

Cooling a macroscopic mechanical oscillator close to its quantum ground state

Leonhard Neuhaus

▶ **To cite this version:**

Leonhard Neuhaus. Cooling a macroscopic mechanical oscillator close to its quantum ground state. Physics [physics]. Université Pierre et Marie Curie - Paris VI, 2016. English. ⟨NNT : 2016PA066555⟩. ⟨tel-01467924v2⟩

HAL Id: tel-01467924

https://theses.hal.science/tel-01467924v2

Submitted on 12 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization













THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PIERRE ET MARIE CURIE

Spécialité : Physique

École doctorale : "Physique en Île-de-France"

réalisée

Laboratoire Kastler Brossel

présentée par

Leonhard NEUHAUS

pour obtenir le grade de :

DOCTEUR DE L'UNIVERSITÉ PIERRE ET MARIE CURIE

Sujet de la thèse :

Cooling a macroscopic mechanical oscillator close to its quantum ground state

devant le jury composé de :

M.	Jean-Marc BERROIR	Examinateur
M.	Pierre-François COHADON	Directeur de thèse
M.	Eddy COLLIN	Rapporteur
M.	Serge GALLIOU	Examinateur
M.	Jack HARRIS	Rapporteur
M.	Antoine HEIDMANN	Membre invité
M.	Tobias KIPPENBERG	Examinateur

Contents

1 Introduction **1**
- 1.1 Historical developments . . . 1
- 1.2 Motivation and principle of the experiment . . . 3
- 1.3 Outline of this thesis . . . 5

2 Theory **7**
- 2.1 Mechanics . . . 7
  - 2.1.1 Vibrational eigenmodes of a solid . . . 8
  - 2.1.2 Eigenmode susceptibility . . . 9
  - 2.1.3 Mechanical damping . . . 9
  - 2.1.4 Thermal motion . . . 10
  - 2.1.5 Effective susceptibility for scalar actuation and measurement . . . 10
  - 2.1.6 Collocation and optomechanical coupling . . . 11
  - 2.1.7 Thermal noise in the scalar model . . . 12
  - 2.1.8 Finite-element method and continuous models . . . 12
  - 2.1.9 Effective mass for special cases . . . 13
- 2.2 Optics . . . 14
  - 2.2.1 Gaussian modes of the electric field . . . 14
  - 2.2.2 Classical field amplitude . . . 15
  - 2.2.3 Optical noise . . . 17
  - 2.2.4 The electric field inside a Fabry-Perot cavity . . . 21
  - 2.2.5 Linear sweep over a resonance . . . 23
  - 2.2.6 Optical susceptibility of a Fabry-Perot cavity . . . 24
  - 2.2.7 Sensitivity of different detection schemes . . . 26
- 2.3 Cavity optomechanics . . . 29
  - 2.3.1 Radiation pressure in a cavity . . . 29
  - 2.3.2 Static effects of radiation-pressure . . . 30
  - 2.3.3 Dynamical effects of radiation-pressure . . . 31

3 Experiment - Samples **39**
- 3.1 Mechanical resonator . . . 39
  - 3.1.1 Choice and design of mechanical oscillator . . . 39
  - 3.1.2 Fabrication . . . 41
  - 3.1.3 Mechanical characterization . . . 50
  - 3.1.4 Limitations of the mechanical quality factor . . . 56

3.2 Coupling mirror . . . 59
3.2.1 Experimental constraints . . . 59
3.2.2 Fabrication of coupling mirrors . . . 61
3.2.3 Optical characterization . . . 65
3.2.4 Loss due to higher-order transverse mode degeneracy . . . 72
3.2.5 Measurement of radius of curvature . . . 74

4 Experiment - Optomechanical cavity inside a cryostat 79
4.1 Cryogenics . . . 79
4.1.1 Dilution refrigeration . . . 79
4.1.2 Horizontal dilution refrigerator from CryoConcept . . . 81
4.1.3 Temperature measurement . . . 85
4.1.4 Electrical connections and noise . . . 86
4.1.5 Inner vacuum chamber (IVC) . . . 88
4.1.6 Gas damping . . . 93
4.1.7 Superfluid ^{4}He damping . . . 96
4.2 Optomechanical cavity . . . 98
4.2.1 Requirements for a cryogenic optomechanical cavity . . . 99
4.2.2 Final cavity design . . . 102
4.2.3 Cavity alignment protocol . . . 102
4.2.4 Electrode actuator . . . 108

5 Experiment - Optical setup 113
5.1 Laser source . . . 114
5.2 Filter cavity . . . 116
5.2.1 Implementation . . . 119
5.2.2 Performance . . . 121
5.2.3 Raman lasing in FPF1 . . . 124
5.3 Modulators . . . 125
5.3.1 Non-resonant fiber amplitude-modulator (AM) . . . 126
5.3.2 Acousto-optic modulator (AOM) . . . 128
5.3.3 Non-resonant fiber phase-modulator (PM) . . . 131
5.4 Cryostat injection optics . . . 133
5.5 Photodetectors . . . 133
5.5.1 Avalanche photodetector . . . 134
5.5.2 Differential photodetector . . . 135
5.6 Homodyne detection . . . 138
5.7 Heterodyne detection . . . 141
5.8 PDH detection . . . 142
5.9 Computer interface of the combined setup . . . 145

6 Feedback control 147
6.1 Brief introduction to feedback . . . 147
6.2 Python RedPitaya Lockbox (PyRPL) . . . 150
6.2.1 Red Pitaya . . . 150
6.2.2 PyRPL overview . . . 152

6.2.3 Simple DSP modules . . . 154
6.2.4 IQ module . . . 156
6.2.5 IIR module . . . 159
6.2.6 Higher abstraction layers . . . 162
6.3 Photo-thermal instabilities . . . 166
6.3.1 Large modulation synchronous detection . . . 166
6.3.2 Photo-thermal instability . . . 169
6.3.3 Feedback-suppression of photo-thermal instabilities . . . 171
6.3.4 Feedback through amplitude-modulated light . . . 173

7 Optomechanics experiments 179
7.1 Room-temperature measurements . . . 179
7.1.1 Brownian motion at 300 K . . . 180
7.1.2 Thermal resistance . . . 186
7.2 Cooling experiment . . . 186
7.2.1 Measurement setup . . . 187
7.2.2 Optomechanically induced transparency . . . 188
7.2.3 Radiation-pressure cooling close to the quantum ground state . . . 191
7.2.4 Heating of the micropillar by absorption . . . 195
7.2.5 Mechanical quality factor . . . 198
7.3 Towards the observation of quantum backaction . . . 202
7.3.1 Calibration of the homodyne detection . . . 202
7.3.2 Observation of back-action . . . 204
7.3.3 Quantum backaction at room temperature? . . . 208
7.4 The noise floor issue . . . 210
7.4.1 Cavity displacement noise floor . . . 211
7.4.2 Phononic crystal mirror substrates . . . 214

8 Conclusion 217
8.1 Summary . . . 217
8.2 Short-term prospects . . . 220
8.3 Long-term outlook . . . 224

Appendices 227

Appendix A Transfer functions of a Fabry-Perot cavity 229
A.1 Mathematica code . . . 229
A.2 Transfer matrix coefficients . . . 231

Appendix B List of functional micropillar samples 233

Appendix C Differential photodiode schematic 235

Appendix D Cryostat wiring schematics 237

Appendix E Operation of the horizontal dilution refrigerator 239

Bibliography **243**

Chapter 1
Introduction

> *Now, mind you, I do not say that I think that quantum mechanics does fail at large distances, I only say that it is not inconsistent with what we do know. If this failure of quantum mechanics is connected with gravity, we might speculatively expect this to happen for masses such that* $GM^2/\hbar c = 1$*, of* M *near* 10^{-5} *grams, which corresponds to some* 10^{18} *particles. Now quantum mechanics gives silly answers for objects of this size; if we calculate the probability that a grain of sand should jump over a wall, we get answers like* $10^{-60,000}$*, which are ridiculous.*
>
> — Richard P. Feynman, *Feynman Lectures on Gravitation*

Optomechanics deals with the mutual coupling of the electromagnetic field with the mechanical degrees of freedom of macroscopic objects [AKM14]. When a light beam is reflected by a movable mirror, the phase of the reflected beam depends on the position of the mirror. Amplifying this phase dependence by the use of high-finesse cavities allows to interferometrically detect the mirror position with ultra-high sensitivities, typically of the order of $10^{-20}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. The most prominent application of this effect can be found in gravitational-wave interferometers, where relative length changes of kilometer-scale cavities of the order of $h \approx 10^{-22}$ are measured in order to observe the tiny deformations of space-time induced by gravitational waves emitted by cataclysmic astrophysical events, as beautifully exemplified by the recent detection of gravitational waves from a binary black hole merger [AAA+16].

The inverse direction of optomechanical coupling can become significant as well: the momentum transferred to the mirror upon the reflection of a light beam can set the mirror into motion. This effect has been familiar since Kepler's observation in the 17th century that comet tails point away from the sun [Kep19]. For about three decades, radiation-pressure has furthermore aroused the interest of the quantum optics community.

1.1 Historical developments

The radiation-pressure force, together with the quantum fluctuations of the light field, can give rise to a fundamental sensitivity limit of interferometric displacement mea-

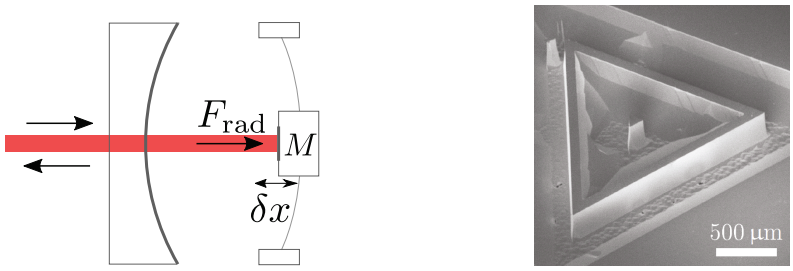


Figure 1.1: Left: The generic optomechanical system, a Fabry-Perot cavity with one movable mirror. The displacement δx modifies the phase of the reflected laser beam, and the radiation-pressure force F_{rad} exerted by the laser modifies the position of the mirror by an amount that is inversely proportional to the effective mirror mass M. Right: Scanning electron microscope (SEM) image of the mechanical resonator studied in this thesis.

surements, which has been studied theoretically since the late 1970s [Cav81]. This *Standard Quantum Limit* (SQL) can however be beaten by using squeezed states of the light field [JR90]. While even today, gravitational-wave interferometers are still not limited by the SQL, efforts to demonstrate and overcome quantum limitations of the measurement sensitivity have greatly motivated the field of quantum optics. This field of research has blossomed since the first demonstration of squeezed light in the 1980s [SHY$^+$85] with nonlinear optical media whose optical index n depends on the laser intensity. Optomechanical systems also display an optical nonlinearity as the physical length of an optomechanical cavity depends on the intensity of the intracavity field (see Fig. 1.1). Quantum optics experiments similar to those performed with atomic gases and non-linear media can therefore be carried out with mechanical systems as well, for example the generation of squeezed light [HR94], QND measurements of the light intensity [HHP97], or entanglement [PDV$^+$05]. This triggered the emergence of a specific subfield within quantum optics devoted to the mechanical degree of freedom. The "Optomechanics and Quantum Measurement" group at Laboratoire Kastler Brossel was (initially as a part of the "Quantum Optics" group) the first experimental group to look into this direction, as early as the 1990s.

Early experiments in optomechanics were however plagued by a much weaker nonlinearity than original quantum optics experiments, as no atomic resonance could be taken advantage of. It proved crucial to boost the interaction between light and matter not only through an optical resonance, but also by designing the optomechanical resonator to have advantageous properties such as low mass and low spring constant, i.e. low frequency. However, in order to avoid technical issues with classical noise that arise predominantly at low frequencies, resonance frequencies of at least 100 kHz are most advantageous in practice. Optomechanics experiments also have to face the specific challenge of thermal noise: if seismic and acoustic noises can be successfully shielded, mechanical motion of room-temperature systems at frequencies up to the GHz range is completely dominated by Brownian motion arising from the oscillators' unavoidable coupling to a non-zero temperature thermal bath. Consequently, much effort in optomechanical research has aimed at the reduction of thermal noise, both by working in a cryogenic environment at technically feasible sub-Kelvin temperatures, and by

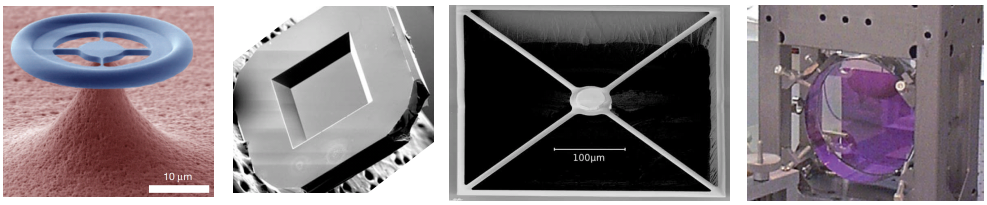


Figure 1.2: Various optomechanical systems at the high end of the mass scale. From left to right: Silica microtoroid ($M \approx 10$ ng) [VDW+12], silicon nitride membrane ($M \approx 50$ ng) [ZSJ+08], dielectric mirror suspended on a silicon nitride trampoline ($M \approx 100$ ng) [KPJ+11], and LIGO test mass ($M \approx 2.7$ kg) [LIG09]. The first two systems have been observed in the quantum regime by now [RDW+11, PYK+15, PPK+16, UML+15].

engineering mechanical oscillators with minimal coupling to their environments, i.e. high mechanical quality factors Q. Employing state-of-the-art microfabrication techniques to implement these design considerations enabled the community to reach a milestone by the beginning of the 2000s: significantly cooling the thermal motion of mechanical oscillator with light, at first by employing an external feedback loop with radiation-pressure actuation of a movable mirror [CHP99], later by replacing the external feedback loop by a Fabry-Perot cavity with photo-thermal [MK04] and eventually radiation-pressure coupling [ACB+06, GBP+06a, SDN+06].

While optomechanical cooling does not reduce the thermal force noise itself, but rather demonstrates the ability to optically modify the mechanical response to external forces, the natural next step is to make the quantum noise of radiation-pressure forces dominate over thermal noise, or equivalently, to reduce the effective thermal occupation number of a mechanical oscillator to its quantum ground state. It is only in this regime that the experiments of quantum optics listed above can be demonstrated with mechanical systems, since excessive thermal noise leads to the immediate destruction of quantum coherence. This quantum regime was reached by a number of groups a few years later, in the order of increasing mass of the mechanical oscillator: first in an ensemble of optically trapped atoms [MMGSK08], next in a 0.3 pg periodically patterned nanostructure [CASN+11], and later in silica microtoroids [RDW+11, VDW+12] and very thin silicon nitride membranes [PPR13, PYK+15, PPK+16, UML+15] with effective masses of about 50 ng (see Fig. 1.2 for a few examples).

1.2 Motivation and principle of the experiment

The work presented in this thesis attempts an exploration of quantum effects in optomechanical coupling with a mechanical resonator mass of $33\,\mu$g, which is three orders of magnitude above the mass of the heaviest mechanical systems for which the observation of quantum effects in optomechanical coupling was reported up to date [RDW+11, PYK+15, PPK+16, UML+15]. An SEM image of our mechanical oscillator is shown in Fig. 1.1. The macroscopic size of the oscillator allows to construct a Fabry-Perot cavity, where one end mirror is the surface of the mechanical oscillator coated with a dielectric mirror. Our system therefore resembles the textbook example of an optomechanical system as closely as possible. The large mass of our optomechanical

system forces us to optimize all other system parameters in order to reach the regime where the quantum fluctuations of radiation-pressure dominate thermal force noise. In particular, we are required to lower the temperature of the thermal environment of the mechanical oscillator in a dilution refrigerator below 1 K. It was expected that a large number of technical problems had to be solved during the construction and optimization of our experiment due to the uncommon combination of parameters. We hope that the discovered solutions that are presented throughout this thesis may find applications inside and outside the optomechanics community.

