 3. Gravitational lensing** in galaxy clusters highlighted a distribution of the total mass that was different from that of the baryonic matter, detected by its X-ray radiation [6].
- 4. Cosmic microwave background radiation (CMB):** At cosmological scales, the existence of dark matter is supported by the observed fluctuations in the cosmic microwave back-

ground. Decoupled from the baryonic matter at the time of recombination, prior to which the two evolved together due to their strong interactions through Compton scattering, the CMB has imprinted in its density distribution the baryonic distribution at recombination. Its striking uniformity thus indicates that the baryonic distribution at early stages of the Universe was similarly uniform, a condition which would have prevented the formation of large scale structure in the Universe. The presence of a gas of cold dark matter can resolve this inconsistency, by having the potential to provide the necessary initial density fluctuations required for the production of the large-scale structures we observe today.

5. Big Bang nucleosynthesis (BBN): BBN refers to the production of nuclei heavier than hydrogen which occurred between around 10 s and 20 min after the Big Bang. Models of BBN that assume the presence of baryonic matter only underestimate the density of deuterium present in the Universe. However, the inclusion of dark matter reconciliates the predictions with the observed values [7].

All the observations above could be utilised to predict the relative proportions of dark matter to baryonic matter in the Universe, rendering ratios consistent with the Λ CDM model. Thus, the evidence for the existence of dark matter is compelling and, while alternative explanations exist in the form of modified theories of gravity (e.g. [8]), none of these theories currently succeeds to adequately explain all the observations, particularly at large scales.

Notwithstanding our ability to measure the dark matter mass on astronomical scales through various independent methods, the mass of its individual constituents is very much unknown. Its Compton wavelength has to be smaller than the scale of the smallest known structures (brown dwarves), otherwise their formation would have been prevented. This sets a lower limit on the dark matter mass of 10^{-22} eV [1]. Model independent upper limits come from the condition of stability against tidal effects on visible structures within dark matter halos, which restrict the dark matter mass below a few solar masses [9]. An example candidate at these scales comes in the form

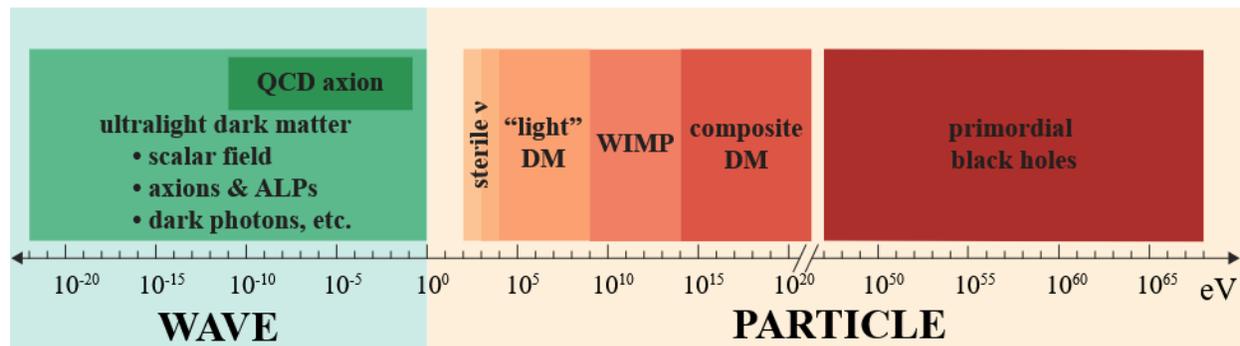


Figure 1.1: Diagram representing some of the main dark matter candidates, spanning more than 80 orders of magnitude in mass

of primordial black holes (black holes formed in the early Universe, before BBN) [10, 11]. Thus, the possible mass range for dark matter covers a tremendous parameter space, extending over more than 80 orders of magnitude, over which many models have been hypothesized (Fig. 1.1).

For a long time, searches for dark matter focused on particle-like candidates, either in the form of individual particles, or as composite objects [12]. Among these, a particularly attractive option is represented by a weakly interacting massive particle (WIMP)[13]. These are particles with masses ranging from about 10 GeV to $\mathcal{O}(\text{TeV})$, produced thermally in the early Universe via freeze-out [14], which can interact with the Standard Model particles via W and Z bosons. Introduced to solve the hierarchy problem [15], the elevated interest that WIMPs have seen has also been motivated by the fact that the predicted amounts of WIMP particles adequately explain the observed dark matter mass. However, years of searches using nuclear recoil experiments and particle colliders (e.g. [16, 17]) have not so far lead to any detection and, as their parameter space is gradually being ruled out, efforts started being redirected towards different classes of DM candidates.

In this context, ultralight dark matter (ULDM) with mass $m_\phi \lesssim 10$ eV stands out as a compelling alternative [18, 19]. Such light particles are necessarily bosonic [20], as otherwise Pauli’s exclusion principle, in conjunction with the existence of an upper bound for the dark matter velocity in order to remain trapped in the galactic gravitational potential, would prevent the accumula-

tion of a large enough density of dark matter to agree with observations (condition known as the Tremaine-Gunn limit [21]). Moreover, if they are to constitute all or a majority of dark matter in the Universe, they have to be characterized by a high mode occupation number ($N_{dB} = n_\phi \lambda_{coh}^3 \gg 1$), which allows them to be treated as a classical field. As such, searches of ultralight dark matter target effects related to their coherent behavior.

Notable ULDM candidates, classified based on their behavior under Lorentz transformations, are scalar fields (e.g. the dilatons and moduli that arise in string theories), pseudo-scalars, such as the axion and axion-like particles (ALPs) and vectors (an example of which would be the so-called "dark photon"). Out of these, an especially promising candidate is the QCD axion, which was theorized as the Goldstone boson of a potential new symmetry (the Peccei-Quinn symmetry) postulated as a solution to the strong CP problem [22], but could also explain the observed dark matter density. Axions and ALPs can be produced non-thermally via the misalignment mechanism [23] and can be detected via their conversion to photons in an electro-magnetic field as the result of its mixing with photons.

For the rest of this work we will focus on an ultralight candidate with a scalar, parity-even, coupling to the Standard Model. The properties of a scalar field as dark matter and its effects on observable physical quantities are explored in the next section.

1.2 Ultralight scalar field dark matter

As a potential dark matter candidate, a scalar field is well-motivated: it is consistent with the Standard Halo Model (SHM) for dark matter and it is a natural product of beyond Standard Model theories such as string theories, in which it can arise as dilatons and moduli. Similar to axions, it can be produced through vacuum misalignment, where the field, displaced from its equilibrium position in its initial state, starts oscillating when the Hubble rate becomes similar to its mass.

1.2.1 Scalar field Lagrangian

Under the assumption of negligible self-interactions, scalar ULDM with a mass m_ϕ can be represented as a classical, non-relativistic, field, which can be incorporated into the Standard Model by the following addition to the Lagrangian density \mathcal{L}_{SM} [24]:

$$\mathcal{L} = \frac{1}{2} \partial_\mu \phi \partial^\mu \phi - V(\phi) + \mathcal{L}_{SM} + \mathcal{L}_{int}. \quad (1.1)$$

The first two terms convey the kinetic and potential energy associated with the scalar field ϕ . The last term represents the interaction of the scalar field with standard model particles.

The potential term $V(\phi)$ can be expressed as a Taylor expansion in ϕ , of which only the first few terms are typically kept:

$$V(\phi) = \sum_{n=0}^{\infty} \frac{c_n}{n!} \phi^n. \quad (1.2)$$

The simplest model, which we will consider here, is a quadratic potential, which gives the field its mass:

$$V(\phi) = \frac{1}{2} \left(\frac{m_\phi c}{\hbar} \right)^2 \phi^2. \quad (1.3)$$

The conclusions we derive remain intact, however, with the inclusion of higher order terms. The Euler-Lagrange equation applied to the Lagrangian of the scalar field in vacuum, renders the Klein-Gordon equation:

$$\left(\partial_\mu \partial^\mu + \left(\frac{m_\phi c}{\hbar} \right)^2 \right) \phi = 0, \quad (1.4)$$

the solution of which describes a field undergoing both temporal and spatial oscillations:

$$\phi(t) = \Phi_0 \cos(\omega_\phi t - \mathbf{k} \cdot \mathbf{r}). \quad (1.5)$$

Here, $\omega_\phi = m_\phi c^2/\hbar$ is the Compton angular frequency.

The amplitude Φ_0 , as will be discussed in section 1.2.4, is stochastic. However, under the assumption that the scalar field constitutes all (or a majority) of the dark matter in the Universe, its average value:

$$\Phi_0 = \frac{\hbar}{m_\phi c} \sqrt{2\rho_{\text{DM}}} \quad (1.6)$$

is related to the local dark matter density ρ_{DM} .

1.2.2 Oscillation of fundamental constants

The simplest possible interaction of scalar field dark matter with the standard model fields can be described by linear terms in the Lagrangian of the form:

$$\mathcal{L}_{int} = \frac{\sqrt{4\pi\hbar c}}{E_P} \phi \left\{ \frac{d_e}{4e^2} F_{\mu\nu} F^{\mu\nu} - d_g \frac{\beta_3}{2g_3} G_{\mu\nu}^A G^{A\mu\nu} - \sum_{q=e,u,d} (d_{m_q} + \gamma_m d_g) m_q \bar{\psi}_q \psi_q \right\}. \quad (1.7)$$

Here, $F_{\mu\nu}$ is the electro-magnetic tensor, G are the gluonic field tensors, β_3 is the beta function for the running of the strong coupling constant g_3 , γ_m is the anomalous dimension giving the energy running of the masses of the strongly-coupled fermions, ψ_q are the fermion spinors and $E_P = \sqrt{\hbar c^5/G}$ is the Planck energy. The dark matter energy scales we are concerned with in this work (\ll keV) allow for interactions with the electron and up and down quarks only. The strengths of these interactions are characterized by 5 dimensionless coupling parameters: d_e acts on the fine structure constant α_0 , d_{m_e} , d_{m_u} , d_{m_d} couple to the masses of the electron and up and down quarks, while d_g is a coupling to the QCD mass scale Λ_3 .

Meanwhile, at these low energies, the Standard Model Lagrangian can be, in turn, written as:

$$\mathcal{L}_{SM} = -\frac{1}{4e^2}F_{\mu\nu}F^{\mu\nu} - \frac{1}{4}G_{\mu\nu}^AG^{A\mu\nu} + \sum_{q=e,u,d} (i\bar{\Psi}_q\mathcal{D}(A, g_3A^A)\Psi_q - m_q\bar{\psi}_q\psi_q). \quad (1.8)$$

The first two terms are related to bosonic interactions for a gauge field A. The third term represents weak and strong interactions with the electron and the u and d quarks. $\mathcal{D}(A)$ is the Dirac operator coupled to the gauge field A.

The addition of \mathcal{L}_{int} to the Standard Model Lagrangian leads to a summation of similar terms that manifests itself as an effective modification of the fundamental constants α_0 , m_e , m_u , m_d and Λ_3 . We hereby focus on the first two, which have direct implications for the work presented here.

For instance, the terms associated with the electro-magnetic interaction add up to:

$$-\frac{1}{4e^2}F_{\mu\nu}F^{\mu\nu} + \frac{\sqrt{4\pi\hbar c}}{E_P}d_e\phi\frac{1}{4e^2}F_{\mu\nu}F^{\mu\nu} = -\frac{1}{4e^2}\left(1 - \frac{\sqrt{4\pi\hbar c}}{E_P}d_e\phi\right)F_{\mu\nu}F^{\mu\nu}, \quad (1.9)$$

equivalent to an effective transformation:

$$\alpha_0 \rightarrow \frac{\alpha_0}{1 - (\sqrt{4\pi\hbar c}/E_P)d_e\phi} \approx \alpha_0\left(1 + \frac{\sqrt{4\pi\hbar c}}{E_P}d_e\phi\right). \quad (1.10)$$

The approximation is valid at linear level, which is the level at which we consider couplings here. Similarly, the terms corresponding to the interactions with the electron field sum up to $-m_e\bar{\psi}_e\psi_e\left(1 + \frac{\sqrt{4\pi\hbar c}}{E_P}d_{m_e}\phi\right)$, leading to an effective electron mass:

$$m_e \rightarrow m_e\left(1 + \frac{\sqrt{4\pi\hbar c}}{E_P}d_{m_e}\phi\right). \quad (1.11)$$

The $\phi(t)$ dependence of these constants implies effective temporal and spatial oscillations at the Compton frequency of the scalar field:

$$\begin{aligned}\frac{\delta\alpha(t, \mathbf{x})}{\alpha} &= d_e \frac{\sqrt{4\pi\hbar c}}{E_P} \phi(t, \mathbf{x}), \\ \frac{\delta m_e(t, \mathbf{x})}{m_e} &= d_{m_e} \frac{\sqrt{4\pi\hbar c}}{E_P} \phi(t, \mathbf{x}).\end{aligned}\tag{1.12}$$

It is worth mentioning that while higher order couplings are possible (for instance, quadratic couplings would cause oscillations at twice the Compton frequency), their effect is suppressed by a factor of $(\phi/M_P)^{n-1}$, where n is the coupling order and M_P the Planck mass. Moreover, in the case of models with quadratic coupling, screening mechanisms in the vicinity of massive bodies are expected to diminish the sensitivity of terrestrial experiments. Such models will therefore not be considered here.

1.2.3 Effect on the Bohr radius

The time and space oscillations of the electron mass m_e and of the fine constant structure α are reflected in effects on the Bohr radius, as follows:

$$a_0(t) = \frac{\hbar}{m_e(t)c\alpha(t)} = \frac{\hbar}{m_e c \alpha} \frac{1 - \frac{\sqrt{4\pi\hbar c}}{E_P} d_e \phi}{1 + \frac{\sqrt{4\pi\hbar c}}{E_P} d_{m_e} \phi},\tag{1.13}$$

or

$$\frac{\delta a_0(t)}{a_0} = -\frac{\delta m_e(t)}{m_e} - \frac{\delta\alpha(t)}{\alpha(t)} = -(d_e + d_{m_e}) \frac{\sqrt{4\pi\hbar c}}{E_P} \phi,\tag{1.14}$$

This modulation affects the size of atoms and of chemical bonds, resulting in, for instance, modulations of the atomic energy levels and modulations of the lengths of solids [25, 26]. The resulting ULDM "strain" driving the length of a solid is given by:

$$h_{\text{DM}}(t) = -(d_e + d_{m_e}) \frac{\sqrt{4\pi\hbar c}}{E_P} \phi.\tag{1.15}$$

This is the effect that the experiment described in this work exploits.

1.2.4 Stochastic properties of the bosonic field

The most common model for the dark matter distribution, dubbed the Standard Halo Model (SHM), stipulates that the dark matter particles in galaxies are trapped and virialized in their gravitational potential, forming a spherical halo which extends far beyond the luminous matter, characterized by an isotropic Maxwell-Boltzmann velocity distribution, with a cut-off determined by the escape velocity. For our galaxy, the dark matter has a local density $\rho_{\text{DM}} \approx 0.4 \text{ GeV/cm}^3$ [20, 27]. As the Solar System moves through the galactic halo towards the Cygnus constellation with a velocity relative to the galactic rest frame $v_g \approx 230 \text{ km/s} \approx 10^{-3}c$, an observer in the Solar System frame measures a spread in the DM velocities $v_{\text{vir}} \approx 166 \text{ km/s}$ [28]. Ignoring the cut-off (at 650 km/h), as well as anisotropies, the velocity distribution can be expressed as:

$$F(\mathbf{v}) = \frac{1}{(2\pi v_{\text{vir}}^2)^{3/2}} \exp\left(-\frac{(\mathbf{v} - \mathbf{v}_g)^2}{2v_{\text{vir}}^2}\right), \quad (1.16)$$

where the bold characters denote vectors.

The spread in velocities results in a spectral broadening of the dark matter frequency, characterized by a corresponding coherence time related to the dark matter angular frequency ω_ϕ :

$$\tau_c = \frac{\hbar}{m_\phi v_{\text{vir}}^2}, \quad (1.17)$$

and by a coherence length:

$$\lambda_c = \hbar/(m_\phi v_{\text{vir}}). \quad (1.18)$$

In order to derive an expression for the frequency distribution, it is convenient to first rewrite the distribution in Eq. (1.16) in terms of the velocity amplitude, by switching to a polar coordinate system (v, θ, ϕ) , with θ being the angle between the directions of \mathbf{v} and \mathbf{v}_g , and then integrating

over the angular coordinates:

$$F(v)dv = \left(\int_0^{2\pi} \int_0^\pi \frac{1}{(2\pi v_{vir}^2)^{3/2}} \exp\left(-\frac{v^2 + v_g^2 - 2vv_g \cos(\theta)}{2v_{vir}^2}\right) v^2 \sin(\theta) d\theta d\phi \right) dv, \quad (1.19)$$

which can be further simplified as:

$$F(v) = \sqrt{\frac{2}{\pi}} \frac{v}{v_{vir} v_g} \exp\left(-\frac{v^2 + v_g^2}{2v_{vir}^2}\right) \sinh\left(\frac{vv_g}{v_{vir}^2}\right). \quad (1.20)$$

The angular frequency of a particle can be related to its distribution via the Doppler shift:

$$\omega = \omega_\phi \left(1 + \frac{v^2}{2c^2}\right), \quad (1.21)$$

which leads to:

$$d\omega = \frac{\omega_\phi}{c^2} v dv = \frac{\omega_\phi}{c} \sqrt{2 \left(\frac{\omega}{\omega_\phi} - 1\right)} dv. \quad (1.22)$$

Using the relation:

$$F(v)dv = F(\omega)d\omega, \quad (1.23)$$

and introducing $\eta \equiv v_g/v_{vir}$, as well as plugging in $\tau_c = c^2/(\omega_\phi v_{vir}^2)$, we obtain:

$$\begin{aligned} F(\omega) &= F(v) \frac{c^2}{v\omega_\phi} = \sqrt{\frac{2}{\pi}} \frac{v}{v_{vir} v_g} \exp\left(-\frac{v^2 + v_g^2}{2v_{vir}^2}\right) \sinh\left(\frac{vv_g}{v_{vir}^2}\right) \frac{c^2}{v\omega_\phi} \\ &= \sqrt{\frac{2}{\pi}} \frac{c^2}{v_{vir} v_g \omega_\phi} \exp\left(-\frac{2c^2(\omega/\omega_\phi - 1) + v_g^2}{2v_{vir}^2}\right) \sinh\left(\frac{c\sqrt{2(\omega/\omega_\phi - 1)}v_g}{v_{vir}^2}\right) \\ &= \sqrt{\frac{2}{\pi}} \frac{\tau_c}{\eta} \exp\left(-\frac{\eta^2}{2}\right) \exp\left((\omega - \omega_\phi)\tau_c\right) \sinh\left(\eta\sqrt{2(\omega - \omega_\phi)\tau_c}\right). \end{aligned} \quad (1.24)$$

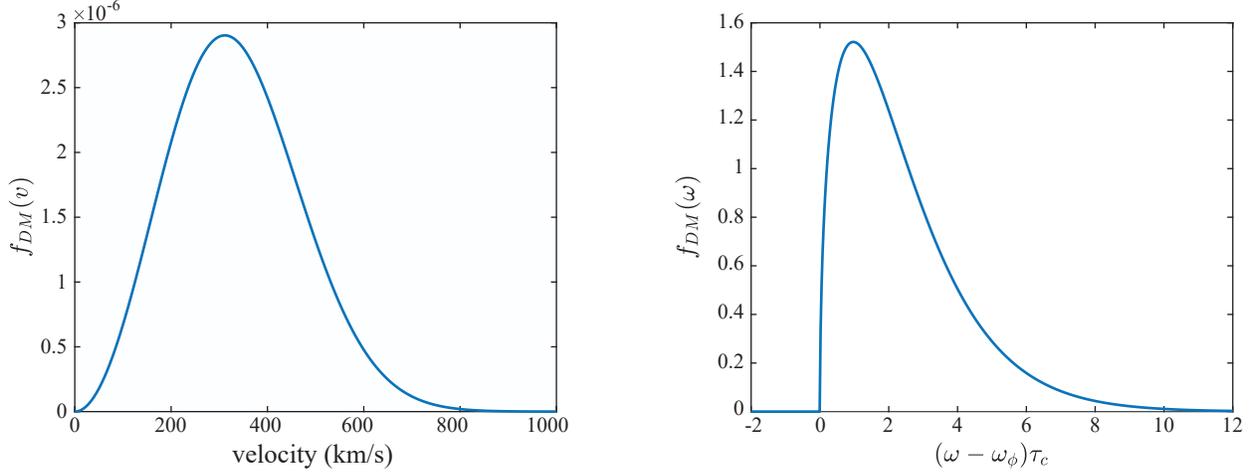


Figure 1.2: Dark matter distribution as a function of the dark matter velocity (left) and angular frequency (right), displaying the characteristic asymmetrical profile

Expressed in terms of the observed (Doppler shifted) Compton frequency $\omega'_\phi = \omega_\phi (1 + v_g^2/(2c^2))$, we can finally rewrite the frequency distribution as:

$$F(\omega) = \sqrt{\frac{2}{\pi}} \frac{\tau_c}{\eta} \exp(-\eta^2) \exp((\omega - \omega'_\phi)\tau_c) \sinh\left(\eta\sqrt{\eta^2 + 2(\omega - \omega'_\phi)\tau_c}\right). \quad (1.25)$$

This equation defines the dark matter characteristic lineshape, which is crucial for making inferences on a dark matter coupling from experimental data (as will become evident in section 7.4). An important feature of the lineshape is its high degree of asymmetry, a signature of a dark matter signal that helps differentiating it from noise peaks of a different origin. Note that during the last step of the previous derivation we assumed the normalization condition $\int_0^\infty (F(\omega)d\omega = 1/2$ instead of 1, which is equivalent to redefining $F(\omega) \rightarrow F(\omega)/2$. The factor of 2 accounts for the fact that experimental frequency-domain data to which the distribution would apply represents the positive part of a two-sided power spectral density. This modified expression is consistent with reference [29].

The consequence is that rather than having a unique frequency, the dark matter field is a sum

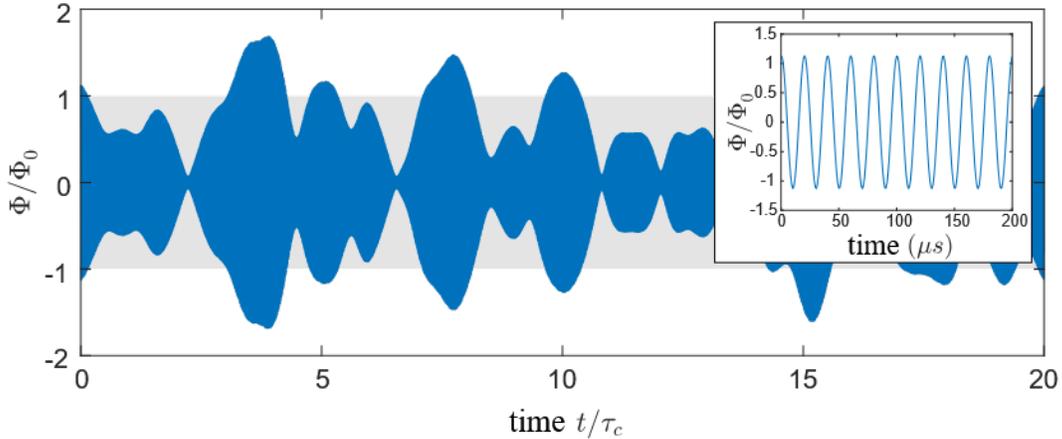


Figure 1.3: Simulation of the amplitude of the dark matter field. Shaded region illustrates the average field amplitude, given by Eq. (1.6) (corresponding to the limit where the total observation time $T_{tot} \gg \tau_c$). Inset shows the field in the $T_{tot} \ll \tau_c$ limit, where the field has a sinusoidal behavior

over all modes of different frequencies distributed according to the probability distribution given by Eq. (1.25). As a result of the interference of the different field modes, the dark matter field will behave stochastically, with an amplitude drawn from a Rayleigh distribution and the phase uniformly distributed between 0 and 2π [30].

1.2.5 Relaxion

A different scenario for a scalar ULDM field is the so-called "relaxion", a ultralight field with a scalar coupling to the Higgs boson. While it was first proposed to solve the electroweak hierarchy problem [31] (essentially the 17 orders of magnitude discrepancy between the Higgs boson mass ≈ 125 GeV and the Planck mass), a relaxion field can also explain the dark matter density in the Universe [32]. Relaxion models allow for the formation of gravitationally bound objects, known as relaxion stars, some of which could be trapped by an external gravitational potential, forming halos [33]. In the eventuality of such a halo having formed in the potential of the Earth or Sun, that would considerably increase the local DM density. Under this scenario, the sensitivity of direct DM detectors, which is dependent on the local DM density, could be augmented by orders

of magnitude.

The possible mass range of such a halo is restricted by astronomical observations. The tightest upper constraints, which come from lunar laser ranging for a halo bound to Earth and from planetary ephemerides for a solar halo, are summarized in reference [33] and illustrated in Fig. 1.4.

The coherence length for a relaxion halo is effectively its radius R_{halo} , determined by the condition that the halo is in gravitational equilibrium. An implication of experimental consequence is that coherence times can be much longer than in the SHM. In the case of a halo bound to Earth, which is more relevant for the current work, and assuming that the radius of the halo is larger than that of the Earth, coherence times can be shown [33] to be approximately:

$$\tau_c = \frac{m_\phi R_{halo}^2}{\hbar} \approx 10^3 \text{ s} \left(\frac{10^{-9} \text{ eV}}{m_\phi} \right)^3 \quad (1.26)$$

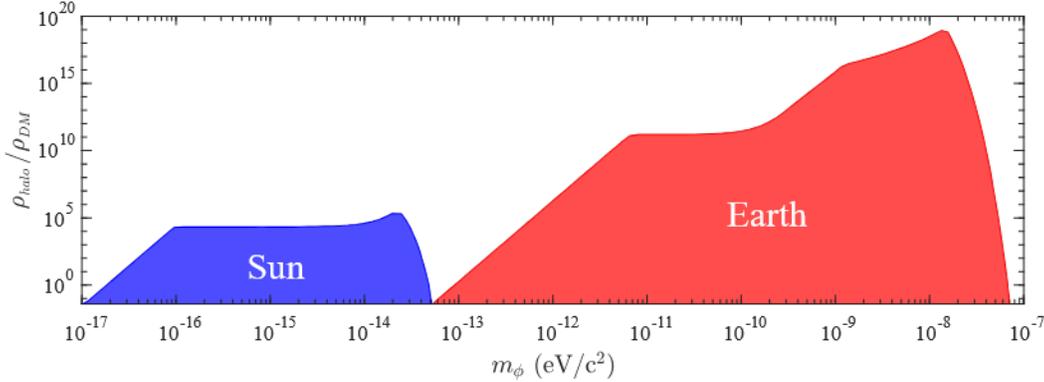


Figure 1.4: Limits on the possible density of a relaxion halo around the Earth (red) and Sun (blue) taken from reference [33], expressed as a ratio to the local dark matter density in the SHM $\rho_{DM} = 0.4 \text{ GeV}/c^2$

1.3 Current bounds on bosonic scalar field DM

We conclude this overview of scalar field ULDM by providing a brief account of some of the various implemented and proposed techniques for the detection of scalar ULDM, as well as current

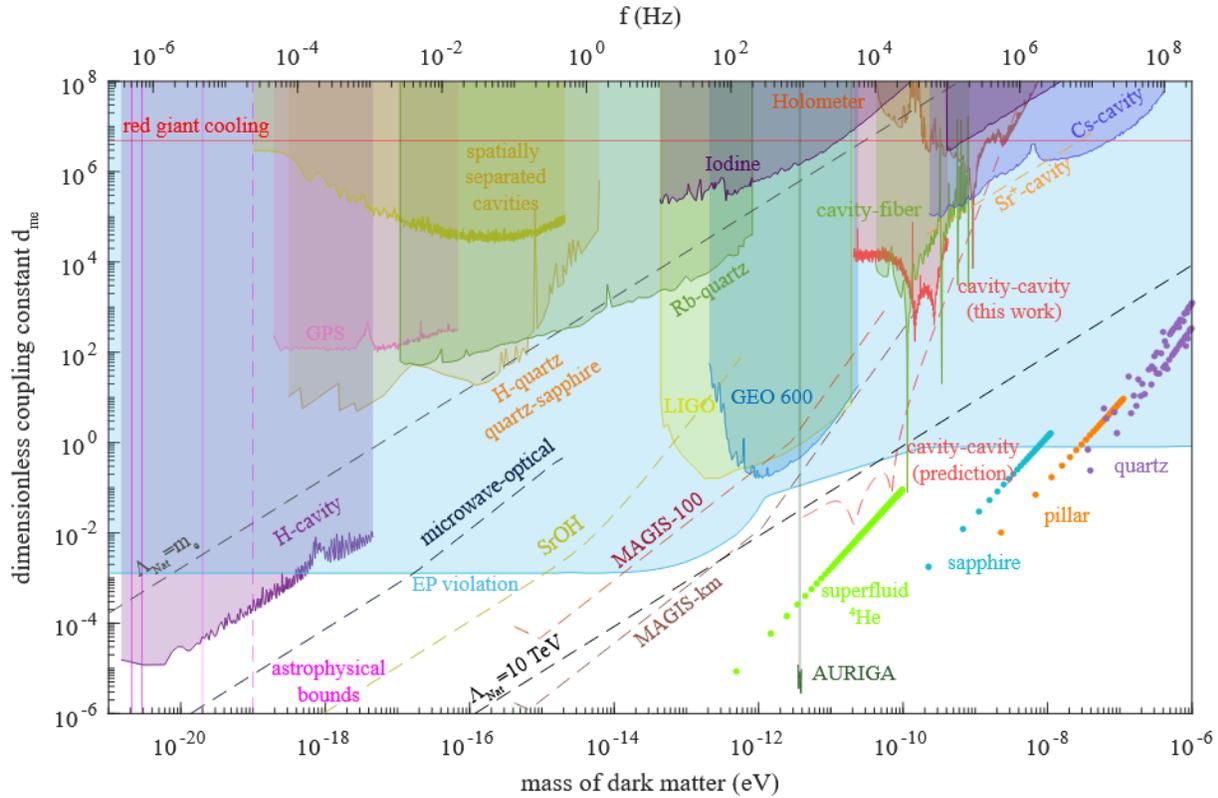


Figure 1.5: Current (solid lines) and projected (dashed lines) bounds on d_{m_e} . Details in text

limits on its coupling. A summary plot, showing both existing and some projected limits, together with theoretical constraints, is provided in Fig. 1.5. Many of these results are also reviewed in more detail in [18].

Some of the most stringent bounds on d_{m_e} over a wide frequency range come from tests of equivalence principle (EP) violation, most notably from the $\ddot{\text{E}}\text{ot-Wash}$ experiment (the torsion balance experiment at University of Washington [34, 35]) and the MICROSCOPE mission, a space-based differential electro-static accelerometer [36]. Scalar couplings to the SM fields lead to a Yukawa-type term to the gravitational force (a "fifth force"), which, in the case of a non-universal coupling, leads to an EP violating acceleration [24]. The combined bound from the aforementioned experiments is represented by the light blue region labeled "EP violation".

A second class of experiments are represented by direct detectors, which search for oscillations of fundamental constants induced by a scalar coupling by looking for their effects on atomic transitions or on the lengths of solids. In contrast with the EP violation tests mentioned before, the sensitivity of these detectors is dependent on the local abundance of dark matter, which makes them particularly favorable in scenarios that predict an increased dark matter local density (such as the relaxion model).

Among the most sensitive detectors for oscillations of fundamental constants are atomic clocks. Optical clocks are primarily sensitive to variations of α , whereas microwave clocks are sensitive to both α and the proton-electron ratio mass m_p/m_e and more weakly to m_g/Λ_{QCD} . Experiments comparing two clocks using different species and hence having different sensitivities to dark matter (e.g. Rb/Cs [37]) place bounds on d_e in the μHz to mHz range. Comparisons of atomic clocks with other frequency references such as ultra-stable optical cavities are sensitive to both d_e and d_{m_e} at higher frequencies (up to the order of 100 MeV). Illustrated in Fig. 1.5 are bounds from comparisons of a Caesium clock [38] (labeled as "Cs-cavity") and a Strontium clock [39] ("Sr-cavity") with optical cavities, as well as a comparison of a Rubidium clock with a quartz oscillator [40] ("Rb/quartz"). Moreover, by using spatially separated atomic clocks and cavities one can attain sensitivity to couplings that typically cancel out in co-located atomic clocks experiments (results from such an experiment [41] are denoted by "GPS" and "Spatially separated cavities"). Prospective sensitivities of future experiments involving microwave and optical atomic clocks [25] and molecular clocks, such as SrOH [42] are also shown in Fig. 1.5.

Bounds from other experiments comparing frequency stable oscillators include frequency comparisons of a H-maser with a Si-cavity [43] and of a bulk acoustic wave quartz oscillator with a H-maser and a cryogenic sapphire oscillator [44].

Spectroscopic measurements in molecules are also sensitive to fundamental constants oscillations. Results obtained from the spectroscopy of molecular iodine (I_2) spectroscopy are also

presented (denoted by "Iodine")[45].

Additionally, the effects of scalar ULDM on atomic energy levels can be exploited in atom interferometers such as MAGIS [46], prospective limits from which are also provided (for the 100 m version of the interferometer, MAGIS-100, as well as a potential 1 km-long version).

Optical interferometers can also constitute scalar ULDM detectors. Michelson interferometers such as the Fermilab Holometer [47] or gravitational-wave detectors (e.g. GEO600 [48]) are sensitive to dark-matter induced oscillations of the thickness and index of refraction of the main beam-splitter, which affect the optical path difference between the arms. A Mach-Zender fiber interferometer with unequal arm lengths was also used to compare the length of a single optical cavity at different times (the DAMNED experiment [49], labeled as "Cavity-fiber").

All the experiments mentioned so far can serve as broadband detectors of a scalar ULDM field. The sensitivity of interferometers and optical cavities can also be enhanced if the dark matter frequency happens to coincide with a relevant mechanical resonance of the detector. Additionally, several current or future experiments focus on the resonant detection of a scalar field. In particular, the ultracryogenic resonant bar AURIGA [50] places tight upper bounds in a narrow frequency range around 900 Hz. Acoustic mechanical resonators consisting of superfluid He or single-crystal solids such as a sapphire cylinder test mass, a sapphire micropillar resonator or a quartz bulk acoustic wave resonator have also been proposed [51] (their expected sensitivities are represented by dots in Fig. 1.5).

Some constraints on the value of d_{m_e} also come from astrophysical considerations. For instance, measurements of the cooling of red giants restrict the value of d_{m_e} below the red line. Values of the mass on the left side of the vertical lines are also excluded by astrophysical observations (see references in [18]).

For reference, Fig. 1.5 also includes naturalness limits for d_{m_e} corresponding to a cutoff of the Yukawa modulus of 10 TeV (the black dashed line) or equal to the electron mass (the grey dashed

line). The regions above these lines are excluded by the scale of radiative corrections to the scalar mass without fine-tuning.

1.4 Overview of work presented

In this work we explore a new method for searching for scalar ULDM that relies on the detection of a differential modulation in the lengths of two Fabry-Pérot cavities, due to the oscillation of the Bohr radius. Such oscillations would emerge as frequency modulations in two beams nearly resonant to the cavities, which can be detected, for instance, in their beat note. The two beams originate from the same laser to eliminate common mode noise. The key element is the use of two cavities which have different response functions to a dark matter drive. Two potential ways of accomplishing this are outlined here.

The first of them involves the use of two ultra-stable cavities with spacers of different lengths and exploits the difference in the response of rigid objects to a mechanical drive above and below their mechanical pole. The implementation of this scheme and its first results (presented in Fig. 1.5 as "cavity-cavity") constitute the main focus of this work.

A second proposed method would be a comparison between two equal-length cavities, one with its mirrors connected by a rigid spacer and one LIGO-like cavity, consisting of two suspended mirrors. In this case, a scalar field would cause oscillations in the length of the rigid cavity below its mechanical pole, while for the non-rigid cavity its effects would be suppressed by many orders of magnitude by the low frequency pendulum they are suspended from. The results of some initial research and development associated with the implementation of this scheme are presented towards the end of this thesis. While its realization has been put on hold due to technical challenges related to the stabilization of the suspended mirrors, this second scheme could provide a path for improving the sensitivity of the experiment in future generations.

I joined this project at its inception and, as such, had the chance to participate in all the de-

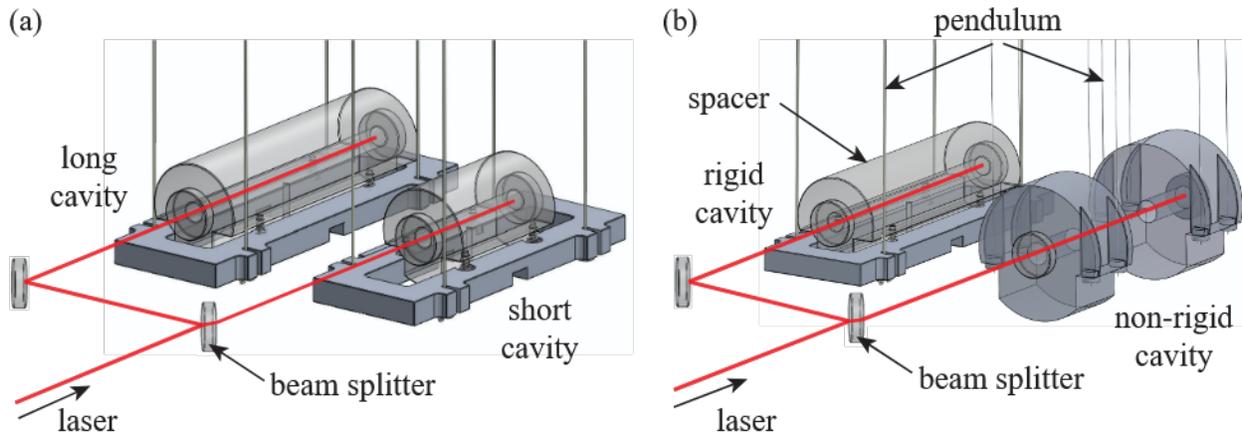


Figure 1.6: Diagram of the experiment for the two versions of the experiment involving either (a) two rigid cavities or (b) a rigid and a non-rigid cavity

velopment stages of its first generation, from the initial experiment planning, to the design, commissioning, construction and testing of the apparatus, concluding with the collection and analysis of the data set that constitutes the result of this experiment cycle. These experiences collectively make the subject of this thesis.

The remainder of the thesis is structured as follows:

Chapter 2 lays the foundations for understanding the experiment. It starts with an overview of optical cavities and an introduction to some of the experimental techniques involved. It follows up with the derivation of the apparatus response and, finally, an evaluation of the prospects of the experiment in relation to the fundamental noise sources.

The following three chapters delve into each of the main subsystems of the apparatus. Chapter 3 addresses the design, construction and assembly of the optical cavities. Chapter 4 discusses the optical and electronic setup for the generation and detection of the beat note. Chapter 5 examines the cryostat inside which the experiment is conducted and the various stages of vibration isolation employed to isolate the cavities from external disturbances, a requirement for achieving high strain sensitivity. A detour into the methods used for measuring the mechanical resonant frequencies, which represent an essential component in the interpretation of our data, constitutes the subject of

chapter 6.

Chapter 7 focuses on the data collection, analysis and results. The Bayesian framework used to interpret the measured beat note as a constraint on d_{m_e} is explained. New bounds on the dark matter interaction and a comparison with other experiments are presented.

Finally, in chapter 8 we showcase some work conducted at the early stages of the experiment towards the construction of a non-rigid cavity, including initial prototypes of suspended mirrors and first steps towards achieving their active position control.

CHAPTER 2

OPTICAL DETECTION OF A DARK MATTER-INDUCED STRAIN

The ability of the two cavities to act as a dark matter detector originates in their unequal response functions to a ULDM field. This response is determined by two effects: a mechanical component, represented by the length change of the cavity under the influence of a dark matter drive and an opto-mechanical component, given by the conversion of the cavity length change to a frequency modulation detectable in the beat note. In this chapter, we establish the groundwork for understanding how optical cavities can serve as strain detectors and develop a mathematical model that captures the response of our detector to a potential dark matter-induced strain. Furthermore, we present theoretical estimates of the fundamental noise for the current two rigid cavities comparison, as well as for a potential future scenario involving a comparison between a cavity with a rigid spacer and a non-rigid cavity made of suspended mirrors.

2.1 Fabry-Pérot cavities as strain detectors

2.1.1 Optical cavities fundamentals

A Fabry-Pérot cavity consists of two highly reflective parallel mirrors facing each other. This configuration acts as a resonator, as light traversing the cavity can form a standing wave, provided a fixed relationship between the distance between the mirrors L and the laser frequency f (the resonant condition) is upheld:

$$f = \frac{nc}{2L}, \quad (2.1)$$

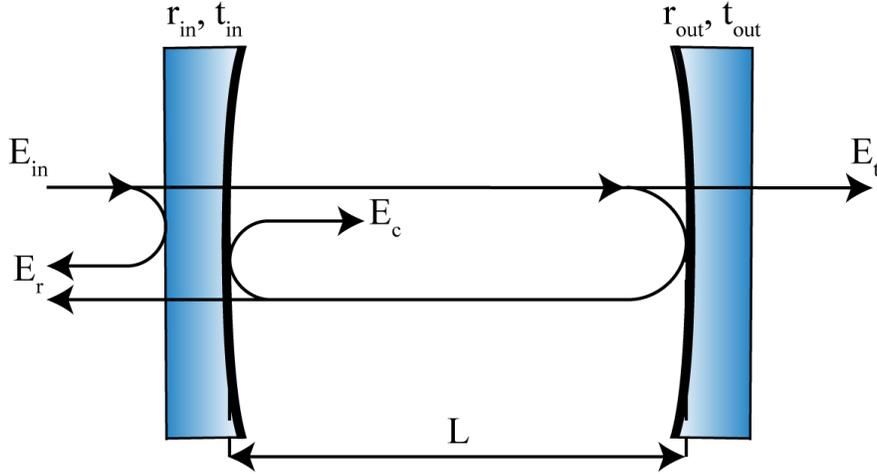


Figure 2.1: Diagram of a Fabry-Pérot cavity illustrating the incident, reflected, transmitted and intra-cavity (circulating) fields

with n an integer denoting the longitudinal mode order and c the speed of light. Hence, optical cavities serve as frequency filters, allowing only a narrow range of nearly-resonant frequencies to pass through, a property which makes them highly effective in laser frequency stabilization.

To analyze the characteristics of the fields inside the cavity, let us consider an optical cavity formed by two mirrors with reflectivities r_{in} and r_{out} and transmissivities t_{in} and t_{out} , placed at a distance L from each other, as depicted in Fig. 2.1. Self-consistency of fields demands that the electric field circulating inside the cavity at a time t , $E_c(t)$, can be formulated as a sum of two components: the fraction of the incident field E_{in} transmitted through the first mirror at time t and the field that entered the cavity at time $t - 2T$ and has undergone a reflection at each mirror (here $T = L/c$ is the one-way photon travel time through the cavity):

$$E_c(t) = t_{in}E_{in}(t) + r_{in}r_{out}e^{-2ikL}E_c(t - 2T), \quad (2.2)$$

where $k = 2\pi/\lambda$ is the wave number and the e^{-2ikL} factor ensures phase continuity between the incident and reflected beams at the second mirror interface.

In the steady state, for a non-varying input field, $E_c(t) = E_c(t - 2T) = E_c$, in which case the

intra-cavity field takes the form:

$$E_c = \frac{t_{in}E_{in}}{1 - r_{in}r_{out}e^{-2ikL}}. \quad (2.3)$$

It is easy to see that the field equation admits an infinite number of equally spaced maxima, corresponding to longitudinal resonant modes and occurring at light wavelengths for which kL is a multiple of π (which is equivalent to Eq. (2.1)). The separation in frequency space between consecutive longitudinal modes is referred to as the cavity **free spectral range (FSR)** and equals:

$$f_{FSR} = \frac{c}{2L}. \quad (2.4)$$

We introduce a few properties of optical cavities that are relevant going forward:

- **Cavity linewidth:** The FWHM (full width at half maximum) of the cavity resonances is commonly referred to as the cavity linewidth and is determined by the reflectivity of the mirrors and the cavity length, as follows:

$$\delta f = \frac{c}{\pi L} \arcsin\left(\frac{1 - r_{in}r_{out}}{2\sqrt{r_{in}r_{out}}}\right) \approx \frac{c}{2\pi L} \frac{1 - r_{in}r_{out}}{\sqrt{r_{in}r_{out}}}; \quad (2.5)$$

- **Finesse:** A parameter used to characterize the sharpness of the resonances and defined as the ratio of the FSR to the cavity linewidth:

$$\mathcal{F} = \frac{f_{FSR}}{\delta f} \approx \frac{\pi\sqrt{r_{in}r_{out}}}{1 - r_{in}r_{out}}; \quad (2.6)$$

- **Light storage time:** The time in which the intra-cavity field decays by a factor of e after the incident light is turned off. It can also be understood as the mean time a photon spends inside the cavity and is proportional to the cavity finesse:

$$\tau_c = -\frac{2L}{c} \frac{1}{\ln(r_{in}r_{out})} \approx \frac{2\mathcal{F}L}{\pi c}; \quad (2.7)$$

- **Cavity pole:** For reasons that will become clear in section 2.1.3, the half-width of the resonance is referred to as the cavity pole, f_p . It represents the frequency above which variations in the intra-cavity field are suppressed and is related to the storage time by:

$$f_p = \frac{1}{2\pi\tau_c}. \quad (2.8)$$

The approximations in Eq. (2.5-2.7) are valid provided that the mirrors have high reflectivity ($r_{in}, r_{out} \approx 1$). This is most often the case for real-world cavities and is applicable to all the cavities featured in the research presented in this thesis.

It follows from Eq. (2.1) that there is a direct relationship between the fluctuations in the cavity length (the displacement noise) and the frequency noise of a beam resonant with the cavity:

$$\frac{\Delta f}{f} = -\frac{\Delta L}{L}. \quad (2.9)$$

From a practical standpoint, this implies that high reflectivity, narrow linewidth, length-stable cavities (such as cavities with a rigid spacer connecting the two mirrors) can be used to reduce laser frequency noise, provided that the frequency of an incoming beam is maintained in resonance with the cavity through a locking technique. Conversely, variations of the distance between the mirrors will be reflected in corresponding shifts in the frequency of the locked laser, allowing such a system to be an exquisitely sensitive detector of changes in the cavity length.

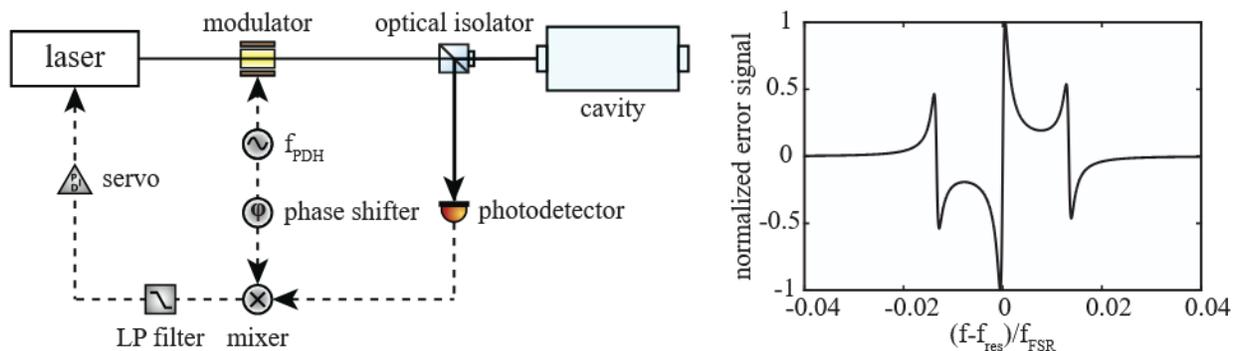


Figure 2.2: Left: Standard configuration for the implementation of PDH locking. Right: Example of PDH error signal as a function of the frequency offset from one of the cavity resonances f_{res} , normalized by the FSR, for a finesse of 1000

2.1.2 Pound-Drever-Hall locking

Locking a laser to the resonance of a cavity involves two key "ingredients": a signal that contains information about the magnitude and direction of the departures of the laser from resonance (i.e. a frequency discriminator) and a feedback loop that acts on the laser frequency to counteract these deviations.

Among the various techniques developed for the stabilization of lasers using optical cavities, one of the most widespread and robust is the Pound-Drever-Hall (PDH) method ([52, 53]), a detailed review - including a full mathematical treatment - of which is provided in [54].

The fundamental idea is that phase modulation at a frequency $f_{PDH} \gg \delta f$ applied to the laser prior to sending it into the cavity (for instance by using an electro-optical modulator) leads to the generation of sidebands around the carrier. When the laser frequency is near-resonant, the heterodyne beat between the carrier, phase shifted after passing through the cavity, and the sidebands, which are directly reflected off the first mirror unaffected by the cavity, measured by a photodetector in reflection, contains phase information encoding the laser detuning from the cavity resonance. Subsequent demodulation with the same local oscillator at f_{PDH} , with a phase adjusted to extract the imaginary part of the signal, and low-pass filtering (to eliminate high frequency components)

give rise to the dispersive error signal described by Eq. (2.10) and depicted in Fig. 2.2:

$$\epsilon = -4\sqrt{P_c P_s} \operatorname{Im} \left(F(\omega_{\text{opt}}) F^*(\omega_{\text{opt}} + \omega_{\text{PDH}}) - F^*(\omega_{\text{opt}}) F(\omega_{\text{opt}} - \omega_{\text{PDH}}) \right), \quad (2.10)$$

with P_c and P_s the power in the carrier and sidebands, respectively, ω_{opt} the angular frequency of the carrier and $F(\omega) = E_{\text{ref}}(\omega)/E_{\text{in}}(\omega)$ the ratio of the reflected and incident fields.

This error signal is antisymmetric and has a zero-crossing on resonance. Moreover, it is linear with the detuning $\delta\omega$ over a frequency range comparable to the cavity linewidth:

$$\epsilon \approx -\frac{4}{\pi} \sqrt{P_c P_s} \frac{\delta\omega}{\delta f}, \quad (2.11)$$

hence being suitable for use in a feedback loop to stabilize the laser frequency.

Significant advantages of this technique are its ability to isolate the effects of frequency fluctuations from laser power fluctuations (because the error signal is 0 on resonance), thus being (to first order) insensitive to the latter, as well as its potential to extend the range of the lock past the cavity bandwidth. However, above the pole, the cavity field cannot track the laser fluctuations anymore, instead capturing the integrated effect of frequency fluctuations over time. In other words, the PDH error signal transitions to being a phase, rather than a frequency, discriminator [53] and the sensitivity decreases as $1/f$, which is usually reflected in a drop in the frequency noise performance of the locked laser.

2.1.3 Dynamic response

Eq. (2.3) describes the field inside a constant length cavity in the presence of a non-varying incident field (the static response). However, to meaningfully interpret the cavity signals in the context of a potential dark matter coupling, as well as to study the impact of noise, it is key to understand the dynamic response of Fabry-Pérot cavities. A comprehensive treatment of the

dynamics of fields in an optical cavity can be found in reference [55]. For the purposes of the present work, it suffices to analyze the behavior of the optical cavity in the presence of laser frequency or amplitude fluctuations and of a cavity length modulation.

Since the field variations we are interested in here are much slower than the optical frequency, for the rest of this section we will, for simplicity, eliminate the fast oscillating component of the electromagnetic fields by performing a transformation: $E \rightarrow Ee^{-i\omega_{\text{opt}}t}$ on the incident, transmitted and circulating fields (with ω_{opt} being the laser angular frequency). Note that Eq. (2.2) still holds after this transformation.

Frequency modulation

The frequency response of the optical cavity near resonance can be obtained by applying the Laplace transform to Eq. (2.2), while imposing the resonant condition $kL = n\pi$ (we will use the notation $\tilde{E}(s)$ to denote the Laplace transform of $E(t)$):

$$\tilde{E}_c(s) = t_{in}\tilde{E}_{in}(s) + r_{in}r_{out}e^{-2sT}\tilde{E}_c(s). \quad (2.12)$$

Given that in an experimental setup it is typically the cavity transmission, rather than the intra-cavity field, that is directly observed, it is more helpful to instead express this relation in terms of the transmitted field, by recognizing that $E_t = t_{out}E_c$. Then, we can define a transfer function relating the cavity incident and transmitted fields:

$$H_f(s) = \frac{\tilde{E}_t(s)}{\tilde{E}_{in}(s)} = \frac{t_{in}t_{out}}{1 - r_{in}r_{out}e^{-2sT}}. \quad (2.13)$$

$H_f(s)$ can be understood as the response of the transmitted field to a frequency modulation of the incident field at frequency f , where $s = 2\pi if$. We can reformulate this transfer function explicitly in terms of f . We will also assume a low-frequency regime, $2\pi fT \ll 1$, which is true for expected

fluctuations in the laser frequency, due, for example, to noise of mechanical or electronic origin.

In this case, we can approximate:

$$H_f(f) = \frac{t_{in}t_{out}}{1 - r_{in}r_{out}(1 - 2(2i\pi f)T)}. \quad (2.14)$$

Combining Eq. (2.6)-(2.8) and noting that $r_{in}r_{out} \approx 1$, the cavity pole f_p can be expressed as:

$$f_p = \frac{1 - r_{in}r_{out}}{4\pi T}. \quad (2.15)$$

Plugging the expression above into Eq. (2.14) and using once again the approximation $r_{in}r_{out} \approx 1$, we can finally write the approximate low-frequency transfer function:

$$H_f(f) \approx \left(\frac{\mathcal{F}t_{in}t_{out}}{\pi} \right) \left(\frac{f_p}{f_p + if} \right). \quad (2.16)$$

It is now easy to recognize this expression as the transfer function of a first order low-pass filter with a pole at f_p , multiplied by a pre-factor corresponding to the total cavity transmission coefficient on resonance. The behavior of this transfer function is displayed in Fig. 2.4.