As stated above, pushing optomechanical systems towards higher mass is disadvantageous for the demonstration of radiation-pressure effects, since heavier oscillators obviously require higher radiation-pressure forces, which are hard to generate without introducing classical noise. The motivation for this thesis can therefore be summarized by the answer to the question: **Why is it nevertheless interesting to perform quantum experiments with mechanical oscillators of masses at the scale of the Planck mass?** We find two answers to this question:

- **Technological aspects.** While a higher mass is unfavorable for the optomechanical coupling, the experimental difficulty does not increase linearly with the mass. A more massive mechanical oscillator involves a larger volume, which allows a number of other optimizations: the optomechanical system can handle higher optical power, the larger volume conducts heat from absorbed laser light more efficiently to its thermal environment, and the mechanical quality factor of a larger oscillator is expected to be more robust with respect to effects from its environment, such as gas damping or other surface-related effects. Furthermore, the larger size enables the use of dielectric mirror coatings and a Fabry-Perot optomechanical cavity, which enable better optomechanical coupling than partially transparent membranes [ZSJ+08] or evanescent coupling [AAU+09]. While recent work with nano-optomechanical systems [GHSN+13] has aimed at the improvement of the quantum efficiency in this kind of system, quantum optics experiments have been performed with Fabry-Perot cavities for several decades. The corresponding maturity of quantum techniques with Fabry-Perot systems is expected to result in larger quantum enhancements with our optomechanical system when sub-SQL schemes are implemented, for example with the planned injection of squeezed light into our optomechanical cavity (see section 8.3). Finally, the resemblance of our table-top system with larger-scale experiments such as LIGO makes it more likely that technology developed for the quantum-enhancement of our experiments may find an application in these larger experiments. In particular, the operation of our experiment at cryogenic temperatures could provide technological hints for the next generation of graviational-wave interferometers, that will operate in a cryogenic environment [Som12].

- **Fundamental aspects.** As early as 1926, attempts to apply the formalism of quantum mechanics to the description of macroscopic point masses were made [Sch26]. Optomechanics is often deemed interesting for tests of quantum mechanics because of the macroscopic size of the mechanical system, composed of 10^{11} to 10^{28} atoms. While quantum effects such as superfluidity and superconductivity are routinely observed in systems of even larger size, the quantum

formalism is not applied to the mechanical degree of freedom in the description of these phenomena. A number of theories devising modifications of quantum mechanics at the macro-scale therefore speculate about the existence of collapse mechanisms that affect quantum superpositions of position states, possibly due to the action of gravity that introduces a universal coupling between different masses involving the position variable [Leg02, AB09, Dió87, Pen96, GRW86]. It is therefore certainly interesting to provide unambiguous evidence for the validity of quantum theory in this particular regime. While the parameters of our optomechanical system are in a regime that should allow for the indirect exclusion of certain collapse models through quantitative measurements of the force noise spectrum acting on the mechanical oscillator in its quantum ground state [BPBU14], the presence of classical noise sources has hindered these observations so far (see section 7.3.2). We expect to provide this evidence by the calibrated measurement of the zero-point fluctuations of our mechanical oscillator once the classical noise sources are removed. At the same time, the solutions of technological challenges with our macroscopic system will help a future generation of macroscopic quantum superposition experiments that are already in the design stage [MERS$^+$08, KAB$^+$16].

1.3 Outline of this thesis

This document is intended to fulfill three purposes. First, it describes the measurement results obtained with our optomechanical system in detail. Second, it reports a large amount of experimental tricks and tools that were accumulated during the work for this thesis. Third, it serves as a manual of the described experimental setup for future generations of PhD students.

After this introduction, Chapter 2 will present the theoretical foundations of optomechanics by modeling the mechanical and optical degrees of freedom and discussing the nonlinear effects arising from radiation-pressure. Chapter 3 continues with the presentation of the mechanical resonator and the coupling mirror samples used throughout the thesis, including design, fabrication and characterization measurements. Chapter 4 deals with the cryostat and the mechanical mounting structure that are employed to create a low-temperature and low-vibration environment for the mechanical oscillator in order to improve the ratio of quantum effects over thermal noise sources. Chapter 5 describes the complex optical setup used to control and detect the motion of the mechanical resonator, together with the corresponding electronics and software enabling fully automated measurements. Chapter 6 provides a short introduction to the theory of feedback control, describes in detail the developed software package PyRPL which implements a versatile feedback controller for quantum optics experiments with an FPGA, and explains how our optomechanical system can be operated in a stable regime despite the presence of parametric instabilities on both sides of the optical resonance. Chapter 7 presents the optomechanics experiments performed during this thesis, including Brownian motion measurements at room temperature and at sub-Kelvin temperatures, optomechanical cooling of our 33 μg mechanical oscillator to a mean thermal occupation number of 20 phonons, the observation of measurement

backaction force noise of quantum and classical origin, and finally a characterization and proposal to remove the classical phase noise source in our optomechanical system. All measurement results are interpreted in light of future experimental developments in chapter 8, which also summarizes the obtained results from all chapters.

Chapter 2

Theory

First we try to get clearly in our minds how far the system of classical mechanics has shown itself adequate to serve as a basis for the whole of physics.

— Albert Einstein, *Physics and Reality*

This chapter will develop the theoretical foundation for the rest of this thesis. The first section discusses the mechanical degree of freedom of the optomechanical system: it is derived how the vibrations of a solid can be modeled as one or few harmonic oscillators. The second section discusses the optical degree of freedom in order to link a measurement of the power of a laser beam that has interacted with an optomechanical system to the properties of the input beam and the displacement of the mechanical oscillator. The last section considers the nonlinearities that arise when the radiation pressure exerted by the light on the mechanical oscillator becomes significant.

2.1 Mechanics

A solid with N atoms has a total of $3N$ degrees of freedom in the three directions of space. The aim of this section is to simplify this situation to only one degree of freedom that we are interested in. To do so, we will first express the displacements of all atoms in the form of normal modes (section 2.1.1). We will then find the susceptibility of the normal modes to an external force vector (2.1.2) and introduce damping to make our model more realistic (section 2.1.3). The damping requires the introduction of a thermal force noise to prevent violation of the equipartition principle (section 2.1.4), a relation also known as Fluctuation-Dissipation theorem (FDT). We will then compute the scalar transfer function for a particular sensor-actuator pair (section 2.1.5), approximate the effect of almost all of the $3N$ normal modes as a mechanical background mode, and introduce the special case of a collocated sensor-actuator pair such as a beam of light (section 2.1.6). We will see that the scalar model for a collocated sensor-actuator pair permits a simple expression for the fluctuation-dissipation theorem (section 2.1.7). We will conclude by commenting on alternative approaches such as the finite-element method and continuum mechanics (section 2.1.8) and apply the developed formalism to two simple examples of mechanical systems (section 2.1.9). The

treatment presented in this section is reminiscent with the ones in [PHH99, Pre11].

2.1.1 Vibrational eigenmodes of a solid

The vibrations of a solid can be viewed as a large number N of atoms which deviate from their equilibrium positions by $x_i(t)$, where the index $i = 3(n-1)+l$ runs from 1 to $3N$ and represents the degree of freedom $l \in \{1,2,3\}$ of atom $n \in \{1,\ldots,N\}$. In a linearized treatment, each atom coordinate x_i is subject to an electrostatic interatomic force $F_{ij} = k_{ij}(x_i - x_j)$ from each of the other atoms with coordinates x_j, with a spring constant k_{ij} which is in practice non-zero only for a small fraction of all pairs of indices when short-ranged interactions are assumed. Newtonian mechanics and an external force G_i on each degree of freedom leads to $3N$ equations of motion of the form

$$M_i\ddot{x}_i(t) + \sum_{j=1}^{3N} k_{ij}(x_i(t) - x_j(t)) = G_i(t) \tag{2.1.1}$$

where M_i is the mass associated with degree of freedom x_i, i.e. the atomic mass. By defining the symmetric stiffness matrix $K_{ij} = -k_{ij}$ for $i \neq j$ and $K_{ii} = \sum_{1\leq j\leq 3N,\, j\neq i} k_{ij}$, the system of equations can be re-written in matrix form as

$$M\ddot{x}(t) + Kx(t) = G(t) \tag{2.1.2}$$

with the mass matrix $M = \mathrm{diag}(M_1,\ldots,M_{3N})$, the position vector $x = (x_1,\ldots,x_{3N})^T$ and the external force vector $G = (G_1,\ldots,G_{3N})^T$. In the absence of an external force, the Ansatz $x(t) = \xi_i e^{i\omega_i t}$ yields the eigenvalue problem

$$(K - \omega_i^2 M)\xi_i = 0 \tag{2.1.3}$$

whose solution results in $3N$ eigenmode shapes ξ_i and mode frequencies ω_i. For convenience, we will number the eigenfrequencies in order of increasing frequency: $\omega_i < \omega_j$ iff $i < j$. Since the the eigenmode shapes can have an arbitrary scalar factor, we will set this normalization with $\xi_i^T\xi_i = 1$ for all modes. Multiplying the equation with ξ_j^T from the left gives

$$\xi_j^T K\xi_i = \omega_i^2 \xi_j^T M\xi_i\,. \tag{2.1.4}$$

Subsequently subtracting the same equation with i and j interchanged yields

$$\xi_j^T M\xi_i(\omega_i^2 - \omega_j^2) = 0\,, \tag{2.1.5}$$

which establishes the orthogonality of the eigenmode shapes of different eigenfrequencies $\omega_i \neq \omega_j$ with respect to the mass matrix, and through the previous equation also orthogonality with respect to the stiffness matrix. By defining the orthogonal matrix of mode shapes $\Xi = (\xi_1,\ldots,\xi_{3N})$ and the mode masses $\mu_i = \xi_i^T M\xi_i$, we can rewrite the orthogonality equations in the form of the diagonal eigenmode mass and frequency matrices μ and ω:

$$\Xi^T M\Xi = \mathrm{diag}(\mu_1,\ldots,\mu_{3N}) = \mu \tag{2.1.6}$$

$$\Xi^T K\Xi = \mathrm{diag}(\mu_1\omega_1^2,\ldots,\mu_{3N}\omega_{3N}^2) = \mu\omega^2\,. \tag{2.1.7}$$

A transformation to the eigenmode basis z defined by $x(t) = \Xi z(t)$ and $G(t) = \Xi F(t)$ allows to rewrite the equation of motion in the simpler form

$$\ddot{z}(t) + \omega^2 z(t) = \mu^{-1}F(t)\,. \tag{2.1.8}$$

2.1.2 Eigenmode susceptibility

The equation of motion is a linear second-order ordinary differential equation, which is commonly solved through the Fourier transform. For a time-dependent function $f(t)$, we define its Fourier transform $f(\Omega) = \int_{-\infty}^{\infty} f(t)e^{i\Omega t}\mathrm{d}t$. The Fourier-transformed equation of motion can then be solved algebraically to read

$$z(\Omega) = \chi(\Omega)F(\Omega) \tag{2.1.9}$$

where the susceptibility matrix

$$\chi(\Omega) = \mathrm{diag}\left(\frac{1}{\mu_i(\omega_i^2 - \Omega^2)}\right) \tag{2.1.10}$$

was introduced.

2.1.3 Mechanical damping

In the absence of an external force, Eq. 2.1.8 predicts that an initial non-zero deflection $z_i(0)$ at rest oscillates perpetually at the corresponding eigenfrequency: $z_i(t) = z_i(0)\cos(\omega_i t)$. In a realistic solid, various damping mechanisms [BMP85] cause the amplitude of an excited mechanical mode to decay exponentially. If the solid is regarded as an isolated system, energy conservation demands that the decayed motional energy leaks into other modes, so we must view mechanical damping as a coupling between normal modes. This coupling cannot be described as simple springs between either atoms or more abstract oscillators, for it would have been already taken into account by the normal mode expansion.

Rather, we must add a force term to the equation of motion that is out-of-phase with the displacement such that the eigenfrequencies can acquire an imaginary component to describe the damped motion. Common models [Sau90] are viscous damping, where the damping term $-\gamma\dot{z}(t)$ is proportional to the velocity, and structural damping, where the damping $-i\omega^2\phi(\Omega)z(\Omega)$ is specified in the form of the frequency-dependent loss angle $\phi_i(\Omega)$ of each mode. While viscous damping is a good model for the damping caused by a gas surrounding the solid, structural damping with constant loss angle yields a better description of intrinsic loss mechanisms for a large range of solids and frequencies. Note that if $\phi(\Omega) = \gamma\omega^{-2}\Omega$, the two models are equivalent. By inserting the structural damping term into the left-hand side of Eq. 2.1.8, we obtain through Fourier transform the modified susceptibility

$$\chi(\Omega) = \mathrm{diag}\left(\frac{1}{\mu_i\omega_i^2(1 - i\phi_i(\Omega) - \Omega^2/\omega_i^2)}\right) \tag{2.1.11}$$

while Eq. 2.1.9 remains unchanged. As the susceptibility of a specific eigenmode with low dissipation is often only measurable near its resonance, where ϕ_i can be assumed to depend only weakly on frequency, the mechanical quality factor

$$Q_i = \left|\frac{\chi_i(\omega_i)}{\chi_i(0)}\right| \approx \phi_i(\omega_i)^{-1} \tag{2.1.12}$$

is a useful experimental quantity.

2.1.4 Thermal motion

From the specific heat of solids [LL80], we know that N atoms in a solid at temperature T store a mean thermal energy of $3Nk_BT$. We thus find that if all $3N$ normal modes are in thermal equilibrium, each should hold on average a thermal energy of k_BT. From the virial theorem [GPS02], we expect this energy to be equally distributed over kinetic and potential energy. The mean potential energy corresponding to the displacement $z = (0, \ldots, 0, z_i, 0, \ldots, 0)^T$ is

$$\langle V \rangle = \frac{1}{2} \left\langle x^T(t) K x(t) \right\rangle = \frac{1}{2} \left\langle z^T \Xi^T K \Xi z \right\rangle = \frac{1}{2} \left\langle z^T \mu\omega^2 z \right\rangle = \frac{1}{2} \mu_i \omega_i^2 \left\langle z_i^2 \right\rangle \tag{2.1.13}$$

and we find for the variance of the thermal motion

$$\left\langle z_i^2 \right\rangle_{\mathrm{th}} = \frac{k_B T}{\mu_i \omega_i^2} \tag{2.1.14}$$

This finite motion is inconsistent with the damped motion derived above, which would predict an asymptotic value $\langle z_i^2 \rangle \to 0$. The insertion of a damping term for each mode in section 2.1.3 was therefore incomplete, and has to go with a fluctuating force term that accounts for the damped energy of other modes leaking into mode i. The thermal force spectrum for viscous damping $S_{FF,i}(\Omega) = 2k_BT\mu_i\gamma_i$ is easily verified to account for the above thermal motion: the integral over all frequencies of the resulting displacement spectrum $S_{zz,i} = |\chi(\Omega)|^2 S_{FF,i}(\Omega)$ indeed evaluates to $k_BT/(\mu_i\omega_i^2)$. In more generality, the fluctuation-dissipation theorem [CW51] predicts that the thermal force spectral density at frequency Ω is given by the product of the mean thermal energy of an oscillator at Ω and the generalized resistance $-\frac{1}{\Omega}\operatorname{Im}(\chi(\Omega)^{-1})$. For structural damping and $k_BT \gg \hbar\Omega$, we find

$$S_{FF,i}(\Omega) = -\frac{2k_BT}{\Omega} \operatorname{Im}(\chi_i(\Omega)^{-1}) = \frac{2k_BT}{\Omega} \mu_i \omega_i^2 \phi_i(\Omega)\,. \tag{2.1.15}$$

Issues with integrability of the force spectrum arise in the low-frequency limit if the loss angle ϕ_i is assumed constant for all frequencies [Sau90]. These problems are however of no practical relevance for the high-frequency and low-loss oscillators employed in this thesis.

2.1.5 Effective susceptibility for scalar actuation and measurement

A scalar measurement of mechanical motion x_{m} is usually performed as a weighted average over the position of many atoms and can be specified by a normalized $3N$-dimensional vector v of the weight factors with $v^Tv = 1$:

$$x_{\mathrm{m}}(\Omega) = v^T x(\Omega) = v^T \Xi \chi(\Omega) \Xi^T G(\Omega)\,. \tag{2.1.16}$$

Similarly, the external force applied with a particular actuator can be written as $G(\Omega) = wf(\Omega)$, the product of a normalized force weight vector w and a scalar force amplitude $f(\Omega)$. By defining the effective mass vector $m = (m_i)$ and the effective

susceptibility χ_m for the actuator-sensor pair with

$$m_i = \mu_i \left(\sum_{j=1}^{3N} \Xi_{ji} v_j \sum_{k=1}^{3N} \Xi_{ki} w_k \right)^{-1} \tag{2.1.17}$$

$$\chi_m(\Omega) = \sum_{i=1}^{3N} \frac{1}{m_i \omega_i^2 (1 - i\phi_i(\Omega) - \Omega^2/\omega_i^2)} \tag{2.1.18}$$

Eq. 2.1.16 takes the simple scalar form

$$x_{\mathrm{m}}(\Omega) = \chi_m(\Omega) f(\Omega)\,. \tag{2.1.19}$$

For $\Omega \ll \omega_{N_0}$, the sum in Eq. 2.1.19 is usually truncated after the first N_0 modes and the remaining part of the sum is approximated by a residual term R which is the total DC-response of all higher-frequency modes

$$R = \sum_{i=N_0+1}^{3N} \frac{1}{m_i \omega_i^2} = v^T K^{-1} w - \sum_{i=1}^{N_0} \frac{1}{m_i \omega_i^2}\,. \tag{2.1.20}$$

The last expression allows to obtain the residual term from the stiffness matrix, removing the need to calculate all $3N$ normal modes. For the description of driven motion, neglecting the loss terms in R as done here is usually a good approximation, as mechanical loss is hard to detect far from resonance. The substitution $K \to K(1 - i\phi(\Omega))$ with an identical loss angle ϕ for all residual modes can be employed when the calculation of the thermal force noise spectrum is desired [Lev98]. If rigid body modes are present in the mechanical system, K is not invertible and the presented treatment must be slightly modified [Pre11].

2.1.6 Collocation and optomechanical coupling

We speak of a collocated sensor-actuator pair when $v = w$. In this special case, all effective masses m_i are positive, which restricts the effective susceptibility to the upper half of the complex plane. The phase of the sensor's response to an actuation $\arg \chi_{\mathrm{m}}(\Omega) \in [0, \pi]$ for all Ω therefore never lags by more than π, which is a very useful property for the design of feedback loops (see chapter 6).

The optomechanical coupling realizes a collocated sensor-actuator pair: the phase-shift imprinted on reflected light by the displacement of the atoms near the surface of a solid is the weighted sum of the individual atom displacements with the weight vector given by the transverse intensity distribution of the light beam multiplied by twice the light's wavevector $2k$. The radiation pressure force follows the same transverse distribution [PHH99] with a scalar amplitude given by $f(t) = \frac{2\hbar\omega_{\mathrm{L}}}{c} I(t)$, where c is the speed of light, ω_{L} its radial frequency and $I(t)$ the instantaneous photon flux. The collocated character of the optomechanical coupling breaks down when absorption of light significantly heats the surface. In this case, the generated heat propagates through the solid and creates thermo-elastic forces which are delayed in time and whose spatial distribution is different from the transverse beam profile where the motion is probed by the light's phase-shift.

2.1.7 Thermal noise in the scalar model

The scalar displacement x_m is also subject to thermal noise. To find the equivalent thermal force noise spectrum of the actuator $S_{ff,\mathrm{th}}(\Omega)$ that results in the expected thermal noise when the scalar force is split into an external and a thermal component $f(\Omega) = f_{\mathrm{ext}}(\Omega) + f_{\mathrm{th}}(\Omega)$, we start by computing the measured thermal displacement noise spectrum:

$$\begin{aligned} S_{ff,\mathrm{th}}(\Omega) :&= \frac{S_{x_m x_m}(\Omega)}{|\chi_m(\Omega)|^2} \\ &= \frac{\langle x_m(\Omega) x_m(\Omega)^* \rangle}{|\chi_m(\Omega)|^2} \\ &= \frac{1}{|\chi_m(\Omega)|^2} v^T \Xi \,\mathrm{diag}\left(S_{zz,i}(\Omega)\right) \Xi^T v \\ &= -\frac{2k_B T}{\Omega} \,\mathrm{Im}\left[\sum_{i=1}^{3N} \frac{\left(\sum_{j=1}^{3N} \Xi_{ji} v_j\right)^2 \chi_i(\Omega)^*}{\chi_m(\Omega)\chi_m(\Omega)^*}\right] . \end{aligned} \tag{2.1.21}$$

In the third line, the covariance matrix S_{zz} of thermal noise is assumed to be diagonal, i.e. thermal noise in different modes is taken to be uncorrelated. This assumption is wrong from a fundamental point of view since for an isolated mechanical system, excess thermal energy in one mode is caused by a lack of energy in another mode, which creates correlations. However, the large number of mechanical modes makes the individual correlations between pairs of modes very small and justifies the assumption in practice. For the collocated case $v = w$, the effective force spectrum simplifies with the help of Eqs. 2.1.17 and 2.1.18 to

$$S_{ff,\mathrm{th}}(\Omega) = -\frac{2k_B T}{\Omega} \,\mathrm{Im}\left(\frac{1}{\chi_m(\Omega)}\right) \tag{2.1.22}$$

and we find that the fluctuation-dissipation theorem holds for the scalar displacement x_m, the scalar force f, and effective susceptibility χ_m.

2.1.8 Finite-element method and continuous models

A solid is fundamentally composed of many atoms, and the atomistic point of view has a number of advantages, such as elucidating the origin of thermal fluctuations. It is however often unrealistic and unnecessary to deal with all $3N$ vibration-modes of N atoms. For many situations, approximating the solid as a continuum yields an excellent approximation, with the advantage that methods of differential calculus may be employed to yield specific solutions for simple geometries. For this, the initial equation of motion 2.1.1 must be replaced by the elastic wave equation [LL58], and the mechanical properties of an isotropic solid can be specified by as few parameters as its elastic modulus, mass density, loss angle and Poisson factor instead of $\mathcal{O}(N^2)$ entries of the stiffness matrix. The principal results are however equivalent between the atomistic and continuum models, and are easily transferred from one domain to the other.

Another advantage of the atomistic presentation above is that its results are easily transferred to mechanical lumped element models: the atoms in the presented model can replaced by larger rigid masses with effective springes attached between them to reduce the number of eigenmodes that must be considered. This yields excellent predictions as long as the lowest eigenfrequency of each element lies well above any frequency of interest such that the element may be approximated as a rigid body. The finite-element method [ZTZ13] is a refinement of this idea, making it practical to evaluate the eigenmode problem for a finite number of lumped elements numerically in reasonable computing time. The formulae of continuum theory are employed by the finite element method to find the effective spring couplings between the different elements.

Throughout this thesis, we will employ the finite-element method to find the lowest-frequency eigenmodes of our mechanical structure, to estimate static stress and theoretical transfer functions of our assemblies, and even to study heat propagation through our structure.

2.1.9 Effective mass for special cases

We will compute the effective mass for two typical scenarios in order to gain intuition and to obtain a theoretical estimate for the effective masses in our optomechanical system.

Mass on a spring

Let a mass M_1 be suspended by a spring k_{12} on a rigid support M_2, which may be modeled by a near-infinite mass $M_2 \to \infty$ whose displacement is nearly zero at all times. The system has two eigenmodes of which one describes the center-of-mass motion of the entire system. The remaining oscillation mode has the mode shape $\xi_1 = (1, 0)^T$, the mode mass $\mu_1 = M$ and the frequency $\omega_1 = \sqrt{k_{12}/M}$. If the scalar motion is detected at mass M_1, the effective mass equals the modal mass: $m_1 = M$.

Compression mode of a pillar

To compute the effective mass of the fundamental compression mode of a pillar of constant cross-section with total mass M and length L, we may approximate the pillar as a one-dimensional chain of $N + 1$ atoms. If atom $j \in \{0, 1, \ldots, N\}$ with mass $M/(N + 1)$ is the atom at an equilibrium distance jL/N from one pillar end, we find for component j of the normalized eigenvector ξ_1 of the fundamental compression mode

$$\xi_{1,j} = \sqrt{\frac{2}{N+1}} \cos\left(\frac{j\pi}{N}\right). \tag{2.1.23}$$

With Eq. 2.1.6 we find for the mode mass $\mu_1 = \frac{M}{N+1}$. To compute the effective mass for displacement actuation and detection at the end face $j = 0$, we use Eq. 2.1.18. We find that the effective mode mass equals half the physical mass:

$$m_1 = \mu_1 \xi_{1,0}^{-2} = \frac{M}{2}. \tag{2.1.24}$$

The same result $m_n = M/2$ holds for harmonics n of the fundamental mode, since the only difference in the calculation is an integer pre-factor in the argument of the cosine in Eq. 2.1.23.

2.2 Optics

This section theoretically derives expressions to model the transformation of an incident laser beam by our experiment until it is finally absorbed by the surface of a photodiode. We begin by defining the electric field of a laser beam (section 2.2.1) and discover that the field can described by its time-dependent complex field amplitude α alone (section 2.2.2). In particular, we will recall the field corresponding to amplitude, phase- and frequency modulations. After a brief reminder of the formalism to describe random signals, we will discuss how both quantum and classical optical noise is expressed by the field amplitude (section 2.2.3) and the optical quadrature spectra. With this mathematical toolbox we are ready to find the transformation of the input laser accomplished by an optomechanical cavity (section 2.2.4). The thus obtained equation of motion for the intracavity field is then either directly solved for the particular case of a linear sweep over the optical resonance (section 2.2.5), or it is expressed in the form of the optical susceptibility for stationary incident field quadratures via the Fourier transform (section 2.2.6). With the analytical expressions of the reflected field, we conclude this section by discussing how information about the fluctuations of the incident laser and the mechanical system are encoded in the measured photocurrent for a number of detection schemes (section 2.2.7).

2.2.1 Gaussian modes of the electric field

The classical electric field $\vec{E}(x, y, z, t)$ is a real-valued three-vector defined at each point in four-dimensional spacetime with coordinates x, y, z, t. The monochromatic light emitted by most lasers occupies only one spatio-temporal mode of the electric field. It oscillates - up to slow modulations $\alpha(t)$ - at the optical frequency ω_L, has the spatial shape of a Gaussian beam $E(x,y,z)$, and is linearly polarized in a particular direction $\hat{p}$ orthogonal to the direction of propagation:

$$\vec{E}(x, y, z, t) = \alpha(t) e^{i\omega_\mathrm{L} t} E(x,y,z)\, \hat{p} + c.c.\,. \tag{2.2.1}$$

All quantities on the r.h.s of Eq. 2.2.1 are complex phasors which are used to express the phase-dependence of the electric field. Since the actual electric field is a real-valued vector, the complex conjugate (*c.c.*) of the r.h.s. must be added to cancel out the imaginary part of the field.

The phasor of the electric field of a TEM_{00} Gaussian beam propagating along the z-direction as a function of the three spatial coordinates x, y, z is given by the expression [Wik16a, KL66]

$$E(x, y, z) = E_0 \frac{w_0}{w(z)} \exp\left(-\frac{x^2 + y^2}{w(z)^2}\right) \exp\left(-ikz - ik\frac{x^2 + y^2}{2R(z)} + i\psi(z)\right) \tag{2.2.2}$$

with the following definitions:

- the wavelength $\lambda = \frac{c}{\omega}$
- the wavenumber $k = \frac{2\pi}{\lambda}$
- the electric field at the origin E_0
- the minimum beam waist w_0
- the Rayleigh range $z_R = \frac{\pi w_0^2}{\lambda}$
- the beam waist along the beam axis $w(z) = w_0\sqrt{1 + \left(\frac{z}{z_R}\right)^2}$
- the radius of curvature $R(z) = z\left[1 + \left(\frac{z_R}{z}\right)^2\right]$
- the Gouy phase $\psi(z) = \arctan\left(\frac{z}{z_r}\right)$.

Higher-order transverse Hermite-Gaussian modes TEM_{lm} have an additional spatial profile given by Hermite-Gauss polynomials H_l and an additional Gouy phase depending on the mode order:

$$E_{l,m}(x,y,z) = E(x,y,z)H_l\left(\frac{\sqrt{2}x}{w(z)}\right)H_m\left(\frac{\sqrt{2}y}{w(z)}\right)e^{i(l+m)\psi(z)} \tag{2.2.3}$$

2.2.2 Classical field amplitude

Unless stated otherwise, we will from now on consider the electric field of a TEM_{00} mode of a Nd:YAG laser at a wavelength $\lambda = 1064\,\mathrm{nm}$, whose waist position, direction and polarization are perfectly matched to the spectrally closest optical mode of the Fabry-Perot cavity under consideration [KL66]. All useful information about the beam is then contained in the complex field amplitude $\alpha(t)$, which varies slowly on a timescale defined by the optical frequency ω_L. Setting

$$E_0 = \sqrt{\frac{2\hbar\omega_\mathrm{L}}{\epsilon_0 c\pi w_0^2}}\ , \tag{2.2.4}$$

where ϵ_0 denotes the vacuum permittivity, results in a normalization of the field amplitude such that

$$I(t) = |\alpha(t)|^2 = \alpha(t)\alpha^*(t) \tag{2.2.5}$$

corresponds to the flux of photons in the beam, i.e. the number of photons flowing through a surface normal to the direction of the laser beam per unit time, each carrying an energy of $\hbar\omega$.

Amplitude modulation

If the magnitude of the mean field amplitude $\bar{\alpha}$ is modulated at a frequency Ω with a relative modulation depth m, we get:

$$\alpha(t) = \bar{\alpha}\left(1 + m\cos(\Omega t)\right) = \bar{\alpha} + \frac{m\bar{\alpha}}{2}e^{i\Omega t} + \frac{m\bar{\alpha}}{2}e^{-i\Omega t}\,. \tag{2.2.6}$$

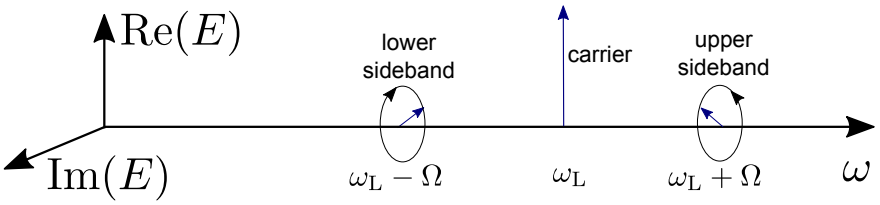


Figure 2.1: Illustration of the sideband picture, showing a carrier with two sideband fields.

Having already factored out the laser frequency phasor $e^{i\omega_L t}$ in Eq. 2.2.1, the DC-term $\bar{\alpha}$ corresponds to the carrier at the optical frequency. The amplitude modulation adds two fields at the frequencies $+\Omega$ and $-\Omega$ with respect to the carrier, corresponding to optical fields at $\omega_L + \Omega$ and $\omega_L - \Omega$. Figure 2.1 visualizes the three optical fields: the carrier field amplitude remains constant, while the upper and lower sidebands rotate as time progresses in the complex plane clockwise and anti-clockwise, respectively. In the case $m \ll 1$, the carrier is the dominant field and the resulting optical flux $I(t)$ corresponds to the projection of the sum of the three phasors onto the carrier, which leads to oscillations of the flux at the modulation frequency Ω:

$$I(t) \approx \bar{I}\,(1 + 2m\cos(\Omega t))\ . \tag{2.2.7}$$

Phase modulation

A phase-modulated laser beam with modulation amplitude ϕ at frequency Ω can be written as the sum of a carrier and a series of sideband fields by employing the Jacobi-Anger expansion

$$\alpha(t) = \bar{\alpha}e^{i\phi\cos(\Omega t)} = \bar{\alpha}J_0(\phi) + \bar{\alpha}\sum_{n=1}^{\infty} i^n J_n(\phi)(e^{in\Omega t} + e^{-in\Omega t}) \approx \bar{\alpha} + \frac{i\phi\bar{\alpha}}{2}e^{i\Omega t} + \frac{i\phi\bar{\alpha}}{2}e^{-i\Omega t} \tag{2.2.8}$$

where J_n denotes the n-th Bessel function of the first kind and the last approximation is only valid for $\phi \ll 1$. As expected for a phase modulation, a constant flux of photons $I(t) = \bar{I}$ results from the modulation since the two sidebands always cancel out when they are collinear with the carrier in the complex plane.

Frequency modulation

A phase-modulated laser beam can be seen in a different way by defining the instantaneous frequency $\omega_i(t) = \omega + \frac{\mathrm{d}(\arg\alpha(t))}{\mathrm{d}t}$, leading to

$$\omega_i(t) = \omega + \phi\Omega\sin(\Omega t)\,. \tag{2.2.9}$$

Thus, a phase modulation at Ω with phase amplitude ϕ is equivalent to a frequency modulation at the same frequency of frequency amplitude $\phi\Omega$. Depending on the optical device probed by a modulated beam, either the phase modulation or the frequency modulation viewpoint can be more advantageous.

Linear optical components

We define an optical component as a device that transforms an optical input field into an output field. If the component is linear, the sum of the output fields resulting from a number of input fields equals the output field resulting from the sum of these input fields. This class of devices is important since it encompasses a large amount of optical devices in the optical signal treatment chain of typical experiments. The only nonlinear optical component considered in this thesis will be the optomechanical cavity, which will however be linearized in order to allow for an easier mathematical treatment.