Amplitude modulation

It is also possible to derive the cavity response to a pure amplitude modulation of the incident field, which we assume to be, in the absence of the perturbation, resonant to the cavity:

$$E_{in}(t) = E_{in}^{(0)} + \delta E_{in}(t). \quad (2.17)$$

In the case of a small modulation, the equation for the intra-cavity field can be solved perturbatively, by looking for a solution of the form:

$$E_c(t) = E_c^{(0)} + \delta E_c(t), \quad (2.18)$$

where $E_c^{(0)}$ is the solution to the field equation in the static case, for an incident field $E_{in}^{(0)}$. In this case, Eq. (2.12) can be rewritten as:

$$\tilde{E}_c^{(0)} + \delta \tilde{E}_c(s) = t_{in} \left(\tilde{E}_{in}^{(0)} + \delta \tilde{E}_{in}(s) \right) + r_{in} r_{out} e^{-2sT} \left(\tilde{E}_c^{(0)} + \delta \tilde{E}_c(s) \right). \quad (2.19)$$

The average and perturbation terms can be equated separately, which allows us to rewrite:

$$\delta \tilde{E}_c(s) (1 - r_{in} r_{out} e^{-2sT}) = t_{in} \delta \tilde{E}_{in}(s). \quad (2.20)$$

The relation above can be converted to a transfer function between an amplitude modulation of the incident light and its effect in the cavity transmission:

$$H_A(s) = \frac{\delta \tilde{E}_t(s)}{\delta \tilde{E}_{in}(s)} = \frac{t_{in} t_{out}}{1 - r_{in} r_{out} e^{-2sT}}. \quad (2.21)$$

This expression is identical with the transfer function previously obtained for a frequency modulation (Eq. (2.13)). Consequently, in the low-frequency limit $sT \ll 1$, we recover the transfer function of a low-pass filter:

$$H_A(f) \approx \left(\frac{\mathcal{F} t_{in} t_{out}}{\pi} \right) \left(\frac{f_p}{f_p + if} \right). \quad (2.22)$$

Length modulation

Lastly, assume the optical cavity length suffers a length perturbation $\delta L(t)$, as depicted in Fig. (2.3). In this case, Eq. (2.2) becomes:

$$E_c(t) = t_{in}E_{in}(t) + r_{in}r_{out}e^{-2ik(L+\delta L(t))}E_c(t-2T). \quad (2.23)$$

We consider the regime where the length perturbation is small, $\delta L(t) \ll \lambda \ll L$. Moreover, we

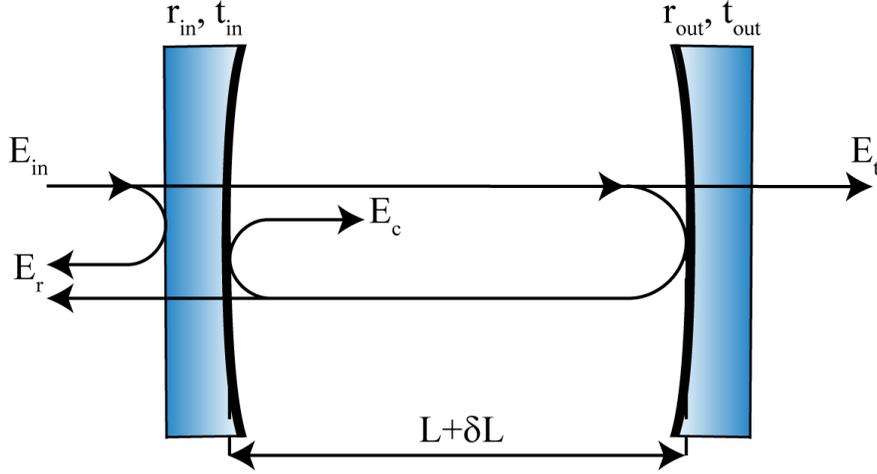


Figure 2.3: Reflected, transmitted and intra-cavity fields

assume that the incident field is constant and that in the absence of this perturbation the laser is in resonance with the cavity, such that $2kL = 2\pi n$, with n an integer. Under these conditions, the effect of the length fluctuation can be expressed as a perturbative term $\delta E_c(t)$ added to the static value of the electric field $E_c^{(0)}(t) = E_c^{(0)}(t-2T)$ corresponding to the unperturbed length, c.f. Eq. (2.2) [56]. In this case:

$$\begin{aligned} E_c(t) &= E_c^{(0)}(t) + \delta E_c(t) \\ &= t_{in}E_{in}(t) + r_{in}r_{out}(1 - 2ik\delta L(t))(E_c^{(0)}(t) + \delta E_c(t-2T)). \end{aligned} \quad (2.24)$$

The static terms can be removed using Eq. (2.2). Then, neglecting the second order perturbation term, we arrive at:

$$\delta E_c(t) = r_{in}r_{out}\delta E_c(t-2T) - 2ikr_{in}r_{out}\delta L(t)E_c^{(0)}(t). \quad (2.25)$$

After again applying the Laplace transform and some rearranging:

$$\frac{\delta \tilde{E}_c(s)}{E_c^{(0)}(s)} = -\frac{2ikr_{in}r_{out}}{1 - r_{in}r_{out}e^{-2sT}}\delta \tilde{L}(s). \quad (2.26)$$

The fluctuations of the intra-cavity field due to the mirror motion may appear as modulations in its amplitude or its phase. As the field derivative $d\tilde{E}_c/df = 0$ on resonance, small modulations of the resonant beam lead to negligible amplitude modulation, making phase modulation the primary effect. We will denote by $\phi(t)$ the contribution to the phase from the length modulation:

$$E_c(t) = E_c^{(0)}(t)e^{i\phi(t)} \approx E_c^{(0)}(1 + i\phi(t)), \quad (2.27)$$

which implies:

$$\delta E_c(t) = iE_c^{(0)}(t)\delta\phi(t). \quad (2.28)$$

After taking the Laplace transform and plugging the result into Eq. (2.26), we arrive at a relation between the length fluctuation and the resulting phase modulation:

$$\delta \tilde{\phi}(s) = -\frac{2kr_{in}r_{out}}{1 - r_{in}r_{out}e^{-2sT}}\delta \tilde{L}(s), \quad (2.29)$$

or, alternatively, expressed in terms of the frequency modulation $\delta f(s) = (s/2\pi)\delta\phi(s)$:

$$\delta \tilde{f}(s) = -\frac{k}{\pi} \frac{r_{in}r_{out}s}{1 - r_{in}r_{out}e^{-2sT}}\delta \tilde{L}(s). \quad (2.30)$$

We can then define a transfer function between the cavity strain $\tilde{h}(s) = \delta \tilde{L}(s)/L$ and the "optical strain", the ratio between the frequency noise and the frequency of the laser traversing the cavity, $\tilde{h}_{\text{opt}}(s) = \delta \tilde{f}(s)/f_{\text{opt}}$. We will refer to this transfer function as the opto-mechanical response of the cavity:

$$H_{\text{OM}}(s) = \frac{\tilde{h}_{\text{opt}}(s)}{\tilde{h}(s)} = -\frac{k}{\pi} \frac{L}{f_{\text{opt}}} \frac{r_{\text{in}} r_{\text{out}} s}{1 - r_{\text{in}} r_{\text{out}} e^{-2sT}}. \quad (2.31)$$

As before, we will look at the low-frequency limit, where $e^{-2sT} \approx 1 - 2sT$. By also using the relation $kL/(\pi f_{\text{opt}}) = 2T$ and plugging in f_p from Eq. (2.15), the transfer function reduces to:

$$H_{\text{OM}}(f) = -\frac{if}{f_p + if}. \quad (2.32)$$

This expression describes a first-order high-pass filter with the pole at f_p .

The cavity transfer functions derived in this section have important implications for the present experiment. The presence of a scalar field, manifested as a cavity length perturbation, is expected to be detected as a frequency modulation in the transmission of the optical cavity, as described by Eq. (2.32). Owing to their high finesse, our science cavities have low-frequency optical poles at 4.7 kHz and 7.4 kHz (as discussed in section 3.5.1 and Appendix A), while the main frequency band over which we report bounds spans 20-90 kHz (see chapter 7 and Fig. 7.6). As a consequence, a dark matter signal can be detected in the cavity transmission virtually unsuppressed. On the other hand, any amplitude or frequency modulation in the measurement band due to laser noise is filtered away in the transmission signal by virtue of the cavities behaving as low-pass filters, c.f. Eqs. (2.16) and (2.22). This dual advantage makes the use of light transmitted through high-finesse cavities particularly appealing. The apparatus we present here retains sensitivity at frequencies around and below the optical poles, as well. In this case the opto-mechanical high-pass filter transfer function needs to be taken into account (Eq. (2.32)).

To better visualize the effect of the cavity on the dark matter signal, as well as on laser noise, Fig. 2.4 presents Bode plots of the normalized low-pass filter transfer function (corresponding to frequency and amplitude modulation) and high-pass filter function (corresponding to length modulation) previously derived.

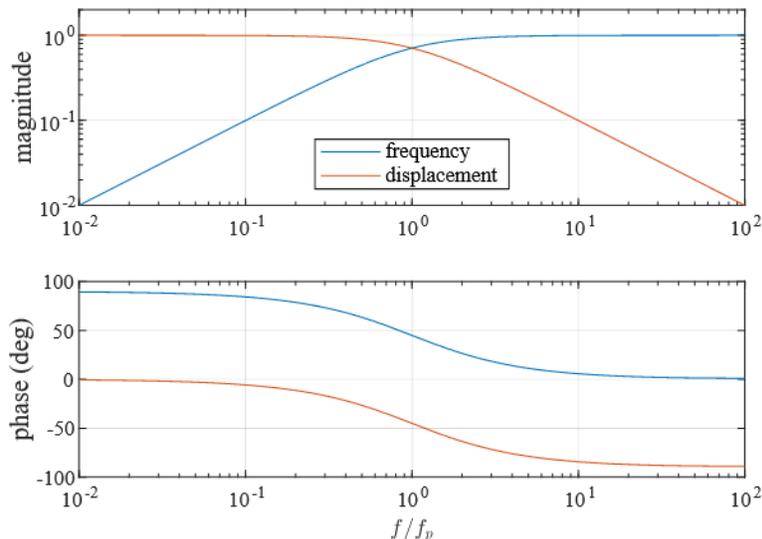


Figure 2.4: Plots of the cavity dynamic responses. Note that the frequency modulation response also applies to amplitude modulation

2.2 Detector response

We will now derive the response of the two-cavity detector. We will start with the study of the frequency dependence of the strain produced in one cavity by the oscillation of fundamental constants. Leveraging insights from this section and the previous one, we will then model how the differential response of the two cavities is imprinted onto the beat note of the two-cavity system.

2.2.1 Mechanical transfer function

We employ a simple "mass-spring" toy model, where we neglect the inner degrees of freedom of the cavity and assume that its whole mass is concentrated at a single point situated at one end [57]. This mass is connected to a reference point placed at the other end of the cavity by a spring of effective spring constant k . Without a dark matter field, the length of the cavity (and spring) is L_0 .

In the presence of temporal oscillations of the Bohr radius, the equilibrium length of the

spring will also oscillate: $L_{eq}(t) = L_0(1 + h_{DM} \cos(\omega_{DM}t))$, with ω_0 the dark matter angular frequency and h_{DM} the nominal strain amplitude, which depends on the strength of the dark matter coupling according to Eq. (1.15). Let us denote the corresponding displacement by $u_{DM}(t) = h_{DM}L_0 \cos(\omega_{DM}t)$.

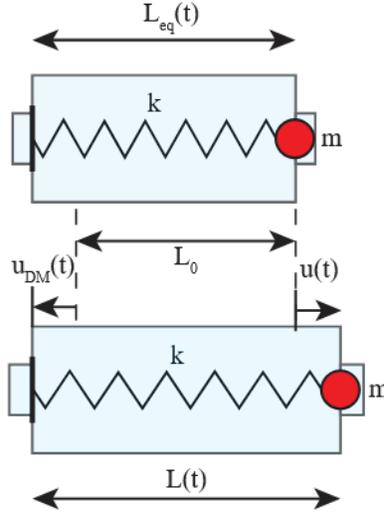


Figure 2.5: Mass-spring toy model for the cavity deformation

With respect to this new equilibrium length, the presence of dark matter emerges as a fictitious driving force acting on the cavity:

$$F_{DM}(t) = -m \left(\frac{d^2}{dt^2} u_{DM}(t) \right). \quad (2.33)$$

The deviation of the cavity from the equilibrium length, $u(t) = L(t) - L_{eq}(t)$ can be described by an equation of motion that includes the dark matter drive, as well as damping. We assume that the latter is dominated by internal losses in the cavity spacer crystal and thus model it as a structural damping term¹ (proportional in magnitude with the amplitude of the displacement, but in phase with its derivative) characterized by the loss factor η . Then, the equation of motion takes the form:

¹However, it was verified that for damping coefficients similar to the experimental conditions ($\eta \approx 10^{-4}$) the quantitative difference between the results obtained assuming structural and viscous damping is negligible.

$$m \frac{d^2 u(t)}{dt^2} + k(1 + i\eta)u(t) = m\omega_{\text{DM}}^2 h_{\text{DM}} L_0 \cos(\omega_{\text{DM}} t) = m\omega_{\text{DM}}^2 h_{\text{DM}} L_0 \text{Re}(e^{i\omega_{\text{DM}} t}). \quad (2.34)$$

Eq. (2.34) admits a solution of the form:

$$u(t) = \text{Re}(Ae^{i\omega_{\text{DM}} t}), \quad (2.35)$$

where

$$m \frac{d^2}{dt^2}(Ae^{i\omega_{\text{DM}} t}) + k(1 + i\eta)Ae^{i\omega_{\text{DM}} t} = m\omega_{\text{DM}}^2 h_{\text{DM}} L_0 e^{i\omega_{\text{DM}} t}. \quad (2.36)$$

Solving Eq. (2.36) for $A(\omega_{\text{DM}})$ renders:

$$A(\omega_{\text{DM}}) = \frac{m\omega_{\text{DM}}^2 h_{\text{DM}} L_0}{-m\omega_{\text{DM}}^2 + k(1 + i\eta)} = \frac{\omega_{\text{DM}}^2 h_{\text{DM}} L_0 ((k/m - \omega_{\text{DM}}^2) - i\eta k/m)}{(k/m - \omega_{\text{DM}}^2)^2 + (\eta k/m)^2}. \quad (2.37)$$

We can introduce $\omega_{\text{res}} = \sqrt{k/m}$, with $f_{\text{res}} = \omega_{\text{res}}/(2\pi)$ the resonant frequency of the cavity. Then, after plugging the above expression into Eq. (2.35) and taking the real part, we are left with:

$$u(t) = h_{\text{DM}} L_0 (\alpha(\omega_{\text{DM}}) \cos(\omega_{\text{DM}} t) + \beta(\omega_{\text{DM}}) \sin(\omega_{\text{DM}} t)), \quad (2.38)$$

where

$$\begin{aligned} \alpha(\omega_{\text{DM}}) &= \frac{\omega_{\text{DM}}^2 (\omega_{\text{res}}^2 - \omega_{\text{DM}}^2)}{(\omega_{\text{res}}^2 - \omega_{\text{DM}}^2)^2 + \eta^2 \omega_{\text{res}}^4}, \\ \beta(\omega_{\text{DM}}) &= \frac{\eta \omega_{\text{DM}}^2 \omega_{\text{res}}^2}{(\omega_{\text{res}}^2 - \omega_{\text{DM}}^2)^2 + \eta^2 \omega_{\text{res}}^4}. \end{aligned} \quad (2.39)$$

It follows that the corresponding strain can be expressed as

$$h(t) = \frac{L(t) - L_0}{L_0} = h_{\text{DM}} \left((1 + \alpha(\omega_{\text{DM}})) \cos(\omega_{\text{DM}}t) + \beta(\omega_{\text{DM}}) \sin(\omega_{\text{DM}}t) \right). \quad (2.40)$$

It is convenient to express the mechanical response of the cavity to the dark matter drive as a transfer function between the nominal dark matter strain and the physical cavity strain:

$$H_{\text{M}}(f) = \frac{\tilde{h}(f)}{h_{\text{DM}}}. \quad (2.41)$$

Then, the magnitude and phase of this transfer function are given by:

$$\begin{aligned} |H_{\text{M}}(f)| &= \sqrt{(1 + \alpha(\omega_{\text{DM}}))^2 + \beta(\omega_{\text{DM}})^2}, \\ \arg(H_{\text{M}}(f)) &= -\arctan\left(\frac{\beta(\omega_{\text{DM}})}{1 + \alpha(\omega_{\text{DM}})}\right). \end{aligned} \quad (2.42)$$

While Eqs. (2.41) and (2.42) were derived using an exact solution to the equation of motion (Eq. (2.34)), for small η the response of a simple harmonic oscillator to a prescribed oscillatory displacement provides a good approximation:

$$H_{\text{M}}(f) \approx \frac{1 + i\eta}{(1 - (\omega_{\text{DM}}/\omega_{\text{res}})^2)^2 + i\eta}. \quad (2.43)$$

In fact, Taylor expansions of the two expressions in terms of η reveal agreement up to $\mathcal{O}(\eta^4)$.

A quick inspection of Eq. (2.43) reveals that the optical cavity is sensitive to dark matter up to the mechanical pole, with a boost in sensitivity around the resonance that is proportional to the mechanical quality factor of the cavity. Above resonance, the signal starts being suppressed by a factor proportional to $(f_{\text{DM}}/f_{\text{res}})^2$.

A more rigorous analysis would incorporate the effect of higher order mechanical modes of the cavity (see, for example, [58]). However, the symmetry properties of the dark matter drive restrict

contributions to odd longitudinal modes only. For the current experiment the next contributing mode ($n = 3$) is expected, based on finite element analysis, to occur at 129 kHz and 230 kHz, respectively, for the two science cavities, which falls outside of the measurement range (capped at 90 kHz). Therefore, the influence of higher order modes is negligible and the simple transfer function derived in this section is adequate.

2.2.2 Beat note frequency modulation and differential transfer function

Consider two optical cavities of different lengths. A resonant laser beam traverses each cavity. For simplicity, suppose that the power is adjusted so that the fields transmitted through the cavities have equal amplitude E_t .

In this case, a dark matter coupling characterized by a nominal strain h_{DM} is revealed as a frequency modulation in the transmitted fields:

$$\begin{aligned} E_{t,X} &= E_0 \exp \left(i\omega_{\text{opt},X} \left(\left(1 + \int^t (H_{\text{OM},X} H_{\text{M},X} h_{\text{DM}} \cos(\omega_{\text{DM}} t)) dt \right) \right) \right) \\ &= E_0 \exp \left(i\omega_{\text{opt},X} \left(1 + \frac{H_{\text{OM},X} H_{\text{M},X} h_{\text{DM}}}{\omega_{\text{DM}}} \sin(\omega_{\text{DM}} t) \right) \right). \end{aligned} \quad (2.44)$$

where X is a stand-in for indices S and L denoting the short and long cavities, respectively and ω_{opt} is the laser angular frequency. The power in the beat note is then:

$$\begin{aligned} P &\approx (E_{t,S} + E_{t,L})(E_{t,S}^* + E_{t,L}^*) \\ &= 2|E_0|^2 + 2|E_0|^2 \cos \left(\omega_{\text{beat}} t + \frac{(H_{\text{OM},S} H_{\text{M},S} \omega_{\text{opt},S} - H_{\text{OM},L} H_{\text{M},L} \omega_{\text{opt},L}) h_{\text{DM}}}{\omega_{\text{DM}}} \sin(\omega_{\text{DM}} t) \right), \end{aligned} \quad (2.45)$$

where $\omega_{\text{beat}} = \omega_{\text{opt},S} - \omega_{\text{opt},L}$ is the beat note angular frequency. Since $\omega_{\text{beat}} \ll \omega_{\text{opt},S}, \omega_{\text{opt},L}$, we can approximate $\omega_{\text{opt},S} \approx \omega_{\text{opt},L} = \omega_{\text{opt}}$. Then, after defining the differential transfer function:

$$H_{\text{diff}} = H_{\text{M},L} H_{\text{OM},L} - H_{\text{M},S} H_{\text{OM},S}, \quad (2.46)$$

and ignoring the DC component:

$$P = 2|E_0|^2 \cos \left(\omega_{\text{beat}} t + \frac{H_{\text{diff}} \omega_{\text{opt}} h_{\text{DM}}}{\omega_{\text{DM}}} \sin(\omega_{\text{DM}} t) \right). \quad (2.47)$$

We recognize this as a frequency modulation with amplitude $f_m = |H_{\text{diff}} f_{\text{opt}} h_{\text{DM}}|$ at the dark matter frequency, which will appear as a sideband in the signal detected by a photodetector. The extraction of f_m using a frequency discriminator thus provides a direct measure of h_{DM} . The transfer function between $h_{\text{opt}} = f_m/f_{\text{opt}}$ and h_{DM} is the detector response:

$$A_{\text{det}} = \frac{h_{\text{opt}}}{h_{\text{DM}}} = |H_{\text{M,L}} H_{\text{OM,L}} - H_{\text{M,S}} H_{\text{OM,S}}|. \quad (2.48)$$

It is worth emphasizing the significance of the fixed phase relationship between the strains of the two cavities, which results from them sharing the same ULDM drive. The ability to subtract the two transfer functions coherently provides a boost to the sensitivity of the experiment. This is in contrast with incoherent noise sources, which add in quadrature. Discarding the phase information by taking the absolute values of the individual transfer functions before their subtraction would lead to an erroneous result.

Fig. 2.6 illustrates the qualitative behavior of the transfer functions of the two cavities and of the resulting detector response. It examines a regime where the optical poles are situated far below the mechanical poles, so that $H_{\text{OM}} \approx 1$ in the region surrounding the latter. In this case, between the longitudinal mechanical resonances, the short cavity reacts to dark matter, while the long cavity, the contribution of which is suppressed, approximately acts as a stable reference. This establishes a range of maximum sensitivity. The sensitivity declines above and below the resonances, as the transfer functions either both drop - above the poles - or start canceling each other out - below the poles (although a differential effect remains beneath the optical poles due to the different opto-mechanical transfer functions).

In the case of a differential measurement between a rigid and a non-rigid cavity of equal lengths, the opto-mechanical transfer functions of the two cavities are identical. The mechanical transfer function of the rigid cavity is represented by the response of its spacer, $H_{M,sp}$. For the non-rigid cavity, however, length oscillations of the pendulum support, described by its response, $H_{M,supp}$, are reduced by each pendulum stage (we can consider N stages, each characterized by a resonant frequency $f_{pend,i}$) in proportion to the square of the frequency. Hence, the overall apparatus response is essentially that of the rigid cavity:

$$A_{det} = \left| H_{M,sp} H_{OM} - \prod_{i=1}^N \left(\frac{f_{pend,i}^2}{f^2} \right) H_{M,supp} H_{OM} \right| \approx |H_{M,sp} H_{OM}|. \quad (2.49)$$

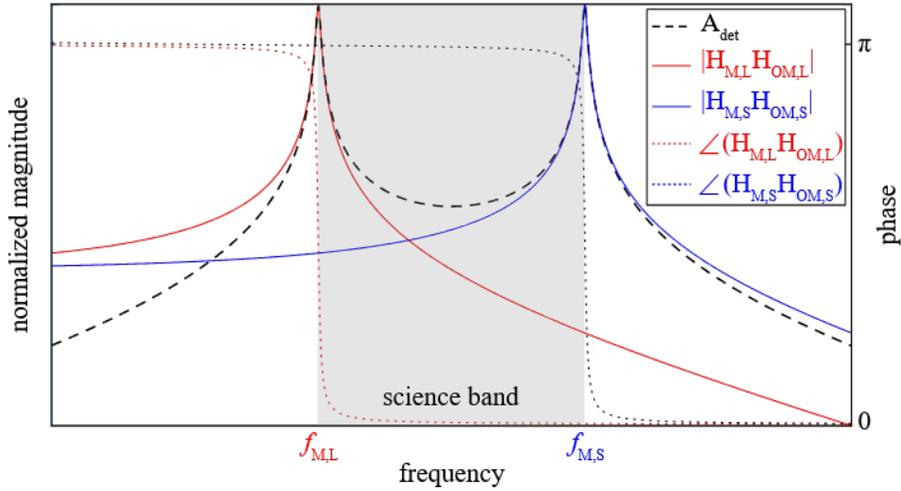


Figure 2.6: Qualitative behavior of the magnitude and phase of the cavity transfer function and overall detector response. Here, the optical poles are much lower in frequency than the mechanical ones and $\eta = 0.01$

2.3 Noise Sources

The sensitivity of the experiment is limited by the aggregate effect of fundamental and technical noise sources. As long as technical noise sources such as laser noise or noise in the detection

chain can be controlled at the relevant level, the ultimate sensitivity is determined primarily by thermal noise in the cavities and shot noise in the detection photodetector, with vibration noise also factoring in at low frequencies. In this section we examine some theoretical estimates of noise sources in order to assess the prospects for the parameter space this experiment can target in each of the two considered scenarios (two rigid cavities or a rigid versus a non-rigid cavity).

2.3.1 Thermal Noise

Thermal vibrations of the atoms in the cavity mirrors, spacers and suspension lead to fluctuations of the axial position and shape of the surface of the mirrors, which are registered by the detector as fluctuations in the cavity length, thereby introducing noise into the system.

The presence of thermal noise is associated with mechanisms that cause energy dissipation, as explained by the fluctuation-dissipation theorem. First formulated by Callen and Welton [59, 60], it posits that any irreversible dissipative process in a system in thermal equilibrium leads to the generation of noise of an amplitude proportional to the mechanical loss corresponding to that process. Quantitatively, given a system characterized by a generalized coordinate $x(t)$, which responds with an impedance $Z(f) = F(f)/\dot{x}(f)$ when driven by an external force $F(f)$, one can express the power spectral density (PSD) of $x(t)$ corresponding to the quiescent state of the system, $S_{xx}(f)$, as:

$$S_{xx}(f) = \frac{k_B T}{\pi^2 f^2} \left| \operatorname{Re} \left[\frac{1}{Z(f)} \right] \right|, \quad (2.50)$$

with k_B being the Boltzmann constant and T the temperature of the system.

In the context of optical cavities, $S_{xx}(f)$ sets a limit on the minimum mirror displacement that can be resolved as a consequence of thermal noise. The dominant type of thermal noise is Brownian noise associated with internal mechanical losses (structural damping) in the material of the mirror coatings and substrates, spacer (for the rigid cavity) and suspension (primarily for the

non-rigid cavity).

Mirror substrate thermal noise:

The thermal noise associated with the mirror substrates can be derived from Eq. (2.50) by applying a cyclical pressure with a profile identical to that of the Gaussian beam incident on the mirror and calculating the power dissipated as a result to find the impedance of the system [61, 62]. The resulting expression for the displacement PSD (in units of m^2/Hz), valid at frequencies below the mechanical resonant modes of the mirror, is:

$$S_{xx,sub}(f) = \frac{2k_B T}{\pi^{3/2} f} \frac{1 - \sigma_m^2}{E_m w} \phi_m(f), \quad (2.51)$$

where σ_m is the substrate Poisson ratio, E_m its Young's modulus, w the beam waist at the surface of the mirror and ϕ_m the loss angle of the substrate material, typically assumed to be frequency independent.

Coating thermal noise:

An approximate expression for the thermal noise in the mirror coatings is given by [63]:

$$S_{xx,coat} = \frac{2k_B T (1 - \sigma_m^2) t_c E_c^2 (1 + \sigma_m)^2 (1 - 2\sigma_m)^2 + E_m^2 (1 + \sigma_c)^2 (1 - 2\sigma_c)}{\pi^2 f E_m w^2 E_m E_c (1 - \sigma_m^2) (1 - \sigma_c^2)} \phi_c. \quad (2.52)$$

Here, in addition to the quantities previously defined, E_c , σ_c and ϕ_c are the Young's modulus, Poisson ratio and loss angle of the coating and t_c its thickness. Eq. (2.52) assumes the properties of the coating are direction-independent, although real-life coatings are anisotropic, with different values of the loss angle in directions parallel with the coating plane and perpendicular to it.

While the substrates are usually made of single-crystal materials which exhibit very low loss

angles, the coatings are primarily designed with their optical properties in mind and are typically made of stacked layers of different materials. As such, thermal noise in coatings tends to be higher, often being the dominant form of thermal noise in optical cavities. Recent advances in the development of crystalline GaAs/AlGaAs coatings [64] enabled the fabrication of mirrors with lower coating thermal noise than substrate noise at room temperature.

Spacer thermal noise:

The spacer Brownian noise can be estimated by treating the spacer as a simple harmonic oscillator with mass m and an effective spring constant associated with its longitudinal mode k , with internal damping characterized by the loss angle ϕ_{sp} . In this case the noise PSD can be expressed as [65]:

$$S_{xx,sp} = \frac{4k_B T k \phi_{sp}}{2\pi f \{ [k - m(2\pi f)^2]^2 + k^2 \phi_{sp}^2 \}}. \quad (2.53)$$

For a ballpark estimate, the spacer can be approximated as a cylindrical tube of length L , outer radius R_{out} , inner radius R_{in} , and Young's modulus E_{sp} . In this case, for the longitudinal oscillation mode (the accordion mode) of the symmetrically supported spacer, one can take m to be half the mass of the spacer and the spring constant $k = E_{sp}\pi(R_{out}^2 - R_{in}^2)/(L/2)$.

Suspension thermal noise:

The suspension thermal noise is associated with position fluctuations of the pendulum the cavities are hanging from. As such, it is more relevant for the non-rigid cavity, where the two mirrors are suspended from separate pendula. The spectral density can be calculated [66, 67] using a modal expansion over the pendulum, torsion and violin modes (the vibrational modes of the pendulum wires), where each individual mode can be described by Eq. (2.53). The result is equivalent with a white noise spectrum corresponding to a thermal environment in equilibrium at temperature T ,

modified by the transfer function of the pendulum. Given that the thermal drive is weaker than that from the vibrations of the pendulum support, in our experiment the suspension thermal noise falls below contributions from vibration noise.

As observed in Eq. (2.51-2.53), the power spectral density of Brownian thermal noise scales linearly with the system temperature and with the loss angle. Thermal noise can be therefore reduced by using materials with high mechanical quality factor and cooling the experiment to cryogenic temperatures.

Brownian motion is not the only mechanism generating thermal noise. Temperature fluctuations lead to thermal gradients that translate into length fluctuations due to the thermal expansion of the material, a phenomenon called thermo-elastic noise. They also cause changes in the refractive index of the mirrors which manifest as phase noise, known as thermo-refractive noise. Collectively referred to as thermo-optic noise, these two effects are characterized by power spectral densities that scale as T^2 and therefore their contribution is significantly smaller than that of Brownian noise in a cryogenic environment. For the purposes of the current work, these thermal noise sources can thus be ignored.

2.3.2 Shot Noise

Shot noise stems from the discrete nature of the photons arriving at the detection photodetector. Since the counted number of photons N follows a Poisson distribution, it is subject to fluctuations of magnitude \sqrt{N} . These fluctuations give rise to noise in the photodetector current I given by $\langle i_{shot} \rangle = \sqrt{2eI}$ (with e the electron charge).

Assuming a detection scheme in which the beat note between the beams resonant to the two cavities is measured, as described in section 2.2.2, with power P in each of the beams reaching the photodetector, the photodetector current is given by:

$$I = \frac{\eta e(2P)}{hf_{\text{opt}}} \left(1 + \cos(\omega_{\text{beat}}t + \Delta\phi_{\text{beat}}) \right), \quad (2.54)$$

where f_{opt} is the laser frequency, ω_{beat} and $\Delta\phi_{\text{beat}}$ are the beat note angular frequency and phase, and η the photodetector quantum efficiency. The corresponding shot noise is related to the DC component of this current and given by a flat power spectral density:

$$S_I(f) = \frac{4e^2\eta P}{hf_{\text{opt}}}. \quad (2.55)$$

The phase noise PSD is then obtained by dividing the shot noise PSD by the signal power:

$$S_\phi(f) = \frac{S_I(f)}{(2\eta eP/hf_{\text{opt}})^2} = \frac{hf_{\text{opt}}}{\eta P}. \quad (2.56)$$

Subsequent conversion to frequency noise by multiplying the phase noise by f^2 and then to strain noise through division by f_{opt}^2 leads to an expression for the strain noise PSD:

$$S_h(f) = \frac{h}{\eta f_{\text{opt}} P} f^2. \quad (2.57)$$

2.3.3 Vibration Noise

Laboratory vibration spectrum

While not a fundamental noise source, seismic motion, laboratory equipment and human activity contribute unavoidable vibration noise to the experiment. Seismic noise is expected to dominate at low frequencies (below 1 Hz), with a prominent peak around 1/7 Hz caused by ocean waves, known as the micro-seismic peak, and fall rapidly above 1 Hz.

The level of vibration noise in the laboratory is a major factor to take into account when designing the vibration isolation for the experiment. A typical spectrum up to 10 kHz, measured on the

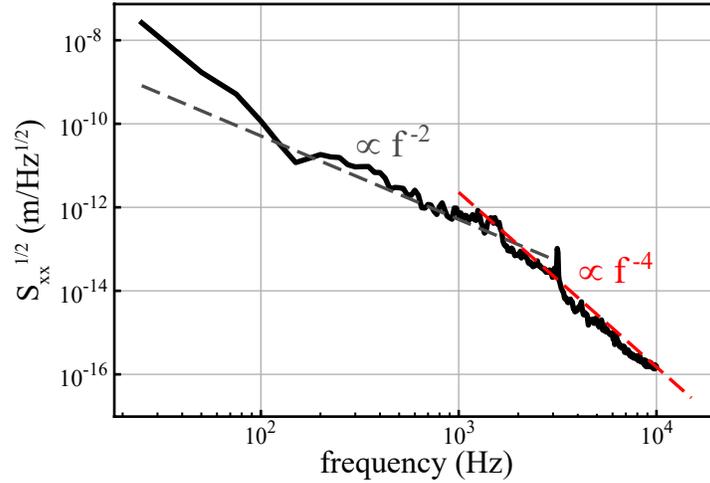


Figure 2.7: Spectrum of vibration noise measured on the optical table. The displacement noise falls approximately as f^{-2} above around 100 Hz and as f^{-4} above around 1000 Hz

optical table using an accelerometer (PCB Piezotronics 393A03), is presented in Fig. 2.7. Below 1 kHz the noise floor exhibits an inverse square dependence on frequency, a behavior that has been reported in previous studies and is a characteristic of vibration noise that is not location-specific [68]. Over this region, the measured data fits to:

$$S_{xx}^{1/2} \approx 5.1 \times 10^{-7} \left(\frac{1\text{Hz}}{f} \right)^2 \left[\text{m}/\sqrt{\text{Hz}} \right]. \quad (2.58)$$

Above around 1 kHz the observed slope changes and the vibration spectrum drops approximately as the fourth power of the frequency:

$$S_{xx}^{1/2} \approx 2.3 \left(\frac{1\text{Hz}}{f} \right)^4 \left[\text{m}/\sqrt{\text{Hz}} \right]. \quad (2.59)$$

Pendula as mechanical filters

The magnitude of vibrations can be lowered beneath the level of the fundamental noises through the use of mechanical filters (an example of which is hanging the cavities from pendula). In this

case, the vibration noise of the environment is multiplied by the transfer function of the vibration isolation setup.

Each filter can be modeled as a simple harmonic oscillator with resonant frequency f_0 and quality factor Q , which is a measure of its damping (some amount of damping is not only inevitable, but often implemented on purpose to dissipate the energy accumulated in its natural modes). Assuming a viscous damping model, the provided vibration isolation is frequency dependent and given by the transfer function of a damped harmonic oscillator:

$$H(f) = \frac{1 + iQ^{-1}(f/f_0)}{1 - (f/f_0)^2 + iQ^{-1}f/f_0}, \quad (2.60)$$

This transfer function behaves as follows:

- In the limit $f \ll f_0$ it is flat and approaches 1 (in other words the motion of the payload follows that of the suspension point)
- At $f = f_0$ (the resonant frequency) it presents a peak of height Q
- For $f \gg f_0$ the transfer function rolls off as $(f_0/f)^2$ up to frequencies $f \approx Qf_0$
- For $f \gg Qf_0$ the vibration isolation improves with frequency as $1/f$. This degradation of performance at high frequencies is a somewhat subtle particularity of the viscous damping model.

Passive mechanical filters can be cascaded, in which case each stage improves the attenuation by a factor proportional to $\sim 1/f^2$ (up to frequencies around Qf_0).

2.3.4 Noise estimates for the current apparatus and future prospects

Fig. 2.8(a) presents the contributions of fundamental noises (thermal and shot noise) and vibration noise for the current apparatus, expressed as strain noise. It assumes that the signal is detected

in the beat note of the two cavities, in transmission. For two unequal cavities, this detection method is advantageous, as their different opto-mechanical transfer functions associated with the conversion of cavity strain noise to frequency noise in the cavity transmission, as given by Eq. (2.32), extend the sensitivity range of the experiment below the mechanical poles.

In evaluating the vibration noise, two estimates of the vibrations inside the cryostat were considered: one involved using the measured spectra of vibrations in the room (Fig. 2.7) and theoretical transfer functions of the cryostat vibration isolation system, while the other relied on cryostat vibration specifications provided by the manufacturing company (see section 5.1). We decided to present the latter, which yielded a more conservative value. The cryostat vibrations spectrum was multiplied by the transfer function of the pendula the cavities hang from (Eq. (2.60)).

The strain noise was also converted to a corresponding limit on the dimensionless coupling constant d_{m_e} . This can be interpreted as the ultimate achievable sensitivity of the apparatus, as long as the technical noise level is lowered below that of fundamental noise. Fig. 2.8(b) presents limits on d_{m_e} , assuming a four-day measurement time (4×10^6 s). Indirect bounds set by EP violation tests (labeled "EP"), as well as the theoretical limit below which the scalar field theory is natural for a 10 TeV cutoff ("naturalness") are provided for reference.

The total noise could be further lowered by using longer cavities and by increasing the optical power going into the system. Prospective noise estimates, corresponding to 30 cm and 1 m cavity lengths and 10 mW of laser power per cavity at the photodetector, are given in Fig. 2.9(a). With these parameters, the experiment can improve on limits from EP violation tests over a wide frequency range and cut into the parameter space below the naturalness line (provided that technical noise can be reduced to comparable levels). Although these length upgrades could not be accommodated in our existing cryostat, they are not on an unreasonable scale to pursue in a tabletop setup.

The limits on the sensitivity to d_{m_e} were also evaluated for the scenario where a rigid cavity

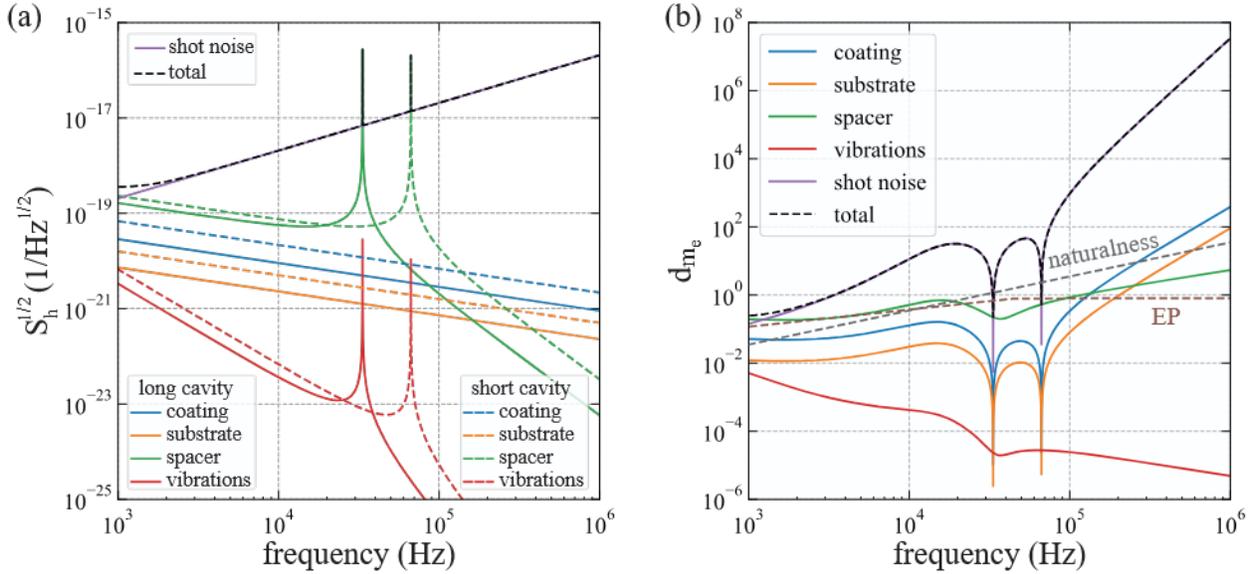


Figure 2.8: (a) Noise analysis for current experiment parameters. Depicted are thermal noise contributions from the mirror coatings and substrates and cavity spacers, the photodetector shot noise and the predicted noise from vibrations. (b) Limits on d_{m_e} set by each of the considered noise sources, for a four-day measurement time. Bounds set by equivalence principle violation tests (“EP”) and a theoretical limit (“naturalness”) are included for comparison

and a non-rigid cavity are used. We assumed a length of 30 cm for both, a limit set by the size of the cryostat, and the same measurement time as before (four days). We also assumed that the non-rigid cavity mirrors hang from two-stage pendula which, according to preliminary calculations, are required to keep the noise from vibrations below the shot noise level. Since in this case the two cavities have equal optical poles, better sensitivity is obtained with this setup if the beat note is measured in reflection. As evidenced by the results presented in Fig. 2.9(b), the upgrade to a non-rigid cavity provides an alternate path to exceeding the naturalness bound, one that is also achievable in our current apparatus.

A detailed list of the parameters used to generate Fig. 2.8 and 2.9 is available in Appendix C.

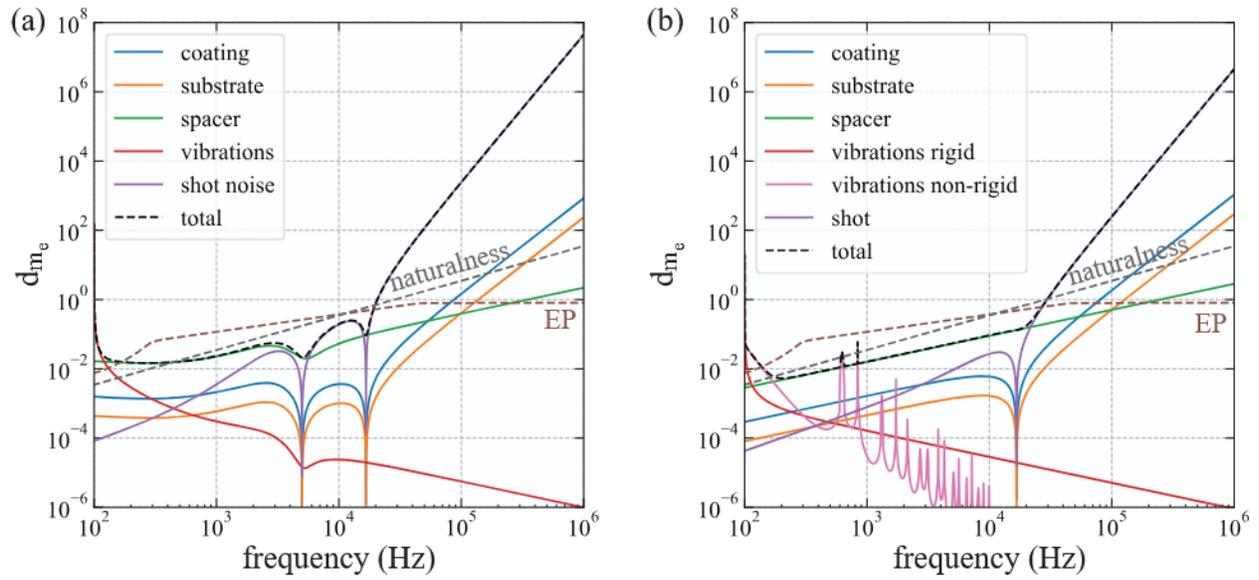


Figure 2.9: Projected limits on d_{m_e} corresponding to fundamental and vibration noise, presented alongside current limits from equivalence principle violation tests and the naturalness bound for (a) two rigid cavities and (b) a rigid and a non-rigid cavity

CHAPTER 3

THE CRYOGENIC CAVITIES

3.1 General design principles and fabrication

Lasers locked to high-finesse optical cavities provide some of the most stable short-term frequency references, fractional frequency variations as low as 4×10^{-17} having been attained [69]. At room temperature, the material of choice for ultra-stable cavities is usually ultra-low expansion (ULE) glass, the low coefficient of thermal expansion (CTE) of which ensures good stability with respect to temperature fluctuations. At cryogenic temperatures, the most common materials are monocrystalline sapphire [70] and silicon [69, 71], which are preferred because of their low CTE at low temperatures.

For the current work, sapphire is used for the mirror substrates and spacers of the cavities. In addition to its low CTE at cryogenic temperatures ($1 \times 10^{-10} \text{ K}^{-1}$ at 4 K [72]), sapphire exhibits low specific heat ($1.1 \times 10^{-2} \text{ J}/(\text{kg} \cdot \text{K})$ at 5 K [73]) and high thermal conductivity (280 W/(m · K) at 4 K [74]), a requirement for effective heat extraction. Moreover, it has a high elastic modulus, which ensures good stability to mechanical perturbations. It is transparent at 1550 nm, the optical wavelength chosen for this experiment, and it is easy to grow into large single crystal boules.

Sapphire is anisotropic and belongs to the trigonal crystal system, displaying a threefold rotation axis, commonly denoted as the C-axis. Its anisotropy causes many of its physical properties (such as elastic moduli, coefficient of thermal expansion, hardness etc.) to be direction dependent. Therefore, the crystal orientation needs to be taken into consideration during the design and construction of the cavities. For optical elements such as the cavity mirrors, the natural choice is to align the optical axis along the C-axis, as it is the only direction in which sapphire does not experience birefringence, as well as the most mechanically symmetric orientation.

The spacers between the mirrors of the optical cavities are made of $\geq 99.996\%$ monocrystalline sapphire, Optical Grade 3. They were designed by us and fabricated by Rubicon Technology (Bensenville, Illinois). They were machined with the crystal C-axis along the optical axis, so that it would be parallel with the C-axis of the mirrors. The alignment of the spacer and mirror C-axes is relevant because stress in the material occurring at interfaces where the crystal structure is dissimilar increases the loss factor and hence contributes thermal noise to the system (which is also the reason why the spacers are single crystal). This orientation is also advantageous because the Young's modulus is the largest along the C-axis. The cavities were designed with the A-axis perpendicular to the optical axis and vertical, so that the crystal structure would have a mirror plane in the vertical plane, which would render the cavity left-right symmetry. However, a manufacturing error led to the orientation of the 15 cm spacer being rotated by 90° around the cavity axis (the A-axis is horizontal). Timeline constraints did not allow for a corrected spacer to be procured for this iteration of the experiment. Nevertheless, the C-axis, which is the direction with the most distinct properties, is still aligned along the cavity axis, as intended, and the effect of this accidental rotation on the mechanical properties was assessed not to limit its performance, as detailed in section 3.2.3.

We chose a cylindrical geometry with symmetric rectangular cutouts at the bottom (sometimes referred to as "notched"), as illustrated in Fig. 3.1, a design choice that has been shown to exhibit a high degree of insensitivity to vibrations [75, 76]. The cavity is supported at four symmetrically placed points in the horizontal plane of the cutouts, marked by 2 mm deep blind holes, 3 mm in diameter. An additional 3 mm diameter vent hole at the bottom of the cavity allows pumping air out of the cavity. An account of how the dimensions of the cut and the positions of the support holes were chosen is provided in section 3.2.2. Details regarding the mounting of the cavities are presented in section 3.3.

Mirror blanks of our specifications were provided by Coastline Optics (Camarillo, California), who also did the polishing of the mirrors and of the flat ends of the spacers. They are

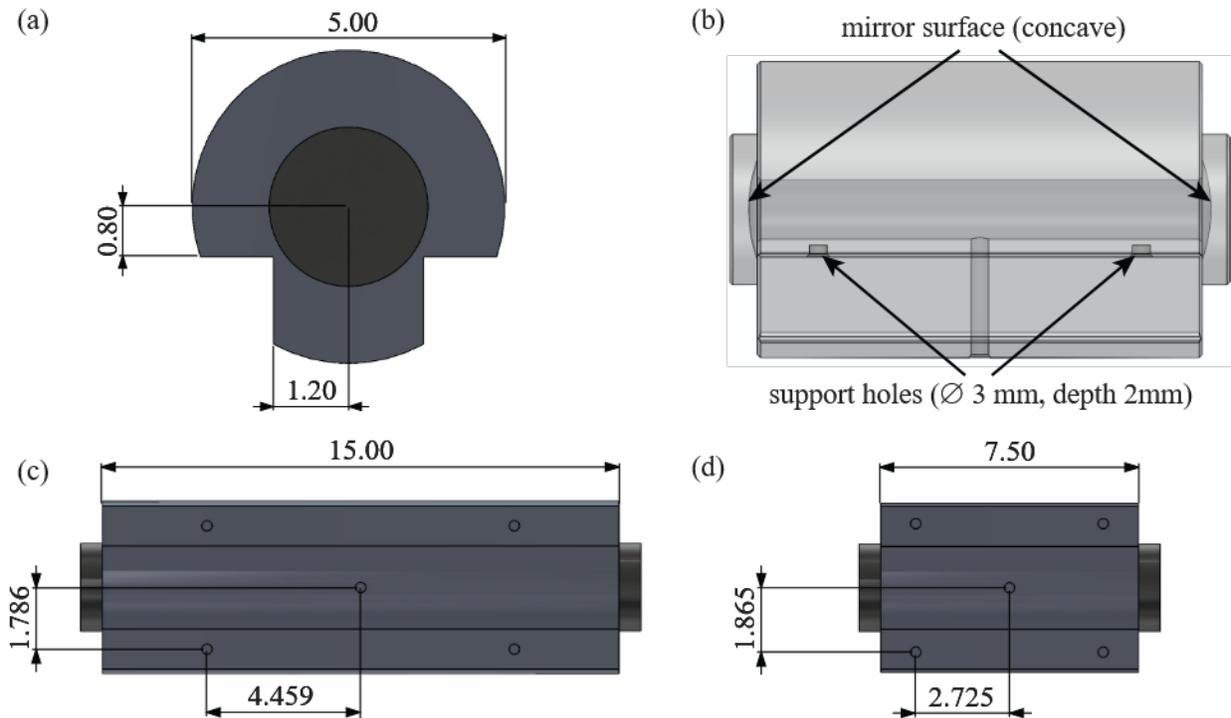


Figure 3.1: (a) CAD model of the 15 cm cavity and 7.5 cm cavity (front view); (b) cavity side view, highlighting the support holes and the mirror surfaces; the mirror curvature is exaggerated for clarity; (c) support holes position for the 15 cm cavity; (d) support holes position for the 7.5 cm cavity

plano-concave, with a 10.2 m radius of curvature and 6.25 mm thick (as depicted in Fig. 3.1 (b)). The mirrors have crystalline coatings made of monocrystalline GaAs/AlGaAs heterostructures, for which loss angles an order of magnitude less than the best dielectric coatings had been demonstrated (2.5×10^{-5} at room temperature [77]). The mirror coating and the optical contacting of the mirrors to the spacer were done by Thorlabs Crystalline Solutions (Santa Barbara, California). The mirrors have a reflectivity of 99.998% (equivalent to a cavity finesse around 150,000).

3.2 Geometry and support position optimization

3.2.1 Vibration sensitivity of the optical cavities

External vibrations that couple into the optical cavities through their supports will degrade the strain sensitivity of the experiment. The sapphire cavities were designed to keep the vibration noise in the system below the fundamental noise sources (thermal and shot noise). To this end, several concurrent strategies are employed. First, the cavities are isolated from their surroundings by multiple layers of passive isolation. They are supported by an inverted pendulum that reduces horizontal vibrations, and by geometric anti-spring filters to target vertical vibrations. These are described in detail in section 5.2. Additional horizontal isolation was provided by hanging the cavities and their supports from a four-wire suspension. Second, careful location of the support points minimize the deformation of the cavities induced by vibrations. Third, the cavities are mounted at the support points on springs and rounded pins, which provide additional isolation while simultaneously ensuring that the support locations are well defined. In the following sections each of these strategies are discussed in more detail.

The frequency stability of the cavities is most affected by two types of vibration-induced deformations: deformations along the optical axis that change the separation between the mirrors (the cavity length) and departures from the parallelism of the mirrors (the relative tilt). Hence, for an acceleration \vec{a} corresponding to the vibration of the supports, the optimal performance of the cavity relies on minimizing the listed quantities. We assume a coordinate system with z along the horizontal optical axis, y in the vertical direction, and x horizontal and transverse to the optical axis:

1. The acceleration sensitivity of the cavity, defined as the ratio of the change in length, in practice computed as the difference between the displacements $u_z^{(i)}$ of the centers of the mirrors along the cavity axis, to an applied acceleration a :

$$k = \frac{(u_z^{(2)} - u_z^{(1)})}{a} \quad (3.1)$$

2. The sensitivity of the cavity tilt around the x-axis to accelerations, expressed as the difference between the tilt angles of the central regions of the mirrors $\theta_x^{(i)}$, divided by an applied acceleration a :

$$k^{(\theta_x)} = \frac{(\theta_x^{(2)} - \theta_x^{(1)})}{a} \quad (3.2)$$

3. The sensitivity of the cavity tilt around the y-axis to accelerations, expressed as the difference between the tilt angles of the central regions of the mirrors $\theta_y^{(i)}$, divided by an applied acceleration a :

$$k^{(\theta_y)} = \frac{(\theta_y^{(2)} - \theta_y^{(1)})}{a}. \quad (3.3)$$

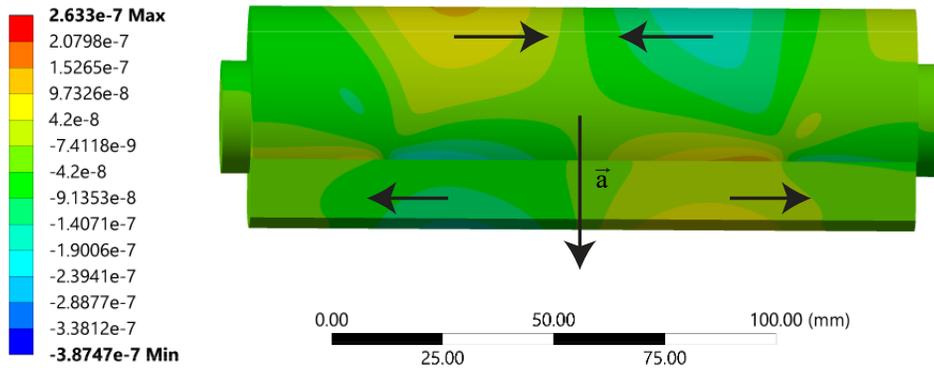


Figure 3.2: Horizontal deformation of the cavity under a vertical force. Arrows illustrate the direction and magnitude of the displacement in a qualitative sense (not drawn to scale)

Since deformations along the cavity axis are the most relevant to its stability, it might be tempting to focus on eliminating the effect of vibrations in this direction. However, transverse vibrations

also affect the cavity length by amounts determined by the Poisson ratio (Fig. 3.2). For this reason, we will address the impact of vibrations in the three directions independently:

1. Longitudinal direction: Elastic deformations due to forces along the cavity axis are considerably suppressed through symmetric placement of the supports. A way to see this is that if the supports are symmetric in the direction of the cavity axis, under a force in the same direction, parts of the cavity will contract, while symmetrically opposite parts will expand, leading to an overall null length change. This effect is illustrated in Fig. 3.3. Since the anisotropy of the material slightly breaks the symmetry along the z -axis, in practice this cancellation is not perfect. However, it is significantly reduced when compared, for instance, to the effect of vertical forces. Furthermore, passive isolation (at frequencies of interest) is easier to achieve in the horizontal direction than vertically with additional pendula. We supplemented the horizontal vibration isolation by hanging the cavities and their cradles from four-wire suspensions, thus relaxing the requirements on their intrinsic vibration stability.
2. Transverse direction: The placement of the support points in a horizontal plane and the left-right symmetry of the cavity reduce effects due to horizontal vibrations orthogonal to the cavity axis. Moreover, the extra attenuation provided by the four-wire pendulum applies to vibrations in this direction as well.
3. Vertical direction: Cavity deformations to vertical accelerations are minimized through an appropriate placement of the support points. For instance, for a given acceleration, the tilt of the cavity depends strongly on the coordinate of the supports along the z -axis (and more weakly on the other directions). It becomes 0, in principle, if the cavity is supported exactly at its Airy points. The cavity length change is dependent on the position of the support in all three directions.