2.2.3 Optical noise

For real lasers, both the modulus and the phase of the field amplitude fluctuate randomly over time. To account for these fluctuations in the modeling of an experiment, the best way is to directly measure their time-trace during the experiment and include these data in the processing of the measurement results. This however puts cumbersome constraints on the design of an experiment and is fundamentally impossible for quantum fluctuations. Therefore, one typically resorts to a statistical description of these random fluctuations by using correlation functions and noise spectra. This approach requires the assumption that the noise process is stationary, i.e. that the statistical properties of the signal are time-independent [CDG$^+$10].

Random signals

Each realization of a stochastic process x is a function $x(t)$ that corresponds to a random sample from a statistical ensemble of such functions [Asp07]. In order to obtain quantities that do not depend on the specific realization, one must resort to ensemble averages of functions $f(x)$ of x denoted by $\langle f(x(t))\rangle$. When the process is stationary, expectation values are independent of time: $\langle f(x(t))\rangle = \langle f(x)\rangle$. A stochastic process is ergodic if the time-average of a single realization converges for long averaging times to an ensemble average: $\langle x\rangle = \lim_{T\to\infty} \frac{1}{2T}\int_{-T}^{T} x(t)\mathrm{d}t$. We assume that all processes we will encounter will obey stationarity and ergodicity.

We define the mean value $\bar{x} = \langle x(t)\rangle$ and the variance $\Delta x^2 = \langle |x(t) - \bar{x}(t)|^2\rangle$ of x. Since the variance neglects the detailed frequency-dependence of the fluctuations, we further define for two random real-valued signals $x(t)$ and $y(t)$ the correlation function

$$C_{xy}(\tau) = \langle x(t)y(t+\tau)\rangle - \bar{x}\bar{y} \tag{2.2.10}$$

and the cross-spectrum as the Fourier transform of the correlation function C_{xy}

$$S_{xy}(\Omega) = \int_{-\infty}^{\infty} C_{xy}(\tau)e^{i\Omega\tau}\mathrm{d}\tau\,. \tag{2.2.11}$$

In the most common case where $x(t) = y(t)$, we speak of the autocorrelation function C_{xx} and the spectrum S_{xx} of x. A straightforward calculation [CDG$^+$10], often termed Wiener-Khinchin theorem, leads to the useful identity

$$S_{xy}(\Omega) = \lim_{T\to\infty} \langle x_T(\Omega)x_T(-\Omega)\rangle \tag{2.2.12}$$

with the windowed Fourier transform[1] defined as

$$x_T(\Omega) = \frac{1}{\sqrt{2T}} \int_{-T}^{T} (x(t) - \bar{x})\, e^{i\Omega t} \mathrm{d}t\,. \tag{2.2.13}$$

Finally, the variance is related to the spectrum via

$$\Delta x^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} S_{xx}(\Omega) \mathrm{d}\Omega\,. \tag{2.2.14}$$

Quantum noise of light

In the quantum description of light [BR04], the complex field amplitude α and its complex conjugate α^* are replaced by the annihilation operator $\hat{a}$ and its hermitian conjugate, the creation operator $\hat{a}^\dagger$, with the commutation relations

$$[\hat{a}(t), \hat{a}(t')] = [\hat{a}^\dagger(t), \hat{a}^\dagger(t')] = 0\,, \quad [\hat{a}(t), \hat{a}^\dagger(t')] = \delta(t - t')\,. \tag{2.2.15}$$

The quantum equivalent of an ensemble average of functions of the field amplitudes is the expectation value of the same functions of the annihilation and creation operators for the quantum state $|\Psi\rangle$ of the system:

$$\left\langle f\left(\alpha(t), \alpha^*(t')\right)\right\rangle \rightarrow \langle\Psi| f\left(\hat{a}(t), \hat{a}^\dagger(t')\right)|\Psi\rangle\,. \tag{2.2.16}$$

The classical picture consisted in the approximation that all quantum commutators are zero. In this case, the "$\rightarrow$" in Eq. 2.2.16 can be replaced by a "=". Otherwise, when there are nonzero commutators as for the creation and annihilation operators of light, the order of terms in products appearing in the function f makes a difference which is not present in the classical picture. This order is dictated by the measurement whose expectation value is to be computed and must be obtained from a physical consideration. The only quantum measurement we will perform on light in this thesis will be photodetection, whose associated quantum operator is the photon flux $\hat{I}(t) = \hat{a}^\dagger\hat{a}$. The lowest optical power measurable with conventional photodetectors lies in the range of a few nW, corresponding to mean photon fluxes $\bar{I} \geq 10^{10}\,\mathrm{s}^{-1}$ in the visible spectrum. The empirically observable smallness of quantum fluctuations suggests to expand the field operators into $\hat{a}(t) = \bar{\alpha} + \delta\hat{a}(t)$ with a classical mean field $\bar{\alpha} = |\bar{\alpha}|e^{i\bar{\phi}}$ and fluctuations $\delta\hat{a}(t)$. The linearization the fluctuations of the photon flux operator is then justified for large mean field amplitudes to a relative error of the order of $|\bar{\alpha}|^{-1} \leq 10^{-5}$:

$$\delta\hat{I}(t) = \hat{I}(t) - \bar{I} \approx \bar{\alpha}^*\hat{a}(t) + \bar{\alpha}\hat{a}^\dagger(t) = |\bar{\alpha}|\hat{X}_{\bar{\phi}}(t) \tag{2.2.17}$$

with the introduction of the quadrature operator $\hat{X}_\phi$ at angle ϕ

$$\hat{X}_\phi = e^{-i\phi}\hat{a} + e^{i\phi}\hat{a}^\dagger\,. \tag{2.2.18}$$

[1] Using the windowed Fourier transform makes the integrand for the majority of physically interesting signals integrable, which is not the case for the ordinary Fourier transform of nonzero stationary stochastic functions. When the index T of the Fourier transform of a stochastic function is omitted, we mean the limit $T \rightarrow \infty$ of the windowed Fourier transform.

In this approximation, the photon flux operator is linear in the field operators, and therefore no commutator appears in the final result. This allows to substitute all quantum operators $\delta\hat{a}$, $\delta\hat{I}$, $\hat{X}_\phi$ by classical random processes $\delta\alpha$, δI, X_ϕ whose statistical ensemble corresponds to the one of the underlying quantum state, and to compute expectation values as classical ensemble averages. In a broader context, the above procedure is known as the semi-classical approximation [RH89, RHGF92].

Coherent, vacuum and squeezed states

Except for the last chapter of this thesis, the input state of our experiment will always be a coherent state $|\alpha\rangle$ of light. This state results from the superposition of the quantum vacuum with a classical field of amplitude α. The resulting mean field is $\bar{\alpha} = \alpha$ and the variance of any quadrature is independent of the quadrature angle ϕ

$$\Delta X_\phi^2 = 1\,. \tag{2.2.19}$$

The fluctuations of a coherent state at different times are uncorrelated, leading to the autocorrelation function $C_{X_\phi X_\phi}(\tau) = \delta(\tau)$ and a white quadrature spectrum. Consequently, the photon flux noise spectrum is $S_{II}(\Omega) = \bar{I}$.

A special case of a coherent state is the vacuum state of light $|0\rangle$, corresponding to a field amplitude of zero. While the fluctuations of this state by itself are undetectable with photodiodes, superpositions of coherent states and the vacuum state are routinely produced at the output ports of beamsplitters illuminated at only one input port with a coherent state from a laser. Since we will model optical loss as a partially transmitting beamsplitter, vacuum entering into bright beams that are eventually measured by a photodiode is unavoidable in real experiments.

As a last example, squeezed states of light $|\alpha, \xi\rangle$ are states for which the quadrature fluctuations are below the level of a coherent state for certain quadrature angles [BR04]. The Heisenberg uncertainty relation for two orthogonal quadratures then requires excess fluctuations with respect to a coherent state at orthogonal quadratures. While the mean field of a squeezed state is essentially α as for a coherent state, the quadrature variances are

$$\Delta X_\phi^2 = \cosh(2r_s) - \sinh(2r_s)\cos(2(\phi - \phi_s)) \tag{2.2.20}$$

with the squeezing factor r_s and squeezing angle ϕ_s defined by $\xi = r_s e^{i2\phi_s}$. If the mean-field $\bar{\alpha}$ is in-phase with the squeezing angle, the photon flux noise spectrum is minimal and amounts to only e^{-2r_s} times the one of a coherent beam, while for mean-fields in quadrature with the squeezing angle, the squeezed light's intensity noise spectrum is e^{2r_s} times higher than a coherent state spectrum.

Classical optical noise

In excess of quantum fluctuations, technical imperfections may cause additional noise of both the phase and the amplitude of light. Sources of technical noise include the variation of material properties related to temperature-fluctuations, finite lifetime effects of the lasing medium, relaxation oscillations of the laser, classical fluctuations of

the effective pathlength crossed by the laser beam, vibrations of the optical components, and electrical noise at the ports of optical modulators. All these sources can be modeled as a phase and amplitude modulator driven as to imprint a random phase $\phi(t)$ and relative amplitude modulation depth $m(t)$ onto the laser beam (see section 2.2.2). Denoting the quantum fluctuations with $\delta\alpha_{\mathrm{q}}$ and assuming all fluctuations to be small compared to the mean field, we can write

$$\alpha(t) = (\bar{\alpha} + \delta\alpha_{\mathrm{q}}(t))\,(1 + m(t))\,e^{i\phi(t)} \approx \bar{\alpha} + \delta\alpha_{\mathrm{q}}(t) + \bar{\alpha}m(t) + i\bar{\alpha}\phi(t)\,. \tag{2.2.21}$$

While the quantum fluctuation term does not depend on the mean field strength, the classical modulations do. Therefore, any nonzero classical laser noise is expected to dominate over quantum noise above a threshold field strength. To characterize classical fluctuations, it is therefore advantageous to perform measurements at high optical power. Contrarily, the experimenter must ensure that experiments probing quantum limits are conducted at powers where classical noise is negligible.

The classical noise can be characterized by the spectra $S_{\phi\phi}(\Omega)$ and $S_{mm}(\Omega)$. Often, the classical noise on the phase and amplitude of the field is correlated, urging us to further specify the classical phase-amplitude correlations with the cross-spectrum $S_{m\phi}(\Omega)$. While the amplitude noise can be estimated simply by measuring the photocurrent spectrum of a laser beam, the other two spectra require a more complicated measurement scheme. In a later section, we will use a Fabry-Perot cavity to obtain estimates of all three noise spectra of our laser source.

Noise in Fourier domain

When writing the fluctuating field of Eq. 2.2.21 in frequency domain through the windowed Fourier transform[2]

$$\alpha(\Omega) = 2\pi\bar{\alpha}\delta(\Omega) + \delta\alpha_{\mathrm{q}}(\Omega) + \bar{\alpha}m(\Omega) + i\bar{\alpha}\phi(\Omega)\,, \tag{2.2.22}$$

we first note the carrier of Fourier amplitude $2\pi\bar{\alpha}$ at Fourier frequency 0, corresponding to the optical frequency ω. Furthermore, each Fourier component of the fluctuations at frequency Ω constitutes an optical sideband at the optical frequency $\omega + \Omega$. When a noisy laser beam is sent through a linear optical device, we can compute the effect of the device on each sideband frequency Ω individually and later sum the output fields together to find the noise transmission of the device. This procedure will allow to write formulae more compactly in the following sections.

To establish a link between phase and amplitude fluctuation spectra and the spectra of the optical quadratures, we define for a beam of mean amplitude $\bar{\alpha} = |\bar{\alpha}|e^{i\bar{\phi}}$ the Fourier-transformed field quadratures[3]

$$X_\phi(\Omega) = e^{-i\phi}\delta\alpha(\Omega) + e^{i\phi}\delta\alpha^*(\Omega) = e^{-i\phi}\delta\alpha(\Omega) + \left(e^{-i\phi}\delta\alpha(-\Omega)\right)^*\,. \tag{2.2.23}$$

[2]By omitting the index T of the windowed Fourier transform, we implicitly form the limit of $T \to \infty$ here and in later sections that make use of the Fourier transform of random functions.

[3]Note that the distinction between the complex conjugate of the Fourier transform and the Fourier transform of the complex conjugate can be confusing. We find $\alpha^*(\Omega) = (\alpha(-\Omega))^*$.

For the field in Eq. 2.2.21, the amplitude and phase quadratures $p = X_{\bar{\phi}}$ and $q = X_{\bar{\phi}+\pi/2}$ read:

$$p(\Omega) = e^{-i\bar{\phi}}\delta\alpha(\Omega) + e^{i\bar{\phi}}\delta\alpha(-\Omega)^* = p_q(\Omega) + 2|\bar{\alpha}|m(\Omega) \tag{2.2.24}$$

$$q(\Omega) = -ie^{-i\bar{\phi}}\delta\alpha(\Omega) + ie^{i\bar{\phi}}\delta\alpha(-\Omega)^* = q_q(\Omega) + 2|\bar{\alpha}|\phi(\Omega) \tag{2.2.25}$$

Using the Wiener-Khinchin theorem from Eq. 2.2.12, we find

$$S_{pp}(\Omega) = 4\bar{I}S_{mm}(\Omega) + 1 \tag{2.2.26}$$

$$S_{qq}(\Omega) = 4\bar{I}S_{\phi\phi}(\Omega) + 1 \tag{2.2.27}$$

$$S_{pq}(\Omega) = S_{qp}^*(\Omega) = 4\bar{I}S_{m\phi}(\Omega)\,. \tag{2.2.28}$$

The number 1 in the first two equations is the quantum noise on the amplitude and phase quadratures. This term is absent from the third equation, since the quantum noise on orthogonal quadratures is uncorrelated for a coherent state. Finally, the photon flux noise spectrum reads

$$S_{II}(\Omega) = \bar{I}S_{pp}(\Omega) = \bar{I} + 4\bar{I}^2S_{mm}(\Omega)\,. \tag{2.2.29}$$

2.2.4 The electric field inside a Fabry-Perot cavity

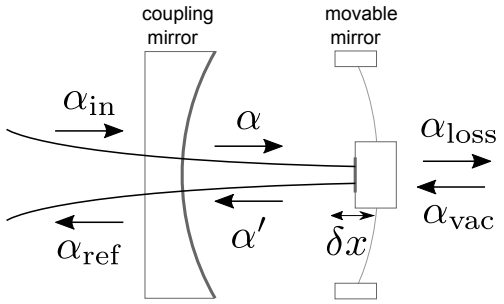


Figure 2.2: Various fields interacting with an optomechanical Fabry-Perot cavity.

In order to compute the electric field inside a Fabry-Perot cavity, we need to introduce the following four electric field amplitudes (cf. Fig. 2.2):

- the incident field α_{in} coming from the laser and oscillating at the frequency $\omega_{\text{L}} = \frac{c}{\lambda}$
- the intracavity field α
- the reflected field α_{ref}
- and the vacuum field α_{vac} which enters the cavity through the loss port.

We model the two cavity mirrors as unitary beam-splitters, which will later enable us to express the reflected field of the Fabry-Perot cavity as a function of the incident

and vacuum field amplitudes only. We assume the input coupler to have a power transmission T and the loss port to lose a fraction P of the power reflected off it.