Based on these considerations, the analysis we carried focused on minimizing the effect of

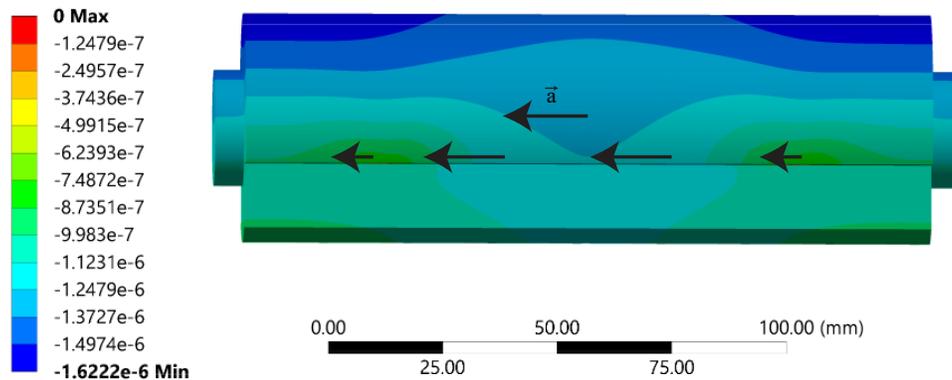


Figure 3.3: Horizontal deformation of the cavity under a horizontal force parallel with the optical axis. Points placed symmetrically with respect to the center are displaced by similar amounts. Arrows illustrate the direction and magnitude of the displacement in a qualitative sense (not drawn to scale)

vertical accelerations. The aim, therefore, was to identify a favorable cut plane in which a set of support points existed for which the length change and relative mirror tilt around the x axis induced by a vertical acceleration (characterized by the acceleration sensitivities k_y and $k_y^{(\theta_x)}$) were simultaneously zero (due to the left-right symmetry of the cavity, we get $k_y^{(\theta_y)} = 0$ automatically).

With the cavity supported at these locations, its stability to vertical vibrations is expected to be determined by a residual acceleration sensitivity stemming from machining tolerances and imperfections in the mounting. Finite element calculations indicate that the horizontal acceleration sensitivity is on the same order of magnitude as this residual vertical acceleration sensitivity for longitudinal vibrations and smaller by a factor of order 1000 for transverse vibrations.

3.2.2 Finite element modeling

The choice of cavity geometry and position of supports was based on finite element analysis, carried out using commercially available software (ANSYS Mechanical [78]).

The quasi-static approach used is similar to that of references [76, 79]. This strategy, much less computationally demanding than examining the response to harmonic perturbations, is justified for

a couple of reasons. First, the effect of vibrations is most pronounced at low frequencies (below 10 Hz), far below the mechanical modes of the cavities (the first of which occur in the kHz range), so forces at these frequencies can effectively be treated as static. Second, as seen in section 2.3.3, each layer of passive isolation adds a degree of attenuation above its resonant frequency scaling as the square of the frequency of the perturbation, so high frequency vibrations are more effectively targeted by the vibration attenuation stages inside the cryostat. Third, while this approach does not account for the behavior of the cavities around their mechanical resonances, it provides information on the broadband "baseline" of the cavity transfer functions, away from the resonances. In other words, minimizing the acceleration sensitivity in the static regime is equivalent to lowering this "baseline" (as illustrated in Fig. 3.4). Rather than applying the perturbation as an acceleration of the supports, it is convenient to work in the reference frame of the supports, in which they are fixed and a gravity-like acceleration is applied to the whole body.

This method was validated for an early iteration of the 15 cm cavity by directly comparing the results of the static analysis carried out in the frame of the supports with those of a full harmonic analysis. For an arbitrary position of the support holes, a sinusoidal displacement with fixed amplitude $x_{support}$ and varying frequency $f = \omega/2\pi$ was applied at their centers and for a few different frequencies the amplitude of the cavity length change, measured as the difference between the displacements at the centers of the mirrors along the optical axis $u_z^{(2)} - u_z^{(1)}$, was computed. This allowed defining a transfer function:

$$H_{cavity}(f) = \frac{u_z^{(2)}(f) - u_z^{(1)}(f)}{x_{support}(f)}. \quad (3.4)$$

This can be related to the acceleration sensitivity k_y obtained from the static stress-strain analysis (Eq. (3.1)) by replacing the static acceleration a with a sinusoidal acceleration at frequency f . In this case, $x_{support} = (2\pi f)^2 a$ and the transfer function can be written as:

$$H_{cavity}(f) = \frac{u_z^{(2)}(f) - u_z^{(1)}(f)}{a/(2\pi f)^2} = 4\pi^2 f^2 k_y. \quad (3.5)$$

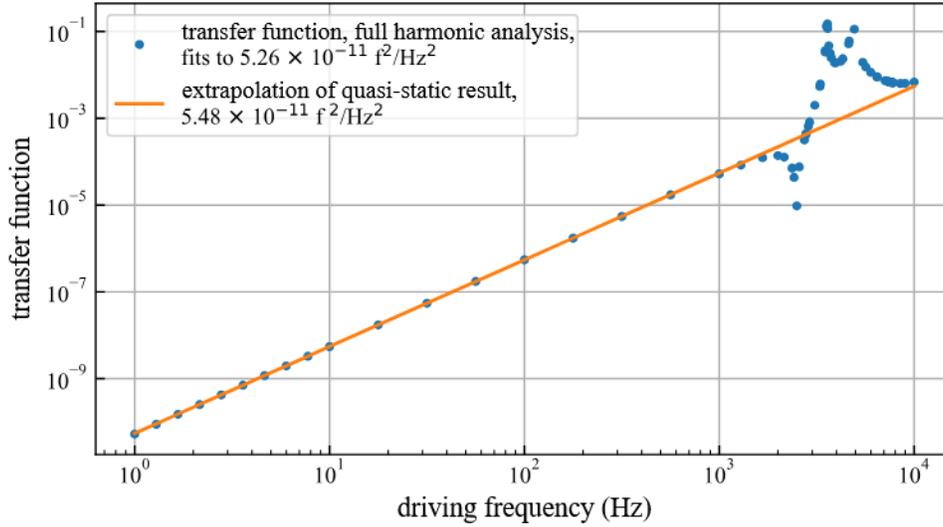


Figure 3.4: Comparison of the harmonic analysis of the cavity deformation with a prediction inferred from the static approximation

It is expected then that away from the effect of the mechanical resonances, the transfer function is proportional to f^2 , with a proportionality constant that can be approximated from the static analysis. As Fig. 3.4 demonstrates, this prediction was confirmed by the results of the finite element simulations.

The supports were modeled as points at the center of the support holes. In order to avoid over-constraining the cavity, the four points were fixed in the vertical direction (by setting the boundary condition $u_y = 0$), but allowed to slide in the horizontal plane. We checked that the results are independent of the size of the supports as long as the contact surface is reasonably close to being point-like by redoing the simulation for one of the cavities with the boundary condition applied instead to 0.3 mm diameter disks placed at the center of the holes. The two sets of boundary conditions (Fig. 3.5) rendered the same results (Fig. 3.6).

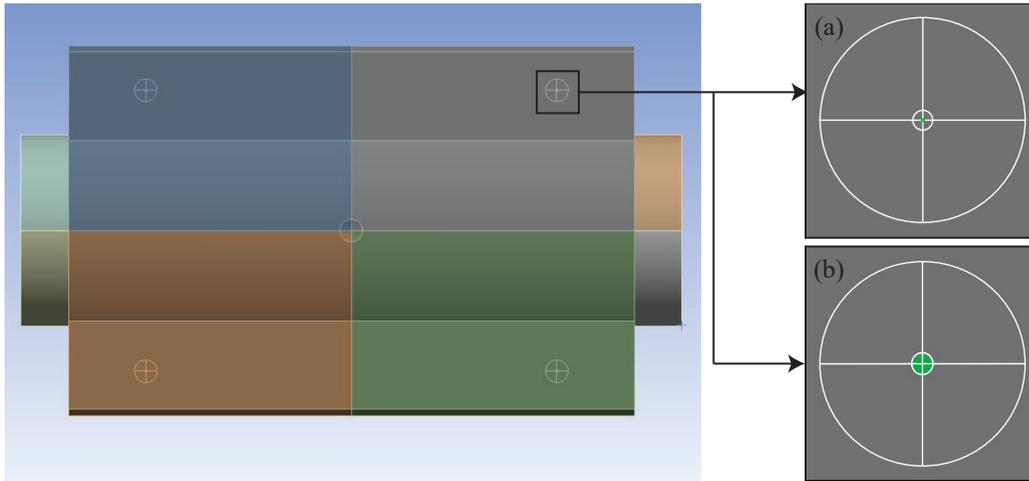


Figure 3.5: Implementation of the supports in the finite element simulation; green shaded regions show the areas over which boundary conditions are applied in the case of (a) point supports and (b) disk supports

For the material properties, provided input parameters were the density of sapphire (3997 kg/m^3) and the anisotropic elasticity matrix at 4 K [80]:

$$D = \begin{bmatrix} 500.0 & 161.6 & 111.4 & 23.26 & 0 & 0 \\ 161.6 & 500.0 & 111.4 & -23.26 & 0 & 0 \\ 111.4 & 111.4 & 502.3 & 0 & 0 & 0 \\ 23.26 & -23.26 & 0 & 151.0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 151.0 & 23.26 \\ 0 & 0 & 0 & 0 & 23.26 & 169.2 \end{bmatrix} \text{ GPa.} \quad (3.6)$$

The locations of the support holes in the horizontal cut plane were varied until coordinates were found for which both the cavity optical length change and the bending of the cavity were null. The dimension of the rectangular cuts provided free parameters to be varied until such positions were possible. The length change and differential tilt were expressed as transfer function as per Eq. (3.4). The zero-crossings of the two transfer functions at the chosen locations are presented in Fig. 3.6 and Fig. 3.7. The simulations were repeated with two different mesh sizes to ensure that mesh

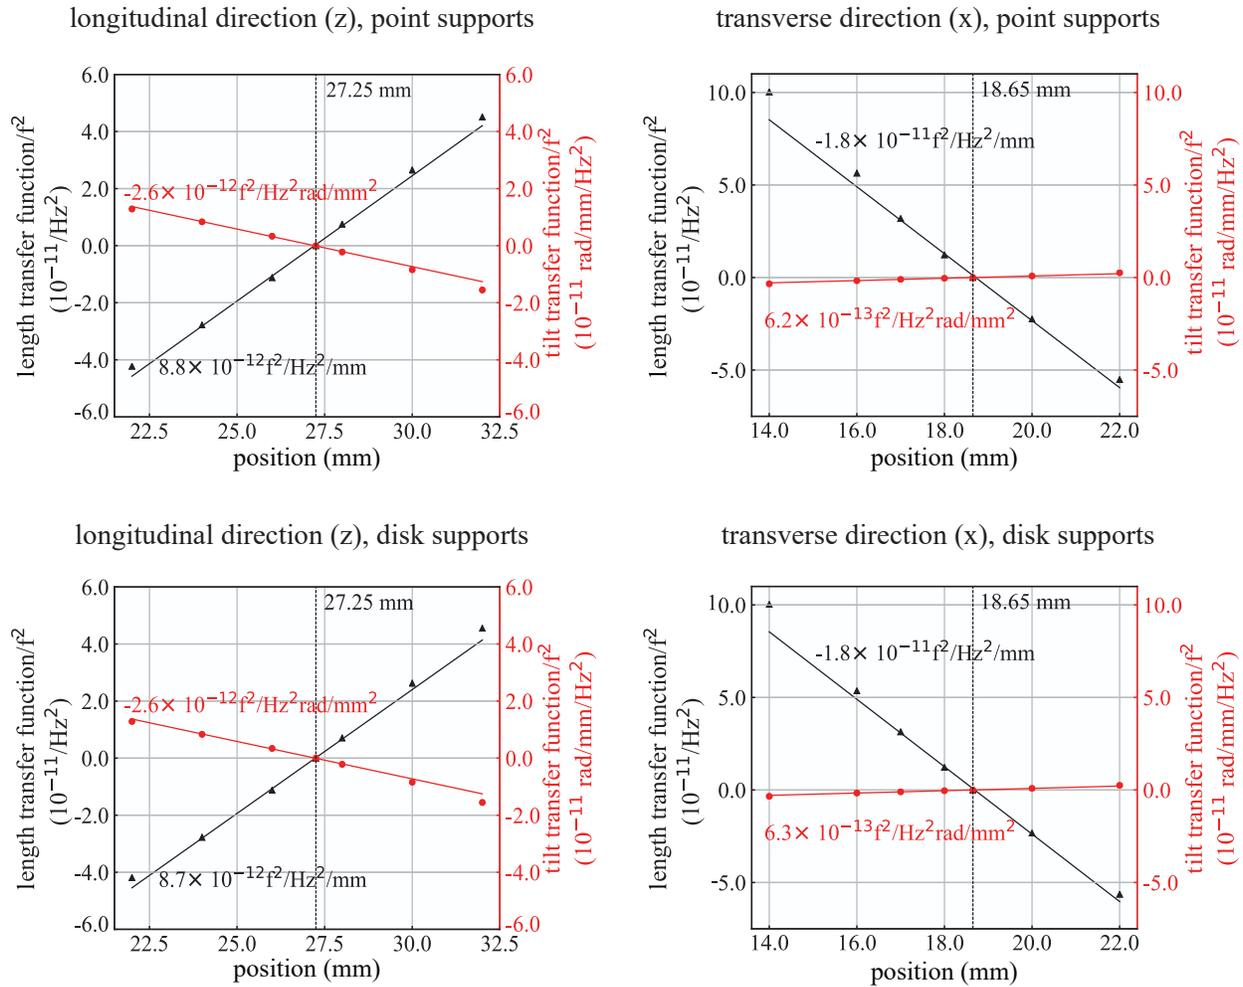


Figure 3.6: Transfer functions corresponding to the length change and mirror tilt of the short cavity as a function of the position of the supports

convergence had been achieved. Moreover, by varying the number of mesh points over which the tilt was calculated, we checked that the simulation results are independent (at the level of the uncertainties caused by machining tolerances) of the size of the region taken into consideration in the calculation up to scales similar to the size of the laser beam waist. The final geometry of each cavity is shown in Fig. 3.1, as well as in Table 3.1.

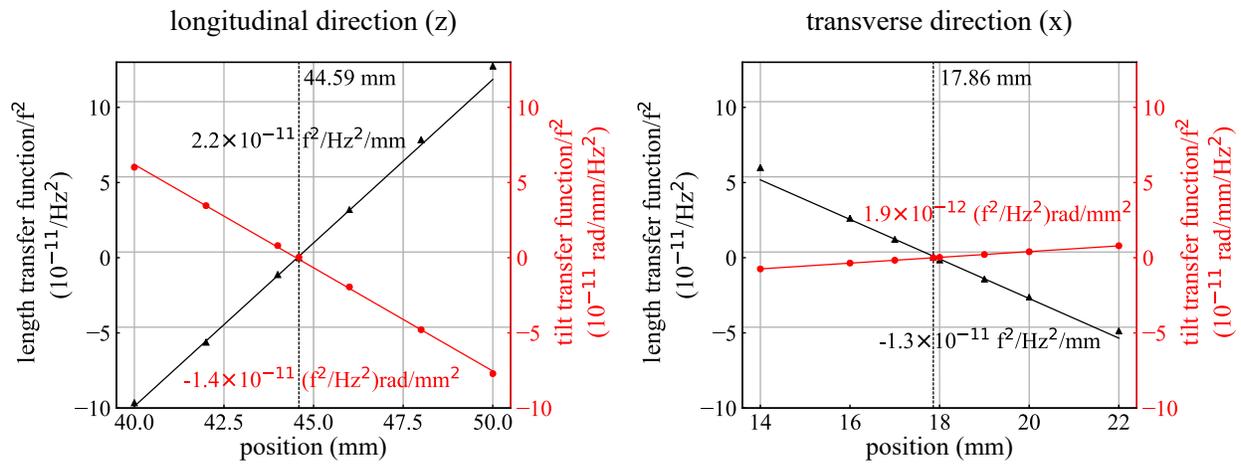


Figure 3.7: Transfer functions corresponding to the length change and mirror tilt of the long cavity as a function of the position of the supports

3.2.3 Effect of machining tolerances and axes misalignment

While in theory with the support holes at their optimal positions the optical lengths of the cavities are perfectly insensitive to vertical accelerations, in practice imperfections in the fabrication of the cavities, as well as the cavity mounting will limit their mechanical performance. The most critical mechanical tolerances are those related to the positions of the support points. The slopes of the plots of the transfer functions against the support position provide information on the effect of inaccuracies in the placement of the support holes. We set an upper limit on this effect by assuming that the holes are displaced in either the longitudinal or the transverse direction by the mechanical tolerance, 0.1 mm. The transfer functions thus estimated for each cavity are compiled in Table 3.2.

The acceleration sensitivity of the cavities is also dependent on the vertical position of the support points, which could be affected by machining imperfections in the position of the horizontal cut plane in which the holes are drilled or in the hole depth. By varying these two parameters in the simulation, we verified that inaccuracies in the vertical position contribute on the same order of magnitude or less as compared to those in the horizontal position. These checks were performed

	Long cavity	Short cavity
Spacer length	15 cm	7.5 cm
Spacer outer diameter	5 cm	5 cm
Inner bore diameter	1 cm	1 cm
Cut-out vertical position below center	0.8 cm	0.8 cm
Cut-out horizontal distance from center	1.2 cm	1.2 cm
Support hole diameter	3 mm	3 mm
Support hole depth	2 mm	2 mm
Support distance from center (z direction)	44.6 mm	27.25 mm
Support distance from center (x direction)	17.9 mm	18.65 mm
Vent hole diameter	3 mm	3 mm
Mirror diameter	2.54 cm	2.54 cm
Mirror radius of curvature	10.2 m	10.2 m
Mirror thickness	6.25 mm	6.25 mm

Table 3.1: Geometrical parameters of the two cryogenic cavities

for the long cavity (Fig. 3.8).

As it was later found out that the A-axis of the 15 cm cavity was accidentally rotated by 90° during fabrication, the acceleration sensitivity was recalculated for the real geometry of the cavity. In this case, the transfer function corresponding to the cavity length change is $1.84 \times 10^{-12} f^2/\text{Hz}^2$ and the support displacement to tilt transfer function is $8.89 \times 10^{-13} (f^2/\text{Hz}^2)\text{rad/mm}$. Reassuringly, these values ended up being on the same order of magnitude as the uncertainties due to machining tolerances, so the rotation of the crystal axes is not expected to significantly degrade the stability of the cavity. This is likely because the orientation of the C-axis was not affected by

	Long Cavity		Short Cavity	
	Length (f^2/Hz^2)	Tilt ($\text{rad}/\text{mm} \times f^2/\text{Hz}^2$)	Length (f^2/Hz^2)	Tilt ($\text{rad}/\text{mm} \times f^2/\text{Hz}^2$)
Longitudinal	2.2×10^{-12}	1.4×10^{-12}	8.8×10^{-13}	2.6×10^{-13}
Transverse	1.3×10^{-12}	1.9×10^{-13}	1.8×10^{-12}	6.2×10^{-14}

Table 3.2: Length and tilt transfer functions for a 0.1 mm longitudinal or transverse displacement of the cavity support holes

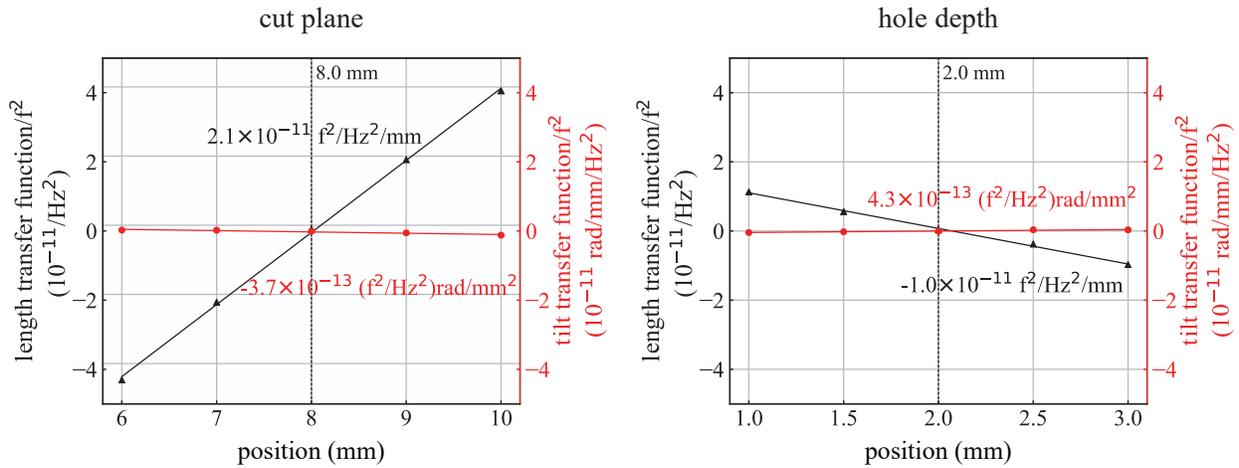


Figure 3.8: Transfer functions corresponding to the length change and mirror tilt of the long cavity as a function of the vertical position of the cut plane and of the hole depth

the fabrication error, and the Young’s modulus is only weakly direction-dependent in the plane perpendicular to it.

3.3 Suspension design and construction

The 7.5 cm and 15 cm cavities lie side by side inside the cryostat, supported on aluminum cradles suspended from the 4 K plate. Each cavity is secured by a safety structure bolted into the 4 K plate, preventing them from falling for more than a few mm in the event of wire failure.

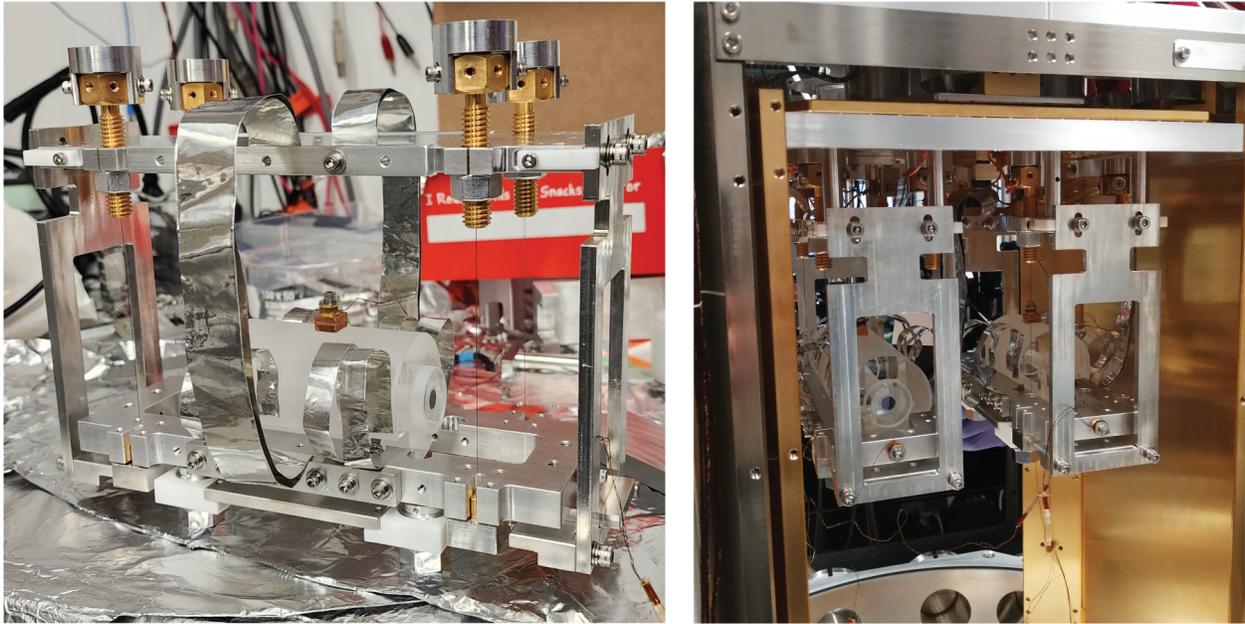


Figure 3.9: Left: horizontal short cavity on its support during assembly. Right: horizontal sapphire cavities side by side inside the cryostat

3.3.1 Mounting

As four support points over-define the support plane, if the contacts of the cavity with its mount were rigid, machining tolerances would make it impossible to ensure that the weight of the cavity is equally distributed. The usual strategy for counteracting this issue is to make the contacts flexible, often in the form of rubber balls [75]. Unfortunately, the realization of a flexible support presents more challenges at cryogenic temperatures, where most materials lose their elasticity.

Hence, in order to preserve well-defined contact points, while also ensuring some degree of compliance of the supports, a novel design was employed, consisting of custom-made stainless steel pins resting on conical steel springs. The rounded-head pins are designed to tightly fit inside the 3 mm holes in the bottom surface of the cavities, ensuring accurate centering of the supports while also maintaining a single contact point at the tip. The steel springs are expected to retain some amount of flexibility even at 4 K.

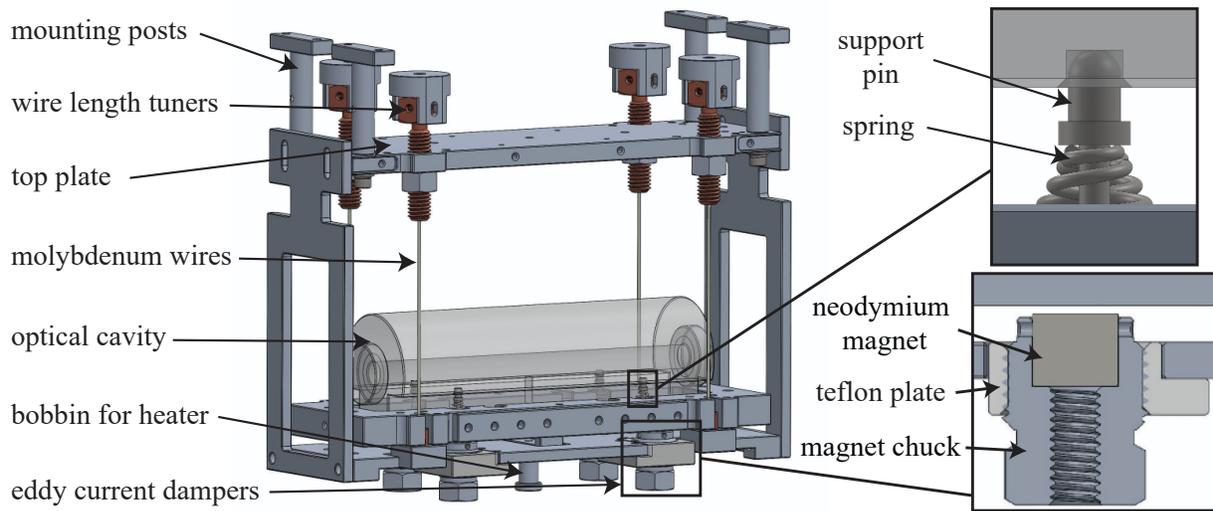


Figure 3.10: Rigid cavity support and suspension. Top inset: support pin and spring. Bottom inset: section through eddy current damper

3.3.2 Four-wire suspension

To add another layer of horizontal vibration isolation, the cradles the sapphire cavities rest on are suspended from four 0.010" thick molybdenum wires, approximately 15 cm long. The structure behaves as a pendulum with the main oscillation mode around 1.5 Hz and is expected to suppress by several orders of magnitude mechanical excitations in the tens of kHz range, which could potentially obscure a dark matter signal. The selection of molybdenum as the material for the wires was motivated by the good trade-off it provides between strength, low loss angle (values in the 10^{-6} range having been reported at 4 K [81]) and reasonable thermal conductivity (57 W/(m · K) at 4 K [82]), as well as being easy to work with (a previous candidate, tungsten, proved to be too brittle to develop a robust assembly procedure).

Each wire goes through a gold-plated copper vented screw that is fastened to the suspension top plate (see Fig. 3.10), which is rigidly attached to the 4 K plate through four posts. This design allows for tuning the vertical position of the attachment points of the wires until the tension in all four of them is equal. The ends of the wires are passed through and wrapped several times around

cylindrical cup-shaped pieces, also made of gold-plated copper (the "wire heads"), which are then filled with melted solder so that the wire is mechanically secured. The wire heads are then firmly clamped to the top plate and cradle, respectively. The use of copper, the immersion of a portion of the wire ends in solder for better thermal contact with the wire heads and the high pressure exerted by the clamps aimed to create a good heat conduction path to cool the cavities through the wires. Moreover, for the same reason, copper braids were attached between the top clamps and the 4 K plate. However, the later realization that the thin molybdenum wires themselves create a bottleneck for heat extraction led to the addition of extra thermal links (more details are provided in Section 3.4.1).

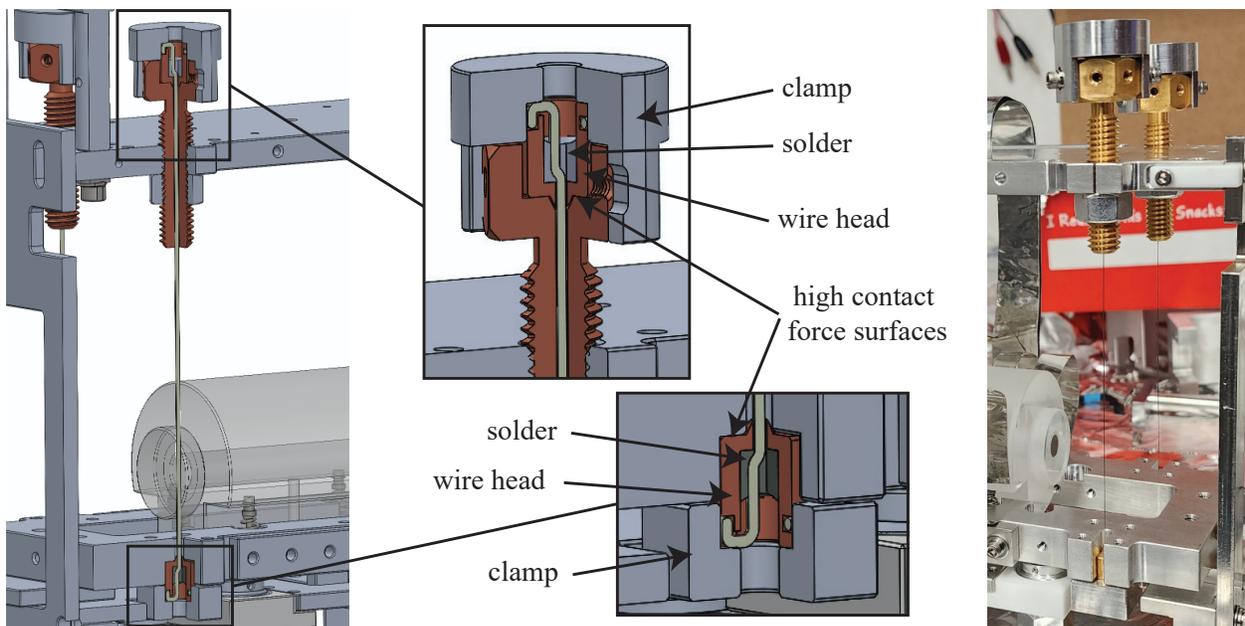


Figure 3.11: Rigid cavity wire assembly

3.3.3 Eddy current damping

The benefit of vibration noise reduction in the measurement frequency band brought by the 4-wire suspension comes at a cost: the resulting pendulum is easily excited and will amplify vibrations close to its resonances. Due to the intrinsic high quality factor of the wires and the high

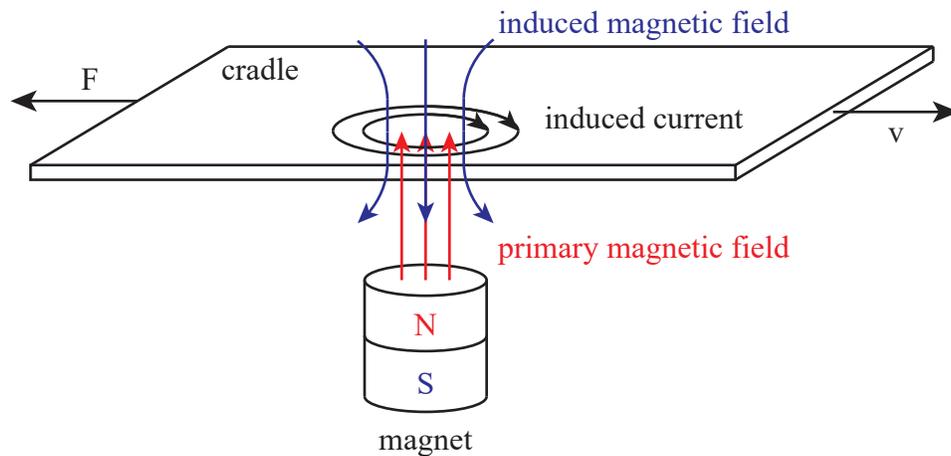


Figure 3.12: Diagram illustrating the principle behind eddy current damping

vacuum environment, this motion takes a long time to settle down on its own. In order to quickly lock the laser to the cavities after a perturbation, we attenuate it using eddy current damping.

The working principle of an eddy current damper is depicted in the diagram in Fig. 3.12. It is based on the induction of eddy currents in the surface of a conducting material when it moves with respect to a magnetic field. These currents generate a magnetic field of their own, with a polarity so as to oppose the motion of the conductor. The resulting behavior is that of a viscous damping mechanism acting on the system, characterized by a damping coefficient that is proportional to the strength of the magnetic field and the electrical conductivity of the conducting object. A notable advantage of this damping method is that it requires no contact between the pendulum and any external structure.

In our system, eddy current damping is implemented using four permanent magnets located on the fixed bottom plate of each cavity's safety structure, beneath the cradle. The magnets, which are neodymium cylinders 7/16" in diameter, 3/8" thick, are placed inside aluminum chucks screwed into the bottom plate and oriented with their poling axis vertical, so that the cradle moves perpendicularly to it. With this mounting method the distance between the magnets and the cradle (and consequently the damping coefficient) is adjustable. Furthermore, their placement allows them to

effectively target all the main oscillation modes of the pendulum (swings in two directions and torsional motion).

3.4 Cooling down the sapphire cavities

3.4.1 Heat extraction and thermal links

One challenge that comes with the stringent vibration noise requirements of the experiment is that there is a general trade-off between heat conduction and vibration isolation in a system. Large surface area connections between different components, which are favorable for heat extraction, also facilitate the coupling of external vibrations.

Since in designing the cavity supports we mostly optimized for vibration isolation, we decided to compensate for the reduced heat conductivity where needed by adding flexible thermal links, made of high thermal conductivity aluminum sheets. This material, which is also used for the heat links inside the cryostat (see section 5.1), is particularly soft, and has been specifically developed for low vibration applications.

At first, heat links were added between the sapphire spacers and the cradles, to overcome the thermal constriction at the point-like contacts between the cavities and the pins they sit on. Initial cool downs of the two cavities, however, pointed towards the thin suspension wires also limiting the cooling of the sapphire cavities.

To gain some intuition about the heat extraction capacity of the molybdenum wires we carried out a theoretical investigation of the time needed to cool down the cavities with and without thermal links. Though rather crude, these estimates proved to be informative. We modeled the evolution of the temperature of the cavities in terms of heat transferred via conduction through the wires and heat links and through radiation to their environment (for simplicity, we neglect any other heat exchange with the outside):

$$\begin{aligned}
& (m_{Al}c_{Al}(T) + m_{Sapph}c_{Sapph}(T)) \frac{dT}{dt} = -8k_{Mo}(T)A_{wire} \frac{d(T - T_{cold})}{dx_{wire}} \\
& - Nk_{link}(T)A_{link} \frac{d(T - T_{cold})}{dx_{link}} - \sigma (e_{Al}A_{Al} + e_{Sapph}A_{Sapph}) (T^4 - T_{cold}^4) .
\end{aligned} \tag{3.7}$$

Here, we ignore thermal gradients across the cavities and treat the two cavities and cradles as having the same temperature $T(t)$ and exchanging heat with the surroundings at temperature $T_{cold}(t)$. In Eq. (3.7), m_{sapph} and A_{sapph} are the total mass and surface area of the two cavities, with $c_{sapph}(T)$ being their temperature dependent heat capacity and e_{sapph} their emissivity. Similarly, m_{Al} and A_{Al} are the mass and surface area of the two cradles, approximated to be made entirely out of aluminum, with thermal conductivity $c_{Al}(T)$ and emissivity e_{Al} . A_{wire} denotes the cross-section area of each of the 8 molybdenum wires supporting the two cavities and A_{link} is the cross-section of a heat link, of which we assume a total number N . k_{Mo} and k_{link} are the temperature-dependent thermal conductivities of molybdenum and of the heat links and σ is the Stefan-Boltzmann constant ($5.67037 \times 10^{-8} \text{ W}/(\text{m}^2\text{K}^4)$).

We further simplify the terms corresponding to the heat conduction by approximating the thermal gradients across the wires and heat links as:

$$\frac{d(T - T_{cold})}{dx} \approx \frac{T - T_{cold}}{L}, \tag{3.8}$$

where L is either the length of the wires or that of the heat links, respectively.

As an estimate of the temperature $T_{cold}(t)$, we used data for the time dependence of the temperature of the cold plate from a previous cool down. Though not expected to be entirely accurate, as the cold plate cooling curve can vary between cool downs depending on the payload attached to it and how it is connected, this model is expected to be more accurate than, for instance, assuming a constant T_{cold} , at the final temperature of the 4 K plate.

The results of the calculations (Fig. 3.13) confirmed the suspicion that the suspension wires act as a bottleneck for heat extraction. In fact, they indicate that without additional heat links, at early

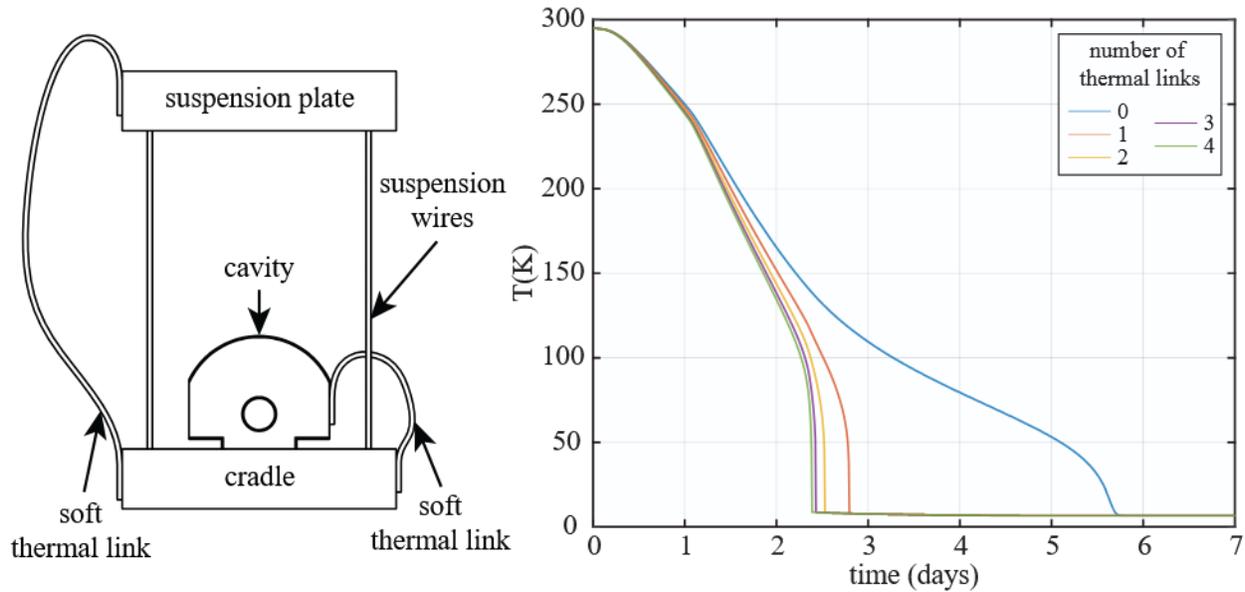


Figure 3.13: Left: thermal links diagram. Right: estimates of the time required to cool the sapphire cavities as a function of the number of thermal links (per cavity) between the suspension plate and the cradle

stages, the cooling of the cavities is predominantly radiative.

Based on these estimates, in the end, four thermal links were attached between the top plates of the suspension and the cradles of each cavity (as displayed in Fig. 3.13). The links were cut out of 0.004" thick sheets of aluminum and are approximately 9" long and 3/4" wide. 1/2" wide strips were also cut for the heat links between the cavity spacers and the cradles (8 of them for the short cavity, 12 for the long cavity). The strips are tightly clamped to the cradles and top plates by flat aluminum plates screwed into their lateral sides. Thermally conductive silver paint (PELCO Conductive Silver Paint from Ted Pella Inc.) was used to attach them to the sapphire spacers.

3.4.2 Temperature sensors and wiring

The temperature of the cryogenic cavities is monitored during cool-down, warm-up and normal operation by 4 temperature sensors, one on each cavity spacer and cradle. The sensors are silicon

diodes (Lakeshore DT-670) with a working range of 1.4-500 K. The reading is done in a 4-wire configuration, where a pair of leads is used to inject a constant current and the voltage across the diode is measured through a second pair. This eliminates contributions due to the resistance of the leads.

Obtaining accurate and reliable temperature measurements at cryogenic temperatures requires special care, so that the sensors maintain good thermal contact with the samples being probed as the cryostat cools down and that heat flow through the leads does not raise the sensor temperature above that to be measured.

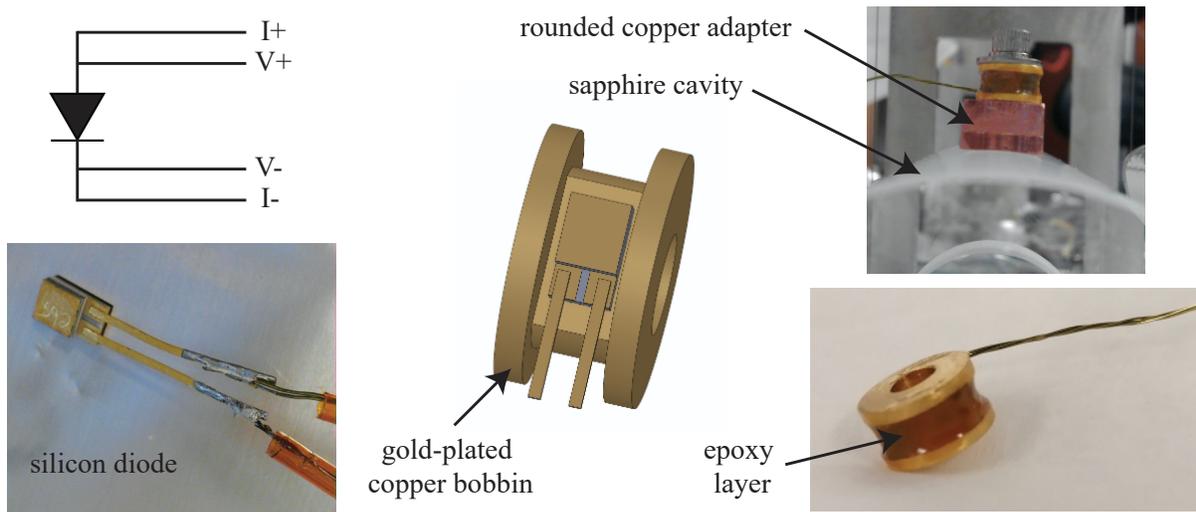


Figure 3.14: Pictures of the temperature sensors and 4-wire measurement diagram

For these reasons, the diodes and their leads are thermally anchored at two places. The first point is at the surface being interrogated and ensures that the sensors and their wires are in thermal equilibrium with the cavity or cradle they are attached to. This is done by gluing each silicon diode to a gold-plated OFHC copper bobbin using a thin layer of thermally conductive epoxy (Epo-Tek H70E). The sensor leads, twisted pairs of 0.005" thick constantan or manganin wire, are then wrapped around the bobbin several times and embedded in a thick epoxy layer, which serves both a mechanical role, fixing the wires in place, and a thermal role, ensuring good heat conduction to

the bobbin. Kapton is used to electrically insulate and protect the leads and diode. The bobbins are then mounted on the cavity cradles and spacers. For the former, the bobbin is directly bolted to the cradle, a high contact force being ensured by the screw. As the same was not possible for the cavity sensors, a copper adapter was designed, with one face rounded to have the same curvature as the spacers. The bobbin is then screwed to the flat face of the adapter and the rounded face is attached using thermally conductive silver paint on the top of the cavities.

The second thermal anchor is at the 4 K plate of the cryostat, where, to avoid heat conduction from the outside, the wires are wrapped around and glued to another copper bobbin affixed to the cold plate. The sensor leads are then connected at the 4 K plate to the cryostat wiring, which feeds the signals outside of the vacuum chamber.

3.4.3 Cool down results

After the addition of the heat links, the cavities took less than two days to reach base temperature, a result not too dissimilar to (and actually better than) the estimates in section 3.4.1. Fig. 3.15 shows the evolution of the temperature of the cavities and their cradles during this time. A sharp temperature drop (which was also theoretically predicted) was observed after around 40 hours. This is consistent with the increase in the thermal conductivity of the heat links below 100 K (the exact temperature of the heat links is not known, but it was approximated in the plot in Fig. 3.15 with the temperature of the short cavity cradle). In the end, the temperature of the cavities during data collection was 6.5 K.

3.5 Optical properties

3.5.1 Cavity parameters

Table 6.1 provides a summary of the important optical cavity parameters. The finesse, reflectivity, transmission, scattering and absorption coefficients were measured by the manufacturer at

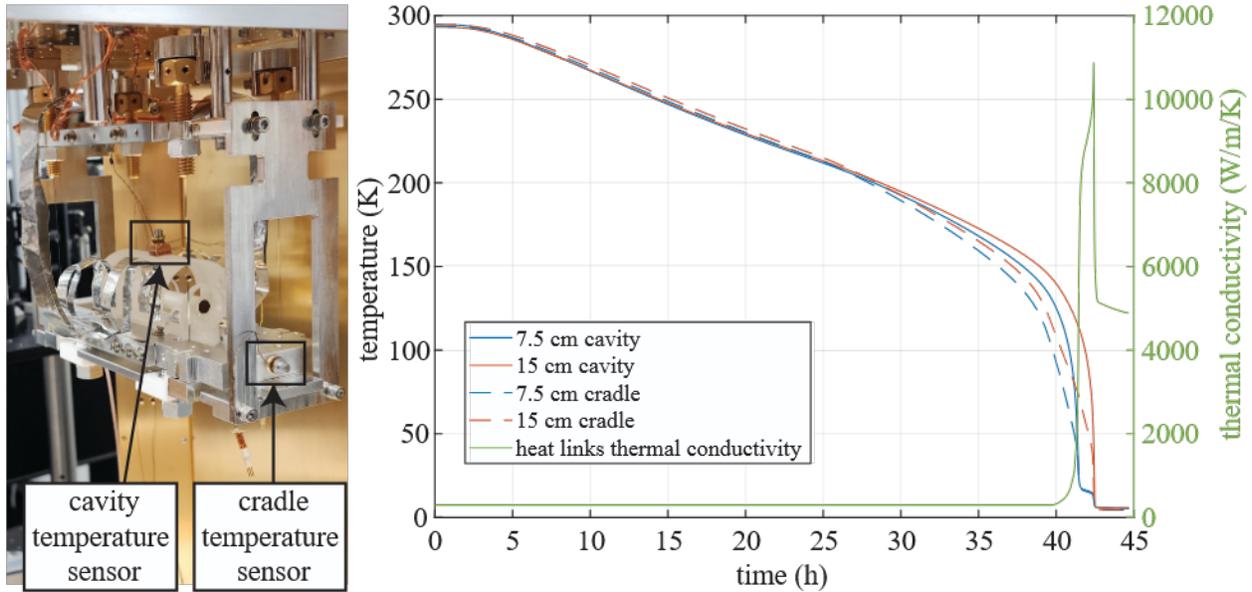


Figure 3.15: Left: position of temperature sensors on the cavities and cradles. Right: cooling curves, plotted together with the estimated thermal conductivity of the heat links

room temperature. The free spectral range, linewidth and cavity pole were calculated based on these specifications and the cavity lengths. The final row provides a value for the cavity pole, measured in situ, after the cavities reached 6.5 K. These values were obtained by measuring the frequency response of the cavities and their optical setups and fitting the data to a low-pass filter function (Fig. 3.16). The detailed procedure is provided in Appendix A.

3.5.2 Crystalline coatings birefringence

The GaAs/AlGaAs coatings of the sapphire cavity mirrors display a high degree of birefringence compared to dielectric coatings. This birefringence is attributed to internal strains in the material induced during the epitaxial growth process or during fabrication [77]. A splitting of the resonant modes was observed and measured in both sapphire cavities by scanning the laser frequency over the resonances. The splitting, which was determined to be around 340 kHz in the 15 cm cavity and around 670 kHz in the 7.5 cm cavity, is much larger than the cavity linewidth

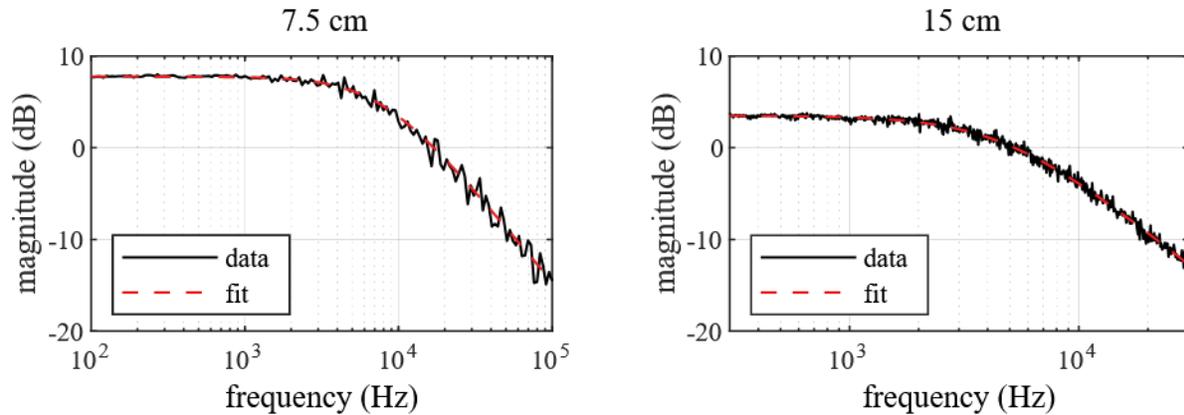


Figure 3.16: Transfer functions of the sapphire cavities, demonstrating their behavior as low-pass filters with corner frequencies at 4.7 kHz (the 15 cm cavity) and 7.4 kHz (the 7.5 cm cavity)

(6.7 kHz and 13.3 kHz FWHM, respectively), allowing for the two polarization eigenmodes to be resolved.

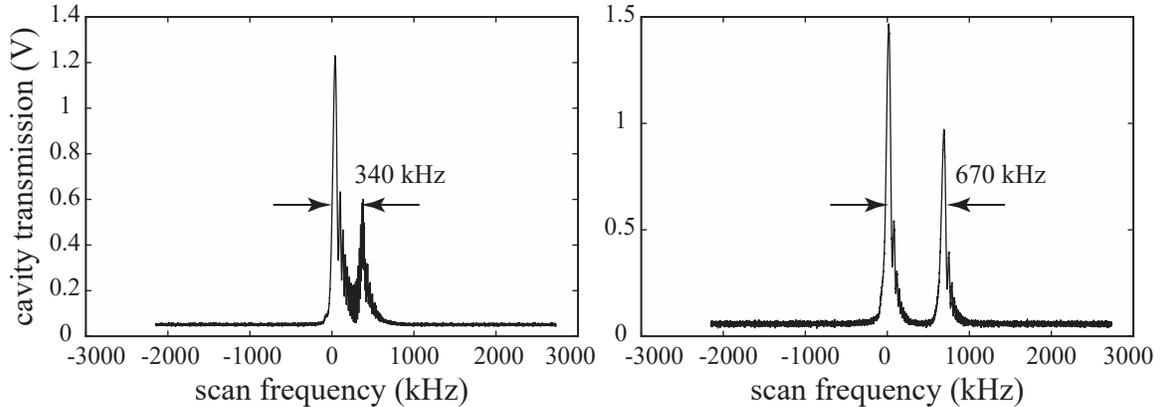


Figure 3.17: Measured birefringence splitting for the 15 cm cavity (left) and 7.5 cm cavity (right)

It is to be noted that excess thermal noise associated with the birefringence of the coatings has been previously observed [83]. As this noise is anticorrelated for the two polarizations, a modified Pound-Drever-Hall locking setup simultaneously interrogating both polarizations has been shown to be efficient in its reduction [83]. However, since so far no evidence of coating thermal noise

	Long cavity	Short cavity
Finesse	148,000	147,000
Transmission	14 ppm	14 ppm
Scattering + absorption	7.3 ppm	7.3 ppm
Reflectivity	99.998	99.998
Free spectral range	1 GHz	2 GHz
Cavity linewidth (FWHM)	6.8 kHz	13.6 kHz
Cavity pole (calculated)	3.4 kHz	6.8 kHz
Cavity pole (measured)	4.7 kHz	7.4 kHz

Table 3.3: Optical properties of the cryogenic cavities

being dominant for the present experiment at this stage has been observed, we decided to simply eliminate one of the polarizations in the incident light as described in section 4.3.

CHAPTER 4

OPTICAL LAYOUT AND DETECTION

Having presented the experiment at a conceptual level and underlined the theoretical considerations underpinning it, we shall now direct our focus to the details of its practical implementation. This chapter starts with an overview of the optical and electronic setup involved in the generation and measurement of the beat note between the frequencies of the two cryogenic cavities. It then explores the laser, the room temperature cavity, the optics and electronics involved in coupling the light into the cavities, locking the laser frequency and generating the beat note, and the detection chain. In the process, it provides an account of the evolution of the setup as it went through several stages of refinement before it reached the sensitivity required to obtain the results presented in this work, outlining the factors motivating these changes and the improvements they contributed to the beat note frequency noise.

4.1 Overview of the experimental setup

Light from a fiber laser is first stabilized to a room temperature reference ULE cavity using the Pound-Drever-Hall technique [54]. The laser frequency is locked at low frequencies (up to 1 kHz) via the laser piezo-electric transducer (PZT). An acousto-optic modulator (AOM) placed after the laser is used to stabilize the frequency of the beam from 1 kHz to 100 kHz. Once locked to the ULE cavity, the beam transmitted through the cavity is picked up and phase modulated using an electro-optic modulator (EOM), which adds symmetrical sidebands 20 MHz above and below the carrier frequency for Pound-Drever-Hall locking. The beam is then split into two branches, each branch being coupled into one of the cryogenic cavities. Acousto-optic modulators placed on each branch shift the frequencies of the two beams onto the resonant frequencies of their respective cavities.

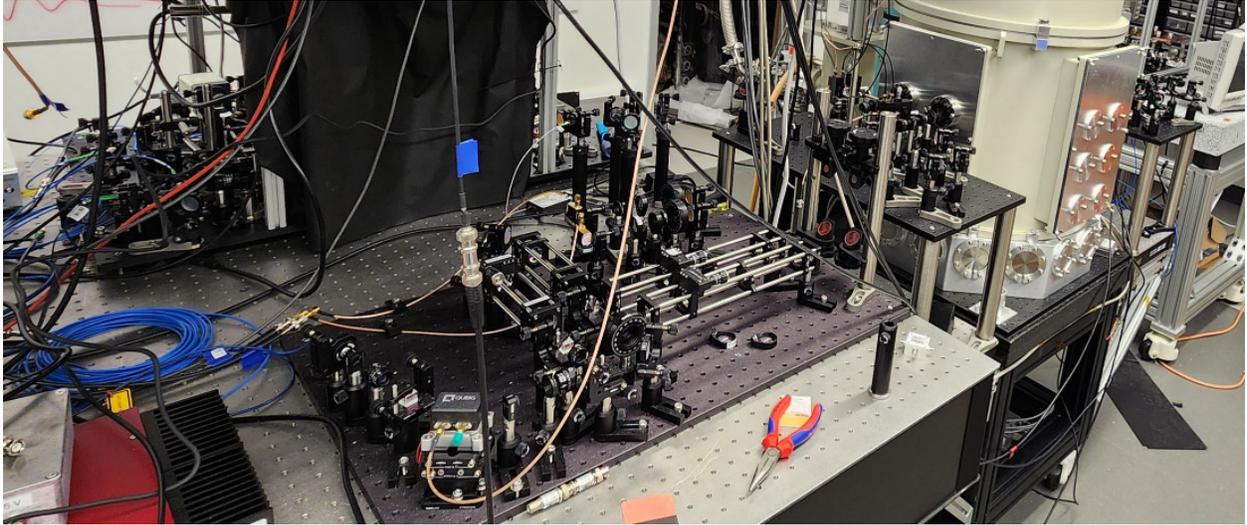


Figure 4.1: Picture of the cryostat and of the optical setup

Pound-Drever-Hall locking is again used to keep the frequencies of the beams in resonance with the cold cavities. This is a low-bandwidth lock (with a bandwidth below 1 kHz), the purpose of which is to compensate for slow drifts. The light transmitted through both cavities is collected on a photodetector, where the two beams beat against each other. The photodetector signal is amplified, sent through a delay line and mixed with its un-delayed version to cancel the carrier (determined by the frequency offset between the resonances of the two sapphire cavities). The resulting signal, which contains the beat note frequency noise information (as explained in section 4.5.2), is read by a data acquisition card (DAQ). A schematic of the entire setup in its final configuration is illustrated in Fig. 4.11.

For the measurement scheme outlined above, the locks to the sapphire cavities operate only at frequencies below the science band (the "passive" approach). In this case, the incident light does not carry information on the length fluctuations of the cavities (and hence on the presence of a dark matter signal) at frequencies of interest. However, as shown in section 2.1.3, the transmitted light would be frequency modulated at the dark matter frequency, which would show up in the beat note. An alternative measurement scheme was tried, where the laser was locked using the AOMs

to the two sapphire cavities over the entire science band (the "active" approach). In this case, the beat note can be generated both with light picked off from the beams going into the cavities or with light picked off from the cavity transmissions, as they are both sensitive to potential dark matter-induced oscillations. We ultimately decided against this approach for several reasons:

- Excess noise caused by the servo was observed in the measurement band when the PDH locks were extended into the science band. The passive approach thus rendered a better noise floor.
- The high bandwidth locking scheme relies on the lock perfectly following the resonant frequencies of the cavities. An imperfect lock would obfuscate the interpretation of the results.
- The beat note measured in transmission was observed to have a lower noise floor than that between the incident light pick-offs (Fig. 4.2). This agrees with the argument presented in section 2.1.3 that optical cavities act in transmission like low-pass filters, cleaning the laser frequency noise above the cavity poles. Since we decided that measuring the transmission beat note was favorable, the high bandwidth lock was not needed.

The optical setup was developed in a modular fashion, with the optics for the cryogenic cavities (referred to as the "two-cavity breadboard"), the room temperature stabilization setup and the detection optics inhabiting separate breadboards. This provided flexibility in making changes to the setup as the experiment went through several development phases. My personal contribution was in setting up the sapphire cavities breadboard and the detection optics and electronics and in the integration of the room temperature stabilization breadboard into the experiment.

4.2 Laser system

The laser used in the experiment is the Koheras ADJUSTIK E15 from NKT Photonics. It is an erbium-doped fiber laser (EDFL) which can generate a maximum power of 40 mW at 1550

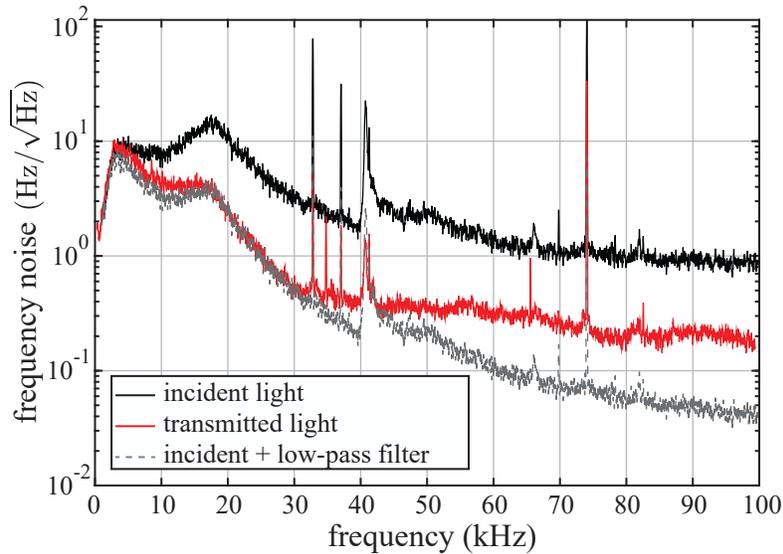


Figure 4.2: Comparison of the frequency noise in the beat note between the two cavity frequencies measured between pick-offs from the incident light and transmitted light. The dashed curve shows a theoretical prediction of the beat note in transmission obtained by applying a digital low-pass filter to the incident data. For simplicity, the pole of the filter was chosen at the 15 cm cavity pole (4.7 kHz). The prediction and measurement show good agreement up to 30 kHz. Despite the extra noise above 30 kHz (potentially caused by optics downstream of the pick-offs, or by the differential filtering by the two cavities, which have different poles), the transmission measurement scheme showcases a notable advantage

nm, a wavelength that is convenient because its use in telecommunications makes optics widely available, and also because both sapphire and silicon (materials that were considered for the cryogenic cavities) are transparent at this wavelength. This laser was selected for its narrow linewidth (0.1 kHz when free-running) and low frequency noise (13 Hz/ $\sqrt{\text{Hz}}$ at 10 kHz). It has a wide thermal tuning range of 1000 pm. A piezo-electric transducer (PZT) allows for faster tuning over a 22 pm range.

The piezo is driven with a triangular waveform to scan the frequency around the resonances of the cavities. Additionally, it can receive a small feedback voltage through its external modulation port to lock the laser to a cavity. The modulation port accepts a differential 2x 5 V peak-to-peak signal with a 2.5 V common mode. Therefore, an op-amp circuit, with an input range of -5 V

to +5 V, is employed to convert the single-ended output of the PID controller to the required differential signal. The input to the modulation port is internally amplified by an amplifier with variable gain. The maximum conversion factor between the input voltage and the frequency shift is 0.8 GHz/V.

In the final configuration of the experiment, the ADJUSTIK E15 is used to seed an ultra-low noise fiber amplifier (NKT Koheras BOOSTIK), which can increase the optical power up to 2 W (in practice, however, the maximum power used for this experiment was 400 mW).

4.3 Two-cavity setup

From the first measured beat note to the final data collection, the optical setup went through several development phases. In order to keep the experimental complexity low during the initial, troubleshooting-heavy, stages, the first incarnation of the experiment did not utilize the room temperature reference ULE cavity.

The cryogenic cavities optical setup contains optics necessary to shift the laser frequency onto those of the sapphire cavities, to couple the light into the cavities and to lock each laser branch to the frequency of the cavity modes. It consists of two parallel nearly identical setups, one for each of the two cavities.

The 40 mW power from the laser is sent to the two-cavity breadboard through a polarization-maintaining fiber. This modular configuration allows for easily switching between the laser and the light stabilized to the room-temperature cavity by changing connections at the input of the fiber. Upon exiting the fiber, the light is collimated and then goes through an optical isolator which prevents back reflections, before being split into two branches.

Acousto-optic modulators

Before sending the laser light into the cryogenic cavities, each branch passes through an acousto-optic modulator, which shifts the laser frequency to resonance with the respective sapphire cavity and serves as a frequency actuator for locking the laser to the cavity.

The acousto-optic modulators are operated in a double-pass configuration [84], where a cat's eye reflector (a lens followed by a mirror placed at its focal plane) ensures that the first-order diffracted beam retraces the trajectory of the original beam on its way back. This eliminates misalignment during locking due to the dependence of the deflection angle on the modulation frequency and doubles the frequency shift (albeit at the cost of decreased efficiency). Telescopes placed in front of each AOM reduce the beam waist in order to obtain optimal diffraction efficiency. A quarter-wave plate located after the AOM rotates the polarization of the light by 90° so that the diffracted beam can be separated from the initial beam with a polarizing beam-splitter (PBS).

The acousto-optic modulators (AMF-150-50-1550 from Brimrose, center frequency: 150 MHz) are used to shift the laser frequency in opposite directions by 184.4 MHz on each pass, aligning the frequencies with modes of the two cavities 737.75 MHz apart. To cover this separation, they had to be operated beyond their 50 MHz bandwidth, which caused a substantial efficiency loss. This issue was resolved in future iterations by using new AOMs with different center frequencies (section 4.4.4). The AOM drives were provided by the two channels of a high power digital RF synthesizer (Moglabs ARF421), which is based on two AD9910 direct digital synthesizer (DDS) sources.

Mode-matching optics

The light is coupled into the TEM₀₀ mode of the cavities. This is achieved using a pair of plano-convex lenses that adjust the beam waist to spatially match the mode of the cavity and two

steering mirrors, which allow aligning the beam to the cavity axis. A typical transmission efficiency of around 20-25% when the laser is locked was realized, with 0.75 mW of power going into each cavity.

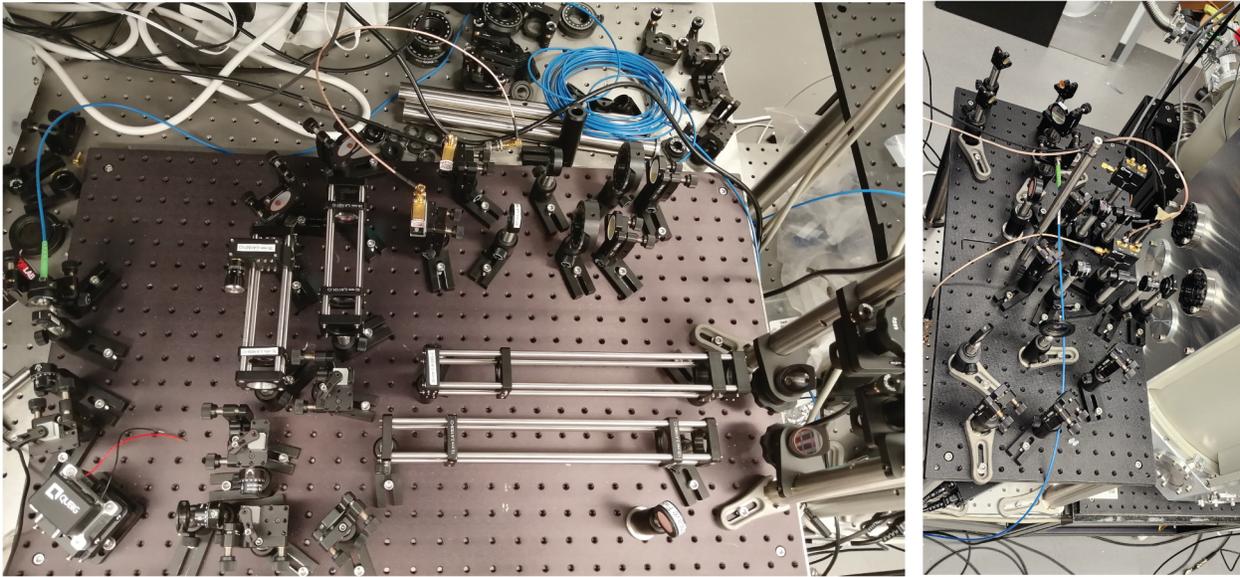


Figure 4.3: Pictures of the cold cavities optical setup (spread between two breadboards). The first breadboard (left) is placed on the optical table and contains the EOM, the AOM double-pass setups and the mode-matching telescopes. The second breadboard (right) is placed on the cryostat stand and contains optics for picking up light for the incident beat note, steering the beam into the cavities and the detection of the cavity reflections

Pound-Drever-Hall locking

To implement Pound-Drever-Hall (PDH) locking, the laser frequency is first phase modulated at 20 MHz by a resonant electro-optic modulator (Qubig PM7-SWIR-1). This modulation takes place before the laser beam is divided into the two paths coupled into the two sapphire cavities. This is to benefit from common mode reduction of any residual amplitude modulation (RAM) caused by the EOM (due to the different transfer functions of the cavities common mode noise is not perfectly canceled out). The cavity reflection, which is used to generate the PDH error signal, is detected by a fast, low-noise photodetector (Qubig PD100-SWIR-1). The photodetector has an

active area of 0.5 mm and a bandwidth of at least 100 MHz.

It is standard to separate the reflected light from the incident beam by placing a polarizing beam-splitter (PBS) and a quarter-wave plate before the optical cavity to form an isolator. Because of the large birefringence of the crystalline coatings, however, we instead use optical isolators consisting of a Faraday rotator and a PBS. In this way, the light entering the cavities is linearly polarized and a half-wave plate can align it with one of the cavity optical axis, so that one polarization mode is eliminated and the shape of the PDH error signal is not affected by the double resonance.

A dedicated driver system (Qubig Q-DRIVER QDR200) applies the 20 MHz drive to the EOM. The driver is capable of locking its output frequency to the resonance of the EOM, for optimal driving efficiency. It also demodulates the photodetector output with the EOM drive to generate the Pound-Drever-Hall error signal.

The PID control loops that lock the laser to the cavities are implemented using a digital laser stabilization module (Toptica DigiLock 110). First, the laser is locked to the 15 cm cavity through its piezo transducer. This is a low bandwidth lock (referred to as the "slow lock"), limited by the 1 kHz bandwidth of the PZT, but with a wider operation range and mainly serves the purpose of addressing slower drifts of the laser with respect to the cavity. Then, locks to each of the sapphire cavities are applied to the respective branches of the beam through the AOMs (the "fast locks"), via feedback applied to the FM modulation ports of the Moglabs ARF driver. These locks were terminated below the science band by applying digital 20 kHz first order low-pass filters on the correction signal. A 1 kHz high-pass filter balanced the induced phase lag. A diagram of the experiment at this stage including the locking scheme is shown in Fig. 4.7(a).

4.4 Integration of the room temperature reference cavity

The room temperature reference cavity stabilizes and filters the laser light before sending it to the cryogenic cavities. The latter is achieved by transmitting the locked laser beam through the

cavity before routing it to the two-cavity breadboard. This approach was motivated by an observed benefit of using the cavity transmission in terms of lowering the frequency noise.

4.4.1 ULE Cavity

For the room temperature stabilization of the laser we use a commercial ultra-stable, high finesse cavity fabricated by Stable Laser Systems. The cavity, which has a plane mirror at one end and a spherical mirror at the other, is made out of ultra low-expansion glass (ULE) and measures 10 cm in length. It rests on a Zerodur mounting block on small viton balls that provide stability against vibrations.

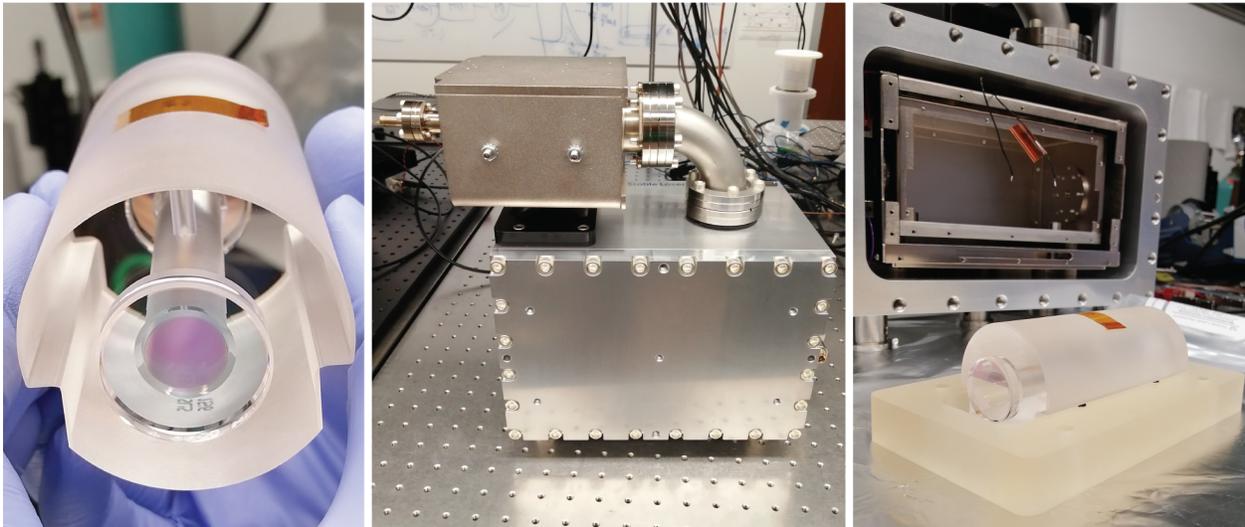


Figure 4.4: Pictures of the ULE cavity and its enclosure

The ULE cavity is placed inside an aluminum vacuum chamber and enclosed by two layers of radiation shields. An ion pump maintains vacuum inside the chamber. The temperature is monitored by thermistors located on the inner radiation shield and stabilized through a heater also placed on the first shield to the zero-crossing temperature of the spacer coefficient of thermal expansion, at 36.21 °C, at a 1 mK/day level.

The optical properties of the ULE cavity are summarized in Table 4.1. Of particular interest

given the filtering strategy we employed is its low cavity pole, 1.5 kHz, which makes it suitable for effectively filtering out noise at frequencies of interest (above 5 kHz).

Spacer length	10 cm
Spacer diameter	5 cm
Spherical mirror radius of curvature	50 cm
Finesse	>500,000
Cavity linewidth (FWHM)¹	3 kHz
Cavity pole¹	1.5 kHz
Free spectral range	1.5 GHz

Table 4.1: Summary of ULE cavity characteristics

4.4.2 Mode alignment

To have the laser in resonance with all three cavities simultaneously one needs to find three modes spaced closely enough in frequency that an AOM drive can bridge the corresponding offsets. We found it useful to measure the wavelengths of the laser locked to the cavities with a wavemeter (Bristol Instruments 671A) and make a map of the available modes to pick a convenient triplet. The distribution of modes is displayed in Fig. 4.5, with the measured wavelengths of the modes of the two cavities and the nearest mode of the ULE cavity labelled. The dotted purple line marks the wavelength of the laser, situated midway between the modes of the two sapphire cavities.

The free spectral ranges of the 15 cm, 7.5 cm and ULE cavities, respectively, are close to 1 GHz, 2 GHz and 1.5 GHz. That means that the mode structure of the three cavities will be nearly periodical, with a period of 6 GHz. In order to find the best cavity modes to use it thus suffices to search within one 6 GHz "unit cell", as outside of it the spacing between the modes will be similar.

¹Corresponding to a finesse of 500,000

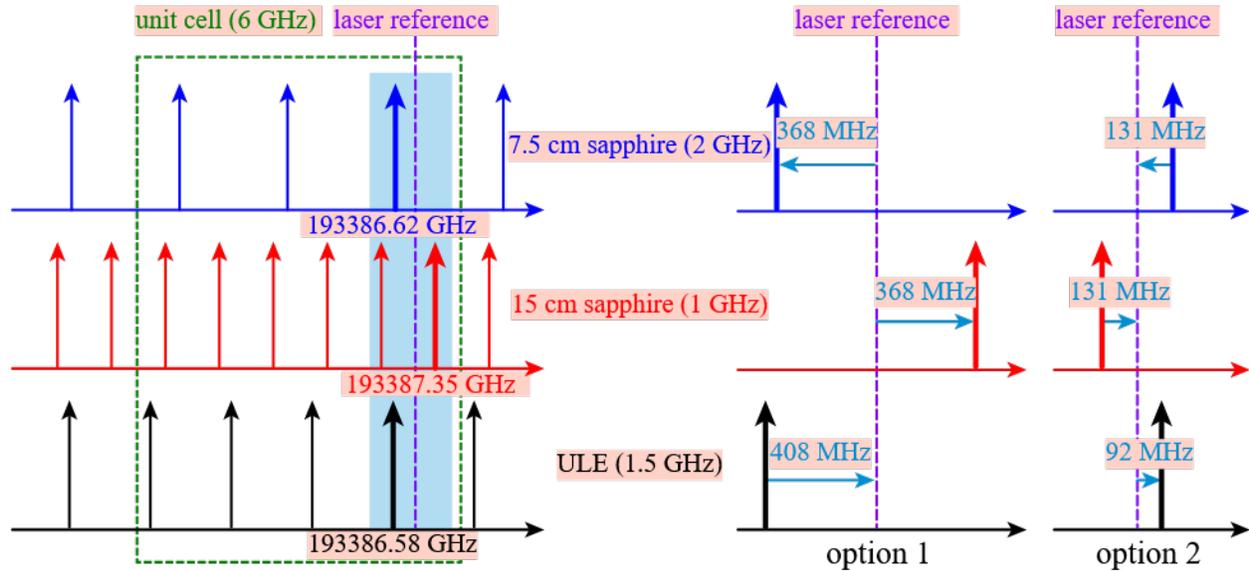


Figure 4.5: Diagram of the alignment of the resonant modes of the three cavities. Left: Alignment of the modes over multiple free spectral ranges. The green box marks the periodicity of the mode structure. Thicker arrows are used to represent the modes the laser was in resonance with when the wavelength was measured (the measured wavelengths are annotated). The light blue box highlights the modes selected for the experiment, with two options for the 15 cm cavity. Right: closer view of the two triplets of modes selected as options and the offsets between them

The nearest three modes, showcased in Fig. 4.5, option 2, were selected for the filtering scheme. However, before acousto-optic modulators at 131 MHz and 92 MHz were available, option 1 was implemented as a temporary solution (section 4.4.3).

4.4.3 Stabilization scheme

While filtering of the laser light is a straightforward solution for reducing the laser noise, it greatly reduces the optical power since the high finesse of the ULE cavity limits the fraction of the light it lets through. At best, we managed to achieve a transmission of 18% of the incoming power. Consequently, a couple of upgrades were necessary before it could be put into practice, including the acquisition of the fiber amplifier mentioned in section 4.2 and an upgrade to more efficient acousto-optic modulators.

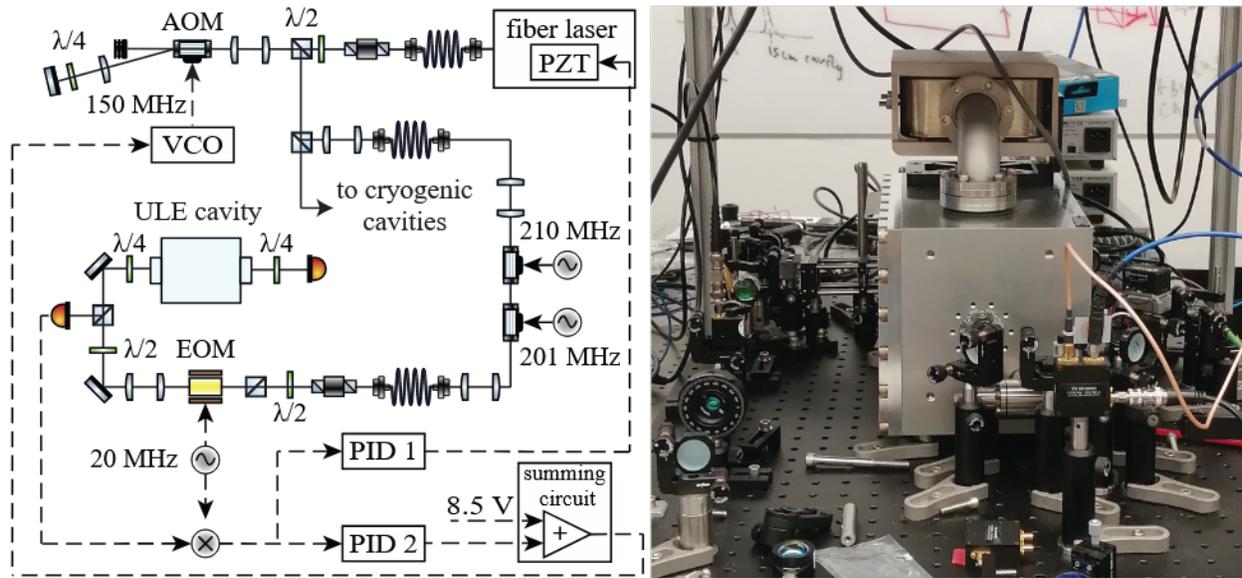


Figure 4.6: Left: diagram of the ULE cavity optical setup in its final configuration. Right: Picture of the ULE cavity inside its vacuum enclosure and the optical setup

For these reasons, an alternative scheme, where the laser was stabilized to the ULE cavity, but not filtered, was implemented as an intermediate stage. In this case, a fraction of the laser power is sent to the ULE cavity for locking, with the majority of it being directed to the two-cavity breadboard.

The optical setup for the ULE cavity is very similar to that used for the cryogenic cavities. The laser light travels through an AOM (AMF-150-50-1550, with a center frequency of 150 MHz) in double-pass, then is coupled into a single-mode polarization maintaining optical fiber to clean the spatial mode, distorted by the AOM. Subsequently, an EOM induces a 20 MHz modulation, before the light is coupled into the TEM₀₀ mode of the ULE cavity. The reflected light is collected for PDH locking.

The same type of driver as for the sapphire cavities setup (Qubig Q-DRIVER QDR200) is used to generate the error signal. The laser is stabilized through a slow piezo lock combined with a fast

lock applied to the AOM, with feedback provided by a Toptica DigiLock 110 controller. Since the ULE cavity lock does not affect the dark matter signal and hence has no bandwidth restrictions, the fast lock was extended into the science band, which allowed for a tighter and more robust lock. Ultimately, its bandwidth was limited by that of the voltage controlled oscillator (VCO) driving the AOM (100 kHz).

The input of the VCO consists of the sum of a constant 8.5 V from a stable DC power source (which drives the VCO at 150 MHz) and the control signal from the DigiLock. A summing circuit was built to add these two contributions.

A PBS and a waveplate split the beam after the locking AOM, directing an adjustable fraction of the stabilized light to the sapphire cavities. Two more acousto-optic modulators (IPF-200-100-1550) with center frequencies at 200 MHz, in single-pass, placed along the ULE cavity path, adjust the laser frequency to match the ULE cavity resonance (in practice, the required total shift was found to be around 411 MHz).

The ULE cavity setup could be integrated with the sapphire cavity optics without any major changes to the latter. The only significant difference is the transfer of the PZT lock to the ULE cavity, leaving the laser locked to the sapphire cavities solely through their AOMs. The increased laser stability achieved by the tight lock to the ULE cavity reduced the demands on the sapphire cavity locks, mostly needed to compensate for low frequency drifts between the cavities, for which a lower gain was sufficient. The result was an improvement in the beat note frequency noise (see Fig. 4.8). However, the lower optical power reaching the cold cavities (about 220 μ W going into each cavity, as measured in front of the cryostat windows) led to the strain sensitivity being limited by the detection electronics noise.

4.4.4 Filtering scheme

The latest version of the experiment, which is shown in Fig. 4.11, builds up on the previous setup, but rather than splitting the stabilized light before the ULE cavity, the filtered light is picked up after it and fiber coupled to the two-cavity breadboard. The two 200 MHz AOMs were also removed and replaced with a single AOM after the ULE cavity.

Some notable modifications were made from the previous stage of the experiment in order to increase the power going into the cryogenic cavities:

- **Laser power:** The NKT BOOSTIK fiber amplifier was added after the laser. This allowed increasing the total optical power to 400 mW. Although in principle the amplifier was capable of providing up to 2 W, the observation of an increase in noise (possibly due to thermal effects resulting from the heating of the ULE cavity) determined us to stop at this value. No effect on the temperatures of the cold cavities was observed after the power increase.
- **Beat note frequency:** Different longitudinal modes of the three cryogenic cavities were picked, such that the nearest cavity modes are now used ("option 2" in Fig. 4.5). This eliminated the need for AOMs with center frequencies above 150 MHz, which typically have lower diffraction efficiency. It also changed the beat note frequency from 738 MHz to 262 MHz, a regime where lower RF losses and electronic noise are expected.
- **Acousto-optic modulators:** New AOMs were acquired, designed specifically so that their center frequencies matched the frequency offsets between the three optical cavities. Specifically, a 92 MHz AOM in single-pass down-shifts the frequency of the ULE cavity transmission before the beam is divided into two, after which two 65 MHz AOMs in double-pass offset the frequency in opposite directions for coupling into the two sapphire cavities. All AOMs use AMTIR as the interaction material, a glass that allows the achievement of high

diffraction efficiencies (75-80%). These changes contributed to a considerable increase in the combined efficiency of the 3 modulators (by a factor of almost 60).

- **Locking scheme:** One takeaway from prior iterations of the experiment is that the cryogenic cavities exhibit excellent relative stability. Therefore, in the spirit of reducing complexity and eliminating servo noise in the science band, the two locks on the cold cavity AOMs were ultimately replaced by a single slow lock to the 15 cm cavity via the 92 MHz AOM. This lock, which has a bandwidth below 1 kHz, mainly aims to correct for the drift between the ULE cavity and the cold cavities and it proved effective at keeping the laser in resonance with both sapphire cavities for many weeks. The 92 MHz AOM was driven by one of the channels of the Moglabs ARF421 synthesizer, which facilitated the application of the correction signal through its modulation port. The second channel and a separate AD9910 DDS, amplified by a high power amplifier (Mini-Circuits TIA-1000-1R8) drive the 65 MHz AOMs.

An immediate improvement in the beat note noise floor, by a factor of 3-5 in the science band, was observed after the implementation of the filtering scheme. The gradual improvement of the frequency noise over the three stages of the experiment is presented in Fig. 4.8.

4.5 Detection

4.5.1 Beat note optics

The light transmitted through the sapphire cavities exits the cryostat, where the two beams are spatially superposed by a 50:50 beam-splitter. Half of the power is collected by a free-space high-speed biased photodetector (Thorlabs DET08C), which is used for the beat note measurement. The other half reaches a second detector which can monitor the beat note power during data collection.

A similar arrangement is used to generate a beat note between the beams before the sapphire cavities. A 90:10 beam sampler inserted in the path of each incident beam picks up 10% of the

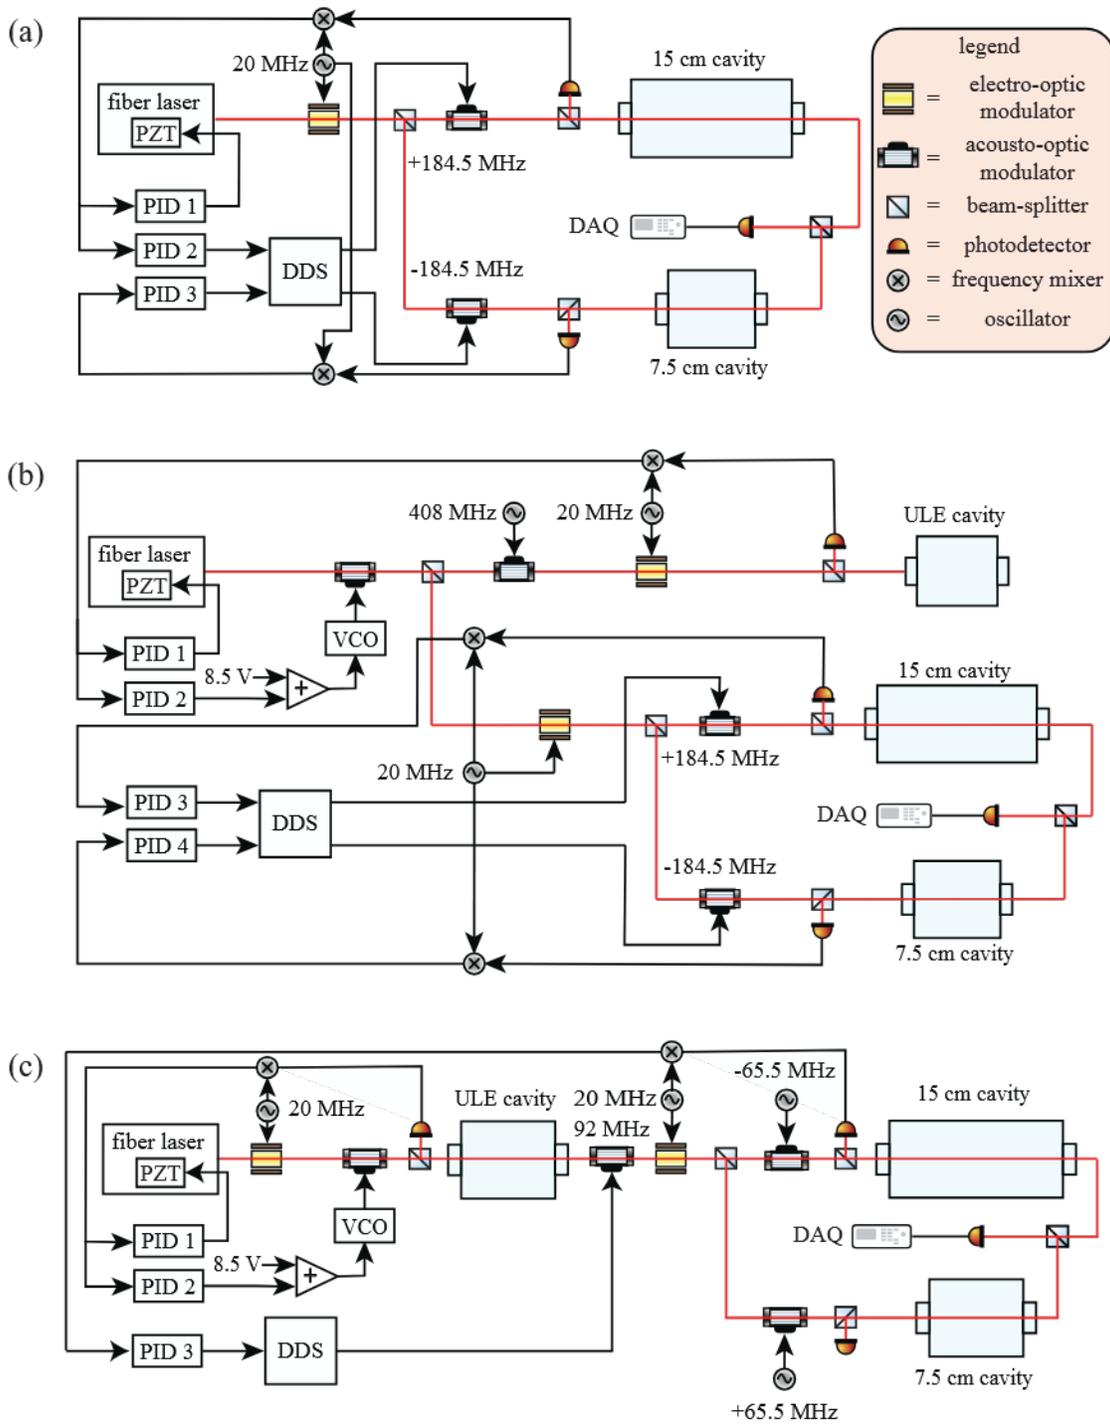


Figure 4.7: Evolution of the optical setup. (a) Without the ULE cavity; (b) With the ULE cavity used only for stabilization; (c) with the ULE cavity filtering all the laser light

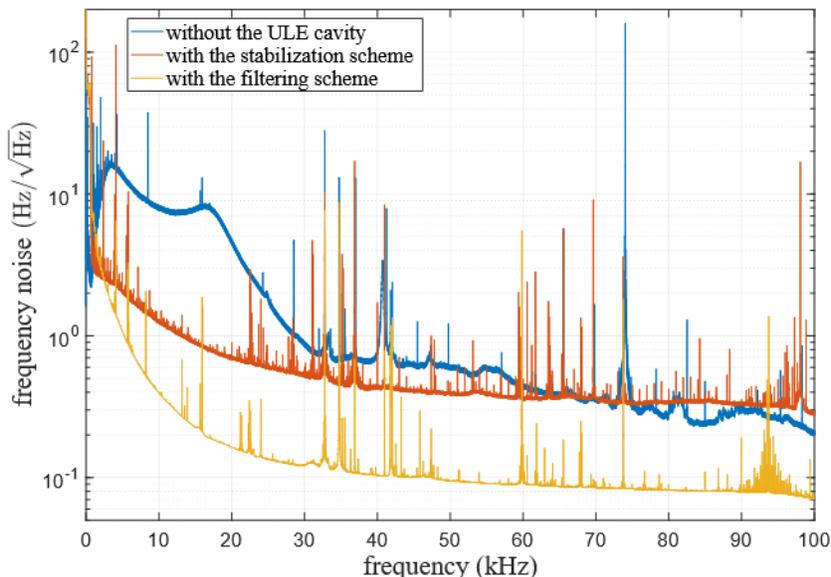


Figure 4.8: Plots of the beat note frequency noise without the ULE cavity, with the stabilization scheme and with the filtering scheme

power, which is sent to a fiber detector (FPD310-FC-NIR from Menlo Systems).

In addition to the beat note detectors, the sapphire cavities transmissions are also individually monitored by two biased detectors (Thorlabs PDA20C2) and two InGaAs cameras (Hamamatsu C14041-10U). These have been useful in the early setup stages, as well as subsequently, for diagnosis. Polarizing beam-splitters preceded by half-wave plates allow to adjust the fraction of the transmitted power going into these, so that no power is wasted during the beat note measurement.

4.5.2 Delay line

A delay line frequency discriminator is used to read the frequency noise off the beat note between the sapphire cavities. This method involves splitting the photodetector signal, after amplification, into two arms, one of which is sent through a delay line. The two branches are recombined using a mixer and the high frequency components are filtered out. When the delayed and un-delayed signals meet at the mixer in quadrature with each other the voltage fluctuations at the output ($\delta V(f)$) are directly related to the frequency noise of the original signal $\delta f(f)$ through the

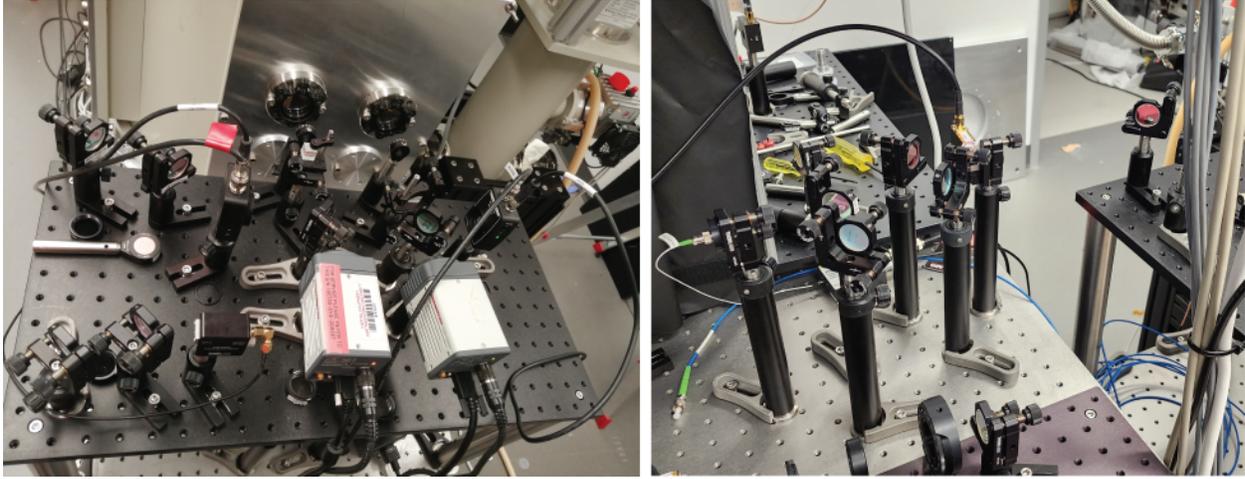


Figure 4.9: Beat note detection setup in transmission (left) and using the pick-up from the incident beams (right)

expression:

$$\delta V(f) = K_L V_R 2\pi\tau_d \frac{\sin \pi f \tau_d}{\pi f \tau_d} \delta f(f). \quad (4.1)$$

In the above, K_L is the mixer efficiency, V_R is the voltage at the RF input of the mixer and τ_d the delay. At low frequencies ($f\tau_d \ll 1$), the voltage noise is linearly proportional to the frequency noise, with a proportionality constant K_d given by:

$$K_d = K_L V_R 2\pi\tau_d. \quad (4.2)$$

More details on the theory and operation of the delay line frequency discriminator are provided in references [85] and [86] and in Appendix B.

This measurement technique presents several advantages:

- It provides a direct conversion of the beat note frequency noise, which is in turn proportional to the differential strain noise, to a voltage noise which can be directly measured with a spectrum analyzer or DAQ;

- It brings the signal down to the frequency range of the DAQ;
- The self-homodyne detection eliminates to first order the drifts in the beat note frequency caused by cavity drifts;
- Demodulation is achieved without the requirement of a second frequency source phase locked to the signal under test.

The delay line consists of a 60 m long coaxial cable (LMR-400 from Times Microwave Systems), corresponding to a 240 ns delay. This part was selected for having low loss at the beat note frequency, attenuating the signal by 5 dB.

Initially, the quadrature condition was realized by mixing down the beat note, at the time at 738 MHz, to 100 MHz. The exact frequency of the mixing signal, coming from a frequency synthesizer, was adjusted until the delayed and un-delayed signals were 90° out of phase with each other. Due to the risk of potential higher electronic noise associated with active components, the intermediate mix-down was later replaced with a mechanical phase shifter (RFPSHT0001W2 from RF Lambda).

The photodetector output undergoes several amplification stages before reaching the delay line. The noise corresponding to the first amplification stage is critical, as it establishes the signal-to-noise ratio at the subsequent components of the detection chain. A low noise amplifier (Femto HSA-X-1-40), with a 40 dB gain and a voltage noise of $310 \text{ pV}/\sqrt{\text{Hz}}$, was selected as the first amplifier in the chain. It is followed by a second amplifier (Minicircuits ZX60-P103LN+, 20 dB gain). The latter was chosen for having a high compression point ($P_{1\text{dB}} = 20 \text{ dBm}$), as it is favorable from a noise performance perspective to send as large a signal as possible into the delay line. This is because the delay line constant K_d is proportional to the input voltage, so increasing the signal strength reduces the frequency noise associated with a fixed level of electronic noise.

After the signal is mixed with its delayed version, a 240 kHz low-pass filter eliminates the

high-frequency components and the demodulated beat note goes to a pre-amplifier (SRS SR560) before being read by the DAQ (Advantech PCIE-1802), which samples it at a rate of 216 kS/s.

Since the frequency range of interest, below 100 kHz, is well within the range of linear operation of the delay line discriminator, the frequency noise can be directly extracted from the demodulated beat note once the delay line constant K_d is determined. The calibration procedure (explained in more detail in Appendix B) entails adding a known sideband either directly to the beat note (by applying a dither to the AOM corresponding to one of the sapphire cavities) or to a signal from a signal generator with the same magnitude as the beat note. The frequency noise corresponding to the sideband δf is found from the sideband-to-carrier ratio before the delay line. The magnitude of the sideband is measured after the delay line and converted to voltage. The delay line constant $K_d = 69.657 \mu\text{V}/\text{Hz}$ is then found as the ratio of the two. As a consistency check, by varying the sideband frequency and magnitude, we verified that the value of K_d is independent of these factors, as theoretically expected. Consistency between the results obtained with two delay lines of two different lengths (30 m and 50 m) was also verified.

Similar to the optical setup, the detection chain was developed in stages, with the delay line length and amplifiers adjusted along the way to accommodate for changes in the beat note power and frequency. These transformations are detailed in Fig. 4.10, with the right-most diagram representing the final configuration.

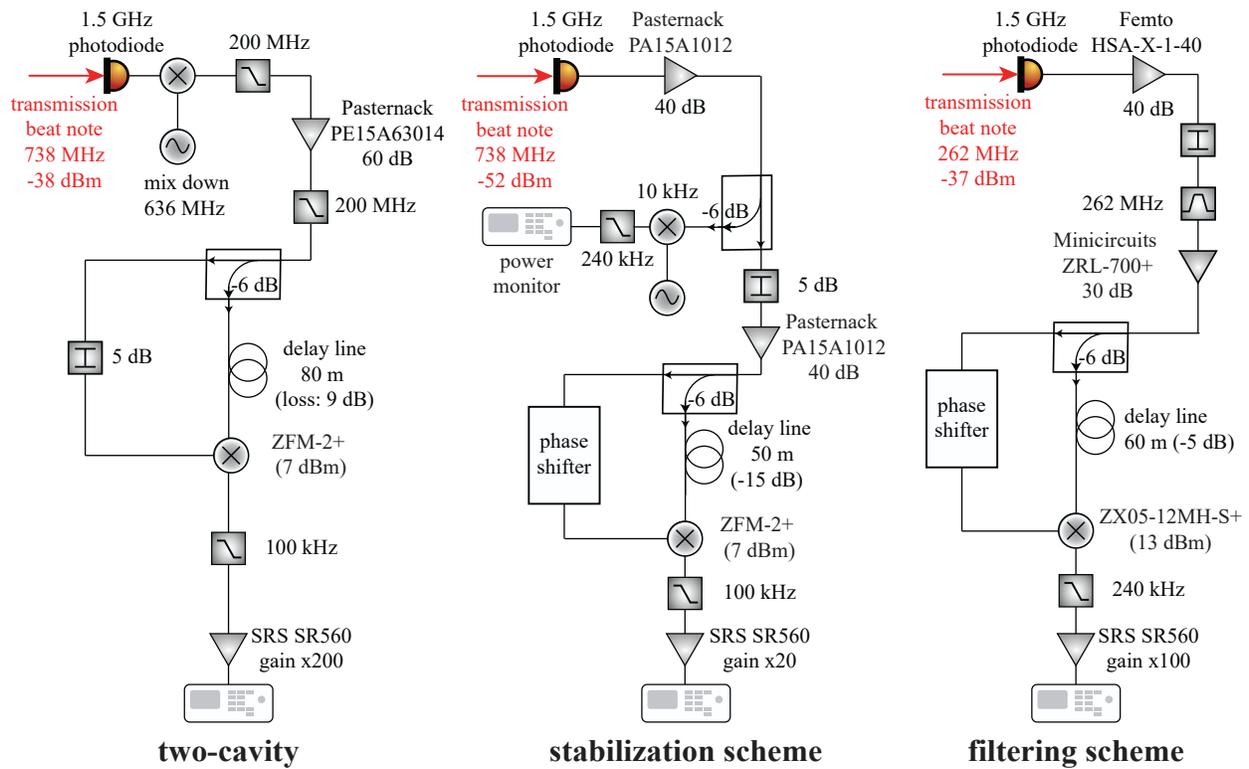


Figure 4.10: Schematic of the evolution of the delay line frequency discriminator setup during each of the three main stages of the experiment

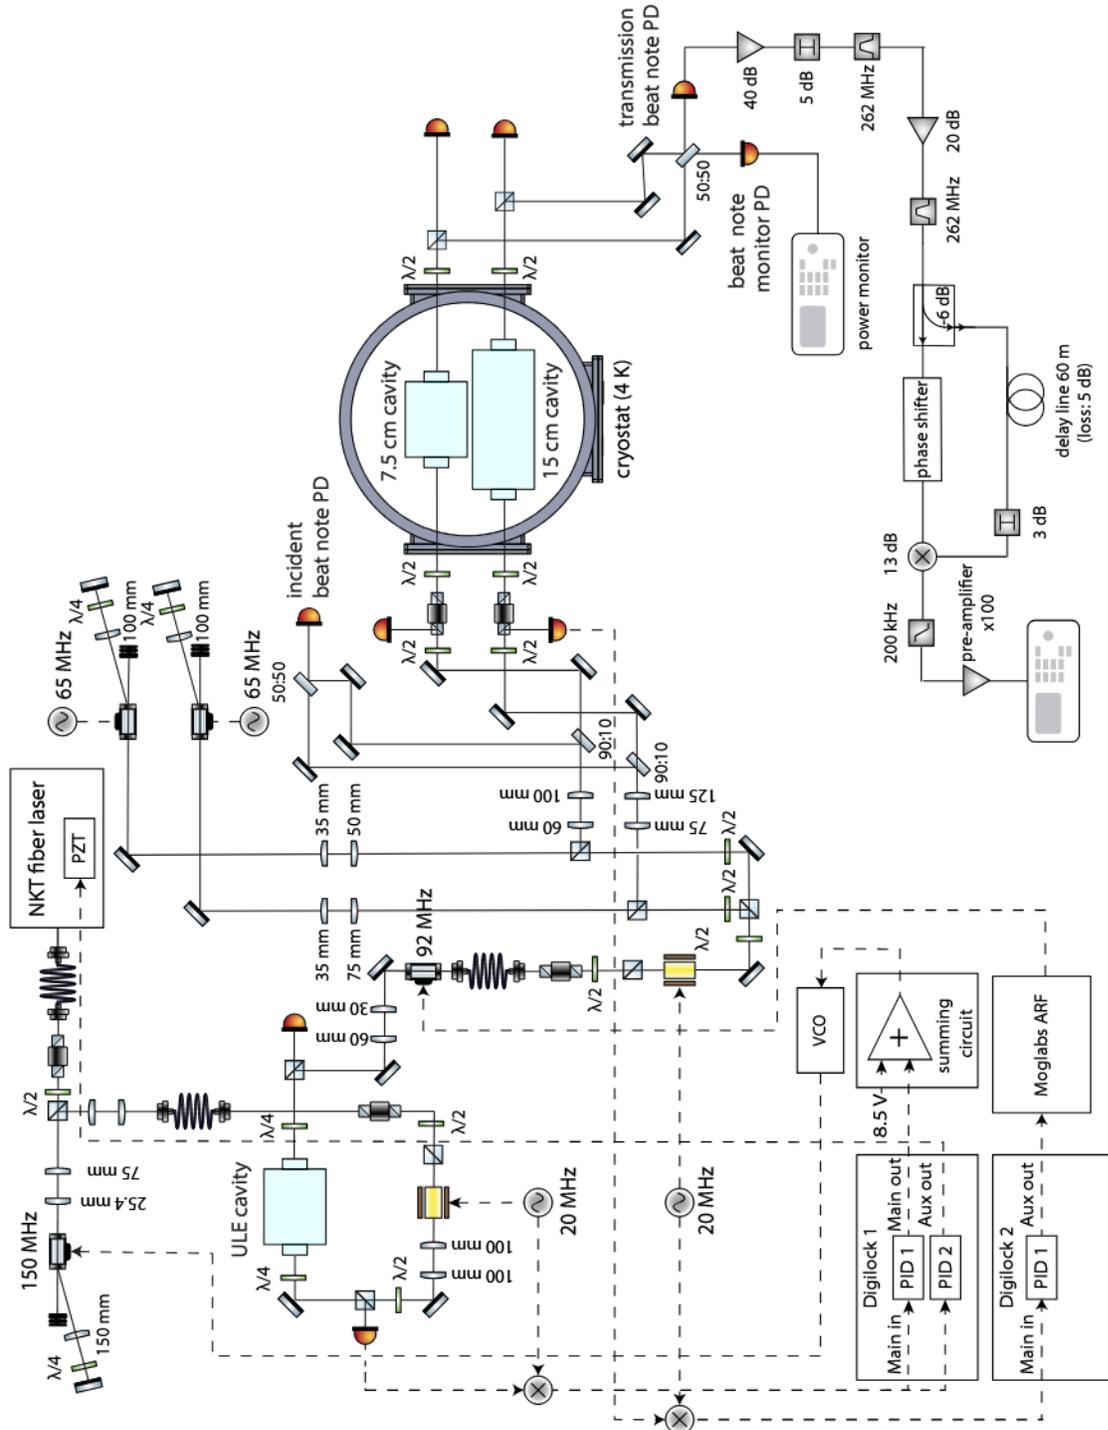


Figure 4.11: Diagram of the optical setup and detection chain

CHAPTER 5

ULTRA-LOW VIBRATION CRYOSTAT

Our measurements were carried out in a unique 4 K cryogenic apparatus that features exquisite vibration isolation, achieved through the implementation of techniques borrowed from gravitational-wave detectors. Conducting the experiment at 4 K leads to a nearly tenfold reduction in thermal noise compared to its room temperature counterpart, enabling sensitivities to displacements as small as 10^{-19} m/ $\sqrt{\text{Hz}}$. However, this is contingent upon technical noise sources, including vibrations, not exceeding this level.

Achieving low vibrations in a cryogenic apparatus is challenging. Wet systems tend to be more quiet, but their regular refills hinder extended continuous periods of data acquisition. Closed-cycle cryocoolers generate vibrations at their main operating frequency (1-2 Hz) and harmonics. These can excite natural resonances of mechanical components at hundreds of Hz and above. Furthermore, vibration isolating structures often restrict heat extraction.

For these reasons, our experiment uses a custom cryostat, engineered with the goal of achieving a high level of vibration isolation, without compromising its thermal performance. Strategies targeting the suppression of vibrations include decoupling the cold plate from vibration sources, as well as a multi-stage vibration isolation chain for both horizontal and vertical vibrations. The cryostat, together with its vibration isolation system, was designed and built by Four Nine Design (Billings, MT), steered by extensive consultations with us to address our scientific requirements.