In a cavity of length $L + \delta x(t)$, the propagating laser beam acquires a phase-shift $\phi = 2kL + 2k\delta x(t)$ per round trip. When the cavity is operated close to a resonance, the phase-shift $\phi = 2n\pi + \delta\phi + 2k\delta x(t)$ lies in the vicinity of an integer multiple of 2π and we can use the approximation $\exp(i(\phi + 2k\delta x(t))) \approx 1 + i\delta\phi + i2k\delta x(t)$. Additionally during each round-trip, the fraction $\sqrt{P}$ of the intracavity field leaves the cavity through the loss port, similar to the fraction $\sqrt{T}$ leaving through the input coupler. In exchange for this, the fraction $\sqrt{T}$ of the incident field and the fraction $\sqrt{P}$ of the vacuum field make their way into the cavity and must be added to the intracavity field during each reflection off the corresponding mirror. If $\tau = 2L/c$ denotes the cavity round-trip time, this leads to the following difference equation for the evolution of the intracavity field:

$$\alpha(t+\tau) = \alpha(t)(1 + i\delta\phi + i2k\delta x)\sqrt{1-T}\sqrt{1-P} + \sqrt{T}\alpha_{\mathrm{in}}(t) + \sqrt{P}\alpha_{\mathrm{vac}}(t). \quad (2.2.30)$$

For a high-finesse cavity, $T + P \ll 1$ and $\sqrt{1-T}\sqrt{1-P} \approx 1 - \frac{T+P}{2}$. With the definitions

$$\kappa = \frac{T+P}{\tau} \quad (2.2.31)$$

$$\eta = \frac{T}{T+P} \quad (2.2.32)$$

$$\kappa_{\mathrm{in}} = \frac{T}{\tau} = \eta\kappa \quad (2.2.33)$$

$$\kappa_0 = \frac{P}{\tau} = (1-\eta)\kappa \quad (2.2.34)$$

$$\Delta = \frac{\delta\phi}{\tau} \quad (2.2.35)$$

$$G = \frac{2k}{\tau} = \frac{\omega_{\mathrm{L}}}{L} \quad (2.2.36)$$

and the approximation that the cavity length fluctuations $\delta x(t)$ and the intracavity field fluctuations around its mean value $\bar{\alpha}$ are small enough to assume $\alpha(t)\delta x(t) \approx \bar{\alpha}\delta x(t)$, the difference equation can be re-written as

$$\frac{\alpha(t+\tau) - \alpha(t)}{\tau} = \left(i\Delta - \frac{\kappa}{2}\right)\alpha(t) + \sqrt{\frac{\kappa_{\mathrm{in}}}{\tau}}\alpha_{\mathrm{in}}(t) + \sqrt{\frac{\kappa_0}{\tau}}\alpha_{\mathrm{vac}}(t) + i\bar{\alpha}G\delta x(t)\,. \quad (2.2.37)$$

Finally, the limit of a high-finesse cavity causes all rates to be small compared to $\frac{1}{\tau}$, which justifies the limit $\tau \to 0$ and yields the differential equation

$$\dot{\alpha}(t) = \left(i\Delta - \frac{\kappa}{2}\right)\alpha(t) + \sqrt{\frac{\kappa_{\mathrm{in}}}{\tau}}\alpha_{\mathrm{in}}(t) + \sqrt{\frac{\kappa_0}{\tau}}\alpha_{\mathrm{vac}}(t) + i\bar{\alpha}G\delta x(t)\,. \quad (2.2.38)$$

The reflected field results from the interference of the fraction $\sqrt{T}$ of the intracavity field and $\sqrt{1-T} \approx 1$ times the incident field that is directly reflected off the input coupler. Unitarity of the input coupler requires these two terms to be added with opposite signs:

$$\alpha_{\mathrm{ref}}(t) = \alpha_{\mathrm{in}}(t) - \sqrt{\kappa_{\mathrm{in}}\tau}\alpha(t). \quad (2.2.39)$$

2.2.5 Linear sweep over a resonance

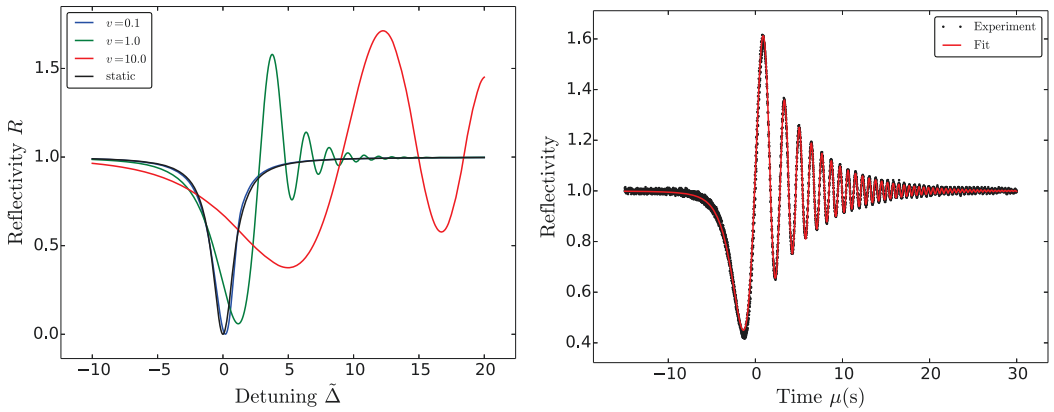


Figure 2.3: Left: Reflectivity of a critically-coupled Fabry-Perot cavity for different normalized sweep velocities v (colored) along with the static reflectivity (black) versus detuning in units of linewidths. Right: Fit of an experimental reflectivity trace with formula 2.2.43.

As an example use-case of the equations from the previous section, we will compute the intensity reflected off a high-finesse cavity for a linear sweep of the laser-cavity detuning $\Delta(t) = \frac{\partial\omega_c}{\partial t}t$ with constant incident field α_{in}. We normalize the sweep velocity in natural units v of cavity bandwidths $\kappa/2$ per lifetime $\frac{1}{\kappa}$ to obtain $\frac{\partial\omega_c}{\partial t} = v\frac{\kappa^2}{2}$. We will neglect the vacuum input field as its amplitude is small and its average zero for multiple sweeps and set $\delta x(t) = 0$. This leaves us to solve the following equation, which can be obtained from Eq. 2.2.38:

$$\dot{\alpha}(t) = \left(iv\frac{\kappa^2}{2}t - \frac{\kappa}{2}\right)\alpha(t) + \sqrt{\frac{\kappa_{\text{in}}}{\tau}}\alpha_{\text{in}}. \tag{2.2.40}$$

Noting that this ordinary differential equation has the form $\dot{\alpha}(t) + P(t)\alpha(t) = Q(t)$, we can multiply both sides by the integrating factor [MF53, Wik16b] $M(t) = e^{\int_{-\infty}^{t} P(s)\mathrm{d}s}$ to allow the resulting equation to be integrated, since the left-hand side becomes the total derivative of $\alpha(t)M(t)$ if $\dot{M}(t) = P(t)M(t)$. Together with the initial condition $\alpha(-\infty) = 0$, this yields the equation

$$\alpha(t) = \sqrt{\frac{\kappa_{\text{in}}}{\tau}}\alpha_{\text{in}}\frac{1}{M(t)}\int_{-\infty}^{t} M(s)\mathrm{d}s \tag{2.2.41}$$

with the integrating factor

$$M(t) = \exp\left(-iv\frac{\kappa^2}{4}t^2 + \frac{\kappa}{2}t\right). \tag{2.2.42}$$

While the solution can be expressed more compactly with the complex error function, the presented form can be integrated numerically, which is found to be faster and still yields a better numerical precision than using the error function. Equations 2.2.5

and 2.2.39 yield an expression for the cavity reflectivity $R = I_{\mathrm{ref}}/I_{\mathrm{in}}$ as a function of the normalized static detuning $\tilde{\Delta} = \frac{\Delta}{\kappa/2} = v\kappa t$:

$$R(\tilde{\Delta}) = \left| 1 - \kappa_{\mathrm{in}} \frac{1}{M(\frac{1}{v\kappa}\tilde{\Delta})} \int_{-\infty}^{\frac{1}{v\kappa}\tilde{\Delta}} M(s)\mathrm{d}s \right|^2 . \tag{2.2.43}$$

Alternatively to the above procedure, we can obtain the well-known Lorentzian dip for the static limit $\dot{\alpha}(t) \to 0$ directly from Eq. 2.2.40.

$$R(\tilde{\Delta}) = 1 - \frac{4\eta(1-\eta)}{1+\tilde{\Delta}^2} \tag{2.2.44}$$

which can become zero only for critical coupling, i.e. if the input coupler transmission T equals the cavity loss P and thus $\eta = 0.5$. Figure 2.3 shows the reflectivity for various sweep velocities and the static approximation. One can see that for slow sweep velocities, the static approximation is sufficient. For normalized sweep velocities close to 1 or even larger, we note a chirped oscillation in the reflected photon flux. This oscillation is the beating between the leaking intracavity field, which has a frequency close to the cavity resonance frequency, and the reflected field whose frequency increases linearly with time as we sweep the laser. While the oscillation decays exponentially with a decay time given by the inverse cavity linewidth, the oscillation phase increases quadratically with time, as expected for a linear increase of the difference frequency. This effect has been discussed extensively in the literature [PBVF97, LWH$^+$99].

2.2.6 Optical susceptibility of a Fabry-Perot cavity

Rarely, the equation of motion of the intracavity field is solved directly as in the previous section. Many times, the physical situation can be well approximated by a sum of stationary input fields. In this case, Fourier-transforming the equation of motion transforms the differential equation for the intracavity field into an algebraic one. This procedure works as long as no nonlinear element is introduced into the cavity, or when a linearized treatment of nonlinear components is feasible.

Susceptibility for the complex field amplitude

By defining the optical susceptibility of a Fabry-Perot cavity as

$$\chi_{\mathrm{c}}(\Omega) = \frac{\kappa/2}{\kappa/2 - i(\Delta + \Omega)} , \tag{2.2.45}$$

we can write the equation of motion of the intracavity field 2.2.38 in Fourier domain

$$\alpha(\Omega) = \sqrt{\frac{2}{\pi}\mathcal{F}\eta}\chi_{\mathrm{c}}(\Omega)\alpha_{\mathrm{in}}(\Omega) + \sqrt{\frac{2}{\pi}\mathcal{F}(1-\eta)}\chi_{\mathrm{c}}(\Omega)\alpha_{\mathrm{vac}}(\Omega) + i\sqrt{\frac{2}{\pi}\mathcal{F}\eta}\chi_{\mathrm{c}}(0)\bar{\alpha}_{\mathrm{in}}\frac{4\mathcal{F}}{\lambda}\chi_{\mathrm{c}}(\Omega)\delta x(\Omega) \tag{2.2.46}$$

where we have assumed that the mean incident field in the optomechanical coupling term is given by the zero-frequency component of the incident field. Apart from the constant pre-factors, the intracavity field equals the input fields multiplied by the cavity susceptibility. The susceptibility is a complex Lorentzian whose squared magnitude

has a FWHM of κ, and which is centered around $\Omega = -\Delta$, which corresponds to the cavity resonance frequency $\omega_c = \omega_\mathrm{L} - \Delta$. As expected for a Lorentzian, the intracavity field undergoes a phase shift of π over the resonance. The finesse $\mathcal{F} = \frac{2\pi}{\tau}/\kappa$ is the ratio between the free spectral range and the FWHM of a cavity resonance. We see that the resonant intracavity photon flux $|\alpha(0)|^2$ is enhanced with respect to the incident one by a factor $\frac{2}{\pi}\eta\mathcal{F}$. The remaining fraction $\frac{2}{\pi}(1-\eta)\mathcal{F}$ comes from the vacuum field. Since a linear cavity does not alter the frequency of the input fields, Eq. 2.2.46 allows us to directly compute the cavity response for all sideband frequencies independently. Importantly, not only the incident mean field is amplified by the finesse, but also the incident classical and quantum fluctuations that lie sufficiently close to the resonance. From the point of view of the cavity, both the incident vacuum fluctuations as well as incident classical laser noise appear similarly as classical drives. In order to enhance the amplitude of quantum fluctuations, it is therefore much more efficient to amplify the field using a cavity than to simply increase the laser power.

For the reflected field, we obtain a similar algebraic equation:

$$\alpha_\mathrm{ref}(\Omega) = (1 - 2\eta\chi_\mathrm{c}(\Omega))\,\alpha_\mathrm{in}(\Omega) - 2\sqrt{\eta(1-\eta)}\chi_\mathrm{c}(\Omega)\alpha_\mathrm{vac}(\Omega) - i\frac{8\eta\mathcal{F}\bar{\alpha}_\mathrm{in}}{\lambda}\chi_\mathrm{c}(0)\chi_\mathrm{c}(\Omega)\delta x(\Omega)\,. \tag{2.2.47}$$

Since the reflected field is the superposition of the input field and Lorentzian intracavity field, its squared magnitude has a Lorentzian dip due to the power lost in the cavity when illuminated near resonance. The reflected field phase changes by 2π over the resonance, since the phase far below and far above resonance, where the incident light cannot interact significantly with the cavity, must coincide modulo 2π. Energy conservation dictates that the reflected field cannot be enhanced by the cavity, thus factors containing the finesse are missing in the reflected field strength. The optomechanical coupling however is enhanced by the finesse, which can be seen as a consequence of the intracavity field acquiring the optomechanical phase-shift as many times as it is reflected off the movable mirror.

Susceptibility for the field quadratures

We have seen in Eq. 2.2.17 that the fluctuations of the photon flux $\delta I(\Omega)$ can be conveniently calculated from the amplitude field quadrature $p(\Omega)$ as defined in Eq. 2.2.24. Furthermore, the input field fluctuations in phase and relative amplitude are naturally decomposed into fluctuations of the amplitude and phase quadratures (cf. Eqs. 2.2.26 and 2.2.27). It is therefore useful to find equations that relate the quadratures of the input fields p_in and q_in as well as the length fluctuations δx to the amplitude and phase quadratures of the reflected field p_ref and q_ref. Since the equations that connect the input and reflected fields are linear, we can write the transformation performed by an optomechanical Fabry-Perot cavity as a matrix multiplication of the vector of input quadratures by a transformation matrix $R(\Omega)$ to obtain the vector of reflected

quadratures[4]:

$$\begin{pmatrix} p_{\mathrm{ref}}(\Omega) \\ q_{\mathrm{ref}}(\Omega) \end{pmatrix} = \begin{pmatrix} R_{11} & R_{12} & R_{13} & R_{14} & R_{15} \\ R_{21} & R_{22} & R_{23} & R_{24} & R_{25} \end{pmatrix} \times \begin{pmatrix} p_{\mathrm{in}}(\Omega) \\ q_{\mathrm{in}}(\Omega) \\ p_{\mathrm{vac}}(\Omega) \\ q_{\mathrm{vac}}(\Omega) \\ \delta x(\Omega) \end{pmatrix} \tag{2.2.48}$$

Eq. 2.2.47 can be used to obtain the ten coefficients of the matrix $R(\Omega)$. As the results become quite lengthy for the most general cases, we use symbolic algebra computer programs such as Mathematica to obtain analytic results. The employed Mathematica code can be found in the Appendix A.

Quadrature spectra

The spectra of the reflected field quadratures are the elements of the quadratures' covariance matrix, which are easily obtained from Eq. 2.2.48:

$$\begin{pmatrix} S_{pp,\mathrm{ref}}(\Omega) & S_{pq,\mathrm{ref}}(\Omega) \\ S_{qp,\mathrm{ref}}(\Omega) & S_{qq,\mathrm{ref}}(\Omega) \end{pmatrix} = \begin{pmatrix} p_{\mathrm{ref}}(\Omega) \\ q_{\mathrm{ref}}(\Omega) \end{pmatrix} \begin{pmatrix} p_{\mathrm{ref}}(-\Omega)^* & q_{\mathrm{ref}}(-\Omega)^* \end{pmatrix} . \tag{2.2.49}$$

The right-hand side evaluates to the covariance matrix of the input spectra, multiplied from the left and right by the Fabry-Perot transformation matrices $R(\Omega)$ and $R(-\Omega)^*$.

2.2.7 Sensitivity of different detection schemes

After its interaction with the cavity, the laser beam is absorbed by a photodiode and generates an electric current $i(t)$ proportional to the instantaneous number of photons $I(t)$. This current is subsequently converted into a voltage by a transimpedance amplifier, possibly further amplified and measured with one of the following instruments:

- an oscilloscope to record a time-trace proportional to $I(t)$,
- a network analyzer (NA) which computes the transfer function from a voltage-driven modulator or opto-mechanical actuator to the output photocurrent, which corresponds for small drive amplitudes to $\delta I(\Omega)/\delta V_{\mathrm{drive}}(\Omega)$,
- an electronic spectrum analyzer (ESA) which performs a low-noise measurement of the photocurrent spectrum $S_{II}(\Omega)$.

The optomechanical coupling imprints the displacement of a mechanical degree of freedom onto the phase of light reflected by it. Since the photocurrent of conventional photodiodes is proportional to the incident photon flux, all information about the phase of the light is lost when it is absorbed by the photodetector. It is therefore the task of the experimenter to design and implement optical components which transfer the phase information of the light beam onto its intensity to make a displacement measurable by photodetection.

[4] All coefficients $R_{ij}(\Omega)$ depend on the frequency Ω. In matrix notation, we omit this dependence for the sake of brevity.

The goal of this section is to derive mathematical expressions for four different detection schemes that link the quantities we wish to measure to one of the three measurable signals listed above. These expressions can be obtained from the coefficients of section 2.2.6. We will discuss the physical meaning of these expressions by considering a number of special cases.