5.1 Cryostat

The cryogenic apparatus (Fig. 5.1 and Fig. 5.2) is cooled by a two-stage pulse tube cryocooler (SHI RP-082B2), with a heat extraction capacity of 40 W at 45 K (the first stage) and 1 W at

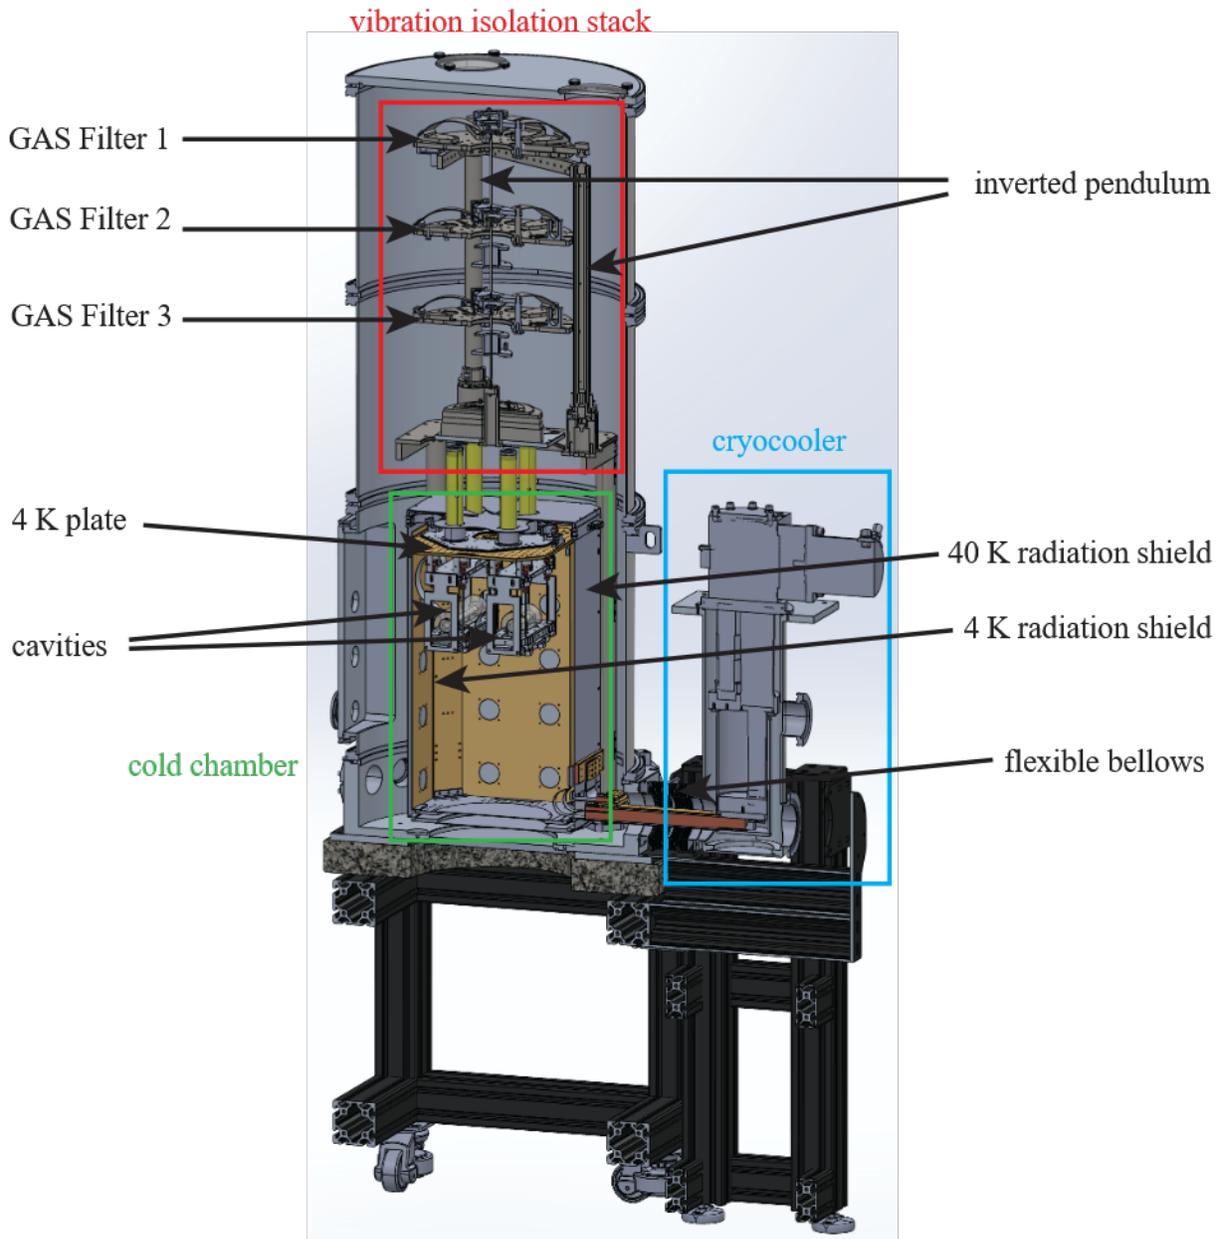


Figure 5.1: CAD drawing of a section through the cryogenic apparatus, illustrating the main components of the system (the cold chamber with the optical cavities, the vibration isolation stack and the pulse tube refrigerator)

4.2 K (the second stage). While pulse tube refrigerators show lower vibrations levels compared to alternatives such as, for instance, Gifford-McMahon coolers, careful design choices are needed to minimize the coupling of noise into the system.

The vacuum space housing the cold chamber and the vibration isolation stack rests on a stand on the floor, on top of a heavy granite plate. The cryocooler is held by a ceiling-mounted support frame (see Fig. 5.2). This configuration establishes a degree of separation between the "noisy" and "quiet" elements, the connection between the two at the vacuum enclosure layer being realized only through flexible bellows. The decision to install the cryostat on the floor and the pulse tube on the ceiling - rather than the other way around - was guided by vibration measurements that determined that the ceiling is the noisier of the two.

The cold chamber is a rectangular box with inner dimensions of 30 cm x 30 cm x 50 cm, allowing sufficient space for the two optical cavities, together with their suspensions. It is confined by two layers of radiation shields: an outer aluminum shield at 40 K and an inner 4 K shield made of gold-plated copper. Both the outer shield and the bottom and side walls of the inner one are supported from the bottom of the vacuum chamber. Meanwhile, the top of the 4 K box consists of a suspended plate (the "4 K plate" or "cold plate") attached to the bottom of the vibration isolation stack, so that the science cavities, which are hanging from it, benefit from multiple stages of vibration isolation.

A copper cold finger (visible in Fig. 5.1) extending through the bellows between the cryocooler and cold chamber serves for heat extraction. In order to prevent the propagation of vibrations from the pulse tube, the sole contact between the cold finger and radiation shields is made through soft, high-conductivity aluminum ribbons (made of the same material as the thermal links used to cool the cavities), as shown in Fig. 5.3. Similar heat links also connect the side and top of the 4 K enclosure to cool the isolated cold plate (these can be observed in Fig. 5.4).

The vacuum can is equipped with electrical feedthroughs in 25-pin D-shaped connectors that

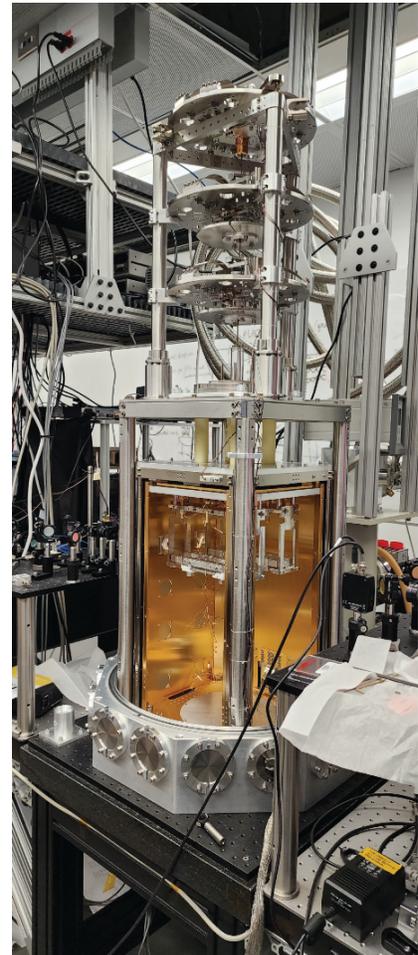
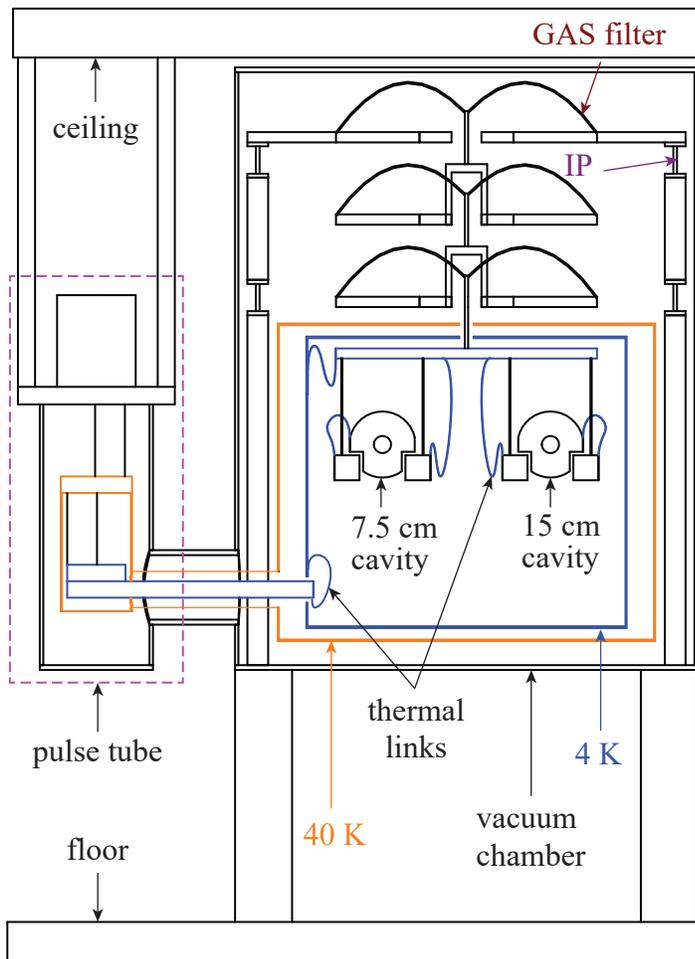


Figure 5.2: Schematic (left) and picture (right) of the cryostat

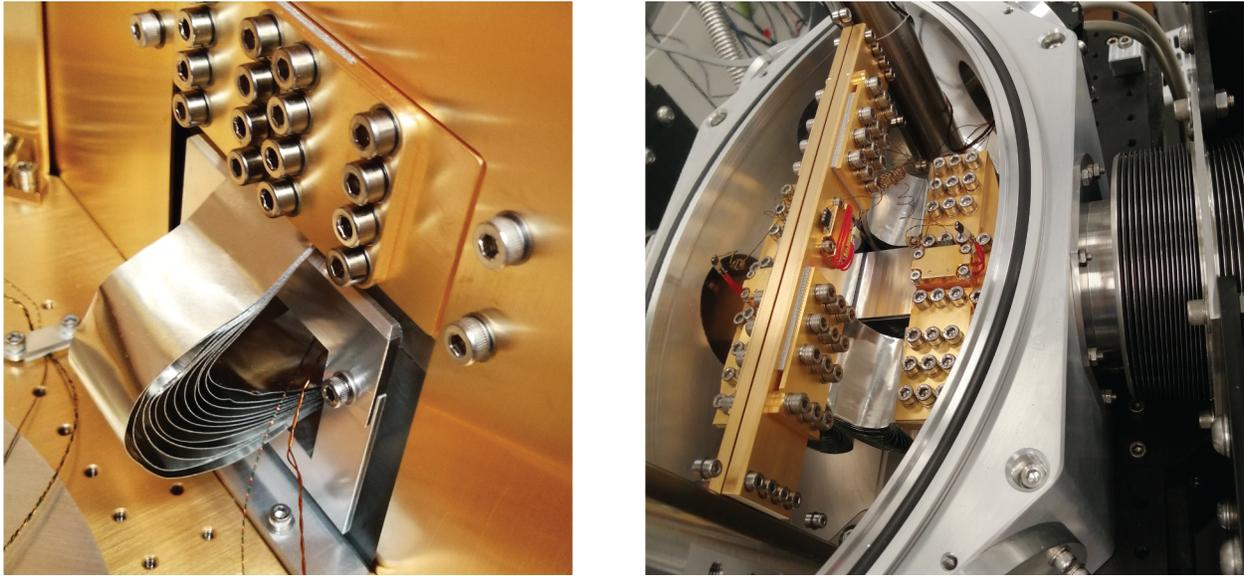


Figure 5.3: Pictures of thermal links inside and outside the cold chamber

connect the cryostat wiring. This include leads for the temperature sensors, heaters, as well as position sensors and actuators for the vibration isolation stages (further details on which are provided in section 5.2) to laboratory electronics. The wires going to the cold chamber are mechanically and thermally anchored at the vibration isolation stack support legs and at the 40 K and 4 K plates. At the first two stages, the wires terminate with Micro DB25 connectors. However, electrical contact issues in the connector on the 4 K plate prompted us to replace it with more robust, cryogenic compatible connectors (part numbers CMR/LF4 and CMR-LF7 from CMR-Direct), mounted into a custom-made mount that is bolted to the 4 K plate (Fig. 5.4). The room temperature leads consist of twisted pairs of thin (32 AWG) copper wire, while phosphor bronze has been used for the leads that enter or are anchored to the cold space in order to reduce heat flow into the system (at 4 K, the thermal conductivity of phosphor bronze is less than 1% that of copper).

In addition to the temperature sensors on the cavities, mentioned in section 3.4.2, the cryostat temperature is monitored at seven locations: at the first and second stage of the cryocooler, on the lateral sides and top plates of the radiation shields and at the room temperature region of



Figure 5.4: Left: anchoring of temperature sensors and heaters leads at 40 K and 4 K. Also visible are the thermal links cooling the 4 K plate. Right: Custom cryogenic connector at the 4 K plate

the vibration isolation payload. Heaters are used during normal operation to prevent the room temperature payload from cooling down, as well as during the warm-up phase, to speed up the process.

When the vacuum cans that enclose the vacuum space are on, access inside the cryostat is possible from three lateral directions through side plates with o-ring seals. Additionally, a removable flange at the bottom allows for access from below. The side plates and radiation shields are equipped with openings located at three different heights, providing alternative positions for installing windows for optical access. The current configuration uses the uppermost apertures on two opposite faces, while all the other ones are closed off to avoid heat transfer between the different thermal regions. All windows have anti-reflection coatings for IR light.

Commercial Conflat (CF) viewports (Thorlabs VC234C), with a 1.5" window opening, were installed on the vacuum can side plates. Custom window mounts were attached to the inner side of the 4 K shield and the outer side of the 40 K shield. These consist of aluminum plates with a circular pocket in which a 0.5" diameter window rests. A ring made of 0.02" thick indium wire forms a soft seal between the optic and the plate that remains malleable even at cryogenic

temperatures, creating a hermetic bond while also ensuring that the window is not pressed on by the thermally contracting plate. A soft teflon washer retains the window in place. The use of teflon screws, which suffer less thermal contraction than metallic ones, prevents clamping issues upon cooling. This design provides sufficient isolation between the different temperature regions on either side of the plate, can withstand thermal cycling and ensures that the window is at no risk of cracking due to excessive pressure from the thermal contraction of the mount.

The total cryostat cool down time is around two days. With the sapphire cavities as payload, the baseline temperature reached by the 4 K plate was around 4.5 K.

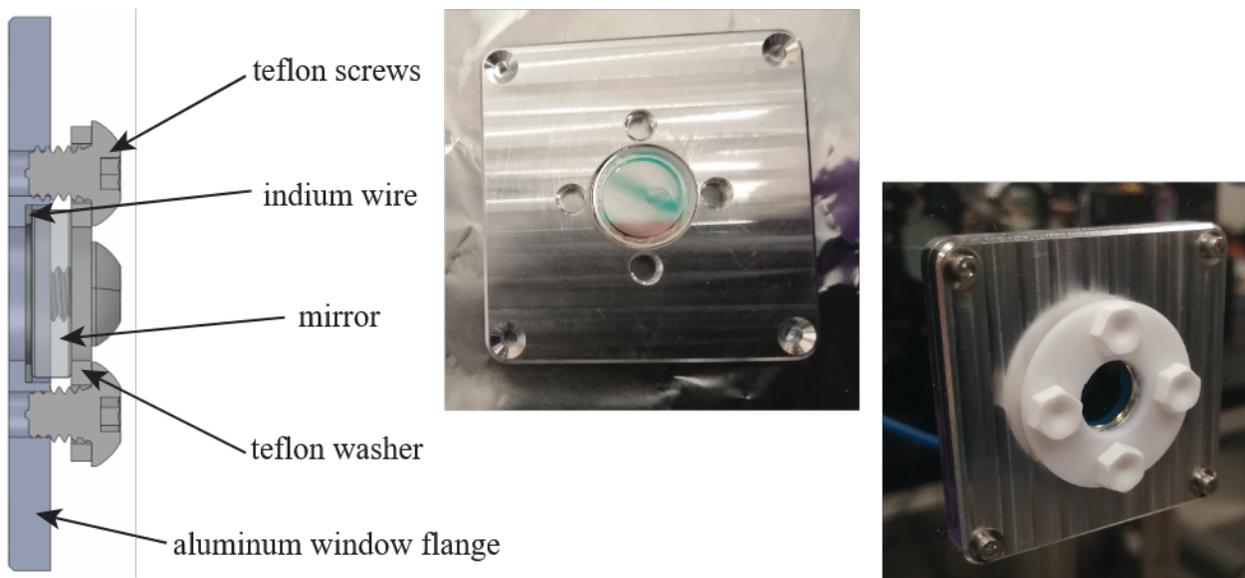


Figure 5.5: Cryogenic-compatible window mounts for the radiation shields

5.2 Vibration Isolation

Vibration isolation methods can be broadly classified into two categories: active and passive. Passive approaches employ materials and devices that absorb or dampen mechanical motion without relying on a control system. Active techniques involve actuators that dynamically adjust the

position of the system to compensate in real time for the motion measured by sensing elements, based on a feedback loop.

Passive vibration isolation (the basics of which are introduced in section 2.3.3) presents a number of advantages, particularly in terms of simplicity, robustness, and reduced need for maintenance once implemented. Moreover, the ability to cascade multiple filters, each contributing an attenuation factor $\sim 1/f^2$, allows for the achievement of very high vibration isolation. Nevertheless, as Eq. (2.60) indicates, such filters are only effective at high frequencies and their effectiveness relies on the ability to design mechanical oscillators with natural frequencies as low as possible.

The relatively high frequency of our science band (above 6 kHz for the data reported here, above around 100 Hz for the prospective comparison between rigid and non-rigid cavities) makes the use of passive vibration isolation favorable. The cryostat is equipped with an inverted pendulum for horizontal isolation and geometric anti-spring (GAS) filters for vertical isolation. In this section, we first introduce each of these devices, and then explain the way they are integrated in our apparatus. At the end, some measurements of the GAS filter performance are presented.

5.2.1 Inverted Pendulum

The inverted pendulum (IP) consists of three legs distributed in a circular pattern around the center, at an angle of 120° with respect to each other. Each leg has a flexure at the bottom that confers flexibility to the assembly. An aluminum plate resting on the three legs supports the following isolation stages. The plate is connected to each of the three legs through another thin flexure at the top of the legs.

The basic idea behind the inverted pendulum is that the flexible elements provide a restoring force to horizontal displacements of the mass supported by the legs. The structure thus acts as a spring with an effective spring constant given by the interplay between gravitational and elastic forces. For a pendulum of length L , assuming a total bending spring constant of the three legs k_θ

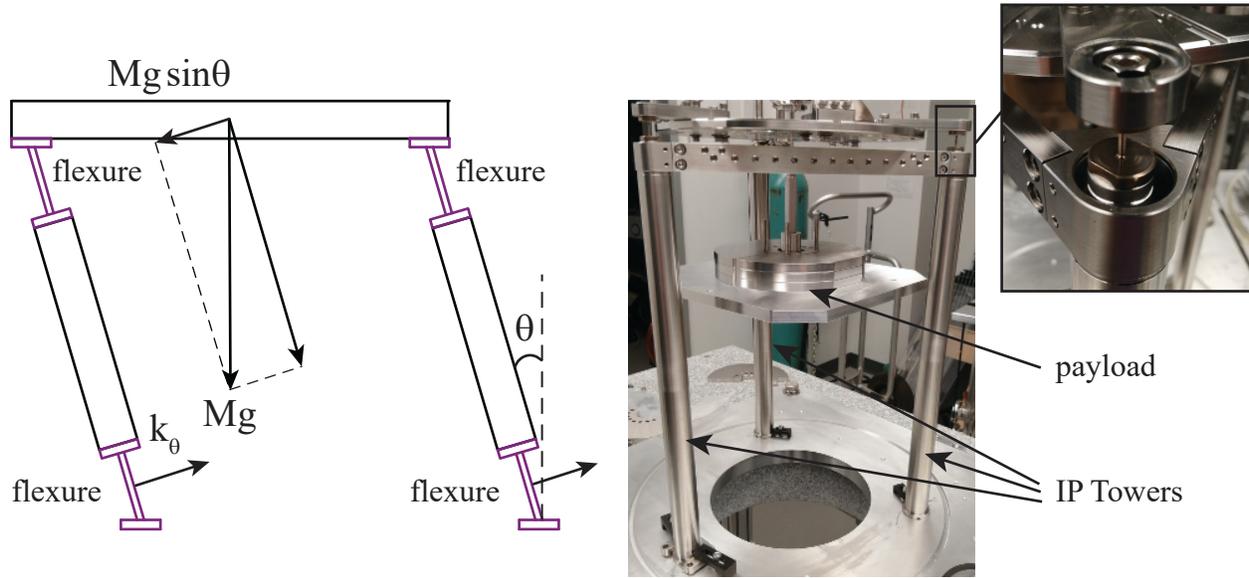


Figure 5.6: Inverted pendulum working principle diagram and pictures of a test inverted pendulum outside of the cryostat. The close-up picture shows the thin flexure at the top. A similar structure is found at the bottom of the leg (inside the inverted pendulum tower)

and a mass M of the load supported by the legs (for simplicity, the mass of the legs is assumed negligible), the effective spring constant k_{eff} is given by:

$$k_{eff} = \frac{k_{\theta}}{L^2} - \frac{Mg}{L}. \quad (5.1)$$

This spring constant can be tuned by varying the load mass and the stiffness of the flexing elements, reaching 0 for a critical mass value $M_c = k_{\theta}/(gL)$. The resonant frequency is then given by:

$$f_{IP} = \frac{1}{2\pi} \sqrt{\frac{g}{L} \left(\frac{M_c - M}{M} \right)}. \quad (5.2)$$

This partial cancellation between the gravitational and elastic forces provides a path toward the realization of suspensions with very low resonant frequencies. For instance, inverted pendula developed for gravitational wave detectors such as LIGO with resonant frequencies as low as 12 mHz were demonstrated [68, 87, 88].

The inverted pendulum in our system has its main resonant frequency around 0.5 Hz. For comparison, with an ideal pendulum a 1 m length would be necessary to achieve a similar performance.

When setting up the inverted pendulum, it is important that the legs are positioned symmetrically. Twists of the inverted pendulum could lead to the legs touching the aluminum covers surrounding them, degrading its performance. Three windows at the top of the vacuum enclosure, situated above each of the inverted pendulum towers, allow visual access to the inverted pendulum. The position of the legs was fine-tuned by leveling the cryostat legs until each of the three legs was reasonably centered inside their corresponding towers and we could verify that the pendulum was able to oscillate freely.

5.2.2 GAS Filter

As seen in chapter 3, despite the horizontal orientation of the science cavities, vertical vibrations are still significant, since they can affect the cavity length through the Poisson ratio. While horizontal vibration isolation can be achieved in reasonably simple ways (such as a pendulum), designing oscillators with low natural frequency (below 1 Hz) in the vertical direction is less trivial. For traditional springs, lowering the resonant frequency requires increasing the length, a route that quickly becomes space inefficient, or reducing the stiffness, which sacrifices yield strength. An elegant and compact solution developed in the context of gravitational wave detectors is the so-called gravitational anti-spring (GAS) filter [89].

A GAS filter consists of symmetrically distributed, radially oriented flexible cantilever blades that are pre-stressed through bending in order to act like springs in compression. The outer ends of the blades are clamped to a base plate, while the inner tips are connected together through a centering structure called a "keystone". The load (which can be the next vibration isolation stage) hangs from this keystone from one wire.

This geometry leads to the generation of nearly horizontal elastic forces concurrently with an

upward vertical restoring force. In a simplified model [68], each cantilever blade can be represented as a compressed spring of undeformed length l_{0x} , which is perfectly horizontal in the equilibrium position ($z = z_{eq}$), together with a vertical spring operating in extension (see Fig. 5.7). We denote by k_z and k_x the respective equivalent springs constants resulting from the collective effect of the three blades in the vertical and nearly horizontal directions and by M the mass of the load. Then, under a vertical displacement $z - z_{eq}$, which results in an elongation of the horizontal springs to a deformed length l and a corresponding tilt by an angle θ , the following equation of motion can be written for the vertical components of the force:

$$M\ddot{z} = -k_z(z - z_{eq}) - k_x(l - l_{0x})\sin(\theta). \quad (5.3)$$

Within the small angle approximation, the expression above can be rewritten as:

$$M\ddot{z} = -\left(k_z - \left(\frac{l_{0x}}{x_0} - 1\right)k_x\right)(z - z_{eq}), \quad (5.4)$$

where x_0 is the horizontal distance between the keystone and the support of the horizontal springs.

This is equivalent with the effect of a spring with effective spring constant

$$k_{eff} = k_z - \left(\frac{l_{0x}}{x_0} - 1\right)k_x. \quad (5.5)$$

As in the case of the inverted pendulum, the simultaneous effect of opposing forces partially canceling each other out generates a low effective spring constant, hence an oscillator with a low mechanical resonance.

The shape of the blades is optimized so that stresses are evenly distributed to avoid stress concentration which could lead to material yielding. For a particular geometry of the cantilever blades, the resonant frequency can be adjusted by radially moving the location of the clamp to the support base or by tuning the mass of the load. That means that for a certain geometry, in order to

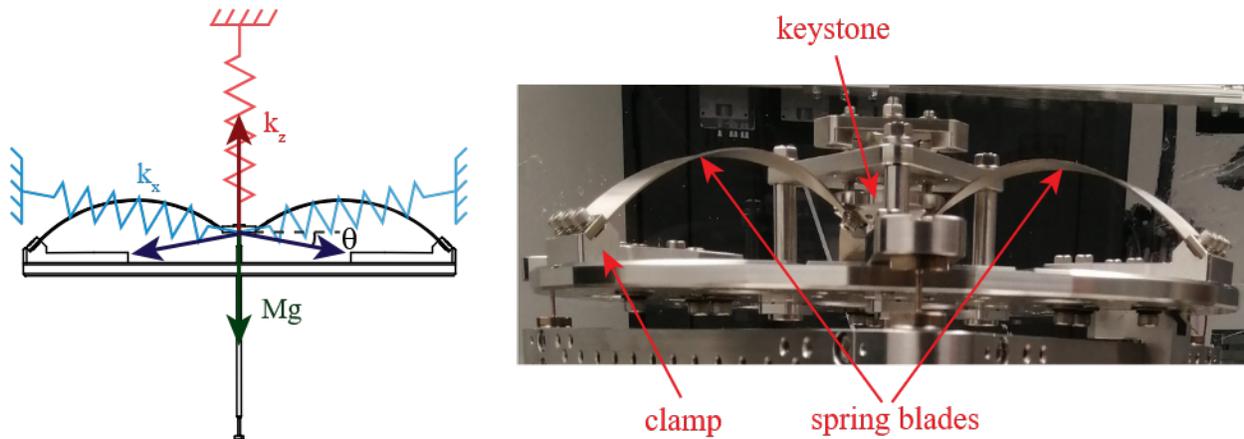


Figure 5.7: GAS Filter

maintain optimal vibration isolation, the payload mass must be set to a fixed value.

5.2.3 Vibration isolation stack

To reduce the effect of lab vibrations below that of shot and thermal noise, a room temperature vibration isolation assembly has been implemented inside the vacuum system, consisting of cascaded alternating filters in the horizontal and vertical direction, which support the cold plate.

A diagram presenting the components of the vibration isolation stack is provided in Fig. 5.8. The first stage of vibration isolation is an inverted pendulum with a resonant frequency of about 0.5 Hz. Three GAS filters with resonances below 0.5 Hz provide vertical isolation. Each of them is expected to reduce vibrations above 100 Hz by at least a factor of 1000, for a total of 9 orders of magnitude of vibration isolation. Each GAS filter hangs from the previous one by a thin rod made of maraging steel, an alloy known for its high strength and low creep. This structure effectively acts as an additional one-wire pendulum. This makes for a total of 7 vibration isolation stages, 4 horizontal and 3 vertical. To these we added the additional cryogenic pendulum supporting each cavity, described in section 3.3, which provides an extra layer of attenuation of horizontal vibrations.

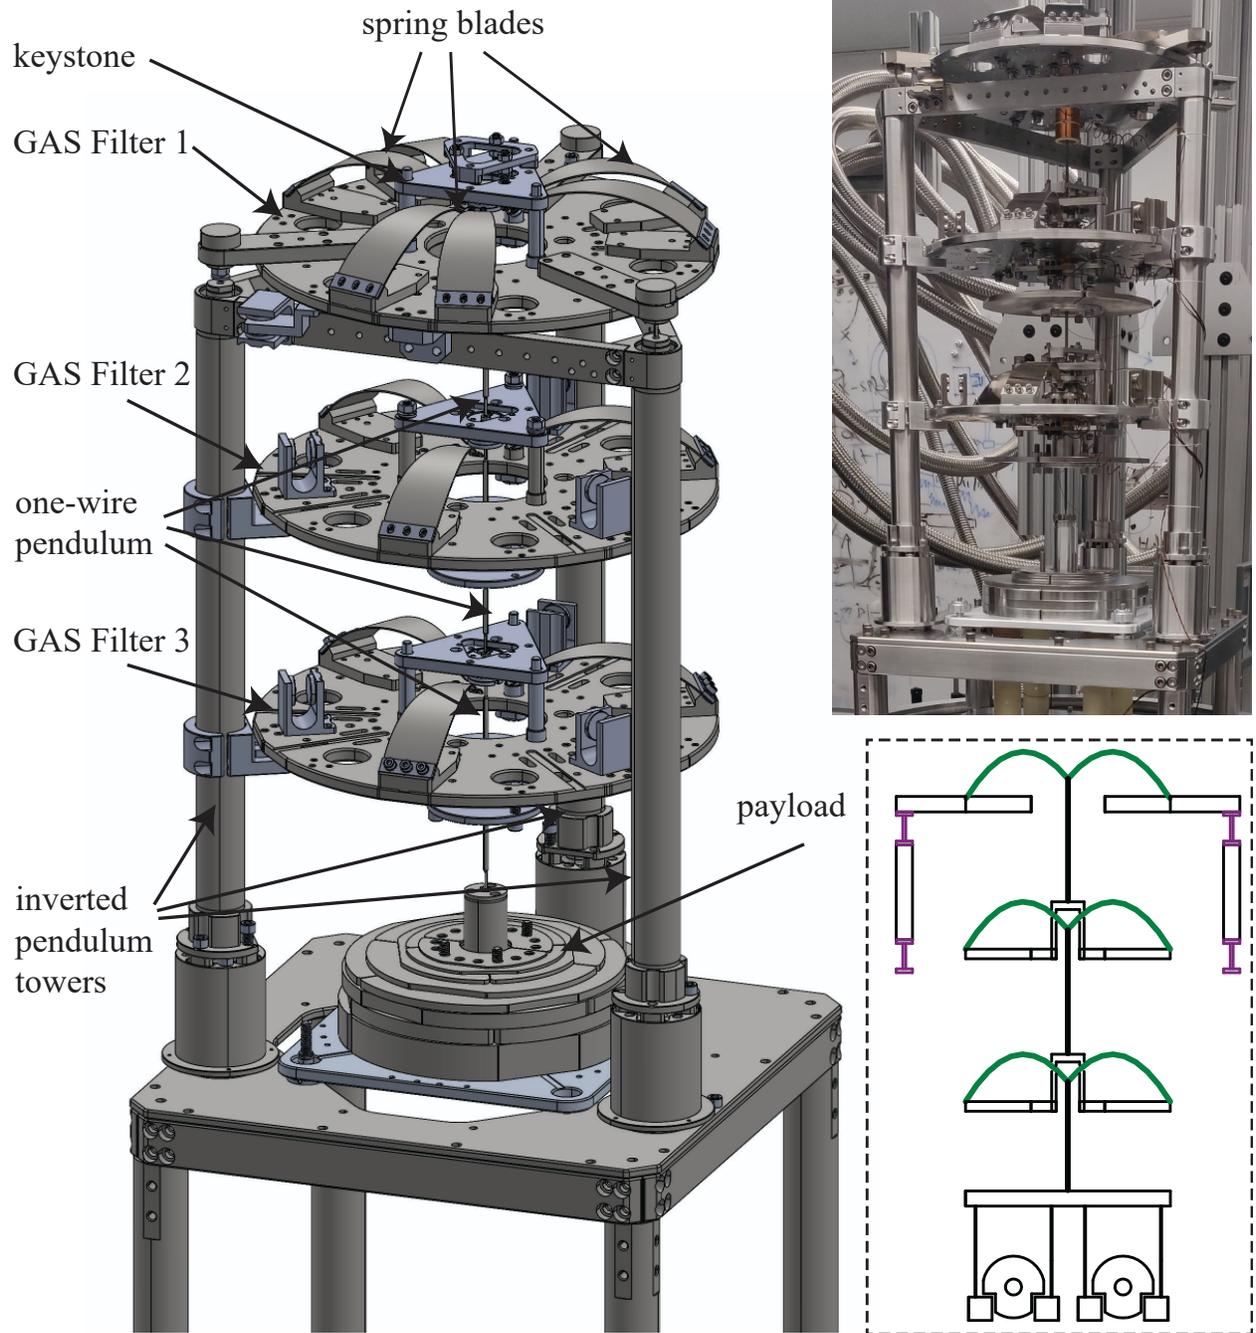


Figure 5.8: CAD drawing of the room temperature vibration isolation stack, showcasing the seven stages of vibration isolation (left), together with a photo of the vibration isolation system (top right) and a simplified diagram (bottom right)

In addition to these passive components, voice coil actuators placed on each vertical filtering stage can be used, in conjunction with LVDT (linear variable differential transformer) sensors installed to monitor the vertical position, for the active damping of the GAS filters' motion. The voice coils have an operating range of 10 mm. While at this stage of the experiment active control was found unnecessary, this feature could prove to be useful in the future, particularly if the implementation of the non-rigid cavity pushes down the science band toward lower frequencies, where the effect of vibrations is more prominent.

The main resonant frequency of the GAS filters stack was found to be 0.326 Hz, as observed from spectra of measurements of the 3 LVDT sensors during operation. For the optimal operation of the GAS filters (i.e. the lowest resonant frequency), the total mass attached to the 4 K plate needs to be 7.6 kg.

5.2.4 GAS filter performance

GAS filter transfer function

Transfer function measurements performed on a fourth nominally identical GAS filter, used for testing and prototyping purposes, confirmed the ability of these devices to reduce vibration levels above 10 Hz by at least 3 orders of magnitude (Fig. 5.9). These tests were done by inducing vibrations at the base plate of the GAS filter with a shaker and measuring the vibration level on the base plate and at the load using two accelerometers. Unfortunately, as the transmissibility tests were conducted in air, coupling of parasitic acoustic noise to the payload accelerometer and cable "shorted" the GAS filter, preventing accurate measurements of the transfer function above around 50 Hz. Moreover, around 70 Hz the drive started to excite resonances of the fixture supporting the GAS filter and shaker.

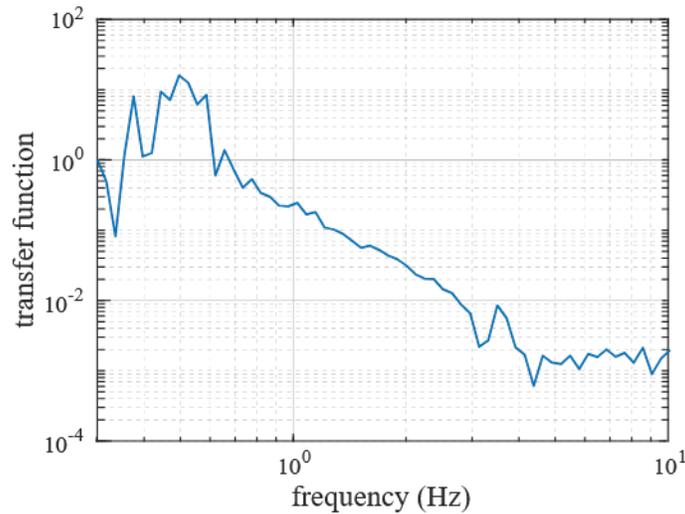


Figure 5.9: Example of transfer function of a test GAS filter

Magic Wand and GAS Filter Tuning

While an ideal oscillator provides vibration isolation that improves above its mechanical resonance as the square of frequency, for the real GAS filter saturation is observed above some critical frequency, an effect associated with the moment of inertia of the blades [68]. This effect can be overcome by changing the mass distribution in a way that optimizes the position of the center of percussion.

A practical solution to this problem is the addition of rods, referred to as "magic wands", between the base plate and keystone that act in parallel with the GAS filter so as to adjust the mass distribution and stiffness of the system [90]. The magic wand is connected to the base plate through a hinge and to the keystone through a thin flexure and has an adjustable counterweight at the end opposing the keystone (see the diagram in Fig. 5.10). The transfer function can be minimized by optimizing the mass m and position l of the counterweight, as defined in Fig. 5.10.

Magic wands for the three cryostat GAS filters (two per GAS filter) were provided by Four Nine Design. While not yet implemented for the vibration isolation system inside the vacuum

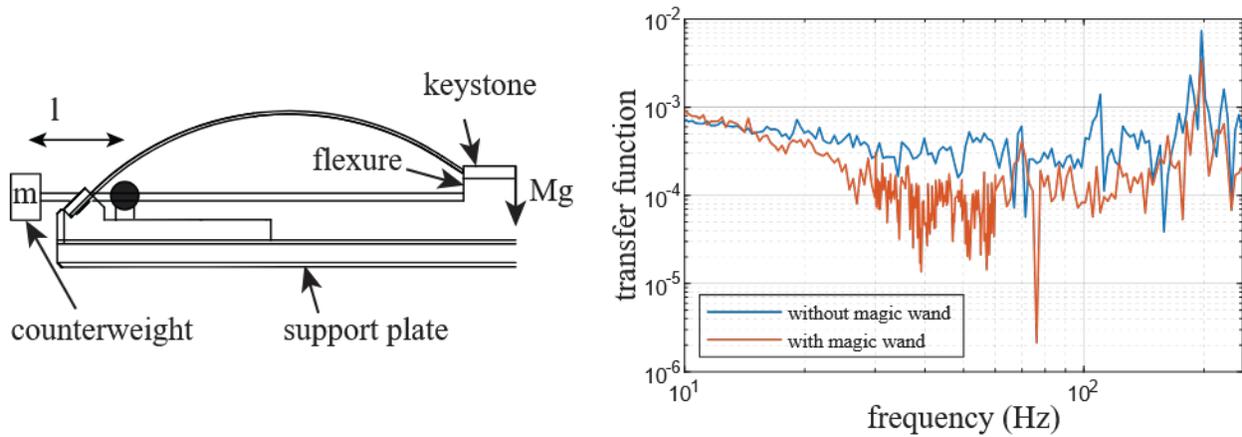


Figure 5.10: Left: magic wand diagram. For simplicity, only one blade and one wand are represented. Right: transfer functions before and after the integration and tuning of the magic wand

space, their performance was evaluated on a test GAS filter outside of the cryostat. Measurements of transfer functions obtained using the same method as for the results in Fig. 5.9 show an improvement in vibration attenuation after the integration and fine tuning of the magic wands. While the measurements were again limited by parasitic coupling of acoustic vibrations, results are promising for opening up the possibility of achieving factors of at least 10^4 of vibration isolation per GAS filter.

CHAPTER 6

MECHANICAL RESONANCES OF THE SAPPHIRE CAVITIES

The model for the detector response developed in section 2.2 takes as parameters the values of the resonant frequencies corresponding to the longitudinal mechanical modes of the cavities and the quality factors of these resonances. The interpretation of the experimental results depends critically on the determination of these quantities.

Finite element analysis offers a means to calculate the cavity resonances. The modal analysis mode of ANSYS Mechanical was used to solve the eigenvalue problem that renders the natural frequencies of the two cavities. The simulation assumed that the cavities were fixed at four points at the centers of the support holes and included the effect of gravity. The calculated frequencies are 35.00 kHz for the long cavity and 66.78 kHz for the short cavity. In addition to the longitudinal modes, the simulations also revealed other resonances inside the science band. While these do not contribute to first order to the results, they can be excited by ambient noise, resulting in peaks in the measured data. Observations of patterns between where peaks corresponding to different modes were expected to occur and where they actually occurred were helpful in the identification of the longitudinal resonances. An inventory of the lowest few resonant modes for both cavities, with their corresponding frequencies, is given in Fig. 6.1.

The simulations took approximately 5 minutes to calculate the first 100 modes for a mesh of around 200000 nodes, using 8 cores and 20 GB of RAM. The dependence on the results of the mesh size was on the order of tens of Hz for the two longitudinal modes.

Ring-down measurements [91, 92] have found high quality factors for bulk sapphire, $Q = 10^8$ – 10^9 . However, the attachment of the thermal links and (to a lesser degree) the optical contacting of the mirrors are likely to have introduced additional damping into the current system. Rather

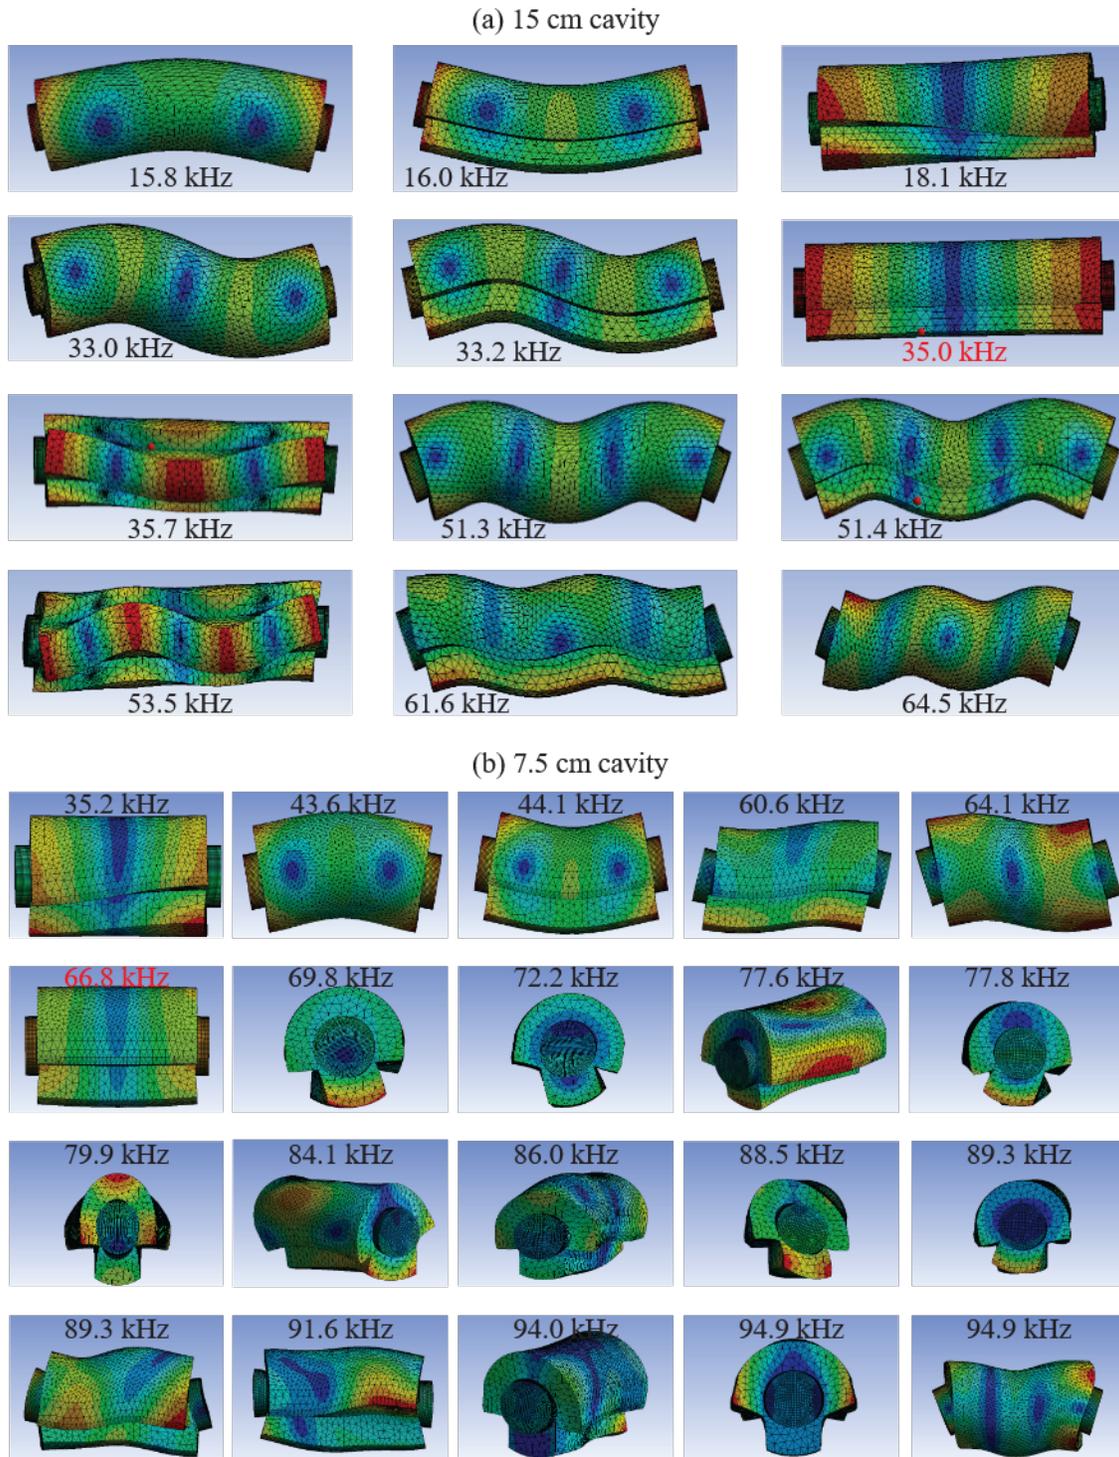


Figure 6.1: Catalog of the lowest resonant modes of the (a) 15 cm and (b) 7.5 cm sapphire cavities, calculated using ANSYS. The longitudinal modes are highlighted in red

than relying on literature values, it is hence much more reliable to measure the quality factors in the actual setup. Measurements of the resonances are also important to determine their frequencies with more accuracy than the simulations can provide, since the latter did not include the heat links and are limited in their capacity to faithfully model the way the cavities are supported and loaded.

A first series of measurements of the resonances was done at room temperature, outside the cryostat, on identical copies of the two sapphire spacers (albeit without the mirrors attached). Although this test setup also did not completely reflect the experimental conditions, these initial measurements served in establishing some procedures for exciting the resonant modes and observing the mechanical response. After having built some confidence in our ability to observe actual resonances, a second set of measurements was conducted in situ, on the cold cavities. The results were used in conjunction with those of the finite element analysis to identify the mechanical resonances of the cavities.

6.1 Room temperature measurements

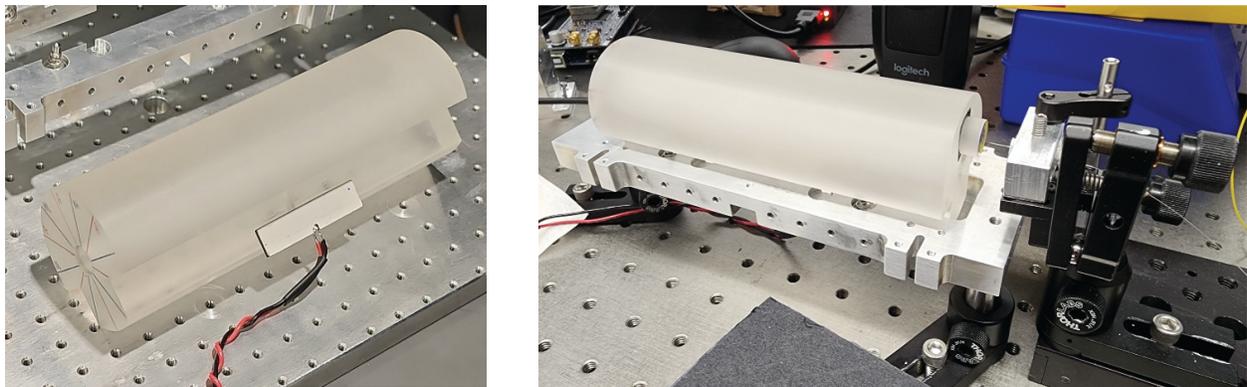


Figure 6.2: Room temperature setup for resonance measurements. Left: piezo glued to the side of a spare spacer. Right: spacer with piezo seated in its cradle; a fiber interferometer detects the motion of a mirror glued to one end of the spacer

Room temperature measurements were conducted on spare copies of the 7.5 cm and 15 cm sapphire spacers. The spacers were seated on nominally identical cradles and supported on conical

springs and rounded pins that replicated the setup used inside the cryostat.

The spacer resonances were excited by applying a drive to a piezoelectric element glued to the side. The piezos were chosen to have resonant frequencies in the vicinity of the expected longitudinal resonances (around 35 kHz for the 15 cm cavity and 67 kHz for the 7.5 cm cavity), after it was found that this enhanced the ability to observe the resonances. Two different detection techniques were employed: one was a purely electrical measurement, by observing the impedance of the piezoelectric element, the second was optical, through interferometric measurements of the motion of the spacer. The results were compared to each other and to those obtained from finite element analysis.

6.1.1 Detection via electro-mechanical coupling

A piezoelectric material can be viewed as an electro-mechanical system, in which electrical energy is converted into mechanical deformation and vice versa. This energy transfer is frequency dependent, with maximum energy being absorbed at the mechanical resonances, where the strong mechanical response is revealed as a large decrease in the equivalent impedance of the piezoelectric.

The electrical coupling of the mechanical behavior of the piezo is conventionally modeled as an equivalent resonant RLC circuit (where the inductance L_m , capacitance C_m and resistance R_m are measures of the mass, spring constant and damping) in parallel with the electrical impedance of the piezo, represented as a capacitive element (with capacitance C_e). The corresponding circuit diagram is depicted in Fig. 6.3(b). The equivalent impedance of this circuit displays a minimum close to the series resonant frequency f_s , which coincides with the mechanical resonant frequency, due to the resonant behavior of the series RLC circuit corresponding to the mechanical part of the system, and a maximum at a slightly higher frequency, corresponding to a good approximation to the resonance f_p of the parallel circuit formed by the electrical and mechanical components (see

Fig. 6.3(c)).

When the piezo is attached to an external system (such as the cavity spacers), the mechanical behavior is determined by the overall mechanical impedance of the piezo and object it is attached to. Therefore, the mechanical resonances of the latter give rise to visible features in the electrical impedance frequency response. An example is presented in Fig. 6.4, where the impedance is measured (by measuring the drive voltage and the resulting current) as a function of drive frequency for the same piezo, first isolated (hanging from a string), then attached to the 15 cm spacer. When the piezo is glued, its internal resonances are strongly suppressed, but additional features with the characteristic dispersive shape expected from a resonance (as depicted in Fig. 6.3(a)) appear. These are close in frequency to where cavity mechanical resonances are expected from the simulations (these frequencies are marked in the picture by dashed vertical lines, with the red line marking the longitudinal mode). These peaks are more prominent when they happen to occur on the wings of the piezo resonances, an observation that guided our choice of piezos.

6.1.2 Optical detection

A second approach to identify the spacer resonances involved searching for peaks in the spectrum of the displacement of a mirror glued to the spacer, measured precisely using a fiber interferometer while the spacer was excited using the attached piezo.

The interferometer consists of a 1310 nm laser, the light of which is sent into one of the two input ports of a fiber coupler. One of the coupler outputs leads into an optical fiber, the other end of which is placed on a translation and tip-tilt stage and brought in close proximity to the mirror on the spacer. When aligned, the light reflected off the mirror surface reenters the fiber and is detected by a photodiode connected to the second input of the coupler. At the photodiode, it interferes with a secondary beam reflected off the cut end of the fiber. The interference signal encodes variations in the distance between the mirror surface and the fiber end.

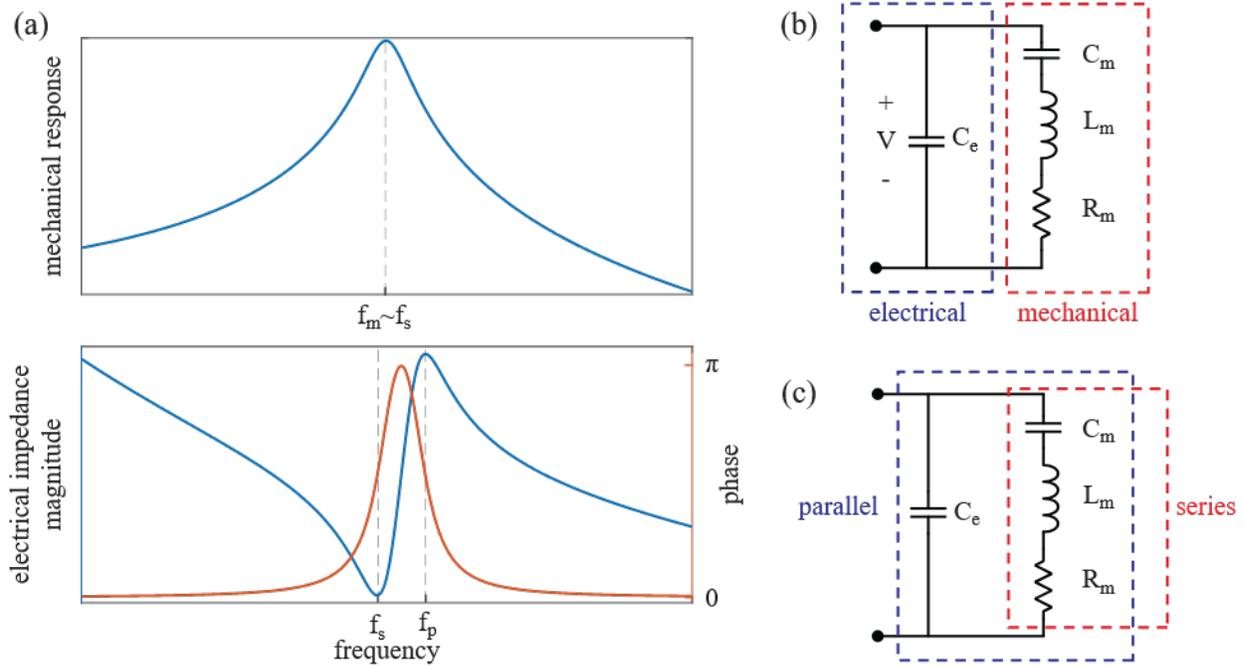


Figure 6.3: (a) Mechanical transfer function and impedance frequency response of a piezoelectric element. (b) Equivalent circuit diagram, highlighting the electrical and mechanical components. (c) Equivalent circuit diagram, highlighting the series and parallel resonant circuits

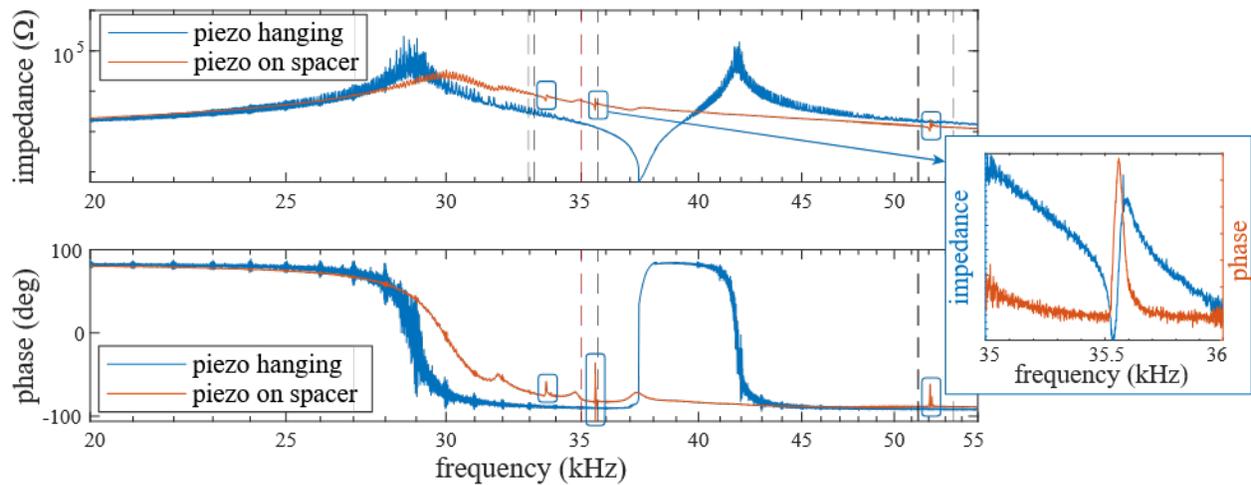


Figure 6.4: Frequency dependence of the impedance magnitude and phase for a piezo hanging from a wire or attached to the 15 cm spacer. Dashed lines mark the expected values of the spacer resonant frequencies (red for the longitudinal mode, black for other modes). The inset shows a close-up of the resonance around 35 kHz

The excitations of the mechanical modes proved to be too weak to be detected directly by a spectrum analyzer while sweeping the drive frequency. However, the feeble resonant response could be extracted by demodulating it with the drive signal using a lock-in amplifier (Zürich Instruments MFLI), as illustrated in the schematic in Fig. 6.5.