Direct detection / Side-of-the-fringe method

Direct detection consists in using a direct photodetector measurement of the photon flux reflected from a Fabry-Perot cavity as the error signal[5]. We have already seen in Eq. 2.2.44 that the mean flux reflected from a cavity equals the incident flux with a Lorentzian dip around the resonance, with a minimum reflectivity on resonance of

$$R_0 = 1 - 4\eta(1-\eta) = \frac{(T-P)^2}{(T+P)^2}\,. \tag{2.2.50}$$

In the limit of negligible loss $P/T \to 0$, the static reflectivity approaches 1 for all detunings, and direct detection of the reflected photon flux does not yield information about the cavity-laser detuning, phase modulation of the input light or cavity length changes at DC-frequencies. Amplitude modulations appear in the reflected photon flux in a similar way as in the incident one. The use of direct detection as an error signal for laser-cavity frequency locking thus relies on the absorption of the cavity, when $T \approx P$, to yield information about the detuning. The transfer functions from p_{in}, q_{in}, and δx to the direct detection error signal are given by the product of the reflected mean field $\sqrt{\bar{I}_{\mathrm{ref}}}$ and the corresponding components R_{11}, R_{21} and R_{51} of the reflected amplitude quadrature p_{ref} (see section A.2).

It is remarkable that without any extra assumptions, the ratio $\frac{R_{51}}{R_{21}}$ does not depend on the detuning Δ. The detailed calculation establishes that a phase modulation of the input laser with modulation depth ϕ_0 at frequency Ω_0 results in the same error signal amplitude as a length modulation of

$$|\delta x_{\mathrm{equivalent}}| = \frac{L}{\nu_{\mathrm{L}}}\Omega_0\phi_0 \tag{2.2.51}$$

where ν_L denotes the laser frequency in units of Hz. The stated conversion factor allows to calibrate the direct detection error signal in units of displacement from a known phase modulation depth of the input beam [GSA+10].

Pound-Drever-Hall detection

In Pound-Drever-Hall (PDH) detection [DHK+83, Bla01], the input laser is strongly phase-modulated at a frequency Ω_{mod} significantly above the cavity linewidth κ, in order to create two optical sidebands at $\pm\Omega_{\mathrm{mod}}$ which practically do not interact with the cavity for small carrier-cavity detunings. The reflected photocurrent is then electronically demodulated at the modulation frequency and low-pass filtered in order

[5]The error signal of known shape is first used to stabilize the laser-cavity detuning to a desired setpoint. In this situation, the fluctuations of the error signal are proportional of the deviation of the photon flux from its mean value, thus indicating the "error".

to detect one quadrature of the beating between the carrier (that has interacted with the cavity) and the two sidebands. When the laser is far-detuned from resonance, the phase modulation does not appear in the reflected photocurrent, and the resulting error signal is zero. The signal has another zero-crossing when the carrier is exactly at the cavity resonance, where both sidebands are symmetric with respect to the π-rotated carrier. For all other detunings, the phase of the carrier (rotated by the cavity) partially transforms the phase modulation into an amplitude modulation and becomes thus detectable by the photodiode. To illustrate this, we consider the photocurrent generated by two phase modulation sidebands that are assumed to remain unaffected upon reflection off the cavity, together with a carrier field with fluctuations $p + iq$ acquired by its interaction with the cavity:

$$\begin{aligned} I(t)/\bar{I} &= \left| J_0(\phi)(1 + p(t) + iq(t)) + iJ_1(\phi)e^{i\Omega_{\mathrm{mod}}t} + iJ_1(\phi)e^{-i\Omega_{\mathrm{mod}}t} \right|^2 \\ &\approx 1 + 4J_0(\phi)J_1(\phi)\cos(\Omega_{\mathrm{mod}}t)q(t) \end{aligned} \tag{2.2.52}$$

Demodulation selects the term at Ω_{mod} which provides an estimate of the acquired phase quadrature of the reflected light. Since the PDH error signal has a simple zero at zero detuning, its dependence on the laser intensity is only a second-order effect which is neglected in the above approximation. Similarly, PDH detection does not require a significant absorption dip as it is the case for direct detection.

Homodyne detection

Homodyne detection consists in sending the signal beam α_{s}, whose fluctuations are to be measured, onto a symmetric beamsplitter, whose second port is fed with a local oscillator beam α_{LO}. The two resulting beams $\alpha_{\mathrm{s}}/\sqrt{2} \pm \alpha_{\mathrm{LO}}/\sqrt{2}$ are sent to two photodetectors whose photocurrents are subtracted electronically. The resulting difference current is easily calculated to be proportional to

$$\delta I_{\mathrm{diff}} = 2|\bar{\alpha}_{\mathrm{s}}\bar{\alpha}_{\mathrm{LO}}| \cos\left(\bar{\phi}_{\mathrm{s}} - \bar{\phi}_{\mathrm{LO}}\right) + |\bar{\alpha}_{\mathrm{LO}}|X_{\mathrm{s},\bar{\phi}_{\mathrm{LO}}} + |\bar{\alpha}_{\mathrm{s}}|X_{\mathrm{LO},\bar{\phi}_{\mathrm{s}}} \tag{2.2.53}$$

Most commonly, a homodyne detector is operated in the limit $|\bar{\alpha}_{\mathrm{s}}| \ll |\bar{\alpha}_{\mathrm{LO}}|$ where the last term of the above equation can be neglected. In this case, the DC-part of the error signal can be used to set the relative phase between local oscillator and signal beam and thus to choose the quadrature $X_{\mathrm{s},\bar{\phi}_{\mathrm{LO}}}$ of the signal beam that is to be detected with the fluctuating part of the photocurrent. Locking the DC-signal to zero results in measuring the signal phase with a gain given by the local oscillator amplitude: $\delta I_{\mathrm{diff}} = |\bar{\alpha}_{\mathrm{LO}}|q_s$. Most of the time, we will employ homodyne detection to measure the phase quadrature of a beam reflected from a Fabry-Perot cavity, and because of the absence of other terms in the difference photocurrent, the homodyne signal will be limited by the signal-to-quantum noise ratio of the reflected quadrature itself (see section 2.2.6).

Heterodyne detection

So far, we have only considered modulations that create symmetric sidebands around the carrier. In more general settings, the sideband spectrum around the carrier may

be asymmetric, for example when the Stokes- and anti-Stokes processes of the optomechanical interaction obey different scattering rates. In this case, heterodyne detection can be employed to obtain an estimate of the sideband spectrum. Heterodyne detection works in a similar way as homodyne detection except that the local oscillator field oscillates at a different optical frequency $\omega_{\mathrm{L}}+\Omega_{\mathrm{het}}$ than the signal carrier at ω_{L}. When the two fields, whose mean fields may be assumed real without loss of generality, are superposed by a beamsplitter, we find

$$\begin{aligned} I(t) &= \left|\bar{\alpha}_{\mathrm{s}} + \delta\alpha_{\mathrm{s}}(t) + (\alpha_{\mathrm{LO}} + \delta\alpha_{\mathrm{LO}}(t))\, e^{i\Omega_{\mathrm{het}}t}\right|^2 \\ &\approx \bar{\alpha}_{\mathrm{s}}^2 + \bar{\alpha}_{\mathrm{LO}}^2 + 2\bar{\alpha}_{\mathrm{s}}\bar{\alpha}_{\mathrm{LO}}\cos(\Omega_{\mathrm{het}}t) \\ &+ \bar{\alpha}_{\mathrm{LO}}\left(\delta\alpha_{\mathrm{s}}(t)e^{-i\Omega_{\mathrm{het}}t} + \delta\alpha_{\mathrm{s}}^*(t)e^{i\Omega_{\mathrm{het}}t} + p_{\mathrm{LO}}(t)\right) \end{aligned} \tag{2.2.54}$$

In the last approximation, we have assumed $\bar{\alpha}_{\mathrm{s}} \ll \bar{\alpha}_{\mathrm{LO}}$. The Fourier transform of the photocurrent fluctuations is thus composed[6] of the beatnote between signal carrier and local oscillator and the beatnote between local oscillator and signal sidebands

$$\begin{aligned} \delta I(\Omega) &= 2\pi\bar{\alpha}_{\mathrm{s}}\bar{\alpha}_{\mathrm{LO}}\left(\delta(\Omega+\Omega_{\mathrm{het}}) + \delta(\Omega-\Omega_{\mathrm{het}}) + p_{\mathrm{LO}}(\Omega)\right) \\ &+ \bar{\alpha}_{\mathrm{LO}}\delta\alpha_{\mathrm{s}}(-\Omega_{\mathrm{het}}+\Omega) + \bar{\alpha}_{\mathrm{LO}}\delta\alpha_{\mathrm{s}}(-\Omega_{\mathrm{het}}-\Omega)^* \end{aligned} \tag{2.2.55}$$

where the two signal sidebands at $\pm\Omega-\Omega_{\mathrm{het}}$ are mapped to the same frequency Ω of the photocurrent. To remove this ambiguity, one usually chooses the heterodyne frequency such that only one of the signal sidebands contains a quantity to be measured, and the other one is in the vacuum state. The measurement at frequency Ω then contains the shot-noise of both signal sideband frequencies along with the signal at only one of the sidebands. For classical modulations, this results in a signal-to-shot-noise ratio that is worse by a factor of two with respect to a homodyne measurement.

2.3 Cavity optomechanics

Through the action of radiation-pressure, light can displace the boundaries by which it is confined. This inherently nonlinear interaction is greatly enhanced in cavity optomechanics by the use of an optical and mechanical resonator. When sufficiently high optical powers and flexible enough boundaries are combined, we expect new phenomena beyond the scope of the equations from the previous section. To understand these, we start this section with a rigorous expression for the radiation pressure in a cavity in section 2.3.1. The first and strongest effect caused by radiation-pressure is due to the mean intracavity field and will be discussed in section 2.3.2. Next, section 2.3.3 shows the different ways how fluctuations can cause a variable radiation-pressure force, and tries to classify the possible effects.

2.3.1 Radiation pressure in a cavity

So far, we have only considered how an oscillation of the cavity length is transduced into a fluctuation of the optical output field. To model the dynamics of the mechanical

[6]The term proportional to p_{LO} can be removed by using a balanced detector such as for the homodyne detection that measures the difference photocurrent the two superpositions of signal and local oscillator obtained with a symmetric beamsplitter.

oscillator more realistically, we identify the cavity length fluctuations δx with the scalar displacement x_m from section 2.1.6, with the weight vector defined by the Gaussian mode shape of the laser beam. As foreshadowed in that section, the optical field also exerts a radiation pressure force with the same transverse profile onto the movable mirror, which is often termed back-action of the light on the mirror to contrast the 'ordinary' effect the mirror position has on the light. For a full reflection off the mirror, the radiation-pressure force is equal to twice the sum of the momentum of all photons reflected per unit time. As a single photon at frequency ω_L has the momentum $\hbar\omega_\mathrm{L}/c$, where $\hbar$ is Planck's constant and c the speed of light, we find for the scalar amplitude of the radiation pressure force from the intracavity field

$$f_\mathrm{RP}(t) = \frac{2\hbar\omega_\mathrm{L}}{c} I(t)\,. \tag{2.3.1}$$

The cavity length fluctuations and the radiation pressure form a collocated sensor-actuator pair, which enables us to use the effective susceptibility $\chi_m(\Omega)$ from Eq. 2.1.18 to describe the dynamics of the mechanical system with a scalar force

$$f(t) = f_\mathrm{RP}(t) + f_\mathrm{ext}(t) \tag{2.3.2}$$

where f_ext can be any external force, for example the thermal force noise with the spectrum given by Eq. 2.1.22. The radiation-pressure coupling in a cavity forms a closed loop: the intracavity power depends on the cavity length, which in turn depends on the mechanical response to the radiation pressure force, proportional to the intracavity power.

2.3.2 Static effects of radiation-pressure

The strongest radiation-pressure force term arises from the mean intracavity field, which we compute from Eq. 2.2.38 by setting the derivative of the field to zero:

$$\bar{I} = \frac{2}{\pi}\mathcal{F}\eta\bar{I}_\mathrm{in}\frac{1}{1+4(\Delta+G\bar{\delta x})^2/\kappa^2}\,. \tag{2.3.3}$$

The radiation-pressure induced static displacement is thus

$$\delta\bar{x} = \chi_m(0)\frac{2\hbar\omega_\mathrm{L}}{c}\bar{I}\,. \tag{2.3.4}$$

To gain some intuition, we insert the second into the first equation and choose natural units:

$$\bar{I} = I_\mathrm{max}\left(1+\left(\frac{\bar{I}-I_\mathrm{center}}{I_\mathrm{bw}}\right)^2\right)^{-1} \tag{2.3.5}$$

with

$$I_\mathrm{max} = \frac{2}{\pi}\mathcal{F}\eta\bar{I}_\mathrm{in} \tag{2.3.6}$$

$$I_\mathrm{bw} = \frac{c\kappa}{4\hbar\omega_\mathrm{L}G\chi_m(0)} \tag{2.3.7}$$

$$I_\mathrm{center} = -\frac{\Delta}{\kappa/2}I_\mathrm{bw}\,. \tag{2.3.8}$$

This equation suggests a graphical solution of the cubic equation for the intracavity photon flux (Fig. 2.4). In a coordinate system with axes $x = \bar{I}$ and $y = f(\bar{I})$, we plot both sides of Eq. 2.3.5 as functions $f(\bar{I})$. The intersections of the two curves are the possible solutions. The right-hand side is a Lorentzian with maximum $I_{\max}$ corresponding to the maximum possible intracavity flux for a given finesse and incident flux. Its width is expressed in units of I_{bw}, which is the necessary intracavity power to optomechanically tune the cavity length by half the width of the optical resonance. The Lorentzian is centered around I_{center}, which is essentially the laser-cavity detuning $\Delta = \omega_{\mathrm{L}} - \omega_c$ in the absence of optomechanical effects.

Let us first consider the typical situation when I_{bw} is positive. In this case, an increase in intracavity power leads to an elongation of the optical cavity, shifting its resonance frequency towards lower values. For a blue-detuned laser, i.e. $\Delta = \omega_{\mathrm{L}} - \omega_{\mathrm{c}} > 0$, an increasing intracavity power pushes the cavity resonance away from the laser frequency and only one stable solution exists. In the opposite case of a red-detuned laser, one to three solutions can exist depending on the initial detuning and the value of $I_{\max}$. Out of the three solutions, only the minimum and maximum possible intracavity powers yield a stable situation. The resulting radiation-pressure bistability causes hysteresis in the intracavity power, as observed experimentally more than 30 years ago [DMM+83].

The parameter I_{bw} can only become negative for $\chi_m(0) < 0$, in the unintuitive case of a spring that statically expands in response to a compressing force. While not caused by an unintuitive spring, the situation of a cavity length that shortens for increasing incident power is however very common in real experiments due to thermal expansion of mirror substrates upon absorption of intracavity light. The only consequence is a reversal of the stable and unstable detunings: Now a red-detuned laser is radiation-pressure DC-stable while a blue-detuned one may enter the regime of optical bistability.

In either situation, when the intracavity power is DC-stable, one can experimentally retrieve an effective detuning $\Delta + G\bar{\delta x}$, the detuning that one would expect from the measured reflected power in absence of static radiation-pressure effects. In the remaining part of this thesis, we will incorporate the static radiation-pressure induced cavity length shift into Δ, since all formulae will only depend on this effective detuning. We can therefore assume δx to have a mean value of zero in the following and may use a simpler formula to compute the intracavity flux:

$$\bar{I} = \frac{2}{\pi}\mathcal{F}\eta\bar{I}_{\mathrm{in}}\frac{1}{1+\left(\frac{\Delta}{\kappa/2}\right)^2}\,. \tag{2.3.9}$$

2.3.3 Dynamical effects of radiation-pressure

Already in early work on static bistability [DMM+83], it was noted that also the spectrum of cavity length fluctuations was altered in presence of radiation-pressure. To understand this effect, we write down the mechanical response to fluctuations $p(\Omega)$ of the intracavity amplitude quadrature:

$$\delta x(\Omega) = \chi_m(\Omega)\left(f_{\mathrm{ext}}(\Omega) + \frac{2\hbar\omega_{\mathrm{L}}}{c}|\bar{\alpha}|p(\Omega)\right)\,. \tag{2.3.10}$$

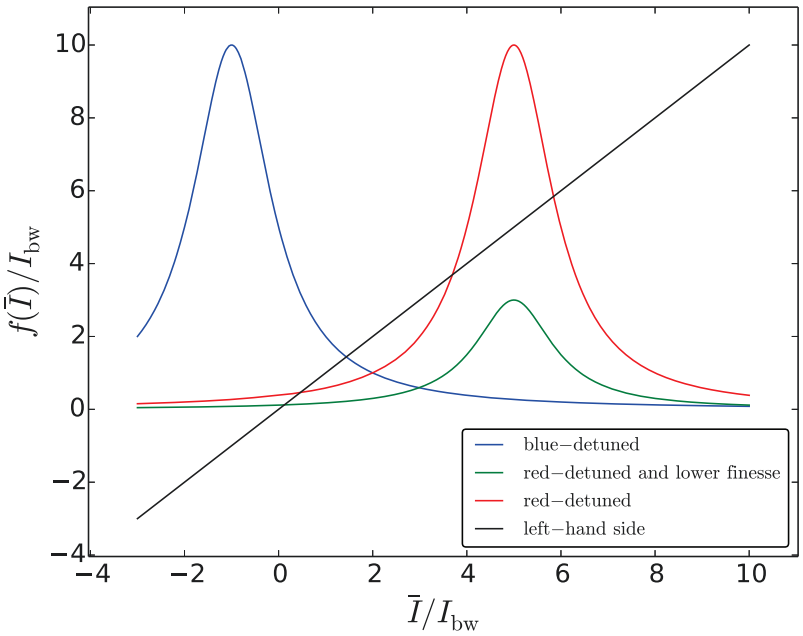


Figure 2.4: Graphical solution of Eq. 2.3.5. While the left-hand side of the equation is the black linear function $f(\bar{I}) = \bar{I}$, the right-hand side is a Lorenzian (colored lines). The solutions in $\bar{I}$ are the abscissa coordinates of the intersections of the black and colored lines in the first quadrant. A change of detuning Δ corresponds to a horizontal shift of the Lorentzian, while a change of incident power or finesse corresponds to a scalar multiplication. The graphical solution nicely illustrates that for the existence of multiple solutions (bistability) at some detuning, a threshold incident power or finesse is necessary such that the maximum slope of the Lorenzian is higher than the slope of the black line.