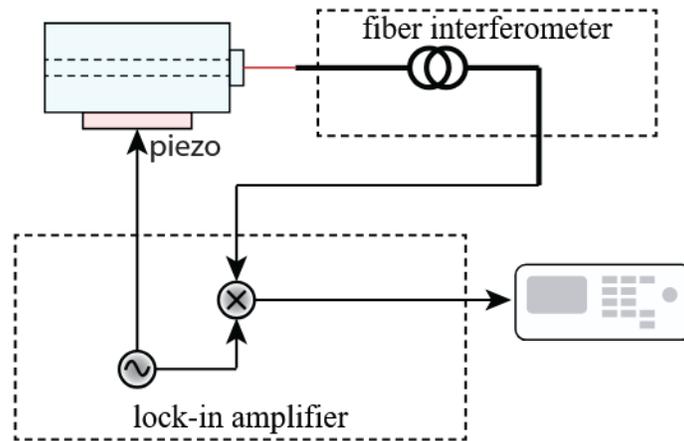


Figure 6.5: Optical detection scheme

6.1.3 Results

A number of resonances for each of the two spacers were detected by both of the previously discussed techniques. The peaks appear slightly shifted from each other when measured with the two methods. This is in part because of the small offset between the impedance minimum and the actual series resonance, but it can also be caused by the sweep speed being too fast compared to the build-up time of the resonance response. Nevertheless, the double detection provides solid evidence that the observed peaks are real.

For the 15 cm spacer, a strong candidate for the longitudinal peak was found at 35.56 kHz. A few less well-defined features potentially indicate other peaks around this frequency (the presence of which was predicted by the finite element analysis), the most prominent of which, at 33.4 kHz, is likely to be one of the two second order bending modes (see Fig. 6.1). Additionally, two double

peaks, one at 16.06 kHz and 16.26 kHz and another around 52.1 kHz, were attributed, based on the finite element simulation results, to other bending modes and a smaller peak at 17.63 kHz looks consistent with the first torsion mode. Overall, compared to the theoretical values, the observed resonances tend to occur about half a kHz higher.

Fig. 6.6 exhibits the resonances identified in the piezo impedance and interferometer signals. In the case of the former, the impedance frequency dependence when the piezo is not attached to the spacer (labeled as "piezo hanging") is also provided for the purpose of highlighting the features that originate from the spacer (the curves corresponding to the isolated piezo sometimes have a different baseline because gluing it to the spacer can dampen the piezo resonances).

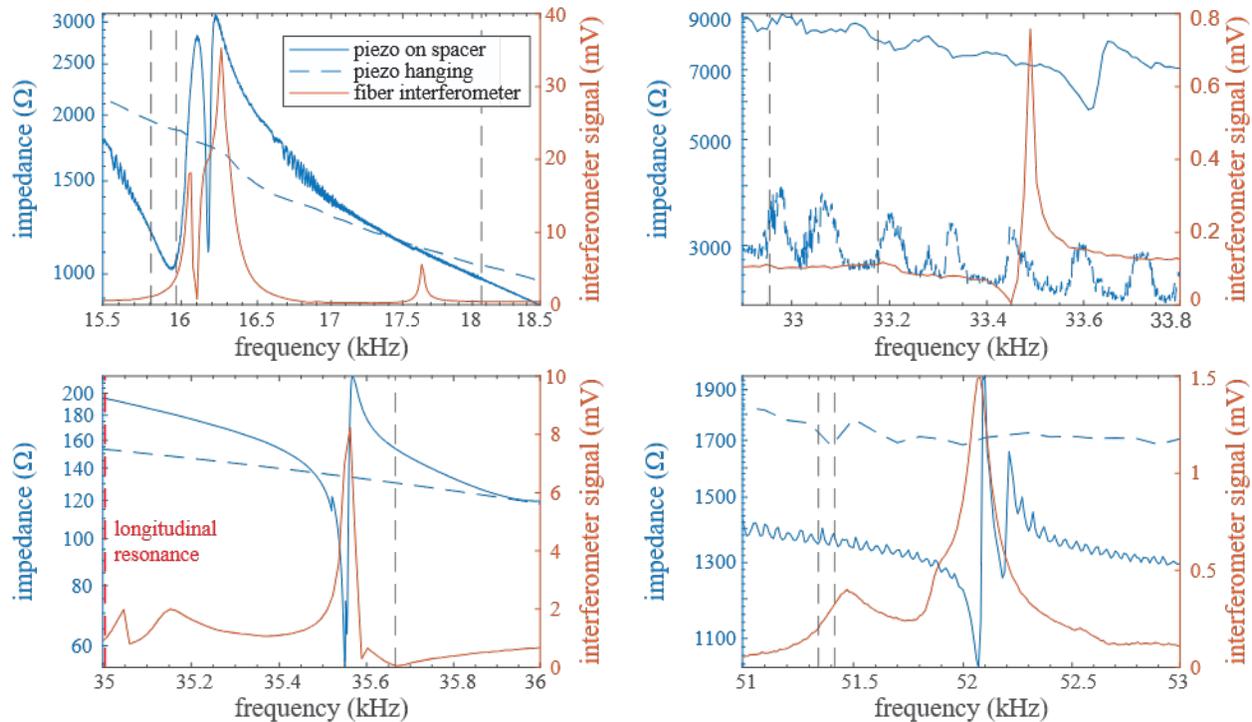


Figure 6.6: Resonances of the 15 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode)

The longitudinal resonance of the 7.5 cm spacer turned out to be more difficult to pin down. We were unable to identify any resonant features in the vicinity of the predicted frequency of 66.8 kHz

and the closest peaks (at 59.96 kHz, 64.08 kHz, 70.04 kHz and 70.98 kHz) align better with other resonances. The most clear observed resonances were two peaks at 46.19 kHz and 46.58 kHz which were attributed to the first order bending modes (expected at 43.6 kHz and 44.1 kHz as per Fig. 6.1). Based on this more than 2 kHz shift between theory and measured values, it is also possible that the discrepancy between the simulations and the real values is bigger in this cavity and one of the observed peaks is, in fact, the longitudinal resonance. Fig. 6.7 presents measurements of the bending modes and a broad scan over the modes between 60 kHz and 80 kHz. The insets of the plot on the right side are narrow scans with finer frequency resolution over the closest observed resonances to 67 kHz.

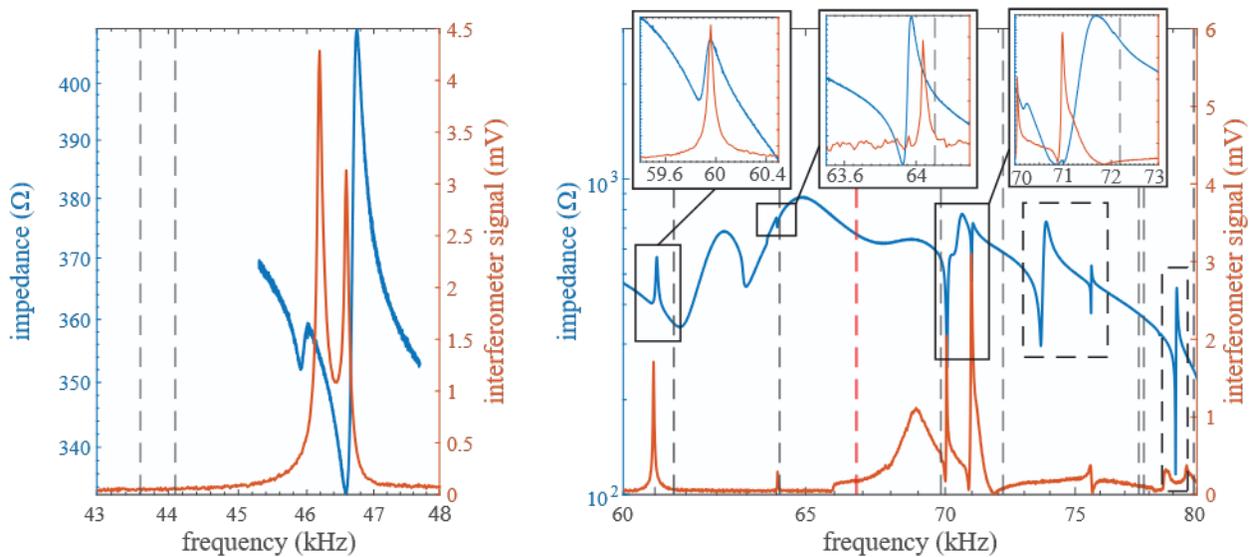


Figure 6.7: Resonances of the 7.5 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode)

All measurements were repeated with two different piezo models, as an additional check that the observed peaks are not an artifact of the specific piezo being used. The method of supporting the cavity was also varied (for instance, by directly seating it on the table or on a damping surface instead of using the pins and springs). The support conditions influenced the quality factor of the

resonances, but did not change their frequency within the measurement resolution. Additionally, clamping the spacer longitudinally was found to massively suppress the peak attributed to the longitudinal resonance of the long cavity, substantiating the conclusion that the observed peak was an internal mode of the cavity.

Ultimately, since the results of the out-of-cryostat measurements are not expected to fully apply to the system inside the cryostat (the spare spacers did not have mirrors and heat links attached, and the glued piezo could potentially modify the resonant frequencies, too), rather than refining this initial study, we decided to proceed with applying the lessons we learned to the actual science cavities.

6.2 In-situ measurements

Compared to the room temperature tests, the main hurdle in measuring the resonances of the cold cavities stemmed from the lack of direct access, as at the time of this investigation opening up the cryostat to alter the setup inside was unfavorable for the overall progress of the experiment. This constrained our options for applying mechanical drives to the system.

Fortunately, the limitations in exciting the resonances most effectively were compensated by having a very sensitive length sensor - the cavities themselves. Rather than requiring an external measurement device, such as the fiber interferometer, or relying on indirect methods, such as the impedance measurements, we were thus able to look for evidence of mechanical excitations in the transmission signals of the cavities.

To identify the frequencies of the resonances, we combined searches of peaks caused by ambient vibrations in the transmission spectra with measurements of the response to driven excitations from a piezo outside of the cryostat and finite element calculations.

6.2.1 Response to ambient vibrations

Excitations of the mechanical resonances due to laboratory vibrations can show up in the transmission spectra of the two cavities, as long as the ambient drive is strong enough to raise the response above the noise floor. Since the cavities are co-located and have identical support structures, many mechanical resonances associated with other components of the system are expected to appear in the two spectra at the same frequencies, which makes them easy to rule out. Peaks present in the signal from one cavity and not the other, of which the analysis of transmission spectra revealed several, were classified as cavity resonance candidates and warranted further investigation. For the long cavity, notable are two peaks around 16 kHz, where the two bending modes of the cavities are expected, and several peaks around 35 kHz (the strongest of which occurred at 34.76 kHz), in the vicinity of the predicted longitudinal mode. For the short cavity, three peaks (at 41.85 kHz, 42.05 kHz and 42.56 kHz) were found in the region where the first order bending modes were anticipated to occur. No peaks were observed close to the longitudinal resonance. All of these are shown, together with the predicted resonant frequencies, in Fig. 6.8.

A retrospective analysis of the beat note measured over several days provided additional insights. Although individual peaks in the beat note cannot be definitively linked to a specific cavity, the long collection time unveiled peaks that were not discernible in the previously presented transmission data. In particular, four main candidates for the 15 cm cavity longitudinal resonance were identified (the frequencies of which are marked in the left panel of Fig. 6.9), as well as several peaks around the longitudinal resonance of the 7.5 cm cavity.

6.2.2 Response to external drive

In addition to the transmission noise analysis, the cavities were also driven externally with a piezo attached to the outer surface of the cryostat vacuum can, with the drive frequency slowly swept over ranges of interest. The purpose was both to try to excite peaks that may have been

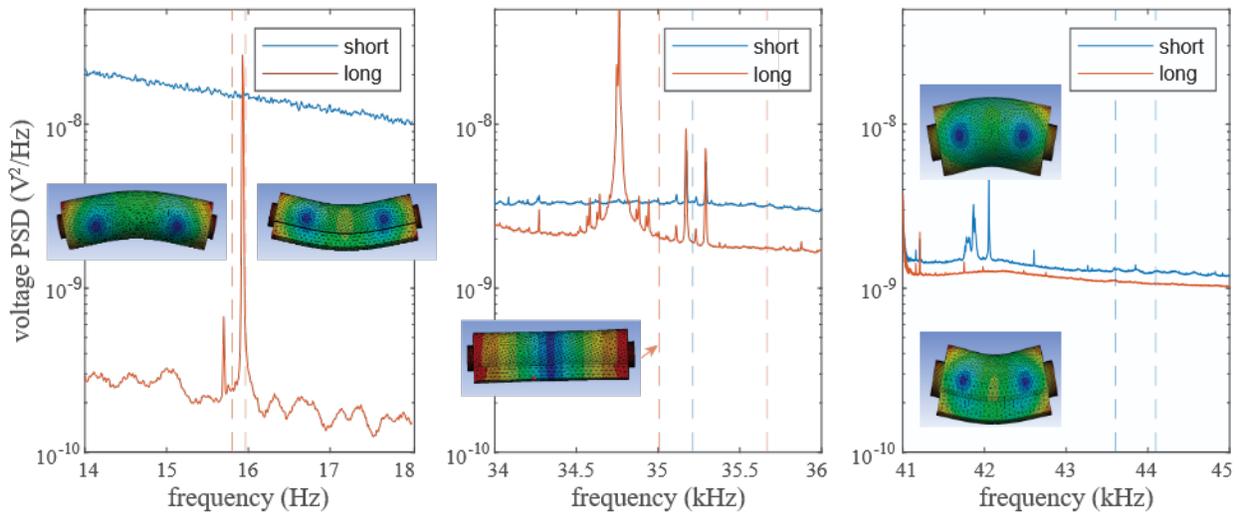


Figure 6.8: Noise peaks in the transmission spectra of the two cavities. Theoretical frequencies of the mechanical resonances of the short cavity (blue) and long cavity (orange) are represented by vertical lines (for the main ones, the mode shapes are also shown)

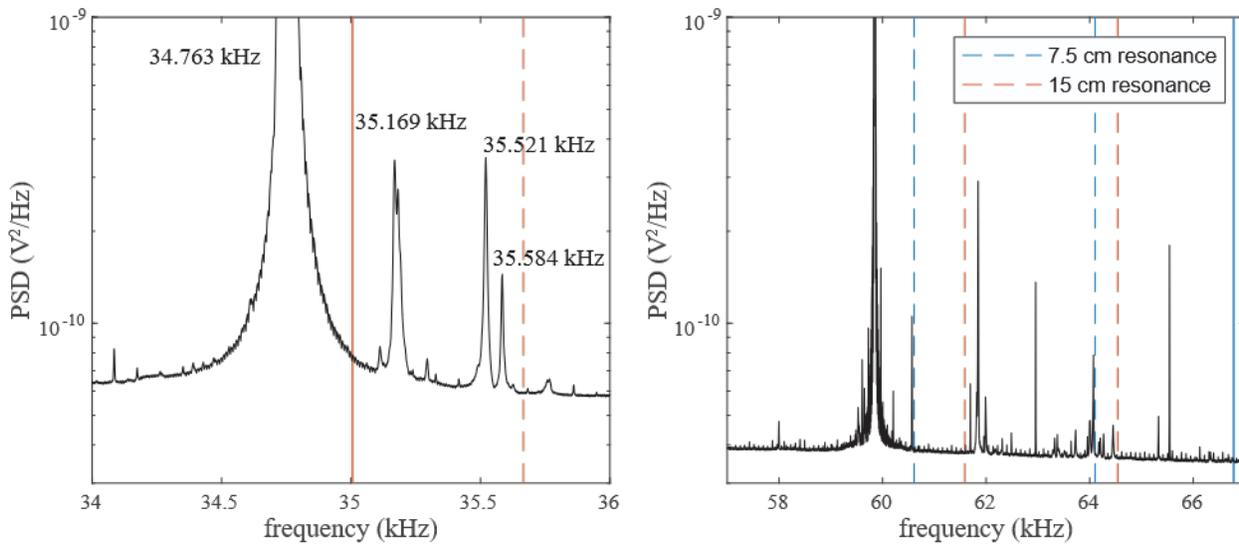


Figure 6.9: Beat note noise peaks around the longitudinal resonant frequencies of the two cavities. The theoretical values of the longitudinal frequencies are shown as solid vertical lines. Other resonances in the plotted frequency range are marked by dashed lines

below the noise level and thus could not be detected in the noise spectra and to confirm, through the observation of resonant enhancements, that the ones previously identified are mechanical in nature.

To enhance sensitivity, a lock-in detection scheme similar to that used in the previous room-temperature investigations was employed. This time, however, the detected signal was the output of the beat note photodetector. Two types of measurements were conducted. In one case, the light from both cavities was allowed to reach the photodetector and the beat note, after amplification, was processed through the delay line frequency discriminator described in section 4.5.2, which removed the carrier to isolate sidebands from mechanical excitations. The resulting demodulated signal served as input to the lock-in amplifier (cf. Fig. 6.10). Alternatively, the light from the cavities was blocked one at a time, and the amplified photodetector output was fed directly into the lock-in amplifier. In this case, we interpret the detected signal as arising from the beating between the carrier and the sidebands generated by the excitation of the mechanical resonances. Although this signal is weaker than the beat note and requires a significantly higher degree of amplification, looking at the individual cavity transmissions allowed us to discriminate which cavity the observed peaks were associated with.

By repeating these measurements with the drive on and off, we confirmed which peaks were mechanical resonances excited by the piezo. In this way, some spurious peaks that were present during the driven scans, but were found to be unrelated to the drive, could be rejected.

Out of the forest of peaks previously observed in the beat note in the 60 kHz-70 kHz range, two peaks, at 61.846 kHz and 64.073 kHz, were confirmed by the lock-in tests. Additionally, one new peak, which was barely visible in the beat note, was found at 66.685 kHz, the closest one to the expected value of the longitudinal resonance (66.8 kHz). All three were confirmed by the light blocking test to be resonances associated with the short cavity. Fig. 6.11 highlights the three resonances measured with the lock-in method, plotted together with the beat note in the absence of

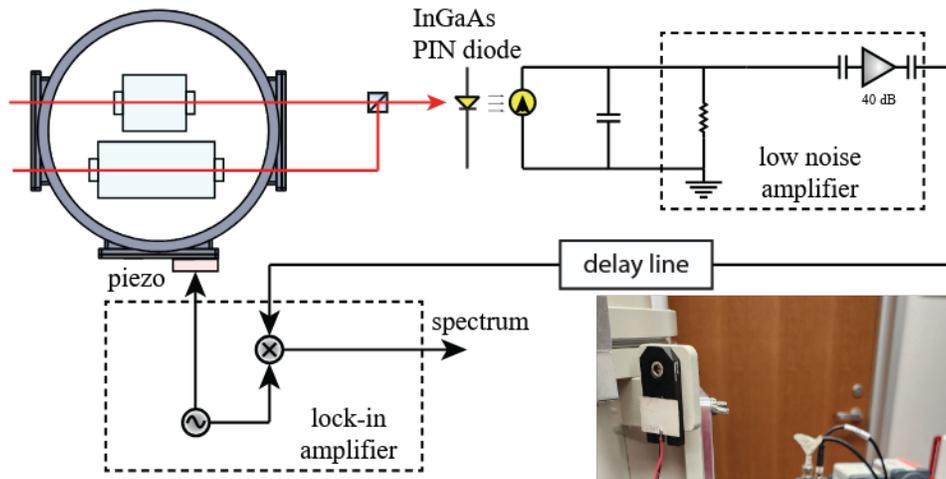


Figure 6.10: Lock-in detection diagram

a drive. The driven response of the 15 cm cavity (showing no peaks) is also presented for reference. The arrows indicate the direction of the frequency sweep (frequency shifts with respect to the peaks in the beat note are consistent with this direction). Additionally, the lock-in measurements also revealed the three peaks around 42 kHz (mentioned in section 6.2.1) as resonances associated with the short cavity.

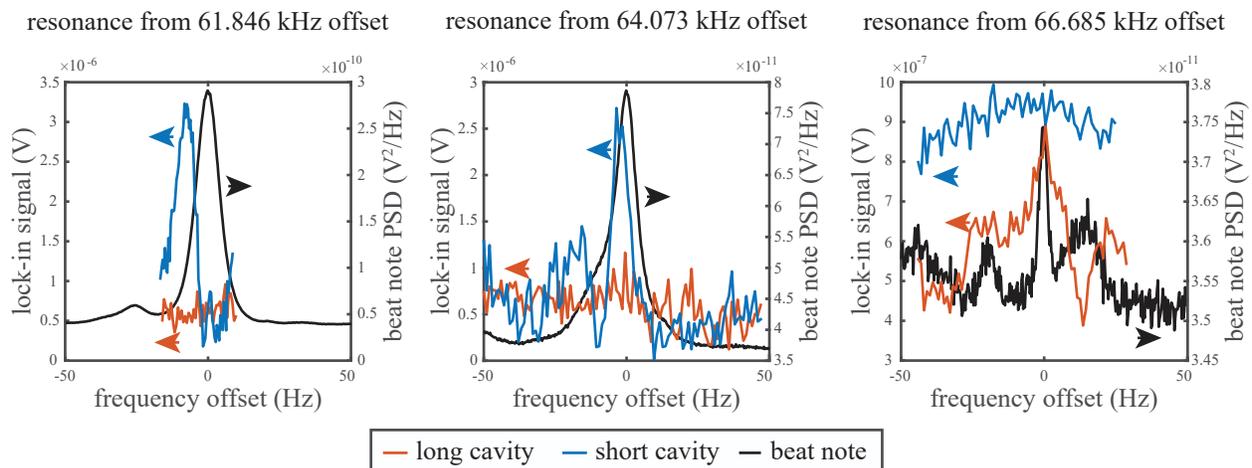


Figure 6.11: Candidates for the short cavity longitudinal mode identified through lock-in measurements

For the long cavity, all four peaks around 35 kHz responded to the applied drive (see Fig. 6.12).

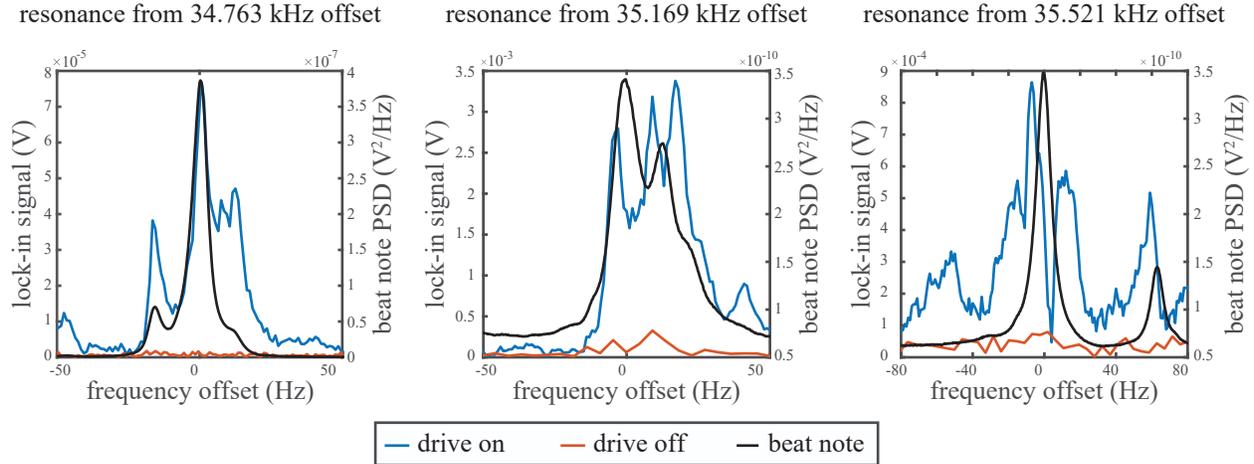


Figure 6.12: Candidates for the long cavity longitudinal mode identified through lock-in measurements. Scans with the drive off are also presented for reference

While the 34.764 kHz resonance appears to be much stronger than the others, the other three peaks could not be definitively ruled out.

	Drive on		Drive off	
	35 kHz peaks visible?	64/66 kHz peaks visible?	35 kHz peaks visible?	64/66 kHz peaks visible?
Beat note	Yes	Yes	No	No
15 cm transmission	Yes	No	No	No
7.5 cm transmission	No	Yes	No	No

Table 6.1: Summary of driven response tests for the longitudinal mode candidates

6.2.3 Corroboration with finite element analysis

During the driven response measurements, a group of peaks of the short cavity emerged most prominently and consistently: the peaks around 42 kHz and 42.5 kHz. This observation led to an idea to hypothesize that these features are first-order bending resonances and, based on this

premise, determine the longitudinal resonance through a one-parameter fit. The fit parameter was implemented as an overall scaling constant α to the stiffness matrix given by Eq. 3.6). Finite element analysis found best alignment with the presumed bending modes for $\alpha = 0.92$. Interestingly, for this case the longitudinal resonances aligned well with the peak observed at 64.073 kHz, while the peak at 61.846 kHz matched one of the bending modes. Fig. 6.13 shows a comparison of a wide driven scan (the scan with the drive off is also provided for reference) against the results of finite element analysis, where the fitted peaks are highlighted as gray vertical lines. This hypothesis is appealing, particularly because the 64.073 kHz resonance was one of the strongest ones observed.

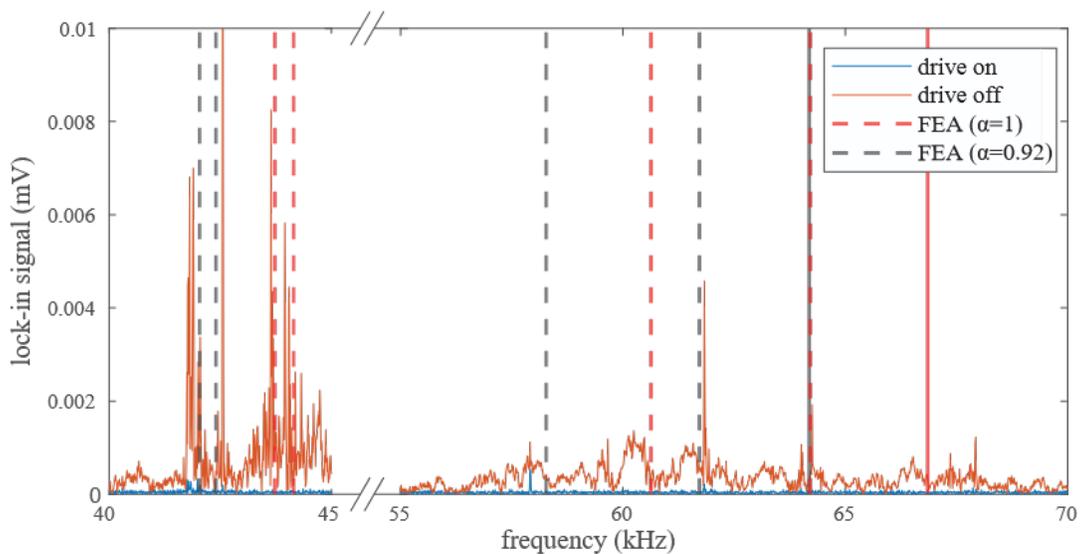


Figure 6.13: Resonances of the 7.5 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode)

On the other hand, without the fitting (in other words for $\alpha = 1$), the much weaker peak at 66.685 kHz (not well evidenced in the wide scan, but visible in Fig. 6.11) emerges as the most likely option for the longitudinal resonance. In this case, the 64.073 kHz and 61.846 kHz peaks can be second order bending modes (the first of the two agreeing much better than the second with

the simulations). This second option is depicted by the red lines in Fig. 6.13)

While this exercise was not entirely conclusive either, it reinforced two peaks (64.073 kHz and 66.685 kHz) as viable candidates for the short cavity longitudinal resonance.

6.2.4 Quality factor estimations

The second purpose of this investigation was to determine the quality factors of the resonances, essentially by measuring their widths.

In the case of the stronger resonances, the width could be easily measured in the transmission or beat note spectrum. However, for the weaker ones, the observation of the full width at half maximum was impeded by the noise floor. In this case, a more reliable approach was to drive these resonances at discrete frequencies around the resonance using the piezo, wait at each frequency long enough for the resonant response to build up, then measure the peak in the spectrum. The width of the resonance was then evaluated as that of the envelope of all of the measured peaks.

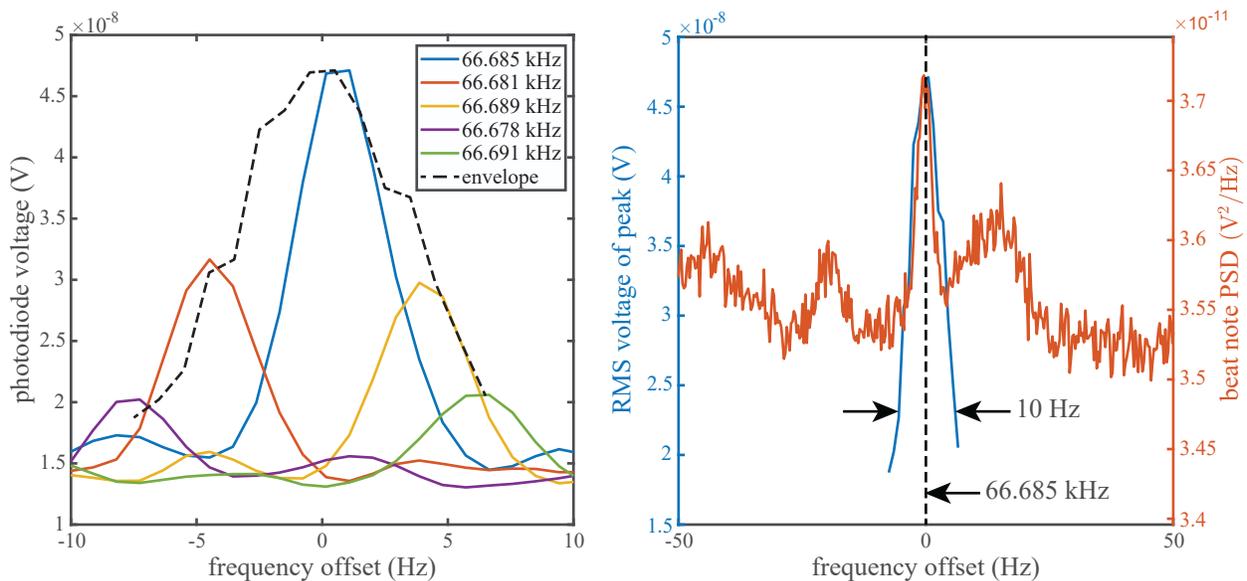


Figure 6.14: Resonances of the 7.5 cm spacer identified through impedance and optical measurements. The vertical lines indicate frequencies at which resonances are expected based on finite element simulations (the red line marks the longitudinal mode)

All the peaks we investigated were found to be around 10 Hz wide, implying quality factors on the order of 10^3 . These are much lower than literature values for bulk sapphire [91]. One possibility is that the quality factors were substantially reduced by the attached thermal links. However, the measurement could also have been limited by the linewidth of the frequency source providing the piezo voltage, as the measured lineshape would be a convolution of the resonance being measured and the drive. For this reason, the measured quality factors should be understood as worst-case values.

While in the end a completely unambiguous identification of the cavity modes was not achieved, these measurements allowed establishing upper and lower bounds on the values of the longitudinal resonances. In the case of the short cavity, these were taken as the frequencies of the two peaks that were identified as viable candidates, 64.073 kHz and 66.685 kHz. For the long cavity, the extreme values of the four prospective peaks in Fig. 6.9 were selected (34.763 kHz and 35.584 kHz). For both cavities, an upper limit of 10 Hz for the widths of these resonances was placed. These values were used to estimate a conservative transfer function of the cavities to a mechanical drive, as explained in section 7.4.1.

These values could be refined with further research, for instance by opening up the cryostat and adding a piezo closer to the cavities. With the means of driving the resonances more efficiently, a measurement of the quality factors using ring-down techniques might also become possible, which would allow overcoming the potential limitation imposed by the drive linewidth.

CHAPTER 7

IMPROVED DIRECT BOUNDS ON ULTRALIGHT DARK MATTER

The apparatus described in chapters 3 to 5 was used to conduct a search for ultralight dark matter in the 5 kHz –100 kHz range. The final data collection was run continuously over a period of approximately 4 days. The results successfully demonstrate the effectiveness of Fabry-Pérot cavities as competitive dark matter detectors in the kHz and tens of kHz domain.

This chapter describes the processes of data acquisition and analysis, detailing the methodology employed in setting upper bounds on the coupling constants d_{m_e} and d_e , defined as:

$$\begin{aligned}\frac{\delta\alpha(t, \mathbf{x})}{\alpha} &= d_e \frac{\sqrt{4\pi\hbar c}}{E_P} \phi(t, \mathbf{x}), \\ \frac{\delta m_e(t, \mathbf{x})}{m_e} &= d_{m_e} \frac{\sqrt{4\pi\hbar c}}{E_P} \phi(t, \mathbf{x}),\end{aligned}\tag{7.1}$$

and culminating with the report of new direct constraints in the aforementioned frequency range.

7.1 Data acquisition

The demodulated beat note between lasers resonant with the cavities was measured using a 24-bit DAQ card (Advantech PCIE-1802L), with a voltage range of ± 2 V and a sampling rate of 216,000 S/s, resulting in a total number of samples $N = 7.48 \times 10^{10}$. The total data collection time is sufficient to properly sample the dark matter linewidth, the FWHM of which in the frequency range considered in this experiment has a minimum of ≈ 5 μ Hz at 5 kHz, in the SHM. The resulting raw data is an array V_k , where k takes integer values from 0 to $N - 1$, sampled at a time interval $\Delta_t = T_{tot}/N$, where $T_{tot} \approx 3.4 \times 10^5$ s is the total measurement time for the entire data set. Some parameters related to the data collection that are relevant for the subsequent analysis are

summarized in table 7.1.

Sampling rate	216,000 S/s
Total collection time	$\approx 3.4 \times 10^5 \text{ s} \approx 4 \text{ days}$
Total number of samples	7.48×10^{10}
Sampling interval	4.63 μs
Frequency resolution	2.9 μHz

Table 7.1: Data acquisition parameters

7.2 Strain spectrum

7.2.1 Discrete Fourier transform review

The first step in data processing is to Fourier transform the time series. For this purpose, we would like to first revise a few relevant definitions.

The discrete Fourier transform (DFT) of the strain noise signal is

$$\tilde{V}_p = \sum_{k=0}^{N-1} \exp(-i \frac{2\pi}{N} pk) V_k, \quad (7.2)$$

where $\tilde{V}_p = \tilde{V}(f_p)$, with $f_p = p/(N\Delta_t) = p\Delta_f$, $p = \overline{0, N-1}$ being the DFT frequencies. The quantities \tilde{V}_p are complex and, for a real signal V_k , obey the relation:

$$\tilde{V}_p = \tilde{V}_{N-p}^*. \quad (7.3)$$

It follows that the DC component \tilde{V}_0 and the Nyquist component $\tilde{V}_{N/2}$ are necessarily real-valued.

The Fourier transform can also be expressed as a linear transformation by a matrix \mathbf{U} , which

is symmetric and unitary:

$$\tilde{\mathbf{V}} = \sqrt{N}\mathbf{U}\mathbf{V}. \quad (7.4)$$

It is useful to define two more quantities: the autocorrelation matrix \mathbf{C} with elements $C_{kk'} = \langle V_k V_{k'} \rangle$, where $\langle x \rangle$ denotes the expectation value of a variable x , and the (two-sided) power spectral density (PSD) matrix of the data series, $\tilde{\mathbf{C}} = N\mathbf{U}\mathbf{C}\mathbf{U}^\dagger$. In terms of the Fourier components \tilde{V}_p , the latter can be expressed as:

$$\tilde{C}_{pp'} = \langle \tilde{V}_p (\tilde{V}_{p'})^* \rangle. \quad (7.5)$$

For a stationary process, the PSD matrix is diagonal and can be reduced to

$$\tilde{C}_p \equiv \tilde{C}_{pp'} = \delta_{pp'} \langle |\tilde{V}_p|^2 \rangle. \quad (7.6)$$

Thus, the PSD is essentially the Fourier transform of the signal autocorrelation and represents a measure of the power distribution of the signal across the frequencies f_p . It is this definition of the power spectral density that is used in the derivations in this chapter (in particular, in section 7.4). This is different from the conventional definition of the PSD, in units of V^2/Hz , which is normalized to be independent of the measurement frequency bandwidth:

$$S_p(f) = \frac{\Delta_t}{N} \tilde{C}_{pp'} = \frac{\Delta_t}{N} |\tilde{V}_p|^2. \quad (7.7)$$

7.2.2 Conversion to strain

A Fast Fourier Transform (FFT) algorithm was employed to compute the DFT of the entire time series (237 GB of data when saved in single precision format). This resource-heavy operation was performed on a 3.2 GHz desktop computer with 16 cores and necessitated 28 hours to complete

with 1.6 TB of RAM.

The frequency-domain voltage noise was then converted to a frequency noise spectrum $\delta\tilde{f}_p$ by dividing the data by the delay line frequency discriminator constant K_d (Eq. 4.2). Subsequent division by the laser frequency yields the differential strain between the two cavities, which represents the sum of the detector noise \tilde{n}_p and any potential dark matter signal \tilde{s}_p :

$$\tilde{d}_p = \frac{\delta\tilde{f}_p}{f_{opt}} = \tilde{s}_p + \tilde{n}_p. \quad (7.8)$$

While it is \tilde{d}_p that will be utilized in the subsequent analysis, all the strain plots that we present here actually depict the strain amplitude spectral density (ASD), which we will denote by $\tilde{h}_p = \sqrt{\Delta_t/N}\tilde{d}_p$. It represents the square root of the conventional PSD defined in Eq. (7.7) and is a quantity that would be more familiar to an experimentalist. Making a clear distinction between the two is important as we move forward.

An averaged spectrum (computed with Welch's method for estimating the PSD, with 10^5 averages) of the resulting strain is presented in Fig. 7.1(a).

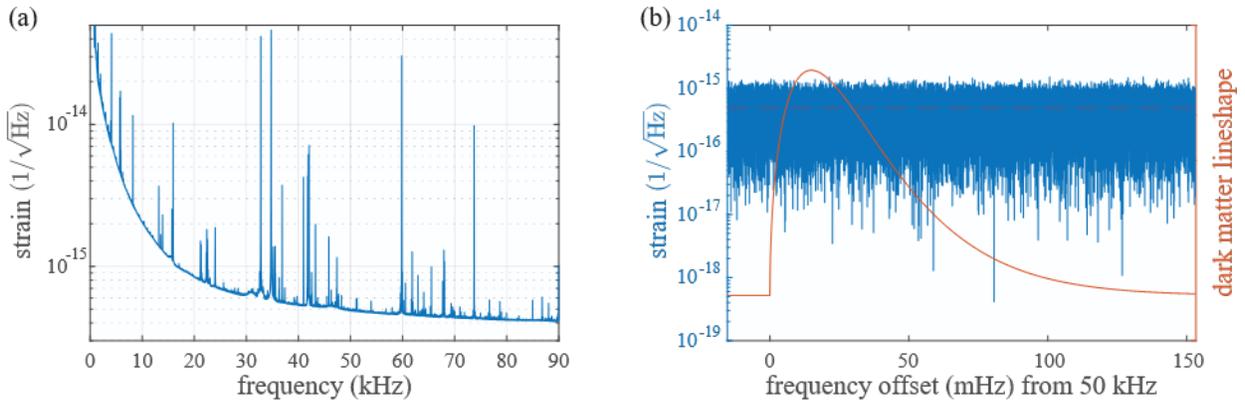


Figure 7.1: (a) Strain noise in the 0-90 kHz frequency range (average over 10^5 FFTs); (b) Strain noise in a narrow window around 50 kHz of a width comparable to the dark matter characteristic linewidth at this frequency; the orange curve shows the dark matter lineshape; the dashed line marks the average value of the noise within the pictured window

7.2.3 Noise floor properties

Before proceeding with detailing the procedure used to place bound on the coupling to dark matter, it is useful to point out some properties of the measured noise floor which will be underlying assumptions in the subsequent analysis.

Stationarity:

The analysis of the measured data in the frequency domain relies on the assumption that the data is stationary (i.e. its probability distribution and consequently its mean, variance and auto-correlation do not change over time). The augmented Dickey-Fuller test [93] performed on the time series confirmed this assumption. The test was implemented using the built-in Matlab function *adftest*.

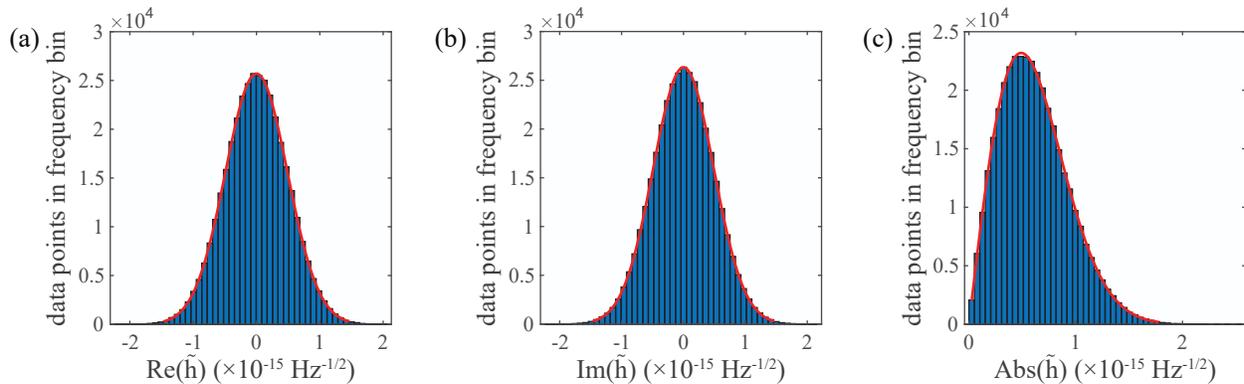


Figure 7.2: Statistical distribution of the strain noise DFT over a window of width $\sim 11/\tau_c$ around 50 kHz. The real and imaginary parts (panels (a) and (b)) fit to a Gaussian distribution. The magnitude (panel (c)) is Rayleigh distributed. The red lines represent fits to the histograms.

Gaussian distribution:

A second important property of the noise floor is that it is normal distributed. Its Gaussian behavior in the time domain leads to both the real and imaginary parts of the (complex) fast Fourier

Transform (FFT) of the strain being Gaussian distributed. As a consequence, the amplitude of the FFT follows a Rayleigh distribution. This is equivalent to a transformation from Cartesian coordinates to polar coordinates (and equivalent to the way in which the Rayleigh distribution arises in the DM lineshape in section 1.2.4). We verified the Gaussian nature of our noise floor by plotting its distribution over narrow frequency windows (of width approximately equal to 11 dark matter linewidths). An example, corresponding to a frequency of 50 kHz, is depicted in Fig. 7.2.

Whiteness:

As evident from figure 7.1(a), the measured strain noise is not white across the measurement frequency range. However, in frequency windows of similar width to that of the dark matter characteristic lineshape, the noise can be appropriately approximated as white (see figure 7.1(b)).

7.3 Investigation of strain spectrum peaks

The strain spectrum was scrutinized for the presence of a potential peak with the characteristic lineshape of dark matter. All the peaks identified by eye were found to have widths of tens of Hz, too broad to be consistent with a dark matter signal and are more likely to be caused by noise of mechanical, electrical or optical origin. A representative example of a peak is displayed in Fig. 7.3(b). Its broad, symmetric profile eliminates it as a possible dark matter signature.

However, since the strain spectrum presented a multitude of features, a more systematic strategy for identifying the presence of a signal was necessary. The procedure was to look for data points in the PSD above the 95% confidence level threshold (defined as the power level P_{th} so that points where the power is greater than P_{th} have a 5% probability to be caused by statistical fluctuations).

The PSD of the data set was split into superbins containing $N = 10^5$ frequency bins each. The width of the superbins was chosen to be significantly larger than the dark matter linewidth across

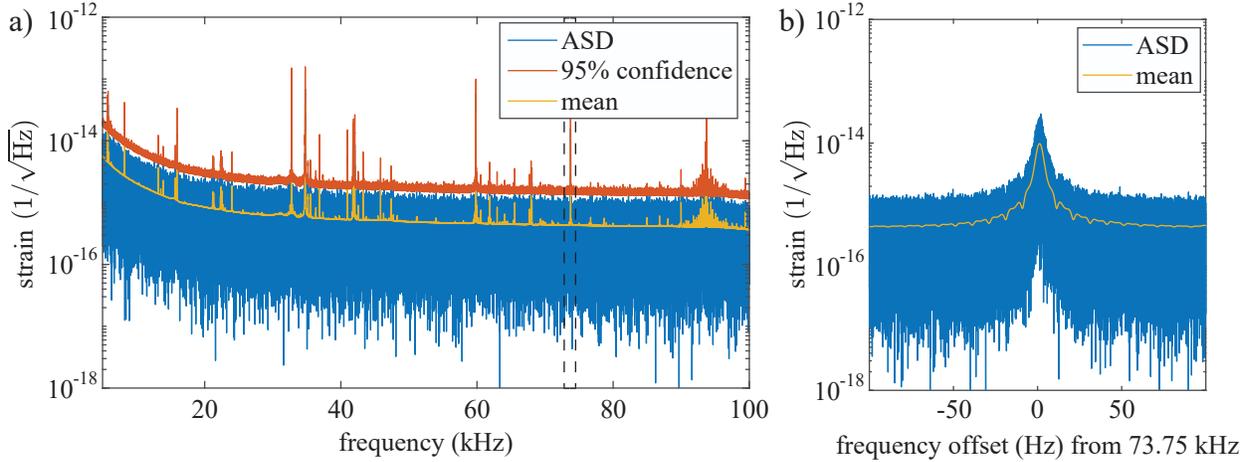


Figure 7.3: (a) Downsampled amplitude spectral density (ASD) of the noise floor, alongside the noise mean value and the 95% confidence level threshold for dark matter detection. The dashed rectangle highlights the peak that is magnified in the adjacent plot; (b) Example of the profile of a noise peak at 73.75 kHz

all frequencies, while still narrow enough to treat the noise within as white. For each superbin, the 95% confidence level threshold was found as the $0.95^{1/N}$ quantile for the frequency bins inside [94]. In other words, we looked for the power spectral density S_{th} , so that the probability that the power spectral density inside a superbin is higher than S_{th} , $\mathcal{P}(S > S_{th}) = 0.95^{1/N}$. The $1/N$ exponent accounts for the effect of searching the peak across N frequencies for each superbin [94].

Based on this criterion, no outliers were identified above the detection threshold. Therefore, we conclude that none of the peaks are statistically significant and they will be treated as part of the noise floor.

The data set, downsampled by a factor of 100,000 for ease of plotting, together with the mean and 95% confidence level threshold are displayed, as amplitude spectral densities, in Fig. 7.3(a).

7.4 Methodology for placing constraints on ultralight dark matter

No peaks indicative of the presence of a ULDM field having been found, we proceed with analyzing the measured noise floor to set limits on a dark matter coupling at 95% confidence level.

The goal of this section is to detail the procedure used to calculate these constraints.

7.4.1 Parameterization in terms of a DM signal

Supposing the detection of a dark matter signal in the frequency data, it is possible to relate it to an underlying ULDM field according to Eq. (1.15):

$$\tilde{s}_p = -(d_e + d_{m_e}) A_p \tilde{\Phi}_p, \quad (7.9)$$

where A_p depends on the discretized apparatus response A_{det} as:

$$A_p = \frac{\sqrt{4\pi\hbar c}}{E_P} A_{det}(f_p), \quad (7.10)$$

and the field Φ_p can be expressed in terms of the dark matter lineshape $F_p = F(\omega_p)$ as

$$\langle |\tilde{\Phi}_p|^2 \rangle = \frac{\pi N}{\Delta T} \Phi_0^2 F_p. \quad (7.11)$$

It is implied by Eq. (7.9) that a coupling of dark matter through the electron mass (through d_{m_e}) is indistinguishable by our apparatus from a coupling through the fine structure constant α (described by d_e). We will consequently adopt the standard approach of setting one parameter to 0 in order to establish an upper bound for the other. In the analysis that follows, we choose to place limits on d_{m_e} , with an understanding that the derived results equally apply to d_e .

Since the dark matter signal and the detector noise are uncorrelated, the PSD $\tilde{\Sigma}_p$ of the frequency domain data \tilde{d}_k (defined in accordance with Eq. (7.6)) can be expressed as the sum of the signal (\tilde{S}_p) and noise ($\tilde{\rho}_p$) power spectral densities:

$$\tilde{\Sigma}_p = \tilde{S}_p + \tilde{\rho}_p = d_{m_e}^2 A_p^2 \frac{\pi N}{\Delta_t} \Phi_0^2 F_p + \tilde{\rho}_p. \quad (7.12)$$

The detector response $A_{det}(f_k)$ is represented for our experiment by the differential transfer func-

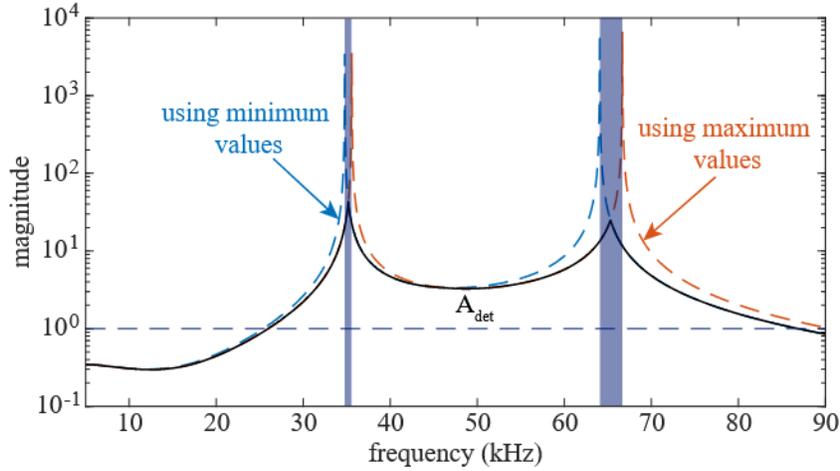


Figure 7.4: Detector response (black), conservatively estimated as the minimum of the transfer functions obtained using lower (blue) and upper (orange) bounds of the mechanical resonances of the cavities. Shaded regions mark the uncertainty intervals for each cavity’s resonance

tion between the two cavities described by Eq. (2.48). We used measured values of the optical poles - 4.7 kHz for the 15 cm cavity and 7.4 kHz for the 7.5 cavity - obtained using the procedure detailed in Appendix A.

As we did not succeed to unambiguously identify the frequencies of the longitudinal resonances of the cavities in our system, we took the conservative approach of computing the transfer function using the extreme values for each resonance, determined through the work presented in chapter 6, and selecting at each frequency the overall minimum of the resulting curves, as illustrated in Fig. 7.4.

7.4.2 Bayesian analysis

In order to determine the 95% confidence level threshold for d_{m_e} we opted for a Bayesian inference approach in the frequency domain [29, 95]. In this framework, one can express the posterior probability of the parameter d_{m_e} given the measured data in terms of the likelihood of measuring a data set $\{\tilde{d}_p\}$ assuming a specific value of d_{m_e} (Bayes’s theorem):

$$\mathcal{P}(d_{m_e}|\{\tilde{d}_p\}) = \frac{\mathcal{P}(d_{m_e})}{\mathcal{P}(\{\tilde{d}_p\})} \mathcal{L}(\{\tilde{d}_p\}|d_{m_e}). \quad (7.13)$$

In the preceding expression, $\mathcal{P}(d_{m_e})$ is the assumed distribution for d_{m_e} , known as the prior, and $\mathcal{P}(\{\tilde{d}_p\})$ is the (unconditional) probability of the results of a measurement being a data set $\{\tilde{d}_p\}$, which can be treated as a normalization constant.

The first step before being able to employ Bayes's theorem is the derivation of the likelihood $\mathcal{L}(\{\tilde{d}_p\}|d_{m_e})$. In general, for stationary, Gaussian noise with mean $\langle n_k \rangle = 0$, the conditional likelihood for measuring a data set $\{\tilde{d}_p\}$ given a model M is represented by a multivariate Gaussian distribution for the noise $\mathbf{n} = \mathbf{d} - \mathbf{s}$ that depends on the autocorrelation matrix \mathbf{C} :

$$\mathcal{L}(\{\tilde{d}_p\}|M) = \frac{1}{\sqrt{\det(2\pi\mathbf{C})}} \exp\left(-\frac{1}{2}\mathbf{n}^\dagger\mathbf{C}^{-1}\mathbf{n}\right). \quad (7.14)$$

Taking into account the relation between autocorrelation and the PSD matrix, the likelihood can be rewritten in terms of the latter [29]:

$$\begin{aligned} \mathcal{L}(\{\tilde{d}_p\}|M) &= \frac{1}{\sqrt{\det(2\pi\mathbf{C})}} \exp\left(-\frac{1}{2}\left(\frac{1}{\sqrt{N}}\tilde{\mathbf{n}}^\dagger\mathbf{U}\right)\left(N\mathbf{U}^\dagger\tilde{\mathbf{C}}^{-1}\mathbf{U}\right)\left(\frac{1}{\sqrt{N}}\mathbf{U}^\dagger\tilde{\mathbf{n}}\right)\right) \\ &\propto \prod_{p=0}^{N-1} \exp\left(-\frac{1}{2}\frac{|\tilde{n}_p|^2}{\tilde{C}_p}\right), \end{aligned} \quad (7.15)$$

where we used both the unitarity of \mathbf{U} and the diagonality of $\tilde{\mathbf{C}}$. The pre-factor can be recovered from the normalization condition that the probability integrates to 1 over the variable space. Then, if we also notice that the values of \tilde{n}_k come in pairs (Eq. (7.3)), we reach the expression [29]:

$$\mathcal{L}(\{\tilde{d}_p\}|M) = \prod_{p=0}^{N/2} \frac{1}{(\beta_p^{-1}\pi\tilde{C}_p)^{\beta_p}} \exp\left(-\beta_p\frac{|\tilde{n}_p|^2}{\tilde{C}_p}\right), \quad (7.16)$$

where $\beta_p = 1/2$ for $p = 0$ and $p = N/2$ and 1 otherwise.