The amplitude quadrature fluctuations $p(\Omega)$ depend not only on the input field fluctuations α_{in}, α_{vac}, but also on the mechanical displacement δx. We can write the dependence in a similar way as in Eq. 2.2.48:

$$\begin{pmatrix} p(\Omega) \\ q(\Omega) \end{pmatrix} = \begin{pmatrix} L_{11} & L_{12} & L_{13} & L_{14} & L_{15} \\ L_{21} & L_{22} & L_{23} & L_{24} & L_{25} \end{pmatrix} \times \begin{pmatrix} p_{\mathrm{in}}(\Omega) \\ q_{\mathrm{in}}(\Omega) \\ p_{\mathrm{vac}}(\Omega) \\ q_{\mathrm{vac}}(\Omega) \\ \delta x(\Omega) \end{pmatrix} \tag{2.3.11}$$

where the coefficients[7] of the matrix $L(\Omega)$ can be obtained from the Mathematica code in Appendix A by omitting the subscript "ref" in the final evaluations. By inserting $p(\Omega)$ from Eq. 2.3.11 into Eq. 2.3.10 and solving for $\delta x(\Omega)$, we obtain a new equation

$$\begin{aligned} \delta x(\Omega) &= \chi_{\mathrm{eff}}(\Omega) f_{\mathrm{ext}}(\Omega) \\ &\quad + \frac{2\hbar\omega_{\mathrm{L}}}{c}|\bar{\alpha}| \left(L_{11}(\Omega)p_{\mathrm{in}}(\Omega) + L_{12}(\Omega)q_{\mathrm{in}}(\Omega) + L_{13}(\Omega)p_{\mathrm{vac}}(\Omega) + L_{14}(\Omega)q_{\mathrm{vac}}(\Omega)\right) \\ &= \chi_{\mathrm{eff}}(\Omega)\left(f_{\mathrm{ext}}(\Omega) + \frac{2\hbar\omega_{\mathrm{L}}}{c}|\bar{\alpha}| p(\Omega)\Big|_{\delta x=0}\right), \end{aligned} \tag{2.3.12}$$

where we have introduced the effective susceptibility $\chi_{\mathrm{eff}}(\Omega)$

$$\chi_{\mathrm{eff}}(\Omega) = \frac{\chi_m(\Omega)}{1 - \frac{2\hbar\omega_{\mathrm{L}}}{c}|\bar{\alpha}| L_{15}(\Omega)\chi_m(\Omega)}\,. \tag{2.3.13}$$

We see that Eq. 2.3.12 has the same form as Eq. 2.3.10, but now the open-loop susceptibility χ_m has been replaced by the closed-loop susceptibility χ_{eff}, which removes the dependence of the right-hand side of the equation on δx. In our linearized treatment, we expectedly find that the effective susceptibility only depends on the mean value of the intracavity field $|\bar{\alpha}|$ and all transfer functions from the different input fluctuations to the displacement are linear.

Furthermore, by combining Eqs. 2.3.11 and 2.3.12 with Eq. 2.2.48, we find for the reflected field quadratures in presence of radiation pressure effects:

$$\begin{aligned} \begin{pmatrix} p_{\mathrm{ref}}(\Omega) \\ q_{\mathrm{ref}}(\Omega) \end{pmatrix} &= \begin{pmatrix} R_{11} & R_{12} & R_{13} & R_{14} \\ R_{21} & R_{22} & R_{23} & R_{24} \end{pmatrix} \times \begin{pmatrix} p_{\mathrm{in}}(\Omega) \\ q_{\mathrm{in}}(\Omega) \\ p_{\mathrm{vac}}(\Omega) \\ q_{\mathrm{vac}}(\Omega) \end{pmatrix} \\ &\quad + |\bar{\alpha}|\chi_{\mathrm{eff}}(\Omega)\frac{2\hbar\omega_{\mathrm{L}}}{c} \begin{pmatrix} R_{15}L_{11} & R_{15}L_{12} & R_{15}L_{13} & R_{15}L_{14} \\ R_{25}L_{11} & R_{25}L_{12} & R_{25}L_{13} & R_{25}L_{14} \end{pmatrix} \times \begin{pmatrix} p_{\mathrm{in}}(\Omega) \\ q_{\mathrm{in}}(\Omega) \\ p_{\mathrm{vac}}(\Omega) \\ q_{\mathrm{vac}}(\Omega) \end{pmatrix} \\ &\quad + \chi_{\mathrm{eff}}(\Omega) \begin{pmatrix} R_{15} \\ R_{25} \end{pmatrix} f_{\mathrm{ext}}(\Omega)\,. \end{aligned} \tag{2.3.14}$$

[7] All coefficients $L_{ij}(\Omega)$ depend on the frequency Ω. In matrix notation, we omit this dependence for the sake of brevity.

The way that Eq. 2.3.14 was written suggests that the quadratures of the reflected field can be understood as the sum of three terms:

1. The first line describes the frequency-dependent rotation of the input quadratures upon reflection off the cavity that occurs in all Fabry-Perot cavities. Neither the mechanical degree of freedom nor radiation pressure appears in these terms.

2. The second line represents a different transformation of the same input quadratures. This transformation cannot be a simple rotation since the "gain" of this line can be tuned by the mean intracavity field $|\bar{\alpha}|$ at the beginning of the line[8]. Furthermore, the frequency dependence is tunable independently from the one of the first line through the effective susceptibility. Since both the first and the second line are added coherently in the output quadratures, both constructive and destructive interference can occur and a rich variety of transformations is implementable with sufficiently compliant mechanical susceptibilities.

3. The last line represents the transduction of mechanical displacement to the output fields that was encountered in section 2.2.6. Here, the mean field-dependent effective susceptibility replaces the intrinsic mechanical susceptibility. The intrinsic susceptibility is only recovered in the limit of low intracavity field strength (see Eq. 2.3.13).

In the situation where two independent lasers (differing by polarization, frequency, or propagation direction etc.) interact with the same mirror, the radiation pressure force of one laser can be modeled as f_{ext} in the equations of the other laser and vice versa. We will often use of this possibility in this thesis, employing a weak laser to detect the radiation pressure effect that the other laser exerts on the mechanical system.

Since the different interference terms can become quite confusing, the optomechanics community has invented names for the most commonly encountered parameter regimes and associated effects, which are summarized in Table 2.1. For effects with sinusoidal input, the transfer function of the optomechanical system from the input modulation to the detected reflected quadrature p_{ref} (e.g. by direct photodetection) or q_{ref} (e.g. by homodyne detection) is directly given by the corresponding coefficients from Eq. 2.3.14. For noise inputs (classical or quantum), Eq. 2.2.49 is used in conjunction with Eq. 2.3.14 to express the detected quadrature spectrum as a function of the spectra of the input quadratures and input force f_{ext}. We notice that the approximation from Eq. 2.2.17 results in the fact that no special treatment must be given to quantum fluctuations, i.e. quantum noise enters just like classical noise into the formalism. While quantum effects undoubtedly are significant in state-of-the-art linear optomechanical systems, all that can be deduced from these experiments about the foundations of quantum mechanics is the existence of a noise source in the optical field quadratures and the mechanical force f_{ext} that we call "quantum noise", and which has the peculiar property that its variance is independent of the mean field/force (cf. Eq. 2.2.26).

[8]The mean intracavity field appears a second time in the denominator of χ_{eff} (see Eq. 2.3.13). However, one easily confirms that the cancellation of the intracavity field is not perfect and a dependence on $|\bar{\alpha}|$ remains.

Name	Δ	$\lvert\bar{\alpha}\chi_m\rvert$	p_{in}	q_{in}	f_{ext}
Quadrature rotation	$\neq 0$	0	mod. or noise	sin or noise	0
Displacement readout	$\neq 0$	0	0	0	sin or noise
RP actuation	0	$\neq 0$	mod. or noise	0	0
OMIT/OMIA	$\neq 0$	$\neq 0$	0	mod. or noise	0
Feedback cooling	0	$\neq 0$	0	0	th. $+ \propto \delta x$
DBA cooling/GSC	< 0	$\neq 0$	0	0	thermal
DBA heating/PI	> 0	$\neq 0$	0	0	thermal
Quantum back-action/RPSN	0	$\neq 0$	vacuum	vacuum	0
Ponderomotive squeezing	$\neq 0$	$\neq 0$	vacuum	vacuum	0
Sideband asymmetry	0	$\neq 0$	vacuum	vacuum	vacuum

Table 2.1: Taxonomy of the effects in linear optomechanics. The parameters determining the regime are: detuning Δ (negative values denote $\omega_{\mathrm{L}} < \omega_c$), strength of radiation pressure effects $|\bar{\alpha}\chi_m|$, incident amplitude and phase quadratures p_{in} and q_{in}, external force f_{ext}. The abbreviations used are: radiation pressure (RP), optomechanically induced transparency/absorption (OMIT/OMIA), dynamical back-action (DBA), ground-state cooling (GSC), parametric instability (PI), radiation pressure shot noise (RPSN). The values in the table are only indicative of the most commonly encountered parameter regimes, e.g. all effects with $\Delta = 0$ also work with a non-zero detuning. A zero input field or force means that this input can be neglected in the description of the effect, a non-zero parameter indicates that the effect disappears if the parameter value is assumed to be zero. "mod." denotes a monochromatic modulation, "thermal" the thermal noise, "noise" any kind of noise and "vacuum" the quantum vacuum state (i.e. quantum noise) at the input. The entry "th. $+ \propto \delta x$" indicates that the mechanical degree of freedom experiences both the thermal force and a force proportional to its displacement.

Dynamical back-action

In order to obtain useful equations for modeling our experiment, we start by modeling the mechanical system as a single mechanical mode with negligible background term. From Eq. 2.1.18, we get

$$\chi_m(\Omega) = \frac{1}{m_0\left(\Omega_0^2 - \Omega^2 - i\gamma_0\Omega_0\right)} \tag{2.3.15}$$

where m_0 is the effective mass, Ω_0 the mode's resonance frequency, and $\gamma_0 = \Omega_0\phi_0 = \Omega_0/Q_0$ its damping rate (assuming structural damping with a loss angle ϕ_0, equivalent to a quality factor of Q_0).

From Table 2.1, we find that we can understand dynamical backaction by neglecting the fluctuations of the optical input fields. The only remaining nonlinear term is χ_{eff} (Eq. 2.3.13), into which we insert the coefficient L_{15} and the susceptibility defined above. By writing the result in an identical form as Eq. 2.3.15

$$\chi_{\text{eff}}(\Omega) = \frac{1}{m_0\left(\Omega_{\text{eff}}^2 - \Omega^2 - i\gamma_{\text{eff}}\Omega_{\text{eff}}\right)}\,, \tag{2.3.16}$$

we have introduced the effective damping rate

$$\gamma_{\text{eff}} = \gamma_0 + 16\bar{I}\mathcal{F}\left(\frac{2\pi x_{\text{zpm}}}{\lambda}\right)^2\left(\frac{1}{1+\left(\frac{\Delta+\Omega}{\kappa/2}\right)^2} - \frac{1}{1+\left(\frac{\Delta-\Omega}{\kappa/2}\right)^2}\right) \tag{2.3.17}$$

and the effective frequency

$$\Omega_{\text{eff}} = \Omega_0 + 8\bar{I}\mathcal{F}\left(\frac{2\pi x_{\text{zpm}}}{\lambda}\right)^2\left(\frac{\frac{\Delta+\Omega_0}{\kappa/2}}{1+\left(\frac{\Delta+\Omega}{\kappa/2}\right)^2} + \frac{\frac{\Delta-\Omega_0}{\kappa/2}}{1+\left(\frac{\Delta-\Omega}{\kappa/2}\right)^2}\right) \tag{2.3.18}$$

with the assumptions $\Omega_{\text{eff}} - \Omega_0 \ll \Omega_0, \kappa$ and $\gamma_{\text{eff}} \ll \Omega_0$. The zero-point rms-motion $x_{\text{zpf}} = \sqrt{\hbar/(2m_0\Omega_0)}$ is commonly used to normalize the mechanical displacement and the intracavity flux $\bar{I}$ can be computed from Eq. 2.3.9. The above equations represent the fact that displacement fluctuations cause a fluctuating radiation pressure force, which can be decomposed into a real part that is added to the restoring force of the mechanical oscillator (optical spring effect), and an imaginary part that contributes to its damping (optical damping). The effect is proportional to both the intracavity intensity and the finesse.

If the effective damping becomes zero or negative, a parametric instability of the optomechanical system occurs and the mechanical oscillator performs self-sustained oscillations. The amplitude of these oscillations grows exponentially and quickly leads to a break-down of the linear approximations for the cavity field modulations. As the displacement becomes comparable to the optical linewidth $\lambda/(2\mathcal{F})$, the displacement-induced modulation of the intracavity-field saturates at the maximum intracavity field for a given incident power, and with the saturation of the corresponding radiation pressure modulation, the displacement amplitude saturates as well. This self-limiting effect is very important for our system since it prevents the damage of the mechanical oscillator.

For radiation pressure enhanced damping rates, one speaks of optomechanical cooling. To see this, we write the thermal force noise spectrum for an enviroment temperature T_{env} near the mechanical resonance frequency from Eq. 2.1.22 for the single-mode mechanical susceptibility from Eq. 2.3.15

$$S_{FF}(\Omega) = 2k_B T_{\mathrm{env}} m_0 \gamma_0 \,. \tag{2.3.19}$$

As long as we are dealing with high-quality factor mechanical oscillators, we may approximate the effective susceptibility as a Lorentzian which leads to easier integrals in the following discussion:

$$\chi_{\mathrm{eff}}(\Omega) \approx i \frac{Q_0}{m_0 \Omega_0^2} \frac{1}{1 - i\frac{\Omega - \Omega_{\mathrm{eff}}}{\gamma_{\mathrm{eff}}/2}} \,. \tag{2.3.20}$$

By computing the mean square thermal displacement from the integrated displacement spectrum

$$\left\langle \delta x^2 \right\rangle = \frac{1}{2\pi} \int_{-\infty}^{\infty} |\chi_{\mathrm{eff}}(\Omega)|^2 \, S_{\mathrm{FF}}(\Omega) \mathrm{d}\Omega = \frac{k_B T_{\mathrm{env}}}{m_0 \Omega_{\mathrm{eff}}^2} \frac{\gamma_0}{\gamma_{\mathrm{eff}}} \,, \tag{2.3.21}$$

we recover the equipartition theorem from Eq. 2.1.14 in the case of negligible optomechanical damping $\gamma_0 = \gamma_{\mathrm{eff}}$. However, if the effective damping rate is larger than the intrinsic one, we can understand the thermal noise spectrum as one arising from an effective temperature

$$T_{\mathrm{eff}} = T_{\mathrm{env}} \frac{\gamma_0}{\gamma_{\mathrm{eff}}} = T_{\mathrm{env}} \frac{Q_{\mathrm{eff}}}{Q_0} \,. \tag{2.3.22}$$

We conclude that optomechanical damping can reduce the effective temperature of the thermal bath that is coupled to mechanical oscillator.