Let us now apply this relation to our data set. In the ensuing analysis we closely follow the reasoning in [29]. Since the Fourier components we look at are far away from the DC and Nyquist frequencies, these components can be neglected. In this case, for a particular possible set of Fourier components that could model the dark matter signal, $\{\tilde{s}_k\}$, and a given value for d_{m_e} , Eq. (7.16) becomes:

$$\mathcal{L}(\{\tilde{d}_p\}|d_{m_e}, \{\tilde{s}_p\}) = \prod_{p=1}^{N/2-1} \frac{1}{\pi \tilde{\rho}_p} \exp\left(-\frac{|\tilde{d}_p - \tilde{s}_p|^2}{\tilde{\rho}_p}\right). \quad (7.17)$$

Moreover, the signal we look for is also Gaussian, which allows us to express its conditional likelihood as:

$$\mathcal{L}(\{\tilde{s}_p\}|d_{m_e}, M) = \prod_{p=1}^{N/2-1} \frac{1}{\pi \tilde{S}_p} \exp\left(-\frac{|\tilde{s}_p|^2}{\tilde{S}_p}\right). \quad (7.18)$$

Since ultimately we are only interested in the posterior probability for d_{m_e} , we would like to eliminate the nuisance parameters \tilde{s}_p by marginalizing over the space of possible sets $\{\tilde{s}_p\}$, which we will denote by $S_{\{\tilde{s}_p\}}$:

$$\mathcal{L}(\{\tilde{d}_p\}|d_{m_e}) = \int_{S_{\{\tilde{s}_p\}}} \mathcal{L}(\{\tilde{d}_p\}|d_{m_e}, \{\tilde{s}_p\}) \mathcal{L}(\{\tilde{s}_p\}|d_{m_e}, M) d\tilde{\mathbf{s}}. \quad (7.19)$$

It is convenient to convert to a coordinate system where the independent variables are the magnitude and phase of the Fourier components $\tilde{s}_p = |\tilde{s}_p|e^{i\theta_p}$. The corresponding volume element is $|\tilde{s}_p|d|\tilde{s}_p|d\theta_p$. The integral in Eq. (7.19) is separable and we can write for one Fourier component:

$$\mathcal{L}_p(\tilde{d}_p|d_{m_e}) = \int_0^\infty \int_0^{2\pi} \frac{1}{\pi \tilde{\rho}_p} \exp\left(-\frac{|\tilde{d}_p - \tilde{s}_p|^2}{\tilde{\rho}_p}\right) \frac{1}{\pi \tilde{S}_p} \exp\left(-\frac{|\tilde{s}_p|^2}{\tilde{S}_p}\right) |\tilde{s}_p|d|\tilde{s}_p|d\theta_p. \quad (7.20)$$

If we also expand $\tilde{s}_p = |\tilde{s}_p|e^{i\varphi_p}$, then:

$$|\tilde{d}_p - \tilde{s}_p|^2 = |\tilde{d}_p|^2 + |\tilde{s}_p|^2 - 2|\tilde{d}_p||\tilde{s}_p| \cos(\theta_p - \varphi_p). \quad (7.21)$$

The integral over the phase can be solved first by noticing that:

$$\int_0^{2\pi} \exp\left(\frac{2|\tilde{d}_p||\tilde{s}_p| \cos(\theta_p - \phi_p)}{\tilde{\rho}_p}\right) d\theta_p = 2\pi I_0\left(\frac{2|\tilde{d}_p||\tilde{s}_p|}{\tilde{\rho}_p}\right), \quad (7.22)$$

where $I_0(x)$ is the modified Bessel function of the first kind. Therefore:

$$\mathcal{L}_p(\tilde{d}_p|d_{m_e}) = \frac{2}{\pi \tilde{\rho}_p \tilde{S}_p} \exp\left(-\frac{|\tilde{d}_p|^2}{\tilde{\rho}_p}\right) \int_0^\infty \exp\left(-\frac{|\tilde{s}_p|^2}{\tilde{\rho}_p} - \frac{|\tilde{s}_p|^2}{\tilde{S}_p}\right) I_0\left(\frac{2|\tilde{d}_p||\tilde{s}_p|}{\tilde{\rho}_p}\right) |\tilde{s}_p| d|\tilde{s}_p|. \quad (7.23)$$

This integral has an analytic solution, since:

$$\int_0^\infty \left(x e^{-Ax^2} I_0(Bx)\right) dx = \frac{e^{B^2/4A}}{2A}. \quad (7.24)$$

Some further simplification, together with Eq. (7.12), yield the final result:

$$\mathcal{L}_p(\tilde{d}_p|d_{m_e}) = \frac{1}{\pi \tilde{\Sigma}_p} \exp\left(-\frac{|\tilde{d}_p|^2}{\tilde{\Sigma}_p}\right), \quad (7.25)$$

Equipped with Eq. (7.25), it is now possible to apply Bayes's theorem. We assume an uninformed prior (a uniform probability distribution for d_{m_e}). In this case, the posterior probability is proportional to $\mathcal{L}(\{\tilde{d}_p\}|d_{m_e}) = \prod_{p=1}^{N/2-1} \mathcal{L}_p(\tilde{d}_p|d_{m_e})$:

$$\mathcal{P}(d_{m_e}|\{\tilde{d}_p\}) \propto \prod_{p=1}^{N/2-1} \frac{1}{\pi \tilde{\Sigma}_p} \exp\left(-\frac{|\tilde{d}_p|^2}{\tilde{\Sigma}_p}\right). \quad (7.26)$$

We deal with the proportionality constant by setting the condition that the probability integrates to 1 over the d_{m_e} space.

The posterior probability having been calculated, placing a bound on the dark matter coupling at the 95% confidence level equates with solving for $d_{m_e}^{95\%}$ the equation:

$$2 \int_0^{d_{m_e}^{95\%}} dd_{m_e} \mathcal{P}(d_{m_e} | \{\tilde{d}_p\}) = 0.95. \quad (7.27)$$

7.4.3 Implementation details

The value of $d_{m_e}^{95\%}$ as defined in Eq. (7.27) was computed over an array of frequencies in the range of interest. In order to ensure that the sampling was done finely enough to probe the ULDM linewidth at every frequency, without sacrificing efficiency, an adaptive grid was implemented, with the separation $\Delta f_{d_{m_e}} = 1/\tau_c$.

At each frequency f_i in the grid, we choose a window $[f_i - 1/\tau_c, f_i + 10/\tau_c]$. We calculate the detector noise model $\tilde{\rho}_i$ as the mean of $|\tilde{d}_p|^2$ over the selected window. Fig. 7.1(b) provides an example of the noise over one such frequency window, with $\tilde{\rho}_i$ displayed as the red dashed line.

The signal PSD \tilde{S}_i and total PSD $\tilde{\Sigma}_i = \tilde{S}_i + \tilde{\rho}_i$ were computed for a grid of values of d_{m_e} , according to Eq. (7.12). Then, for each value of d_{m_e} , the posterior was calculated using Eq. (7.26). In practice, it was easier to compute its logarithm and subsequently exponentiate. Finally, $d_{m_e}^{95\%}$ was obtained by integrating the posterior over the d_{m_e} grid up to 95%.

Since running the algorithm described above on our large data set was computationally expensive, these operations were partially performed on the high-performance computing cluster at Northwestern University (QUEST), where the computation could run on a node with 128 cores.

7.4.4 Injection of a simulated dark matter signal

The analysis algorithm was tested by injecting a digital dark matter signal with a Compton frequency of 40 kHz and $d_{m_e} = 10^6$ into the real data set and verifying that the expected d_{m_e} is extracted back after running the script.

An artificial signal that correctly incorporates the stochastic behavior of dark matter was generated by scaling Gaussian-distributed white noise by the DM characteristic lineshape. The white noise was simulated by drawing random values from a Gaussian distribution with mean 0 and standard deviation 1. The DFT of the time series thus obtained was multiplied by $F(\omega)$. The inverse Fourier transform then represents a series in time domain mimicking a dark matter signal. The result, averaged over 10 periodograms, is shown in the inset of Fig. 7.5(a).

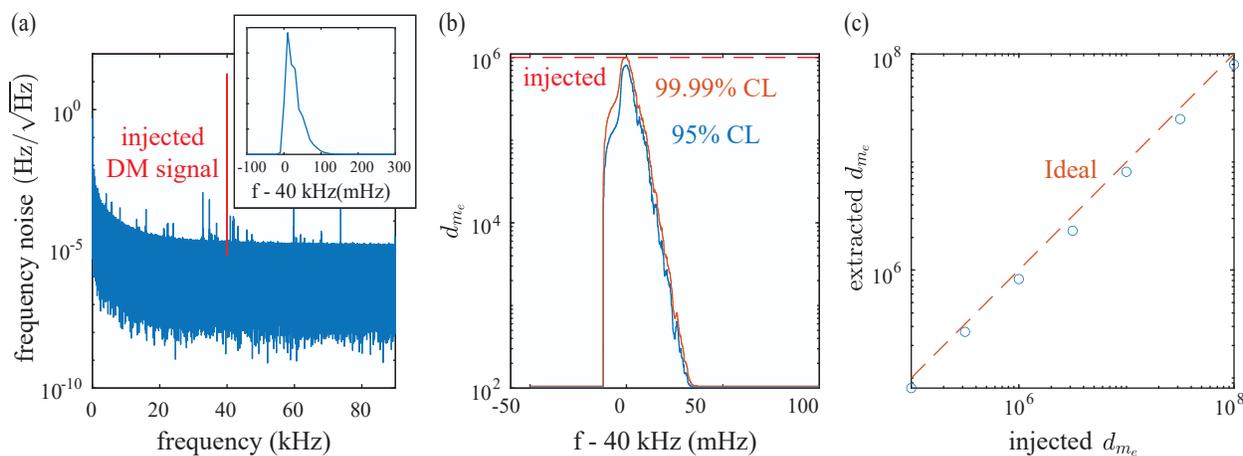


Figure 7.5: (a) Plot of the frequency noise PSD with an added artificial dark matter-like signal (the peak marked with red) corresponding to $d_{m_e} = 10^6$. Top right inset: plot of the simulated dark matter signal, displaying the expected asymmetric lineshape; (b) Recovered dark matter signal after passing the injected data through the analysis algorithm, with 95% and 99.5% confidence level (CL); (c) Scan over a range of d_{m_e} values (computed at the 95% CL), with the dashed line marking the expected values, to which the extracted result is consistently close

The simulated dark matter signal is multiplied by the detector response transfer function and added to a subset of the (time domain) experimental data (the beat note frequency noise) corresponding to a measurement duration of 1000 s.

The FFT of the sum of signal and noise is fed to the d_{m_e} computation algorithm. The extracted dark matter signal, computed over a frequency grid of 5000 points uniformly distributed over a range starting 0.5 Hz below and ending 1 Hz above the injected signal frequency is shown in figure 7.5(b), computed at confidence level 95% (blue) and 99.99% (orange). The results are close to the

expected value $d_{m_e} = 10^6$.

We scanned the value of d_{m_e} corresponding to the injected signal over several orders of magnitude and confirmed that for each data point the extracted d_{m_e} is close to the injected d_{m_e} (figure 7.5(c)).

7.5 Results and discussion

The data analysis procedure described in the previous sections was carried out on the measured strain for two different models for the dark matter distribution: the standard halo model (SHM), the properties of which were described in section 1.2.4, and a relaxion star gravitationally bound to Earth (discussed in section 1.2.5).

For the SHM we assumed a local density $\rho_{\text{DM}} \approx 0.4 \text{ GeV}/\text{cm}^3$, based on astronomical observations [27], and coherence time $\tau_c = \hbar/m_\phi v_{\text{vir}}^2$. In the case of a relaxion star, the local density is frequency dependent and, for the frequency range considered here (20–90 kHz), it ranges from $6.1 \times 10^{10} \text{ GeV}/\text{cm}^3$ to $1.4 \times 10^{13} \text{ GeV}/\text{cm}^3$, a factor of 10^{11} to 10^{13} greater than for the SHM.

Two regimes can be identified, depending on how the measurement time T_{tot} compares to the dark matter coherence time τ_c . In the case where $T_{\text{tot}} \gg \tau_c$ the amplitude and phase values are averaged over the distributions presented in section 1.2.4 and the amplitude takes the value Φ_0 given by Eq. (1.6). This is referred to as the deterministic regime. On the other hand, when $T_{\text{tot}} < \tau_c$ (the stochastic case) the amplitude and phase of the dark matter field must be treated as random variables, drawn, respectively, from Rayleigh and uniform distributions, as outlined in section 1.2.4. The field stochasticity has been shown to cause a reduction in the sensitivity of dark matter by a factor of 3 [96].

For the work described here, the coherence time in the SHM varies from 5 s to 100 s over the frequency range we report on (5-90 kHz), many orders of magnitude smaller than the measurement time $T_{\text{tot}} = 3.4 \times 10^5$, firmly placing the experiment within the deterministic case. However, in

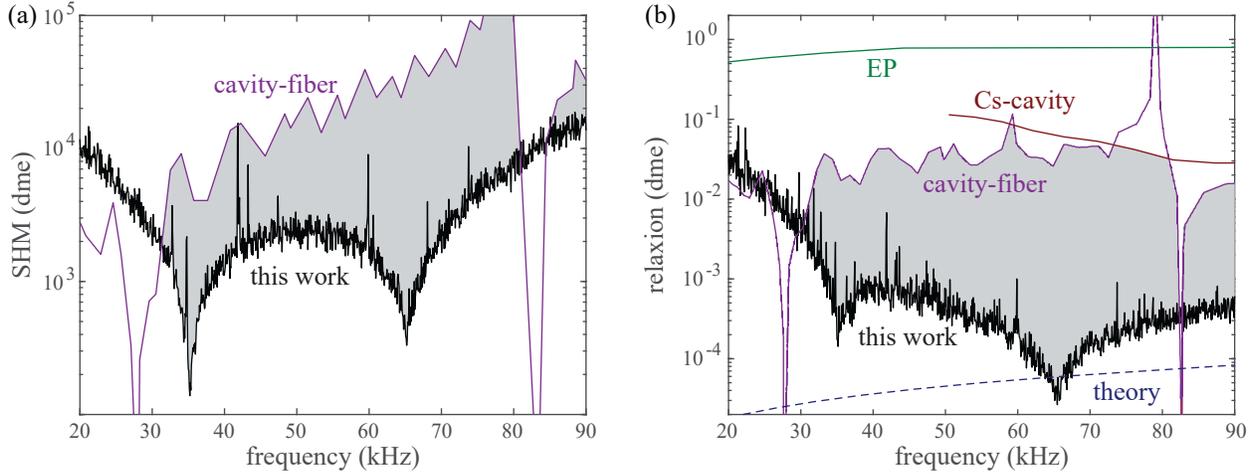


Figure 7.6: New bounds on d_{m_e} in the context of a) the Standard Halo Model and b) a relaxation star bound to Earth. Also shown are limits from two other direct ULDM detectors, cavity-fiber [49] and Cs-cavity [38], as well as indirect limits from EP violating tests [18] and a theoretically motivated target [97]

the relaxation model, coherence times, given by:

$$\tau_\phi \approx 10^3 \text{ s} \left(\frac{10^{-9} \text{ eV}}{m_\phi} \right)^3, \quad (7.28)$$

are much longer, ranging from $\tau_c = 2 \times 10^4 \text{ s}$ to $\tau_c = 2 \times 10^6 \text{ s}$ for Compton frequencies between 20 kHz and 90 kHz, thus spanning both regimes.

An advantage of the Bayesian analysis described in section 7.4.2 is that it incorporates the stochasticity of the scalar field, so it appropriately accounts for both cases.

Fig. 7.6 presents the inferred limits on d_{m_e} at the 95% confidence level for the standard halo and relaxation models in the 20 kHz-90 kHz frequency range. The results improve by one to two orders of magnitude over a majority of this range previous bounds, which came from length comparisons of an optical cavity with itself at a different time using an optical delay line (labeled "cavity-fiber" in Fig. 7.6 and 7.7). Also shown are results coming from comparisons of an optical cavity with an atomic clock ("Cs-cavity" [38]).

The amplitude of the signal searched for is proportional to that of the dark matter field, which

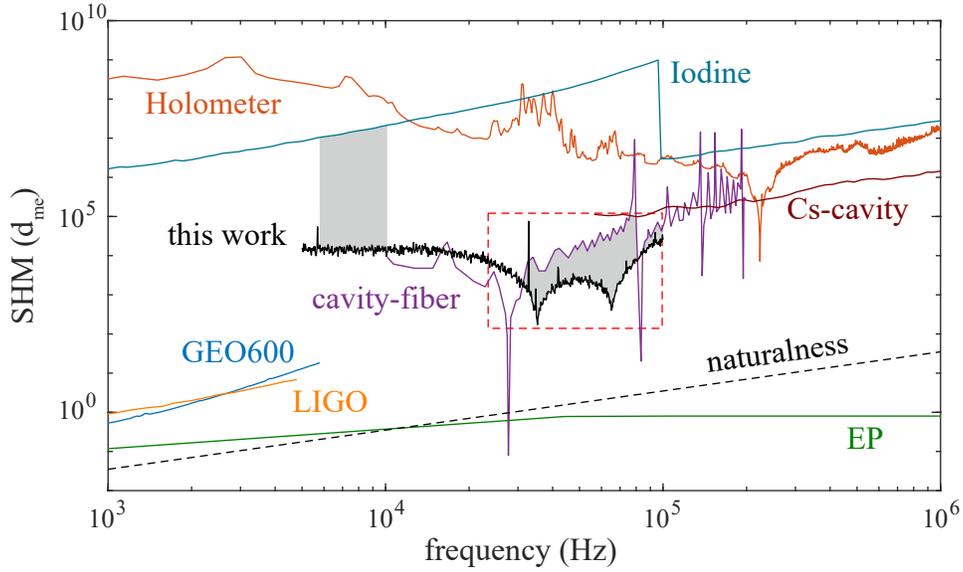


Figure 7.7: Comparison of bounds on d_{m_e} in the SHM over a wider range with other direct and indirect observations and theoretical limits. Details in text

scales as $\sqrt{\rho_{\text{DM}}}$, according to Eq. (1.6). As a consequence of the increase in local dark matter density, the limits we place in the relaxion scenario are 6 orders of magnitude lower than in the SHM. This is an advantage that direct detectors have over indirect detection methods, such as Equivalent Principle (EP) violation tests ([18, 34, 35]), which are insensitive to the local DM density. The enhanced sensitivity of our detector allows us to surpass limits from EP violation by a large margin in the relaxion model. Between 64 kHz and 66 kHz, our bound reaches into a theoretically motivated region for which Higgs-relaxion mixing can occur [97], marked with a dashed line in Fig. 7.6(b). However, the sensitivity starts dropping below 34 kHz, where the measurement time becomes shorter than the coherence time.

Fig. 7.7 presents our results over an extended frequency range, in conjunction with bounds set by other experiments. Other limits shown in this figure, on top of those already described, come from gravitational wave detectors (GEO600 [48] and LIGO [98]), from the Holometer [99] and from molecular Iodine spectroscopy [45]. The line marked as "naturalness" represents a limit imposed by radiative corrections to the mass of scalar particles above which d_{m_e} is excluded without

significant fine-tuning [100].

These first results demonstrate the potential of our apparatus to act as a sensitive ULDM detector. The limits we set exceed by up to two orders of magnitude previous direct bounds for both the SHM and for a relaxion star gravitationally bound to Earth over a majority of the frequency range between 20 kHz and 90 kHz. In addition to the main science band, the work presented here sets new bounds, two to three orders in magnitude lower than previous, in the less explored frequency band between 5 kHz and 10 kHz.

CHAPTER 8

TOWARD THE REALIZATION OF A NON-RIGID OPTICAL CAVITY

Ideally we would compare the variations in resonant frequencies of a cavity with a rigid spacer and a LIGO-style cavity, with its mirrors suspended from two separate pendula. Building and operating the non-rigid cavity, however, poses significant technical challenges. First, both passive vibration attenuation and active position control are required to maintain the mirrors aligned and the cavity length stable. Second, to achieve low thermal noise, it is optimal that the entire bottom stage of the suspension, to which the mirror is attached, as well as the fibers that hold it, are made of the mirror material (in our case sapphire). This eliminates the presence of interfaces between dissimilar materials at the mirror stage and keeps the mirror separated from thermal noise contributions due to higher-loss materials at the upper stages by one layer of isolation. Because of the substantial research involved in addressing these requirements, we instead compared two unequal length rigid cavities for the initial iteration of the experiment. Nevertheless, some steps toward the realization of the non-rigid cavity were taken. In this chapter we summarize these early efforts, which include the construction and characterization of a suspended mirror prototype and the development of actuators to control the position of the mirror. These results and the experience we gathered could constitute a starting point for realizing the non-rigid cavity in future generations.

8.1 Vibration isolation

A first decision that was made was the number of stages required to reduce the level of vibrations below that of thermal and shot noise.

Finite element calculations were performed for a simple one-wire model of a pendulum, to get a sense of the order of magnitude of the vibration attenuation for a single-stage pendulum and

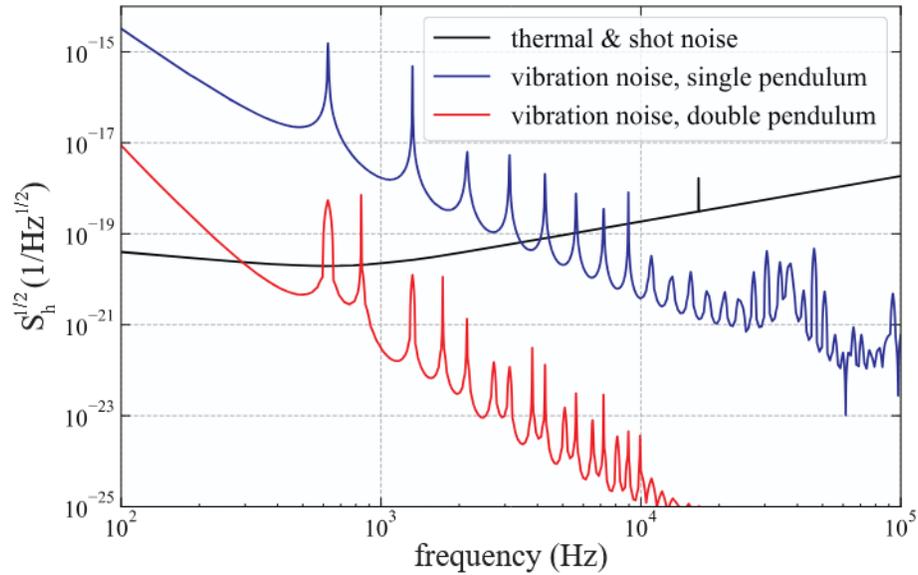


Figure 8.1: Predicted vibration noise for a single-stage and a double-stage pendulum, compared to the expected total sum of thermal and shot noise

for a double pendulum (two stacked pendula). The simulations computed the transfer function between a harmonic displacement of the support and the resulting displacement at the center of the mirror. The obtained transfer function was multiplied by the predicted level of vibrations at the cold plate (corresponding to an acceleration of $10^{-7} g/\sqrt{\text{Hz}}$). Based on the results, which are exhibited in Fig. 8.1, the double pendulum is sufficient to reduce the vibration noise below the sum of thermal and shot noise above 200 Hz (except for a finite number of peaks due to resonances of the structure). While more layers of isolation could extend the science band towards lower frequencies, this comes with the cost of increased complexity of the system. In view of these results, we decided to aim for a double-stage pendulum for this experiment. The estimates in Fig. 8.1 are obtained using the parameters in Table C.1.

8.2 Pendulum oscillation modes and transfer function

To understand the dynamic response of the pendulum, it is necessary to take into account all its relevant resonant modes. In addition to the primary swinging modes, a simple pendulum exhibits a resonance of the wires in the vertical direction, as well as the three rotational modes of the mirrors (often denoted by "pitch", "yaw" and "roll"). Vibrational modes of the suspension wires (the "violin" modes) also affect the motion of the mirror. These modes are illustrated in Fig. 8.2(c). Although analytical approaches to model the behavior of single and multiple-stage pendula exist (see, for instance, [101]), here we opt instead to use finite element modeling to analyze the pendulum resonant modes and transfer functions.

We developed a toy model of an all-sapphire two-stage pendulum (Fig. 8.2(a)), with each stage supported by four 15 cm long fibers, 310 μm in diameter. This thickness lies within the range that offers the best trade-off between vibration isolation (which worsens as the wire diameter increases) and heat conduction (which is facilitated by the use of thicker wires).

The resonant frequencies of the pendulum and its transfer function to a horizontal oscillation in the direction of the laser beam were computed using ANSYS. A horizontal sinusoidal displacement with amplitude $x_0 = 1$ mm was applied to the top support structure, and the displacement of the center of the mirror along the same axis is measured. The ratio of the two, x/x_0 , defines the transfer function. The result is plotted in Fig. 8.2(b). It shows a double peak around 1 Hz (corresponding to the swinging modes of the two pendulum stages), as well as peaks due to the mirror pitch modes and wire violin modes. The transfer function displays the expected f^{-4} behavior over approximately a decade in frequency above the pendulum modes. However, the presence of the pitch modes at a few tens of Hz disrupts this roll-off, and the forest of violin modes above 100 Hz causes a saturation effect in the vibration isolation performance. Nonetheless, about 5 orders of magnitude of vibration attenuation are expected at 100 Hz and over 8 orders of magnitude above 1000 Hz (except at the resonant frequencies).

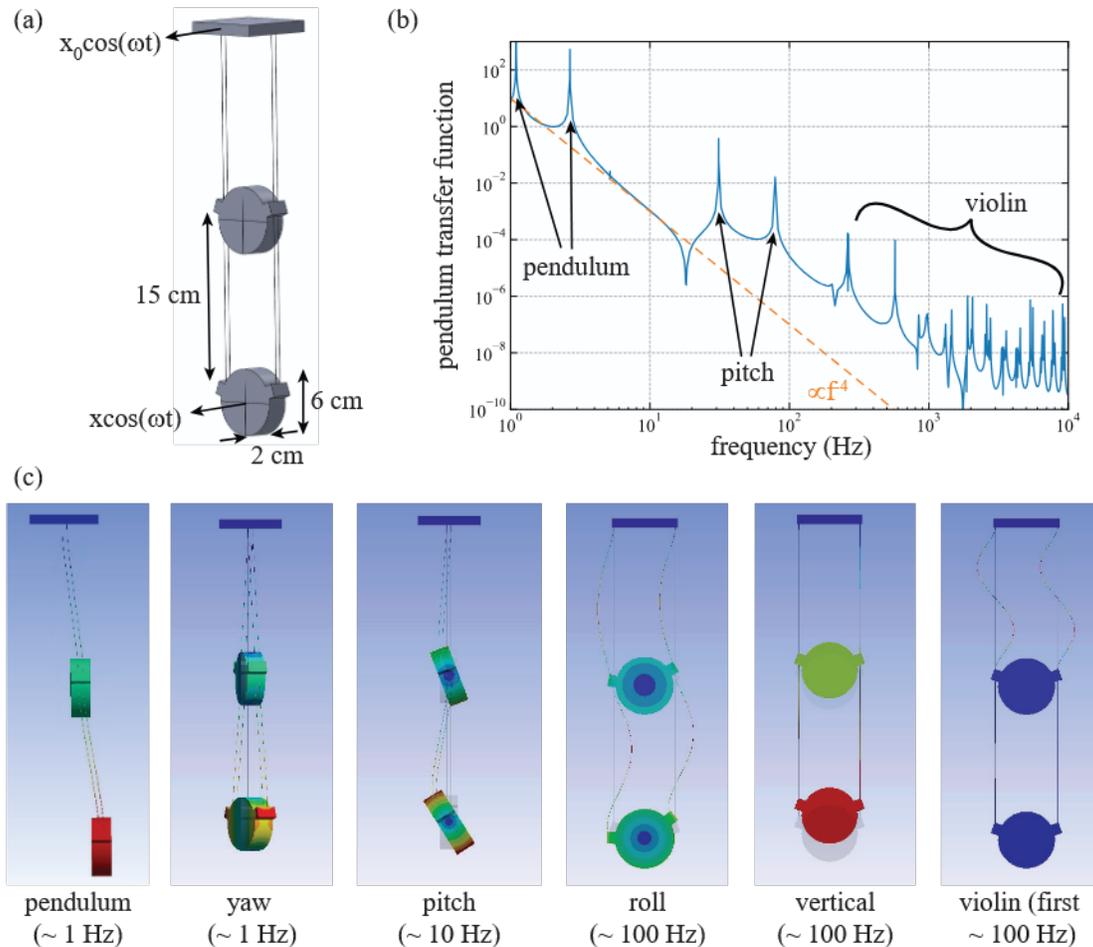


Figure 8.2: a) Toy model for the double pendulum; b) Transfer function for a horizontal displacement; the dashed line indicates the expected slope for a f^{-4} roll-off; c) Pendulum resonant modes and the approximate frequencies where they can be expected to occur

8.3 Development and characterization of a suspended mirror

8.3.1 Pendulum prototype

Several iterations of prototypes for the suspended mirrors have been developed, the latest of which is shown in Fig. 8.3. For ease of fabrication and reduced cost, aluminum was used at the prototyping stage even for the bottom mass, although the final version would be expected to be made of sapphire. Molybdenum wires are used for the suspension. They are attached at the top to

vented screws, which allow for the vertical tuning of the wire attachment points.

The biggest challenge in making a non-rigid cavity is the control of the mirror motion, necessary for alignment and stabilization. Drives from the environment perturb the position of the mirror and excite its mechanical resonances. Therefore, it is key to have a damping mechanism, as well as to be able to control the position and orientation of the mirrors so that they maintain alignment with each other. Moreover, the system the pendulum is damped against needs to be vibration isolated to a similar degree as the main pendulum, otherwise vibrations are introduced into the system by the dampers and motion actuators. For this reason, a second two-stage pendulum is built alongside the main one, with each stage of the pendulum (referred to as recoil masses) enclosing the corresponding stages of the main pendulum (the test mass, which holds the mirror, and the intermediate mass). A schematic representation and pictures of the test mass and recoil mass pair are provided in Fig. 8.3.

In our system we implement both passive and active damping. The former is achieved using two 0.5" diameter magnets on the sides of the bottom recoil mass, which act as eddy current dampers for the test mass. The latter is provided by electro-magnetic actuators, described in more detail in section 8.4.2. These enable control over all degrees of freedom: four actuators at the bottom stage enable motion along the cavity axis, as well as angular motion corresponding to yaw and pitch. Three actuators on the top of the intermediate stage allow for vertical motion and pitch adjustments. Two more actuators on the sides of the intermediate stage provide motion in the two horizontal directions, as well as yaw. The intention is that the intermediate stage actuators will be used for coarse adjustments, while those acting on the test mass will only provide fine adjustments, to keep magnetic noise and heat generation associated with magnetic fields and coil currents to a minimum.

While ultimately the goal is to implement the double-stage suspension described so far, for simplicity we started by testing the bottom stage of the pendulum only. All the work presented in

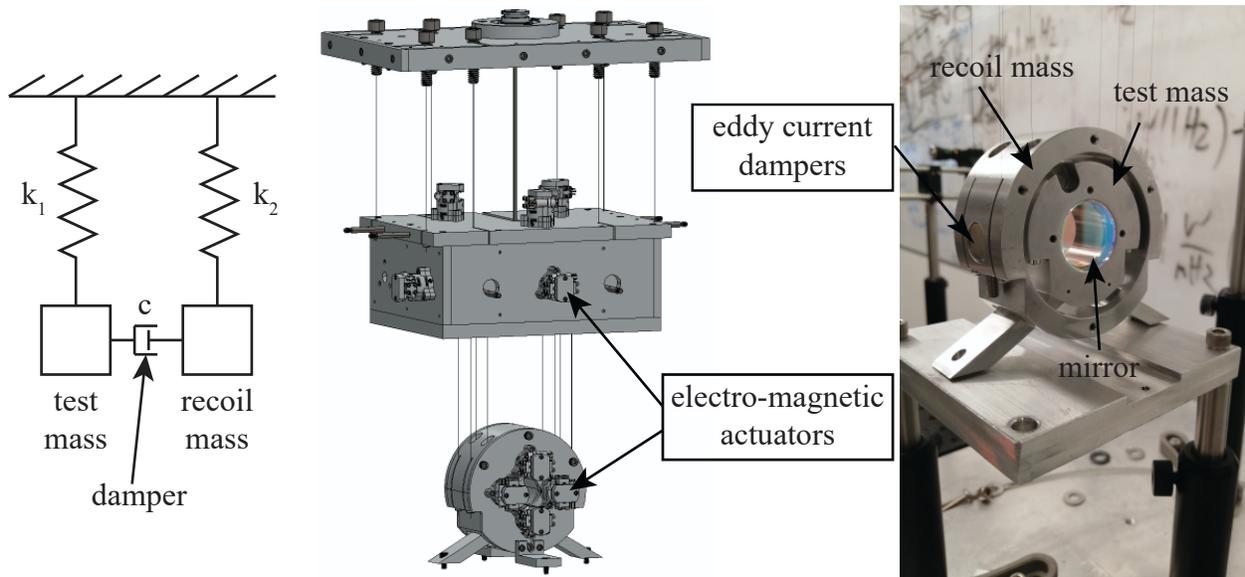


Figure 8.3: Left: diagram of a test mass and a recoil mass with a damping mechanism. Center: CAD model of the double-stage pendulum. Right: view of the test mass and recoil mass with the front cover removed

the following sections was performed using the single-stage pendulum.

8.3.2 Measurements of the pendulum motion

It is informative to know, in view of understanding the operation range needed for the position actuators, how much the pendulum is expected to move in a quiet environment. For this reason, optical measurements of the mirror motion were conducted.

The technique used was an optical lever: a laser beam was directed toward the mirror at an angle and the reflected beam was sent to a quadrant detector. Displacements of the mirror perpendicular to its surface and tilts result in motion of the beam hitting the detector, as illustrated in Fig. 8.4.

For these tests, the pendulum was suspended from a vibration isolation setup consisting of an inverted pendulum and a GAS filter similar to the ones inside the cryostat. The measurements were conducted inside an acrylic enclosure to minimize the effect of environmental perturbations, such as air currents.

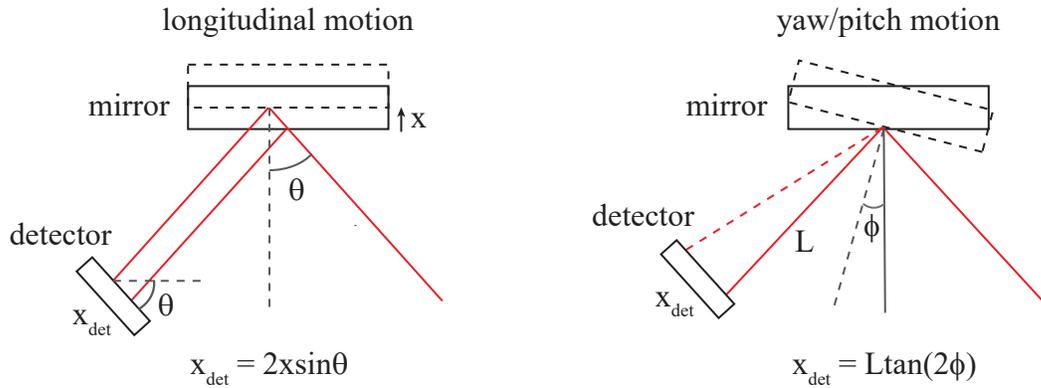


Figure 8.4: Schematic representation of the optical lever, showing the conversion of longitudinal and angular displacements of the mirror to displacements of the beam at the photodetector

Amplitudes of the horizontal motion of the beam reaching the quadrant detector of up to around 10 μm have been observed after the pendulum was left to settle to its most quiet state (an example of a measured trace is provided in Fig. 8.5). Given the distance between the mirror and the photodetector and the incident angle of the beam on the mirror, these translate into 10 μm of pure longitudinal motion or 30 μrad of pure yaw (although likely a combination of the two).

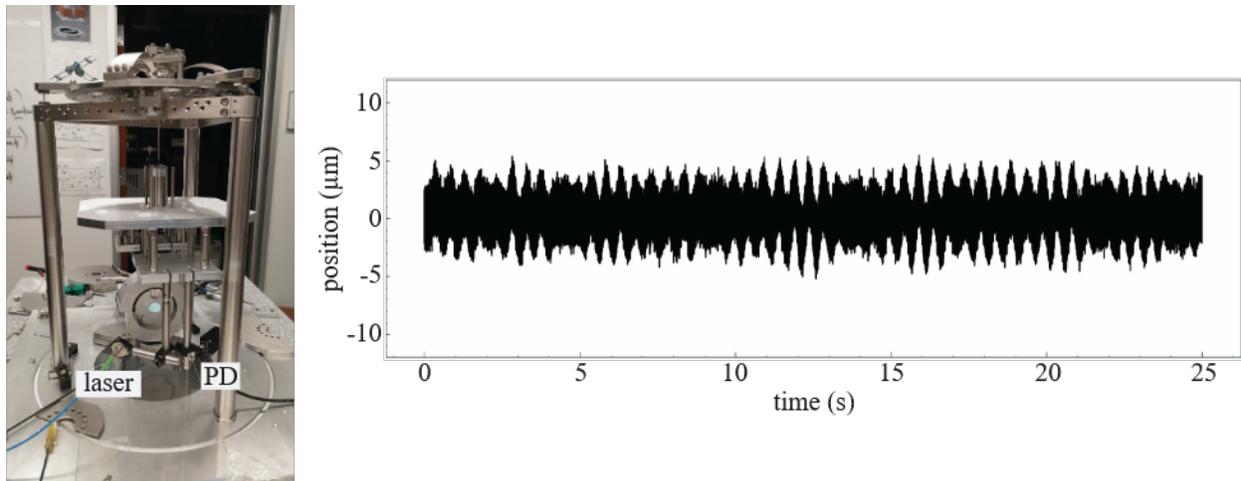


Figure 8.5: Picture of the setup used for measurements of the pendulum motion and example of a measured trace

8.3.3 Cavity with one suspended mirror

To gain some experience with the operation of non-rigid cavities, we also built an optical cavity featuring one suspended mirror. To enable alignment, the second mirror was mounted on a tip-tilt stage fixed to the table. The mirrors had the same radius of curvature as the sapphire mirrors used in the experiment (10.2 m). However, for easier coupling into the cavity, we opted for a lower reflectivity of 97% (equivalent to a cavity finesse of 100).

We coupled light into the TEM₀₀ mode of the cavity. Even without position control of the suspended mirror, coupling to higher-order transverse modes was low, indicating that mirror tilts had minimal impact on the coupling efficiency. However, the signal from a photodetector measuring the cavity transmission (Fig. 8.6) revealed that the longitudinal motion of the pendulum spanned several free spectral ranges, corresponding to oscillations with an amplitude of around 8 μm . This value is consistent with the optical lever measurements. Locking the laser beam frequency was not feasible under these conditions, providing additional confirmation of the critical importance of active position stabilization.

8.4 Active position control

Actuators for the control of the mirror position and orientation are required in order to align the two suspended mirrors so that light can be coupled into the cavity. They are also necessary in order to adjust the length of the cavity, so that the resonance condition for the laser beam is achieved. Moreover, used in closed loop, they compensate for fluctuations in the position of the mirror caused, for instance, by vibrations, and provide damping of the pendulum resonances, for the stable operation of the cavity.

Two types of actuators were considered: piezoelectric and electro-magnetic. Despite some preliminary experiments employing piezo actuators - placed either between the mirror and the test mass or between the intermediate mass and its recoil mass - ultimately we favored the use of

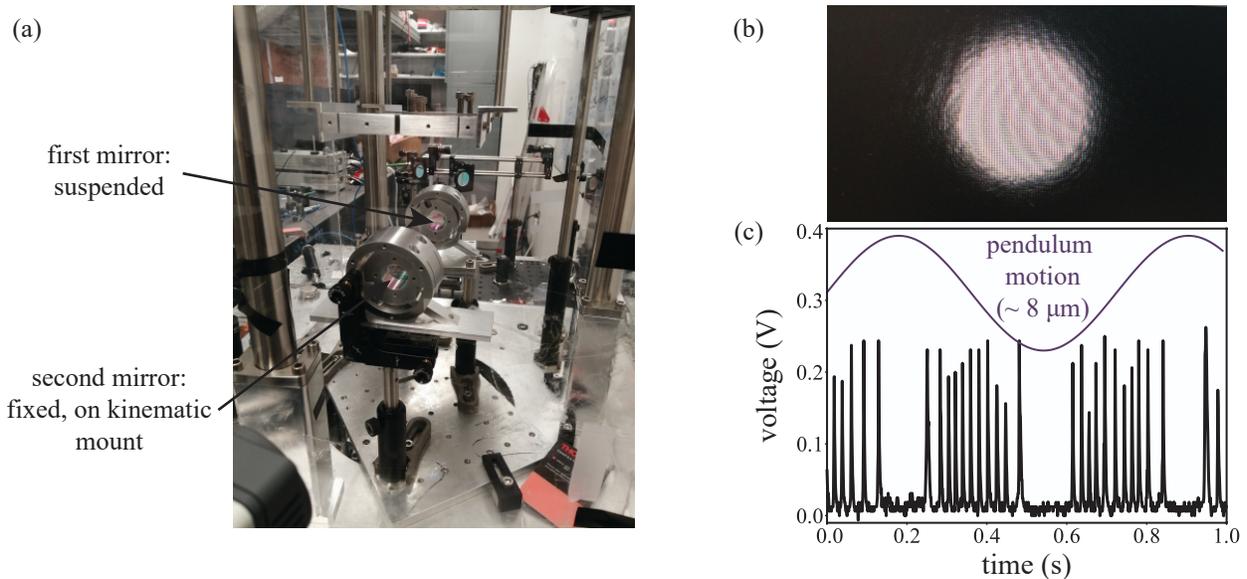


Figure 8.6: Setup of the low-finesse optical cavity with a suspended mirror, as well as a picture of light coupled in the TEM00 mode, detected by an InGaAs camera, and an example of the transmission signal detected by a photodetector

electro-magnetic actuators. This decision was driven by the fact that the latter require no contact between the test mass and recoil mass, and can be made significantly more compact than their piezo counterparts for motion ranges in the tens of μm and above.

8.4.1 Actuation requirements

Certain specifications for the range and resolution of the position actuators need to be met in order to make the laser beam resonant with the cavity and to minimize coupling into higher-order transverse modes due to misalignment of the mirrors with respect to each other.

As far as longitudinal motion is concerned, the condition of being able to adjust the cavity length so that it becomes resonant with the laser wavelength sets a motion range requirement of $0.75 \mu\text{m}$ (half the laser wavelength). However, in order to compensate for the motion of the pendulum due to ambient drives, the actuators need a motion range at least as big as the observed amplitude of pendulum oscillations ($10 \mu\text{m}$).

The required resolution for longitudinal motion is dictated by the need to tune the resonant wavelength to within the operating bandwidth of the acousto-optic modulators used for locking the laser. This corresponds to a resolution of around 100 nm.

Actuators that provide transverse displacements and tilts correct for alignment inaccuracies of the mirrors, which result in power being coupled into higher-order modes. These couplings can be calculated from the mode overlap integrals for a given degree of misalignment [102]. We set a limit on the actuation resolution based on the condition that at most 2% of the laser power is coupled into higher-order modes. For our cavity mirrors, this translates to a transverse displacement resolution of 100 μm and an angular displacement (pitch and yaw) resolution of 10 μrad . It is noteworthy that the demands for transverse displacements are less stringent. This is a consequence of the large radius of curvature of the mirrors used.

The resolution requirements for each degree of freedom are summarized in Table 8.1.

Degree of freedom	Resolution
Longitudinal	100 nm
Transverse	100 μm
Pitch & yaw	10 μrad

Table 8.1: Actuation requirements

8.4.2 Optical Sensors and Electro-Magnetic actuators (OSEMs)

Low-noise electro-magnetic actuation for the active control of suspended mirrors has been developed for gravitational wave detectors, such as LIGO [103] and KAGRA [104]. The ability to make such actuators compatible with high vacuum and, in the case of KAGRA, also cryogenic-compatible, has thus been demonstrated. For this experiment, we built and tested devices that

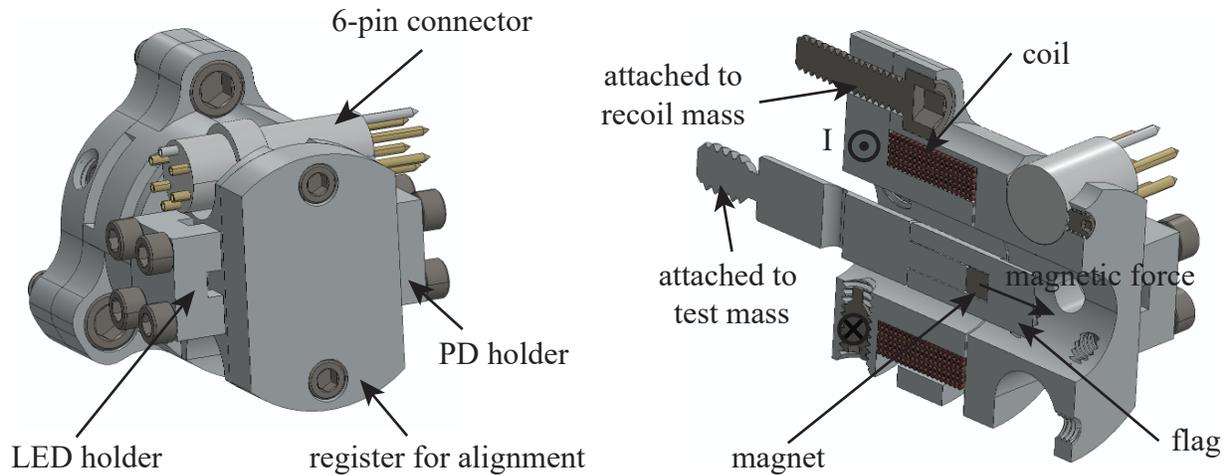


Figure 8.7: Diagram of an OSEM showing the main components

operate on the same principle as those used by LIGO and KAGRA, but miniaturized for the much smaller scale of our cavities.

The Optical Sensors and Electro-Magnetic actuators (OSEMs) are compact devices that integrate a position sensor and an actuator. The sensing system consists of an LED and a photodiode (PD) located on opposite sides of the main body of the OSEM, which is mounted on the recoil mass. A flag attached to the test mass and situated inside the OSEM main body, concentric with it, partially shadows the active area of the photodiode from the light emitted by the LED (see Fig. 8.7). Relative motion of the two pendula changes the surface area of the shadowed region, resulting in a variation in the voltage measured over a resistor in parallel with the PD. The actuator is a copper coil concentrically aligned with the flag. Currents through the coil generate an electro-magnetic force which pushes on a magnet inside the flag. With multiple OSEMs placed appropriately (see Fig. 8.3) motion over all degrees of freedom of the mirror can be measured and controlled.

The coils, made of 85-90 turns of 30 AWG copper arranged in 5 layers, were wound by hand, with each layer glued with epoxy to prevent both unwinding and field fluctuations due to movement of the wire. The PD (OSRAM SFH 2400-Z) and LED (Rohm Semiconductor SCM-013RT)

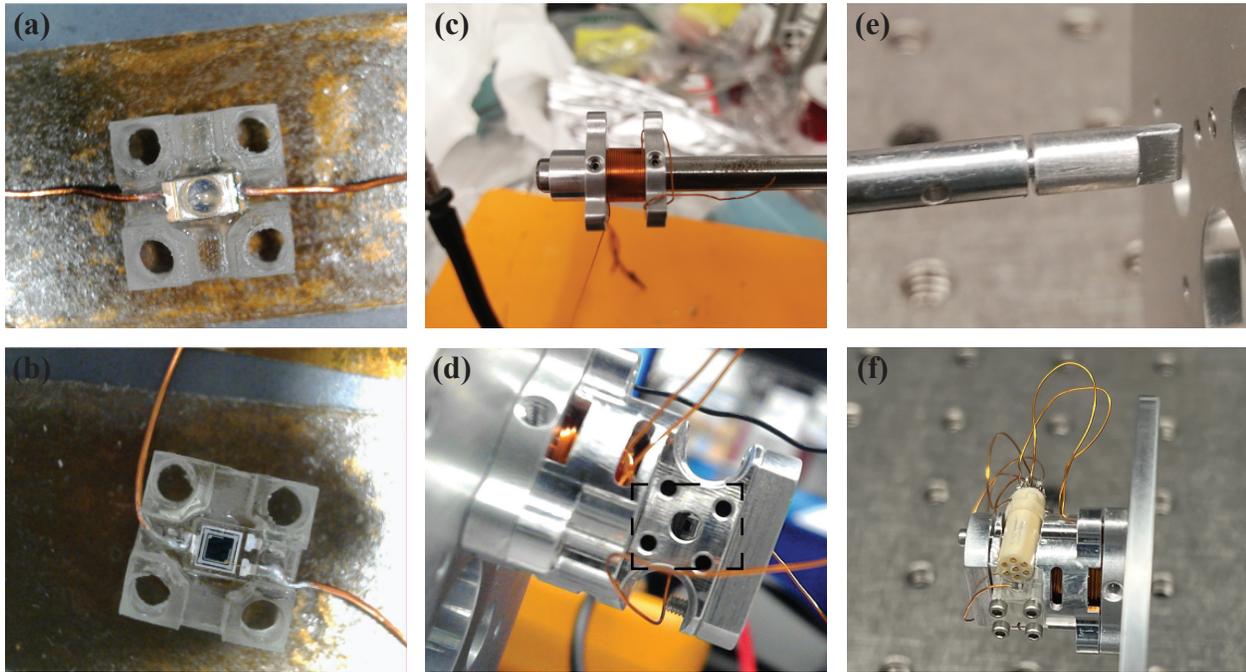


Figure 8.8: Pictures of OSEM components during assembly: (a) LED glued to its holder; (b) photodiode glued to its holder; (c) coil during winding process; (d) photodiode viewed from the LED side; (e) flag with magnet inside; (f) fully assembled OSEM

are glued to holders that screw into the main body of the OSEM. These holders are electrically insulating in order to eliminate the risk of shorts. For speed purposes, we decided to fabricate them using 3D printing for the prototyping stage. However, it was found that because of the insufficient resolution of the printer only about a half of the fabricated parts were adequate for use. In the future, it could be better to use a milling machine to fabricate these parts. A cryogenic-compatible connector (CMR-LM7 from CMR-Direct) is used to connect the leads of the PD and coil. A register at the end of the OSEM is used to align the edge of the flag with the center of the PD holder. This register is removed during operation. The entire device has a length of around 2 cm (0.82”).

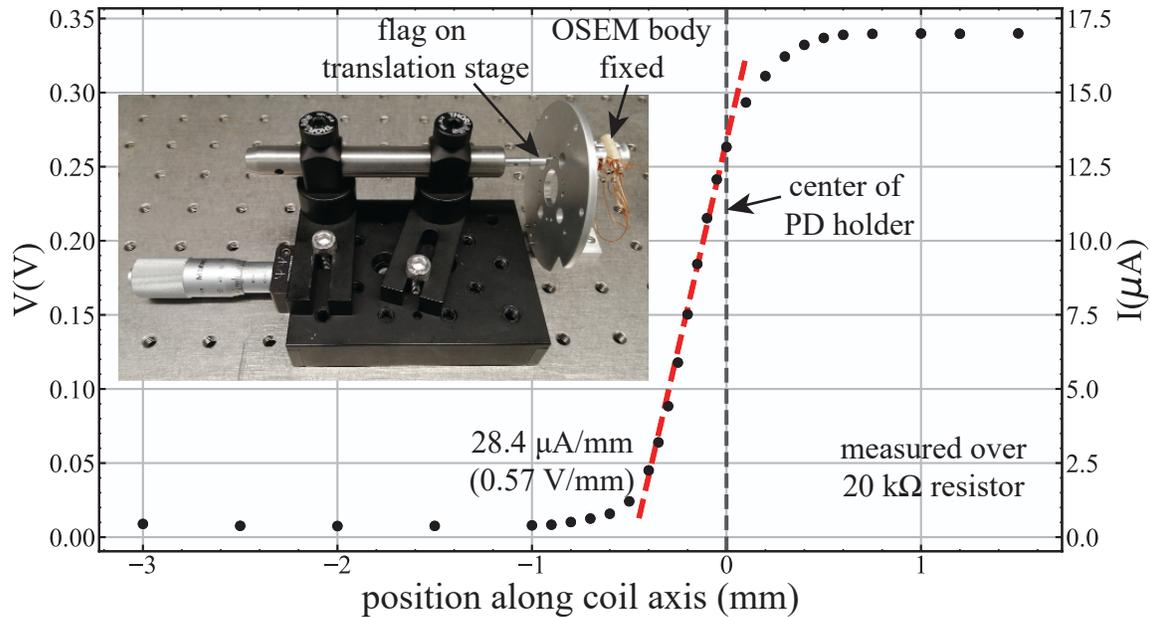


Figure 8.9: Dependence of the PD voltage on the flag position. Inset: setup for the OSEM sensor tests

8.4.3 OSEM characterization

Tests of the sensor and actuator were conducted both to assess their performance and for calibration purposes.

The sensor was tested by measuring the photodiode signal across a $20\text{ k}\Omega$ resistor as the flag, mounted on a translation stage, was moved with respect to the OSEM main body. The test setup can be seen in the inset of Fig. 8.9. The dependence of the measured voltage on the flag position presents a linear region of around 0.4 mm , which determines the OSEM operation range. Given the slope of that region, $28.4\text{ }\mu\text{A}/\text{mm}$, and the typical dark current of the photodiode, 1 nA , the minimum detectable position change is 35 nm . These range and resolution are sufficient according to the requirements considered in section 8.4.1. The voltage against position curve in Fig. 8.9 serves as calibration for interpreting the sensor signal.

The actuator test was performed with the OSEM on the pendulum (Fig. 8.10) and involved

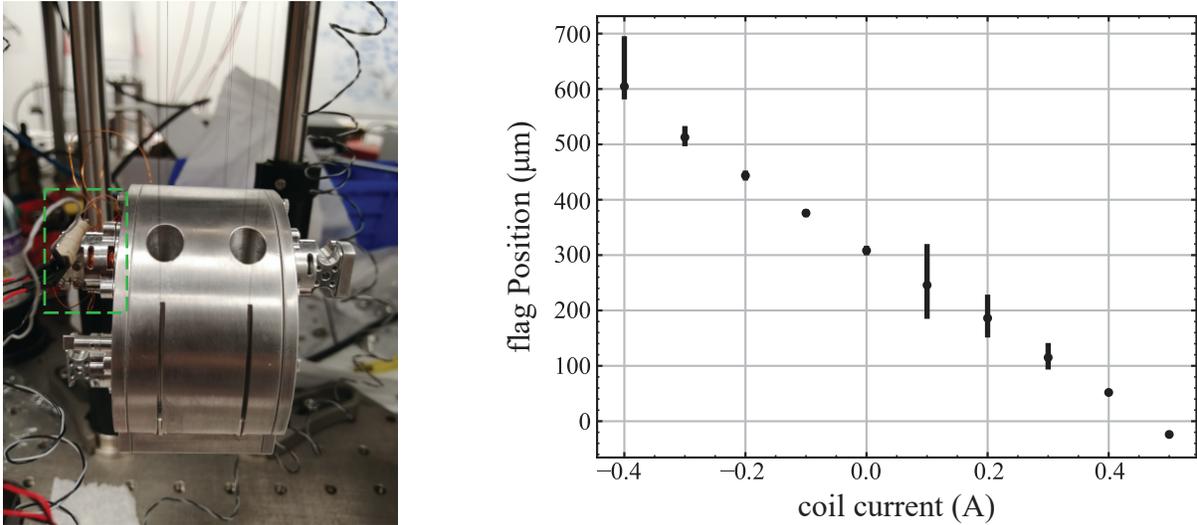


Figure 8.10: Left: picture of the OSEM on the pendulum during the actuation tests (the top OSEM, which was used for the test, is highlighted in green). Right: plot of the measured flag position against the coil current

measuring the flag position using the OSEM photodiode output for different values of the DC current going through the coil. The pendulum position motion varies linearly with the coil current, which is important for close-loop operation. The error bars represent the rms value of the measured position. They become larger at the lower limit of the current as the flag position nears the limit of the sensitivity range of the sensors. The large error bars of the data points corresponding to currents of 0.1-0.3 A are the result of an accidental excitation of the pendulum.

8.4.4 First steps towards active control

We end this chapter by reporting on some preliminary attempts at active control of the mirror. To start with, a feedback loop was implemented for one OSEM only (the top OSEM shown in Fig. 8.10). The OSEM PD signal is used as the input of the PID loop, the output of which is sent to the coil. While the active control reduces the motion of the flag, it only does so to an amplitude of around 5 μm. Moreover, a simultaneous measurement using the optical lever (the

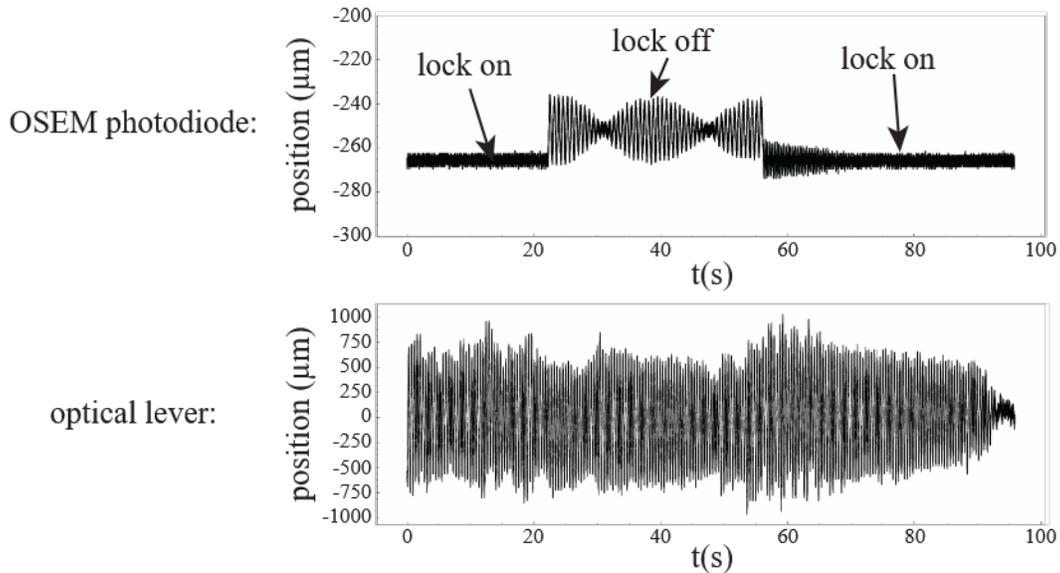


Figure 8.11: Output of the OSEM photodiode and of the optical lever quadrant detector as the PID loop for stabilizing the pendulum is turned off and on

bottom of Fig. 8.11) shows no reduction in the motion of the beam reaching the quadrant detector. We interpret this as a result of the interplay of more than one degree of freedom in the pendulum motion. In particular, this is likely to be the effect of the yaw component, which the OSEM cannot correct for (and which gets amplified by the optical lever by a factor equal to the lever arm, as shown in Fig. 8.4).

This hypothesis was confirmed through the analysis in frequency domain of traces measured with the active control on and off (Fig. 8.12). While the provided feedback was successful in significantly reducing the longitudinal oscillations of the pendulum (the peak at 1.4 Hz), it excited the yaw mode at 1.9 Hz.

The next step, therefore, is to establish control loops to simultaneously target multiple modes of motion of the pendulum. The development of this procedure constitutes the subject of further research.

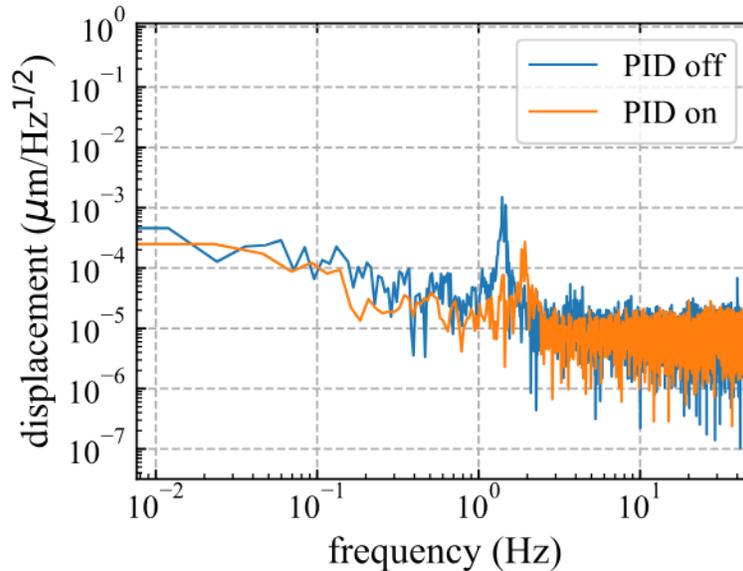


Figure 8.12: Comparison of the amplitude spectral density of the flag displacement with the pendulum active control on and off

8.5 Challenges associated with the non-rigid cavity

A few milestones need to be reached before a non-rigid cavity can be built.

First, as discussed in the previous section, simultaneous active control of all the pendulum degrees of freedom is required. This is not a trivial endeavor because of the coupling between these degrees of freedom. If independent feedback loops are implemented for each OSEM, with the PD signal as input for the PID loop and the corresponding coil current as output, cross-coupling leads to instability. The solution is to have each control loop correspond to a mode of motion of the pendulum. A diagonalization step between the OSEM output signals and the physical coordinates of the pendulum is then required.

Second, the intermediate stage of the pendulum needs to be added. The procedures developed for the alignment and stabilization of the single pendulum will have to be extended to the double pendulum.

Third, as the project transitions from the aluminum prototype to the sapphire pendulum, some

fabrication techniques have to be established. Hydroxide catalysis bonding [105], a method that forms strong chemical bonds between similar flat surfaces using a hydroxide solution, can be used to bond mirrors to the pendulum mass. Indium and gallium bonding have been used to attach the sapphire fibers of the KAGRA gravitational wave interferometer [106], a method that we could borrow. These methods will need to be explored and experimented with.

CHAPTER 9

CONCLUSIONS

This thesis reports on the first results from a new broadband ultralight scalar field dark matter detector. The new method compares variations in the lengths of two Fabry-Pérot cavities that are simultaneously probed by measuring the beat note between light from a common laser resonant to them. The design and construction of the apparatus, as well as the detection and the analysis of its first data run are reported, along with future prospects after apparatus upgrades.

Substantial technical challenges had to be overcome to lower the noise floor to a level where new regions of the dark matter parameter space could be explored. A new cryogenic apparatus, carefully designed to reduce thermal, shot, and vibration noise in the system, was built. Custom sapphire optical cavities with low-loss crystalline coatings achieved high stability at cryogenic temperatures. They were suspended using multiple stages of passive isolation from environmental disturbances side by side inside a 4 K custom cryostat, which reduced the thermal noise nearly tenfold.

Two of my most significant research contributions were designing the sapphire cavity spacers and setting up the room temperature optics and detection electronics. The first involved extensive finite element simulations conducted to minimize the sensitivity of the cavities to vibrations. The second consisted of assembling the optics for locking and generating the beat note, as well as setting up and testing the delay line discriminator that canceled the beat note carrier. It also included the integration of the room-temperature ULE cavity, which filtered the laser to achieve a factor of 3–5 reduction in frequency noise.

The apparatus was used to collect data over a period of 4 days (3.4×10^5 s). Although no evidence of dark matter was detected in the beat note spectrum, Bayesian analysis of the data

places new upper limits on the coupling of scalar field ULDM to ordinary matter. Two different models were assumed: the Standard Halo Model and a relaxion star bound to the Earth. In both cases, the new limits exceed previous bounds from direct detection experiments by up to two orders of magnitude. Additionally, new direct bounds are set in the less explored region of the frequency space between 5 kHz and 10 kHz in the Standard Halo Model.

The sensitivity of the apparatus is currently limited by technical noise, which is about two orders of magnitude above the estimated level of fundamental (shot and thermal) noise. Further research is necessary to conclusively identify the bottleneck. A thorough investigation of the nature of technical noise in the apparatus is thus an immediate next step for the experiment. We suspect, however, that the excess noise could be residual amplitude and/or frequency noise on the laser. In this case, potential avenues for improvement include implementing additional filtering stages or amplitude stabilization. Up to a hundredfold improvement over the current results could be achieved after these upgrades.

Two future paths to increase the sensitivity of the apparatus by several additional orders of magnitude are examined. The first approach proposes the use of longer optical cavities (up to 1 m in length). The second entails comparing the lengths of a cavity with a rigid spacer and one comprised of two suspended mirrors. First steps toward the building of such a cavity, including tests of a prototype pendulum and the development and testing of sensors and actuators for position control, rendered promising results. Both scenarios have the potential to lower the level of shot and thermal noise by approximately two orders of magnitude, while also extending the science band down to a few hundreds of Hz. These upgrades would enable surpassing existing indirect experimental bounds from EP violation tests, and explore regions within the natural parameter space.

REFERENCES

- [1] Particle Data Group Collaboration, S. Navas *et al.*, “Review of particle physics,” *Phys. Rev. D*, vol. 110, p. 030001, 3 2024.
- [2] D. Harvey, R. Massey, T. Kitching, A. Taylor, and E. Tittley, “The nongravitational interactions of dark matter in colliding galaxy clusters,” *Science*, vol. 347, no. 6229, pp. 1462–1465, 2015.
- [3] F. Zwicky, “On the Masses of Nebulae and of Clusters of Nebulae,” *The Astrophysical Journal*, vol. 86, p. 217, Oct. 1937.
- [4] V. C. Rubin and W. K. Ford Jr., “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions,” *The Astrophysical Journal*, vol. 159, p. 379, Feb. 1970.
- [5] V. C. Rubin, W. K. Ford Jr., and N. Thonnard, “Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa \rightarrow Sc.,” *The Astrophysical Journal*, vol. 225, pp. L107–L111, Nov. 1978.
- [6] D. Clowe *et al.*, “A Direct Empirical Proof of the Existence of Dark Matter,” *The Astrophysical Journal*, vol. 648, no. 2, pp. L109–L113, Sep. 2006.
- [7] S. Burles, K. M. Nollett, and M. S. Turner, “Big Bang Nucleosynthesis Predictions for Precision Cosmology,” *The Astrophysical Journal*, vol. 552, no. 1, pp. L1–L5, May 2001.
- [8] M. Milgrom, “A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis.,” *The Astrophysical Journal*, vol. 270, pp. 365–370, Jul. 1983.
- [9] M. A. Monroy-Rodríguez and C. Allen, “The end of the macho era, revisited: New limits on macho masses from halo wide binaries,” *The Astrophysical journal*, vol. 790, no. 2, pp. 159–, 2014, ISSN: 0004-637X.
- [10] B. J. Carr and S. W. Hawking, “Black holes in the early universe,” *Monthly Notices of the Royal Astronomical Society*, vol. 168, no. 2, pp. 399–415, Aug. 1974, ISSN: 0035-8711.
- [11] B. Carr and F. Kühnel, “Primordial black holes as dark matter: Recent developments,” *Annual Review of Nuclear and Particle Science*, vol. 70, no. Volume 70, 2020, pp. 355–394, 2020, ISSN: 1545-4134.

- [12] D. M. Jacobs, G. D. Starkman, and B. W. Lynn, “Macro dark matter,” *Monthly Notices of the Royal Astronomical Society*, vol. 450, no. 4, pp. 3418–3430, May 2015, ISSN: 0035-8711.
- [13] J. L. Feng, “Dark Matter Candidates from Particle Physics and Methods of Detection,” *Ann. Rev. Astron. Astrophys.*, vol. 48, pp. 495–545, 2010.
- [14] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter,” *Physics Reports*, vol. 267, no. 5, pp. 195–373, 1996, ISSN: 0370-1573.
- [15] S. Dimopoulos, S. Raby, and F. Wilczek, “Supersymmetry and the scale of unification,” *Phys. Rev. D*, vol. 24, pp. 1681–1683, 6 1981.
- [16] E. Aprile *et al.*, “First Dark Matter Search Results from the XENON1T Experiment,” *Phys. Rev. Lett.*, vol. 119, p. 181 301, 18 2017.
- [17] R. Agnese *et al.*, “Results from the super cryogenic dark matter search experiment at soudan,” *Phys. Rev. Lett.*, vol. 120, p. 061 802, 6 2018.
- [18] D Antypas *et al.*, “New horizons: Scalar and vector ultralight dark matter,” *arXiv preprint arXiv:2203.14915*, 2022.
- [19] E. G. M. Ferreira, “Ultra-light dark matter,” *The Astronomy and Astrophysics Review*, vol. 29, no. 7, 2021.
- [20] D. F. J. Kimball and K. van Bibber, *The Search for Ultralight Bosonic Dark Matter*. Springer, 2022.
- [21] S. Tremaine and J. E. Gunn, “Dynamical role of light neutral leptons in cosmology,” *Phys. Rev. Lett.*, vol. 42, pp. 407–410, 6 1979.
- [22] J. E. Kim and G. Carosi, “Axions and the strong CP problem,” *Rev. Mod. Phys.*, vol. 82, pp. 557–601, 1 2010.
- [23] J. Preskill, M. B. Wise, and F. Wilczek, “Cosmology of the invisible axion,” *Physics Letters B*, vol. 120, no. 1, pp. 127–132, 1983, ISSN: 0370-2693.
- [24] T. Damour and J. F. Donoghue, “Equivalence principle violations and couplings of a light dilaton,” *Phys. Rev. D*, vol. 82, no. 8, p. 084 033, 2010.
- [25] A. Arvanitaki, J. Huang, and K. Van Tilburg, “Searching for dilaton dark matter with atomic clocks,” *Phys. Rev. D*, vol. 91, p. 015 015, 1 2015.

- [26] Y. V. Stadnik and V. V. Flambaum, “Searching for dark matter and variation of fundamental constants with laser and maser interferometry,” *Phys. Rev. Lett.*, vol. 114, p. 161 301, 16 2015.
- [27] P. F. de Salas and A. Widmark, “Dark matter local density determination: recent observations and future prospects,” *Rept. Prog. Phys.*, vol. 84, no. 10, p. 104 901, 2021.
- [28] K. Freese, M. Lisanti, and C. Savage, “Colloquium: Annual modulation of dark matter,” *Rev. Mod. Phys.*, vol. 85, no. 4, pp. 1561–1581, 2013.
- [29] A. Derevianko, “Detecting dark-matter waves with a network of precision-measurement tools,” *Phys. Rev. A*, vol. 97, no. 4, p. 042 506, Apr. 2018.
- [30] J. W. Foster, N. L. Rodd, and B. R. Safdi, “Revealing the dark matter halo with axion direct detection,” *Phys. Rev. D*, vol. 97, no. 12, p. 123 006, 2018.
- [31] P. W. Graham, D. E. Kaplan, and S. Rajendran, “Cosmological relaxation of the electroweak scale,” *Phys. Rev. Lett.*, vol. 115, no. 22, p. 221 801, 2015.
- [32] A. Banerjee, H. Kim, and G. Perez, “Coherent relaxion dark matter,” *Phys. Rev. D*, vol. 100, no. 11, p. 115 026, 2019.
- [33] A. Banerjee, D. Budker, J. Eby, H. Kim, and G. Perez, “Relaxion stars and their detection via atomic physics,” *Communications Physics*, vol. 3, no. 1, 2020.
- [34] S. Schlamminger, K.-Y. Choi, T. A. Wagner, J. H. Gundlach, and E. G. Adelberger, “Test of the equivalence principle using a rotating torsion balance,” *Phys. Rev. Lett.*, vol. 100, p. 041 101, 4 2008.
- [35] T. A. Wagner, S. Schlamminger, J. H. Gundlach, and E. G. Adelberger, “Torsion-balance tests of the weak equivalence principle,” *Classical and Quantum Gravity*, vol. 29, no. 18, p. 184 002, 2012.
- [36] J. Bergé, P. Brax, G. Métris, M. Pernot-Borràs, P. Touboul, and J.-P. Uzan, “Microscope mission: First constraints on the violation of the weak equivalence principle by a light scalar dilaton,” *Phys. Rev. Lett.*, vol. 120, p. 141 101, 14 2018.
- [37] A. Hees, J. Guéna, M. Abgrall, S. Bize, and P. Wolf, “Searching for an oscillating massive scalar field as a dark matter candidate using atomic hyperfine frequency comparisons,” *Phys. Rev. Lett.*, vol. 117, p. 061 301, 6 2016.

- [38] O. Tretiak *et al.*, “Improved bounds on ultralight scalar dark matter in the radio-frequency range,” *Phys. Rev. Lett.*, vol. 129, p. 031 301, 3 2022.
- [39] S. Aharony, N. Akerman, R. Ozeri, G. Perez, I. Savoray, and R. Shaniv, “Constraining rapidly oscillating scalar dark matter using dynamic decoupling,” *Phys. Rev. D*, vol. 103, p. 075 017, 7 2021.
- [40] X. Zhang *et al.*, “Search for ultralight dark matter with spectroscopy of radio-frequency atomic transitions,” *Phys. Rev. Lett.*, vol. 130, p. 251 002, 25 2023.
- [41] M. Filzinger *et al.*, “Ultralight dark matter search with space-time separated atomic clocks and cavities,” *Phys. Rev. Lett.*, vol. 134, p. 031 001, 3 2025.
- [42] I. Kozyryev, Z. Lasner, and J. M. Doyle, “Enhanced sensitivity to ultralight bosonic dark matter in the spectra of the linear radical sroh,” *Phys. Rev. A*, vol. 103, p. 043 313, 4 2021.
- [43] C. J. Kennedy *et al.*, “Precision Metrology Meets Cosmology: Improved Constraints on Ultralight Dark Matter from Atom-Cavity Frequency Comparisons,” *Phys. Rev. Lett.*, vol. 125, p. 201 302, 20 2020.
- [44] W. M. Campbell, B. T. McAllister, M. Goryachev, E. N. Ivanov, and M. E. Tobar, “Searching for Scalar Dark Matter via Coupling to Fundamental Constants with Photonic, Atomic, and Mechanical Oscillators,” *Phys. Rev. Lett.*, vol. 126, p. 071 301, 7 2021.
- [45] R. Oswald *et al.*, “Search for dark-matter-induced oscillations of fundamental constants using molecular spectroscopy,” *Physical Review Letters*, vol. 129, no. 3, p. 031 302, 2022.
- [46] M. Abe, P. Adamson, M. Borcean, D. Bortoletto, K. Bridges, S. P. Carman, *et al.*, “Matter-wave atomic gradiometer interferometric sensor (magis-100),” *Quantum Science and Technology*, vol. 6, no. 4, p. 044 003, 2021.
- [47] L. Aiello *et al.*, “Constraints on Scalar Field Dark Matter from Colocated Michelson Interferometers,” *Phys. Rev. Lett.*, vol. 128, no. 12, p. 121 101, 2022.
- [48] S. M. Vermeulen *et al.*, “Direct limits for scalar field dark matter from a gravitational-wave detector,” *Nature*, vol. 600, no. 7889, pp. 424–428, 2021.
- [49] E. Savalle *et al.*, “Searching for dark matter with an optical cavity and an unequal-delay interferometer,” *Phys. Rev. Lett.*, vol. 126, no. 5, p. 051 301, 2021.
- [50] A. Branca *et al.*, “Search for an Ultralight Scalar Dark Matter Candidate with the AURIGA Detector,” *Phys. Rev. Lett.*, vol. 118, p. 021 302, 2 2017.

- [51] J. Manley, D. J. Wilson, R. Stump, D. Grin, and S. Singh, “Searching for Scalar Dark Matter with Compact Mechanical Resonators,” *Phys. Rev. Lett.*, vol. 124, p. 151 301, 15 2020.
- [52] R. V. Pound, “Electronic frequency stabilization of microwave oscillators,” *Review of Scientific Instruments*, vol. 17, no. 11, pp. 490–505, Nov. 1946, ISSN: 0034-6748.
- [53] R. W. P. Drever *et al.*, “Laser phase and frequency stabilization using an optical resonator,” *Appl. Phys. B*, vol. 31, pp. 97–105, Jun. 1983.
- [54] E. D. Black, “An introduction to Pound–Drever–Hall laser frequency stabilization,” *American Journal of Physics*, vol. 69, no. 1, pp. 79–87, Jan. 2001, ISSN: 0002-9505.
- [55] M. Rakhmanov, “Dynamics of laser interferometric gravitational wave detectors,” Ph.D. dissertation, California Institute of Technology, Nov. 2000.
- [56] K. Izumi, *Cavity response in transmission*, Japanese Gravitational Wave Document Database, JGW-T1201121-v2, 2012.
- [57] A. Arvanitaki, S. Dimopoulos, and K. Van Tilburg, “Sound of dark matter: Searching for light scalars with resonant-mass detectors,” *Phys. Rev. Lett.*, vol. 116, p. 031 102, 3 2016.
- [58] E. Savalle, “Téster la relativité générale avec des horloges dans l’espace, et explorer les possibilités de détection de matière noire avec des atomes froids dans l’espace et au sol,” Ph.D. dissertation, Observatoire de Paris - SYRTE, Nov. 2020.
- [59] H. B. Callen and T. A. Welton, “Irreversibility and generalized noise,” *Phys. Rev.*, vol. 83, pp. 34–40, 1 1951.
- [60] H. B. Callen and R. F. Greene, “On a theorem of irreversible thermodynamics,” *Phys. Rev.*, vol. 86, pp. 702–710, 5 1952.
- [61] Y. Levin, “Internal thermal noise in the LIGO test masses: A direct approach,” *Phys. Rev. D*, vol. 57, pp. 659–663, 2 1998.
- [62] T. Chalermongsak, F. Seifert, E. D. Hall, K. Arai, E. K. Gustafson, and R. X. Adhikari, “Broadband measurement of coating thermal noise in rigid Fabry–Pérot cavities,” *Metrologia*, vol. 52, no. 1, p. 17, 2014.
- [63] G. M. Harry *et al.*, “Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings,” *Class. Quant. Grav.*, vol. 19, pp. 897–918, 2002.

- [64] T. Chalermsoongsak *et al.*, “Coherent cancellation of photothermal noise in GaAs/Al_{0.92}Ga_{0.08}As Bragg mirrors,” *Metrologia*, vol. 53, no. 2, p. 860, 2016.
- [65] P. R. Saulson, “Thermal noise in mechanical experiments,” *Phys. Rev. D*, vol. 42, pp. 2437–2445, 8 1990.
- [66] G. I. González and P. R. Saulson, “Brownian motion of a mass suspended by an anelastic wire,” *The Journal of the Acoustical Society of America*, vol. 96, no. 1, pp. 207–212, Jul. 1994.
- [67] G. González, “Suspensions thermal noise in the LIGO gravitational wave detector,” *Classical and Quantum Gravity*, vol. 17, no. 21, p. 4409, 2000.
- [68] T. Sekiguchi, “A study of low frequency vibration isolation system for large scale gravitational wave detectors,” Ph.D. dissertation, University of Tokyo, Jan. 2016.
- [69] D. G. Matei *et al.*, “1.5 μm Lasers with Sub-10 mHz Linewidth,” *Phys. Rev. Lett.*, vol. 118, p. 263 202, 26 Jun. 2017.
- [70] S. Pugla, “Ultrastable high finesse cavities for laser frequency stabilization,” Ph.D. dissertation, Imperial College London, May 2007.
- [71] W. Zhang *et al.*, “Ultrastable Silicon Cavity in a Continuously Operating Closed-Cycle Cryostat at 4 K,” *Physical Review Letters*, vol. 119, p. 243 601, Dec. 2017, ISSN: 0031-9007.
- [72] C. Taylor, M. Notcutt, E. Wong, A. Mann, and D. Blair, “Measurement of the coefficient of thermal expansion of a cryogenic, all-sapphire, Fabry-Perot optical cavity,” *Optics Communications*, vol. 131, no. 4, pp. 311–314, 1996, ISSN: 0030-4018.
- [73] T. Akutsu *et al.*, “Overview of KAGRA: Detector design and construction history,” *Progress of Theoretical and Experimental Physics*, vol. 2021, no. 5, Art. No. 05A101, May 2021, Funding by Ministry of Education, Culture, Sports, Science and Technology (MEXT), ISSN: 2050-3911.
- [74] M. Wolfmeyer and J. Dillinger, “The thermal conductivity of sapphire between 0.4 and 4° K,” *Physics Letters A*, vol. 34, no. 4, pp. 247–248, 1971, ISSN: 0375-9601.
- [75] S. A. Webster, M. Oxborrow, and P. Gill, “Vibration insensitive optical cavity,” *Phys. Rev. A*, vol. 75, p. 011 801, 1 Jan. 2007.

- [76] L. Chen, J. L. Hall, J. Ye, T. Yang, E. Zang, and T. Li, “Vibration-induced elastic deformation of Fabry-Perot cavities,” *Phys. Rev. A*, vol. 74, p. 053 801, 5 Nov. 2006.
- [77] G. D. Cole, W. Zhang, M. J. Martin, J. Ye, and M. Aspelmeyer, “Tenfold reduction of Brownian noise in high-reflectivity optical coatings,” *Nature Photonics*, vol. 7, no. 8, 644–650, 2013.
- [78] ANSYS, Inc., *ANSYS Workbench 2020 R2*.
- [79] T. Nazarova, F. Riehle, and U. Sterr, “Vibration-insensitive reference cavity for an ultranarrow-linewidth laser,” *Applied Physics B: Lasers and Optics*, vol. 83, no. 4, pp. 531–536, Jun. 2006.
- [80] J. B. Wachtman Jr., W. E. Tefft, D. G. Lam Jr., and R. P. Stinchfield, “Elastic Constants of Synthetic Single Crystal Corundum at Room Temperature,” *Journal of research of the National Bureau of Standards. Section A, Physics and chemistry*, vol. 64A, no. 3, 213–228, 1960.
- [81] W. Duffy, “Acoustic quality factor of polycrystalline molybdenum from 50 mK to 300 K,” *Physica B: Condensed Matter*, vol. 169, no. 1, pp. 463–464, 1991, ISSN: 0921-4526.
- [82] Y. S. Touloukian, R. W. Powell, C. Y. Ho, and P. G. Klemens, “Thermophysical properties of matter - the TPRC data series. Volume 1. Thermal conductivity - metallic elements and alloys,” Jan. 1970.
- [83] J. Yu *et al.*, “Excess Noise and Photoinduced Effects in Highly Reflective Crystalline Mirror Coatings,” *Phys. Rev. X*, vol. 13, p. 041 002, 4 Oct. 2023.
- [84] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts, “Double-pass acousto-optic modulator system,” *Review of Scientific Instruments*, vol. 76, no. 6, p. 063 112, Jun. 2005, ISSN: 0034-6748.
- [85] *Phase noise characterization of microwave oscillators*, Product Note 11729C-2, Hewlett Packard Company, 1985.
- [86] A. L. Lance, W. D. Seal, and F. Labaar, “Phase noise and AM noise measurements in the frequency domain,” in *IN: Infrared and millimeter waves. Volume 11 (A85-47951 23-33)*. Orlando, K. J. Button, Ed., vol. 11, Jan. 1984, pp. 239–289.
- [87] A. Takamori *et al.*, “Inverted pendulum as low-frequency pre-isolation for advanced gravitational wave detectors,” *Nuclear Instruments and Methods in Physics Research Section*

- A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 582, no. 2, pp. 683–692, 2007.
- [88] A. Takamori, “Low frequency seismic isolation for gravitational wave detectors,” Ph.D. dissertation, University of Tokyo, 2002.
- [89] G Cella, R DeSalvo, V Sannibale, H Tariq, N Viboud, and A Takamori, “Seismic attenuation performance of the first prototype of a geometric anti-spring filter,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 487, no. 3, pp. 652–660, 2002.
- [90] A. Stochino, R. DeSalvo, Y. Huang, and V. Sannibale, “Improvement of the seismic noise attenuation performance of the monolithic geometric anti-spring filters for gravitational wave interferometric detectors,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 580, no. 3, pp. 1559–1564, 2007.
- [91] T. Uchiyama *et al.*, “Mechanical quality factor of a cryogenic sapphire test mass for gravitational wave detectors,” *Phys. Lett. A*, vol. 261, pp. 5–11, 1999.
- [92] J. Bourhill, E. Ivanov, and M. E. Tobar, “Precision measurement of a low-loss cylindrical dumbbell-shaped sapphire mechanical oscillator using radiation pressure,” *Phys. Rev. A*, vol. 92, p. 023 817, 2 2015.
- [93] D. Dickey and W. Fuller, “Distribution of the estimators for autoregressive time series with a unit root,” *JASA. Journal of the American Statistical Association*, vol. 74, Jun. 1979.
- [94] J. Scargle, “Studies in astronomical time series analysis. ii - statistical aspects of spectral analysis of unevenly spaced data,” *The Astrophysical Journal*, vol. 263, Jan. 1983.
- [95] J. Romano and N. Cornish, “Detection methods for stochastic gravitational-wave backgrounds: a unified treatment,” *Living Reviews in Relativity*, vol. 20, no. 2, Apr. 2017.
- [96] G. P. Centers *et al.*, “Stochastic fluctuations of bosonic dark matter,” *Nature Communications*, vol. 12, no. 1, p. 7321, 2021.
- [97] Y.-D. Tsai, J. Eby, and M. S. Safronova, “Direct detection of ultralight dark matter bound to the sun with space quantum sensors,” *Nature Astronomy*, vol. 7, pp. 113–121, 2023.
- [98] A. S. Göttel *et al.*, “Searching for scalar field dark matter with ligo,” *Physical Review Letters*, vol. 133, no. 10, p. 101 001, 2024.

- [99] L. Aiello *et al.*, “Constraints on scalar field dark matter from colocated michelson interferometers,” *Physical Review Letters*, vol. 128, no. 12, p. 121 101, 2022.
- [100] A. A. Geraci, C. Bradley, D. Gao, J. Weinstein, and A. Derevianko, “Searching for ultra-light dark matter with optical cavities,” *Physical Review Letters*, vol. 123, no. 3, p. 031 304, 2019.
- [101] C. I. E. Torrie, “Development of suspensions for the geo 600 gravitational wave detector,” Ph.D. dissertation, University of Glasgow, Nov. 1999.
- [102] D. Z. Anderson, “Alignment of resonant optical cavities,” *Applied Optics*, vol. 23, no. 17, 2944–2949, 1984.
- [103] S. J. Cooper *et al.*, “Sensors and actuators for the advanced LIGO A+ upgrade,” *Review of Scientific Instruments*, vol. 94, no. 1, p. 014 502, Jan. 2023.
- [104] T. Akutsu *et al.*, “Compact integrated optical sensors and electromagnetic actuators for vibration isolation systems in the gravitational-wave detector KAGRA,” *Review of Scientific Instruments*, vol. 91, no. 11, p. 115 001, Nov. 2020.
- [105] K Haughian *et al.*, “The effect of crystal orientation on the cryogenic strength of hydroxide catalysis bonded sapphire,” *Classical and Quantum Gravity*, vol. 32, no. 7, p. 075 013, 2015.
- [106] D. Chen, “Study of a cryogenic suspension system for the gravitational wave telescope KAGRA,” Ph.D. dissertation, Tokyo U., Mar. 2015.

APPENDIX A

OPEN LOOP MEASUREMENTS OF CAVITY TRANSFER FUNCTIONS

Understanding the frequency response of the optical setup is crucial for the design of the control loop that locks the laser to the cavities. Additionally, it provides useful information about the behavior of individual components of the system. We present here a method for directly measuring the open loop transfer function of the optical cavity together with the elements of the PDH setup (the photodetector, the mix-down electronics and filters employed in the generation of the error signal, the acousto-optic modulator and its driver). We also show how these measurements were used as a way to measure the poles of the sapphire cavities.

The transfer functions of the optical cavities and Pound-Drever-Hall optics and electronics are typically measured by injecting a small modulation to the laser frequency and measuring the response as the modulation frequency is swept. Because the laser needs to remain in resonance with the optical cavity for the duration of the measurement, they are usually measured in closed loop, with the laser locked to the cavity. However, the good relative stability of our two cryogenic cavities allows us to directly measure the open loop transfer functions of each cavity, together with its PDH locking setup, with the laser stabilized by using the other cavity as a frequency reference.

The measurement scheme, presented in Fig. A.1, is similar to the two-cavity setup which was used for the first beat note measurement (section 4.3), with the exception that a network analyzer (Moku:Lab from Liquid Instruments) replaced the locking electronics for one of the cavities.

The laser is locked to one of the cavities through its piezo-electric transducer. Though a low bandwidth (1 kHz) lock, this is enough to maintain the laser in resonance with the second cavity for the duration of the measurement. The network analyzer adds a sinusoidal perturbation to the signal driving the AOM of the second cavity through the frequency modulation port of the frequency

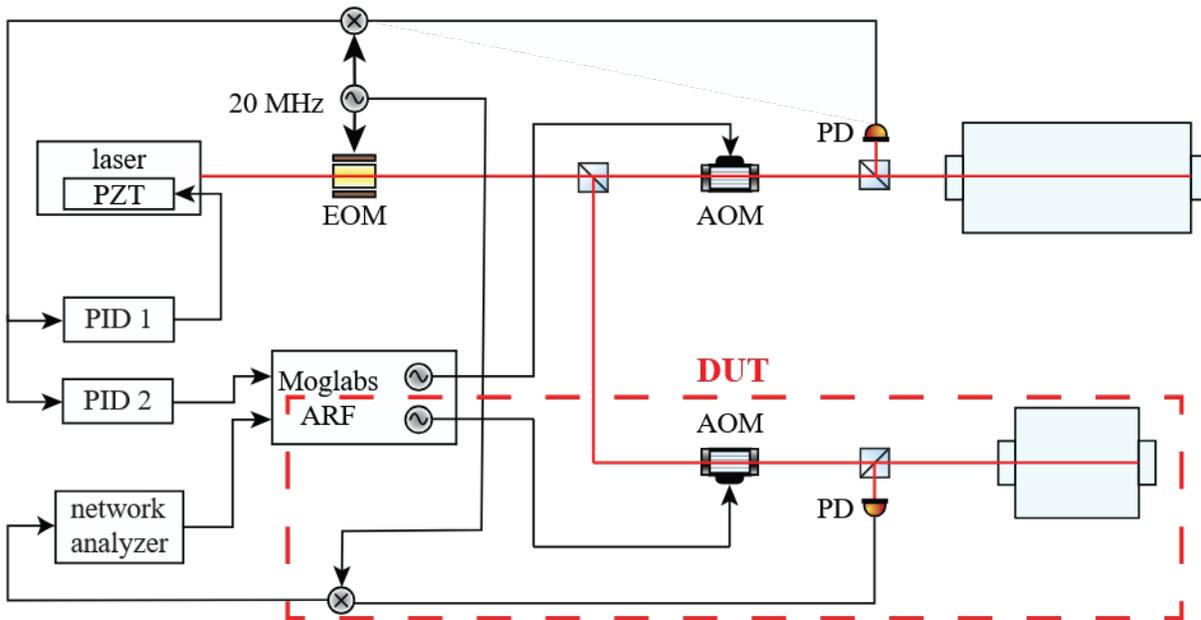


Figure A.1: Measurement scheme for the open loop transfer function. The response of the optical cavity, together with the optics and electronics involved in the Pound-Drever-Hall locking setup is measured for one cavity while the laser is locked to the other cavity. The red box marks the device under test (DUT)

synthesizer driving the AOM. The cavity reflection is measured by the PDH photodiode. The signal is mixed down with the 20 MHz signal driving the EOM to generate a PDH error signal, which is then fed into the input port of the network analyzer. As the frequency of the injected perturbation is varied, the magnitude and phase of the response are recorded, thus rendering the transfer function of the entire control loop with the exception of the control electronics.

The transfer functions are dominated by two features (Fig. A.2): a pole at a few kHz and a steep drop close to 1 MHz, corresponding to the bandwidth limit of the frequency synthesizer driving the AOM. The former is of particular interest, as it is consistent with the low-pass filter response expected from the Fabry-Pérot cavities. The lack of any other features suggests that the rest of the system has an approximately flat response up to around 1 MHz.

The magnitudes of the measured transfer functions up to a few tens of kHz were fitted to the functional form of a first order pole:

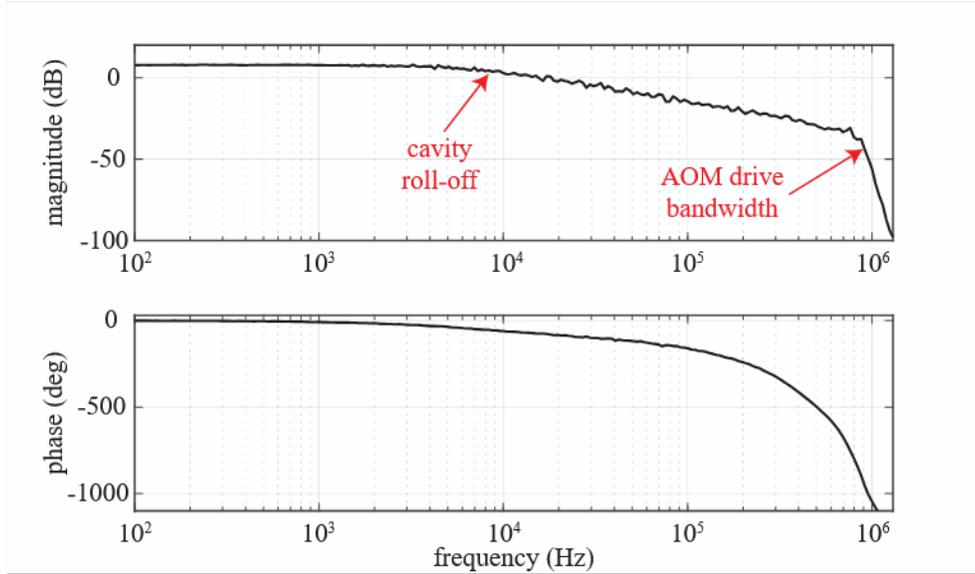


Figure A.2: Bode plot of the measured open-loop transfer function of the 7.5 cm cavity, optical setup, AOM, photodetector

$$y(f) = 20\log_{10} \left(\frac{A}{\sqrt{f^2 + f_p^2}} \right) \quad (\text{A.1})$$

where $y(f)$ is the magnitude in dB, f the frequency, A is a fitting constant associated with the overall gain of the system and f_p a fitting constant corresponding to the cavity pole.

The fits reproduced the data well (Fig. A.3). The pole corners were determined to be 4.7 kHz for the long cavity and 7.6 kHz for the short cavity. These values are similar to, though not in complete agreement with the frequencies of the cavity poles calculated from the finesse: 3.4 kHz for the long cavity and 6.8 kHz for the short cavity. It is worth mentioning that the finesse values were provided by the manufacturer of the cavities and were obtained at room temperature (our measurements were performed at 6.5 K), through a different (ring-down) technique.

As a note, the phase information of the transfer function could not be fitted to an order one pole. This is not surprising, since the huge roll-off due to the bandwidth limit of the AOM drive (which can be understood as a high order low-pass filter), starts a lot earlier than the magnitude

roll-off, already affecting the phase behavior in the region considered for fitting).

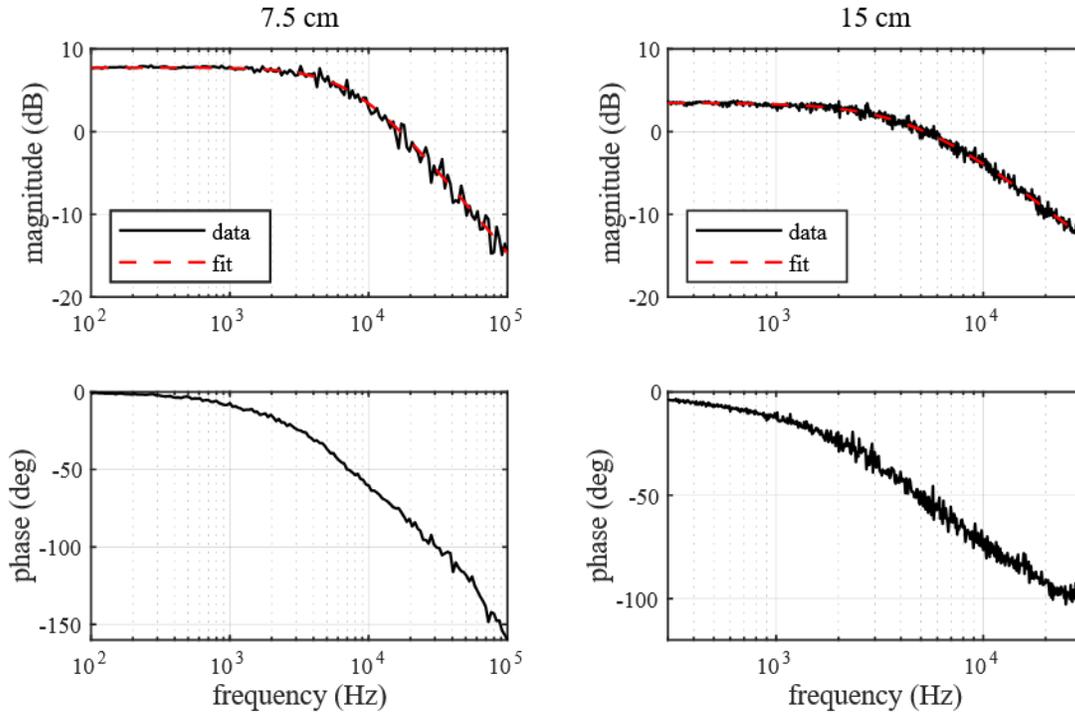


Figure A.3: Measured transfer functions for the 7.5 cm (left) and 15 cm (right) cavities. The magnitudes were fitted to a single pole with corner frequencies at 7.4 kHz for the 7.5 cm cavity and 4.7 kHz for the 15 cm cavity

APPENDIX B

DELAY LINE FREQUENCY DISCRIMINATOR

Operation principle

The delay line frequency discriminator (Fig. B.1) converts the frequency noise of a signal to voltage noise. This is done by splitting the signal into two branches, one of which is delayed before they are recombined using a mixer, and low-pass filtering the output. The conversion takes place in two stages. First, the delay line transforms the frequency fluctuations of the signal into phase fluctuations of the delayed signal relative to the undelayed signal. Second, the mixer translates the phase shift between the two paths into a proportional voltage.

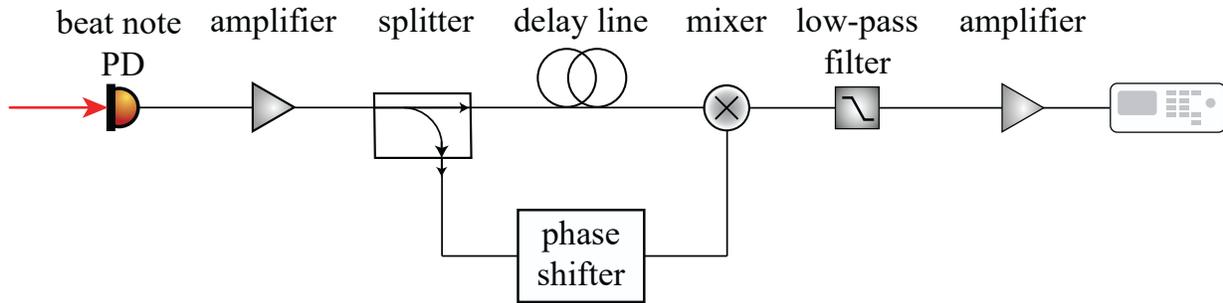


Figure B.1: Schematic of the delay line frequency discriminator

To understand its operation, we will consider a frequency modulation with amplitude δf , at a frequency f_m , applied to an input signal of amplitude V_0 and frequency f_0 :

$$V_{in}(t) = V_0(t) \cos \left(2\pi f_0 t + \frac{\delta f}{f_m} \cos(2\pi f_m t) \right). \quad (\text{B.1})$$

Assuming a delay time τ_d and an amplitude V_R of the signal at the R input of the mixer, which is

characterized by a mixer constant K_L , the mixer output is given by:

$$\begin{aligned}
 V_{out}(t) = & K_L V_R \cos \left(2\pi f_0 t + \frac{\delta f}{f_m} \cos(2\pi f_m t) + 2\pi f_0(t - \tau_d) + \frac{\delta f}{f_m} \cos(2\pi f_m(t - \tau_d)) \right) \\
 & + K_L V_R \cos \left(2\pi f_0 t + \frac{\delta f}{f_m} \cos(2\pi f_m t) - 2\pi f_0(t - \tau_d) - \frac{\delta f}{f_m} \cos(2\pi f_m(t - \tau_d)) \right) \\
 & + \text{higher order terms.}
 \end{aligned} \tag{B.2}$$

A low-pass filter at the output eliminates the high-frequency components. The filtered signal is:

$$\begin{aligned}
 V(t) = & K_L V_R \cos \left(2\pi f_0 t + \frac{\delta f}{f_m} \cos(2\pi f_m t) - 2\pi f_0(t - \tau_d) - \frac{\delta f}{f_m} \cos(2\pi f_m(t - \tau_d)) \right) \\
 = & K_L V_R \cos \left(2\pi f_0 \tau_d - 2 \frac{\delta f}{f_m} \sin(\pi f_m \tau_d) \sin \left(2\pi f_m \left(t - \frac{\tau_d}{2} \right) \right) \right).
 \end{aligned} \tag{B.3}$$

The phase shifter on the undelayed path can be tuned so that the signals on the two branches meet at the mixer in quadrature, $2\pi f_0 \tau_d = (2n + 1)\pi/2$, with n integer. In this case,

$$V(t) = K_L V_R \sin \left(2 \frac{\delta f}{f_m} \sin(\pi f_m \tau_d) \sin \left(2\pi f_m \left(t - \frac{\tau_d}{2} \right) \right) \right). \tag{B.4}$$

Under the assumption that the frequency modulation is small, $\delta f \ll f_m$, the amplitude of the output voltage δV is given by:

$$\delta V \approx K_L V_R \left(2 \frac{\delta f}{f_m} \sin(\pi f_m \tau_d) \right) = 2\pi \tau_d K_L V_R \delta f \frac{\sin(\pi f_m \tau_d)}{\pi f_m \tau_d}. \tag{B.5}$$

In this case, for modulation frequencies small enough that $f_m \tau_d \ll 1$, the output voltage is directly proportional to the modulation amplitude, with a proportionality constant:

$$K_d = 2\pi \tau_d K_L V_R. \tag{B.6}$$

It is important to note that, although it is called a "constant", K_d actually depends on the power of the input signal.

Calibration procedure

The delay line frequency discriminator K_d is determined by measuring the response of the system to a known signal. This was done either by applying a dither to the beat note, or with a frequency modulated signal from a signal generator, with the same power as the beat note. The calibration procedure is illustrated in Fig. B.2. In this case, a 60 kHz modulation was applied to the beat note using the acousto-optic modulator that shifts the frequency of the light going to the short cavity.

For small frequency modulations (corresponding to total phase deviations much smaller than 1 rad), the phase spectral density S_ϕ (in units of rad^2/Hz) is related to the single sideband power density P_{ssb} (in units of W/Hz) and the total signal power P_s (in units of W) as [85]:

$$S_\phi = 2 \frac{P_{ssb}}{P_s}. \quad (\text{B.7})$$

This is converted to the frequency noise spectral density S_f (in units of Hz^2/Hz) using the relation:

$$S_f = S_\phi f_m^2 = 2 f_m^2 \frac{P_{ssb}}{P_s}. \quad (\text{B.8})$$

We calculate the total frequency noise in the sideband (in units of Hz) from the magnitudes of the carrier and sideband peaks observed in the two-sided spectrum of the beat note, measured on a spectrum analyzer (left panel of Fig. B.2):

$$\delta f = \sqrt{2} f_m \times 10^{(P_{ssb}^{(\text{dBm})} - P_c^{(\text{dBm})})/20}. \quad (\text{B.9})$$

Here, $P_{ssb}^{(\text{dBm})}$ and $P_c^{(\text{dBm})}$ are the magnitudes of the sideband and carrier expressed in dBm .

The corresponding voltage noise at the output of the delay line frequency discriminator is determined as the root-mean-square voltage V_{rms} corresponding to the sideband (the square root of the power spectral density of the signal integrated over the peak width). Then, the delay line frequency discriminator constant is calculated as:

$$K_d = V_{rms}/\delta f = 69.657 \mu\text{V}/\text{Hz} . \quad (\text{B.10})$$

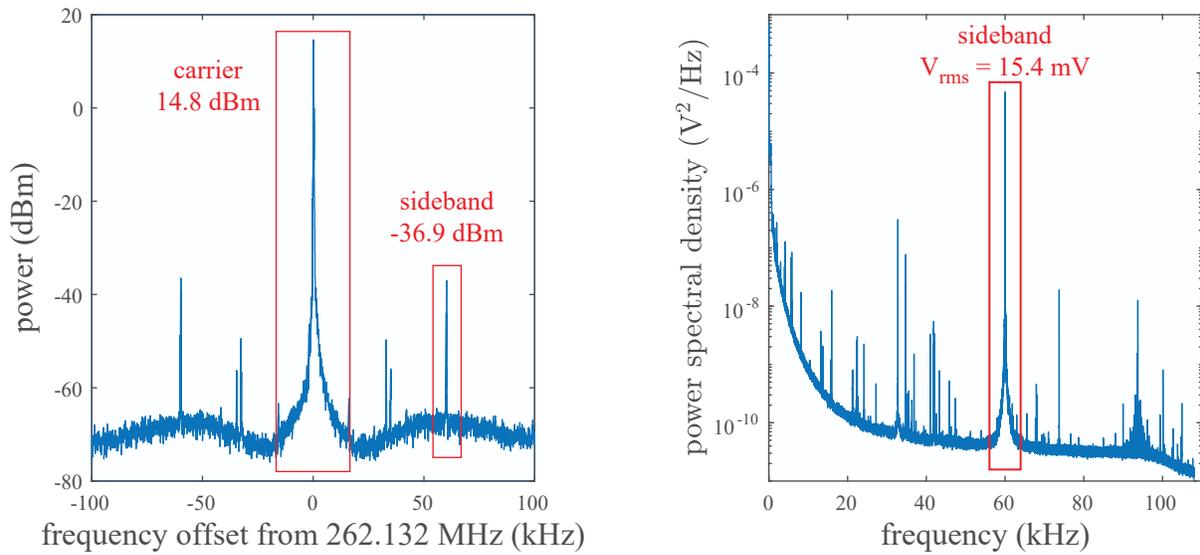


Figure B.2: Left: magnitudes of the beat note carrier and of the sideband from an applied 60 kHz dither before demodulation. Right: 60 kHz sideband after passing through the delay line frequency discriminator

APPENDIX C
LIST OF PARAMETERS USED FOR NOISE PROJECTIONS

	Parameter	Value	Description
General properties	λ	1550 nm	Laser wavelength
	\mathcal{F}	150,000	Cavity finesse
	T	4 K	Cavity temperature
	T_{tot}	4×10^6 s	Measurement time
	η	0.8	PD quantum efficiency
Substrate & Spacer	ρ_m, ρ_{sp}	3997 kg/m ³	Density
	E_m, E_{sp}	464 GPa	Young's modulus
	σ_m, σ_{sp}	0.29	Poisson ratio
	ϕ_m	10^{-8}	Substrate loss angle
	ϕ_{sp}	10^{-4}	Spacer loss angle
	R_{out}	2.5 cm	Spacer outer radius
	R_{in}	0.5 cm	Spacer inner radius
	R_m	1.27 cm	Mirror radius
	t_m	0.625 cm	Mirror thickness
	R_c	10.2 m	Mirror radius of curvature
Mirror coating	E_c	100 GPa	Coating Young's modulus
	σ_c	0.32	Coating Poisson ratio
	ϕ_c	5×10^{-6}	Coating loss angle (at 4 K)
	t_c	10 μ m	Coating thickness

Non-rigid cavity suspension	E_{susp}	70 GPa	Wire Young's modulus
	σ_{susp}	0.25	Wire Poisson ratio
	d_{susp}	310 μm	Wire diameter
	L_{susp}	15 cm	Wire length
Rigid cavity vibrations	$f_{0,z}$	1.5 Hz	Pendulum resonant frequency
	$f_{0,y}$	100 Hz	Wire vertical resonant frequency
	H_z	6×10^{-12}	Horizontal vibration sensitivity
	H_y	2×10^{-12}	Vertical vibration sensitivity
	a_{vib}	$10^{-7} g/\sqrt{\text{Hz}}$	Cryostat vibration acceleration
Rigid-rigid (current apparatus)	L_L	15 cm	Long cavity length
	L_S	7.5 cm	Short cavity length
	w_L	0.66 mm	Long cavity beam waist
	w_S	0.55 mm	Short cavity beam waist
	P	0.1 mW	Power at PD from one cavity
Rigid-rigid (projection)	L_L	100 cm	Long cavity length
	L_S	30 cm	Short cavity length
	w_L	1.00 mm	Long cavity beam waist
	w_S	0.78 mm	Short cavity beam waist
	P	10 mW	Power at PD from one cavity
Rigid-non-rigid (projection)	L	30 cm	Cavity length
	w	0.78 mm	Beam waist
	P_{in}	10 mW	Incident power for each cavity

Table C.1: Experimental parameters used in the noise projections in Fig. 2.8 and Fig. 2.9