Quantum back-action

The effective damping rate introduced in the previous section motivates the question whether a fluctuating force is associated with the optical damping, as one would expect from the fluctuation-dissipation theorem (FDT). The FDT also holds in this case, but a more general version than the one section 2.1.4 is required. The original paper on the fluctuation-dissipation theorem [CW51] states that the variance associated with a damping mechanism of an oscillator at frequency ω is related to "the mean energy at the temperature T of an oscillator of natural frequency ω", which was approximated in section 2.1.4 by $k_B T$ for low-frequency mechanical oscillators. For the optical field at frequency ω_{L}, this energy is at the relevant frequencies $\omega_{\mathrm{L}} \pm \Omega_0 \approx \omega_{\mathrm{L}}$

$$E(\omega_{\mathrm{L}}) = \hbar\omega_{\mathrm{L}} \left(\frac{1}{2} + \frac{1}{e^{\frac{\hbar\omega_{\mathrm{L}}}{k_B T}} - 1} \right) \approx \hbar\omega_{\mathrm{L}} \left(\frac{1}{2} + 7 \times 10^{-20} \right) \,. \tag{2.3.23}$$

The mean energy of the optical degree of freedom is therefore compeletly dominated by the quantum fluctuations $\hbar\omega_{\mathrm{L}}/2$. We therefore have to take the vacuum fluctuations of the optical field into account in our calculation to understand the force fluctuations associated with optomechanical coupling.

The fluctuations of the amplitude quadrature of the intracavity field can be directly read-off from Eq. 2.2.46 by setting α_{in} equal to vacuum fluctuations independent from the vacuum α_{vac} entering through the loss port. With the definition from Eq. 2.2.24, we find by choosing a real intracavity field

$$\begin{aligned} p(\Omega) = &\sqrt{\frac{2}{\pi}\mathcal{F}}\sqrt{\eta}\left(\chi_{\mathrm{c}}(\Omega)\alpha_{\mathrm{in}}(\Omega) + \chi_{\mathrm{c}}(-\Omega)^*\alpha_{\mathrm{in}}(-\Omega)^*\right) \\ &+ \sqrt{\frac{2}{\pi}\mathcal{F}}\sqrt{1-\eta}\left(\chi_{\mathrm{c}}(\Omega)\alpha_{\mathrm{vac}}(\Omega) + \chi_{\mathrm{c}}(-\Omega)^*\alpha_{\mathrm{vac}}(-\Omega)^*\right) . \end{aligned} \tag{2.3.24}$$

Together with Eqs. 2.3.9, 2.2.26 and 2.2.45 we get

$$S_{II}(\Omega) = \frac{1}{\pi}\mathcal{F}\bar{I}\left(\frac{1}{1+\left(\frac{\Delta+\Omega}{\kappa/2}\right)^2} + \frac{1}{1+\left(\frac{\Delta-\Omega}{\kappa/2}\right)^2}\right) . \tag{2.3.25}$$

By multiplying this equation with $(2\hbar\omega_{\mathrm{L}}/c)^2$, we obtain the radiation pressure force noise spectrum. It is clear that the force spectrum can be expressed in the form of an effective temperature of the radiation pressure force noise. In this way, the fluctuation-dissipation theorem also holds for the radiation pressure damping. The associated temperature is usually extremely low, but it grows linearly with the intracavity power. One goal of this thesis is to experimentally reach the regime where this radiation pressure force noise is the dominant force acting on the mechanical oscillator. Entering this regime would allow to study the associated limitations for displacement sensing and possible schemes to overcome these limits.

Clearly, if the mean energy of the optical field is above the vacuum level, as for example in the presence of classical noise, the force fluctuations are expected to increase in proportion to the ratio of excess noise to quantum noise.

Chapter 3

Experiment - Samples

> *This is why we regard the struggle to isolate one degree of freedom from all others – i.e. the struggle for high Q – as a noble task for experimenters.*
>
> — V. B. Braginsky, V. P. Mitrofanov, V. I. Panov, *Systems with Small Dissipation*

Armed with a theoretical understanding of the optomechanical interaction from the previous chapter, we proceed to discuss the experimental implementation of both the mechanical and optical resonators in our experiment. It will become clear that our initial decision on how to solve the problem of obtaining a high mechanical quality factor Q, together with the desire to implement a high-finesse Fabry-Perot cavity for the optical readout, makes most design choices of the mechanical and optical system unambiguous. We will then discuss in detail the employed methods for the mechanical and optical characterization of our samples. Finally, we will conclude that both the mechanical and optical properties of our system are severely degraded if other undesired modes become degenerate with the respective mode of interest, a situation which can fortunately be avoided in practice to a large extent.

3.1 Mechanical resonator

3.1.1 Choice and design of mechanical oscillator

The choice of the opto-mechanical oscillator presented here had been made before the beginning of this thesis work [Kuh13]. The general aim of the experiment to obtain experimental evidence for quantum effects with an optomechanical system is synonymous in state-of-the-art systems to a maximization of the ratio of the optomechanical measurement rate to the thermal decoherence rate. In the following we present a number of design considerations along with the implementation imperatives derived from them:

1. Existing expertise and collaborations as well as the conceptual simplicity suggested to employ a high-finesse Fabry-Perot cavity for the optical part of the

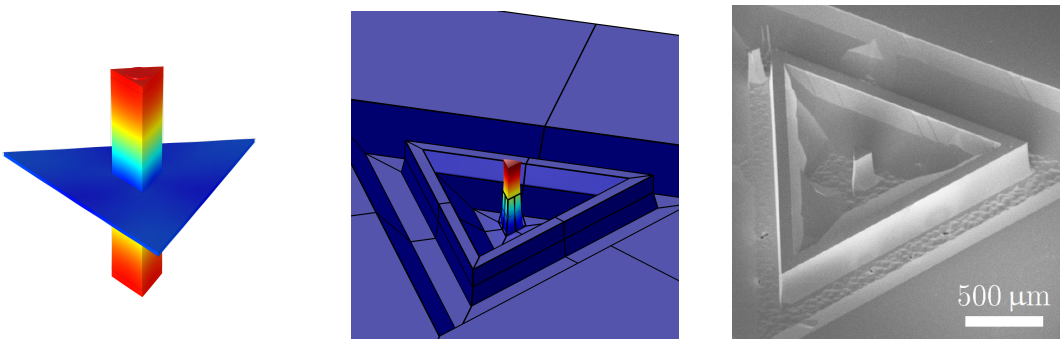


Figure 3.1: Left: Schematic quartz micropillar geometry. Color indicates total mechanical strain of the fundamental compression mode, increasing from blue to red. The thin black lines indicate the un-elongated geometry, showing that the strain points predominantly along the pillar axis. The pillar is symmetrically suspended on a thin membrane, which is almost unstrained. Middle: COMSOL finite-element model for the resonator. Around the central pillar, one can see the suspending membrane, the decoupling shield, an outer suspending membrane, and finally a part of the surrounding quartz wafer which are all included in the simulation. Right: Scanning electron microscope image of a realized micropillar structure.

system. This implies the coating of a dielectric Bragg mirror onto the mechanical resonator. This in turn constrains the lateral size of the resonator to exceed the optical beam waist size by at least a factor of 5 to avoid excessive optical clipping loss.

2. The desire to minimize thermal decoherence of the mechanical oscillator translates into optimizing the mechanical quality factor of the mechanical mode of interest.

   (a) To minimize the effect of structural damping, monocrystalline quartz was chosen as resonator material since it has the lowest known mechanical loss angle ($\phi < 10^{-9}$) at cryogenic temperatures [GCI+12].

   (b) Since the mechanical quality factor of the mirror coating material is known [CCF+04] to be below 10^4, a structure which stores a mimimum amount of elastic energy in the mirror coating was searched for. This motivated to use a compression mode of a high aspect-ratio geometry such as a pillar, with the mirror on one end, where the mode's displacement is maximum, but the mechcanical stress minimum.

   (c) To keep mechanical surface losses on the sides of the pillar negligible [El 93], wet etching along the principal crystal axes was chosen in favor of higher-roughness methods such as deep reactive ion etching.

   (d) To confine the mechanical energy to the resonator, i.e. to keep clamping loss at a minimum, a symmetric suspension of the pillar through a thin membrane was employed. The minimum realizable membrane thickness was imposed by the fabrication process to about 10 μm. After initial experiments leading to quality factors of the order of only 10^4, an additional symmetric decoupling shield was inserted between the quartz wafer and the central pillar, which eventually resulted in quality factors above 10^6.

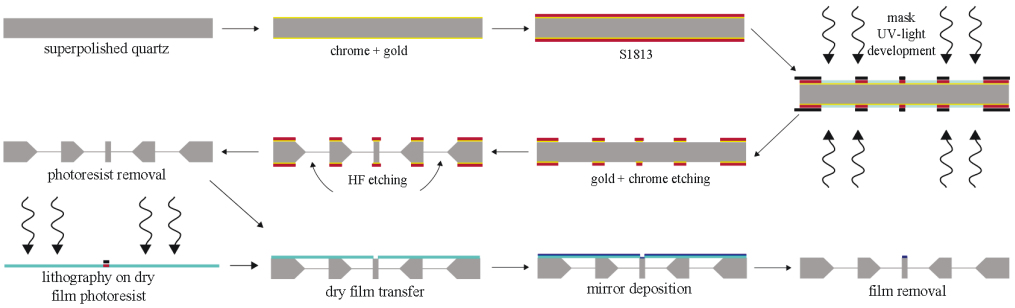


Figure 3.2: Fabrication procedure of the micropillar resonators.

3. The pillar length is proportional to the effective mass m_{eff} and roughly inversely proportional to the frequency Ω_m of the compression mode. A figure of merit for ground-state cooling is the product of peak effective susceptibility and mode frequency, where the susceptibilty describes the effectiveness of radiation-pressure actuation while higher frequencies result in a lower initial thermal phonon number to start with at a given temperature. With Eq. 2.1.18, we easily see that the effect of pillar length cancels out in this figure of merit. Therefore, the pillar length could be optimized based on other considerations. Firstly, a high ratio between pillar length and membrane thickness is beneficial for keeping clamping losses low. Secondly, a mode frequency above 1 MHz makes it much easier to keep technical noise sources, such as seismic and acoustic excitations or classical laser noise below the level of the expected quantum effects. A compromise between these two effects resulted in a mode frequency of roughly 4 MHz, or a pillar length of 1 mm.

The final geometry is shown in Figs. 3.1 and 3.3. The design and fabrication were carried out in collaboration with the physics department of ONERA, Châtillon, France. Simulations by ONERA were used to find out the geometries obtainable through the wet-etching process. The computer program SAMCEF was employed by ONERA to carry out a finite-element simulation to optimize the dimensions in order to obtain a robust geometry that yields minimum mechanical clamping loss.

3.1.2 Fabrication

The fabrication process of the micropillar resonator is illustrated in Fig. 3.2. More detailed descriptions of the various steps can be found in the references [KBD+11, ZNT+17, Kuh13, Zer17]. Fabrication can be divided into two parts: the creation of the mechanical resonator, and the coating of a dielectric mirror on top of it. We will conclude this section with a discussion of the effects observed in the fabrication.

Fabrication of the mechanical structure

The mechanical structure was entirely fabricated in the cleanroom of ONERA, Châtillon. The fabrication steps are:

1. We purchase square quartz bars of dimensions 1.5" × 1.5" × 17 mm from the company Gemma-Quartz. These bars are sawed into 1.8 mm thick wafers by the company IEV. After initial attempts, only quartz referenced as "very high quality" was used. This specification guarantees less than 10 crystal defects per cm^3. Such lattice imperfections would cause defects in the final structure [Kuh13] and are observed in about 10 % of the samples with this wafer quality. The optical axis of the quartz is oriented normal to the wafer surface, while an angular corner marks which of the wafer sides is oriented along the electrical axis of the crystal (see description in Fig 3.3). The wafers come with a specified surface roughness below 10 μm.

2. Wafers are then superpolished on both sides by the company SESO from Aix-en-Provence, France, down to a thickness of 1 mm. The company guarantees a final surface microroughness of less than 2 nm rms. The superpolishing step is necessary since the high-reflectivity mirror will eventually be coated on one wafer surface, and the mirror roughness essentially reproduces the substrate properties. To keep diffraction loss below 10 ppm, a microroughness below 0.3 nm rms is necessary.

3. An etching mask consisting of 15 nm of chromium and 200 nm of gold is coated on both wafer sides at ONERA. To improve the masks's resistance against the etchant, the coating is preceeded by an optimized cleaning process.

4. A 1.4-μm thick layer of the positive photoresist S1813 is spin-coated on both sides of the wafer. The double-sided mask aligner EVG 620NT (earlier samples: EVG AL-6) is employed to transfer the pattern of the photolithography mask as shown in Fig. 3.3 onto both sides of the pillar. Both theoretical expectations and fabrication experience suggest that a good alignment of the pattern on both sides of the wafer is crucial to obtain final structures with high mechanical quality factors. The specified accuracy of the mask aligner is about 1 μm, and alignment is in practice limited by the ability to discern the alignment marks on the first insulated side of the wafer and the photomask. For the best samples, we believe that top and bottom structures lie less than 3 μm apart.

5. The exposed resist is developed following the standard procedure recommended for the photoresist and the gold and chromium is removed in the areas that are not protected by the resist. For gold etching, we use "Gold Etch TVA", while the chromium is etched with "Cr Etch 18" from ROHM. The good alignment of top and bottom masks is verified with an optical microscope after this step. The remaining photoresist is usually not removed until after the next step.

6. Anisotropic etching of the structure is performed in a mixture of hydrofluoric acid (HF) and ammonium fluorure (NH_4) at room temperature under steady agitation. With an etching rate of roughly 20 μm/h from each side, a nominal etching time of slightly less than 25 h is required for a final membrane thickness below 20 microns. To avoid over-etching and destruction of the samples, the structure is first etched only for about 20 h, after which the membrane thickness is estimated with an optical microscope that is equipped with an interferometric

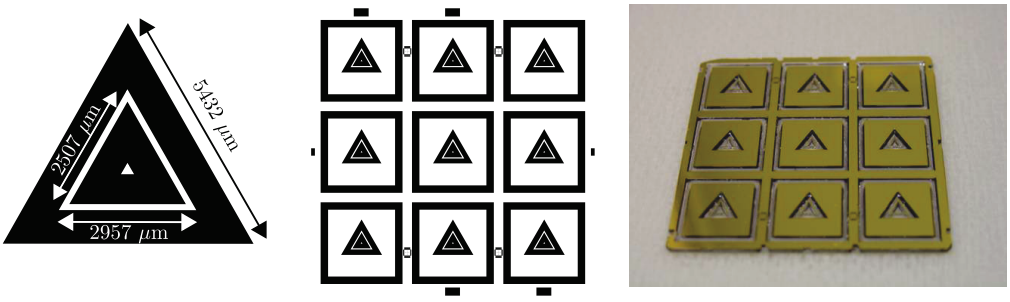


Figure 3.3: Left and middle: Final geometry of the photomask design. The same photomask is used on both sides of the pillar. Black zones are not coated with chromium, and will be etched into the quartz when a positive photoresist is used. The central equilateral triangle defines the cross-section of the micropillar with side-length of 240 μm. Right: Gold coated quartz wafer after roughly 24 h of etching in HF. The nine samples per wafer are clearly discernible. The upper left corner of the wafer is cut at a small angle with respect to the horizontal direction in the picture. This indicates that the electrical crystal axis points along the horizontal edge of the wafer, such that it is parallel to the bottom edge of each pillar's surface. The top row of samples is named C1-3 from left to right, the middle one C4-6, and the bottom one C7-9.

measurement of the height of the focal plane. From this, a more precise estimate of the etching rate of the specific solution can be deduced to finalize the etching down to a roughness of about 20 μm in one to three more steps.

7. The thin quartz membrane holding each of the nine 1 cm × 1 cm samples to the wafer frame is broken with a scalpel, while the central parts of all samples are protected from debris with a polymer form screwed on both sides of the wafer [Kuh13]. The resulting nine samples are named according to the pattern "year-month-day-Cn", where n denotes the position of the sample on the wafer (see Fig. 3.3).

8. Each sample is subsequently mounted in a vacuum chamber in the cleanroom. A Michelson interferometer is used to measure the frequency and quality factor of the pillar compression mode near 3.6 MHz. Once the resonance frequency is found through a swept excitation of a piezoelectric actuator underneath the sample support, the quality factor is measured with a mechanical ring-down measurement using the same piezo.

9. Each sample is individually optimized by iterations of 3 to 15 minutes of additional etching and subsequent measurement of its mechanical quality factor. This very time-consuming iteration is ended either when the quality factor exceeds one million, or when an inspection with the microscope suggests that the membrane holding the pillar or decoupling shield has too many holes to resist another few minutes of etching. Fig. 3.4 shows the observed evolution of the quality factors for two different wafers during optimization.

10. Once a sample is optimized, the remaining gold and chromium is removed in the same way as after the photolithography step. Sometimes, two iterations of the

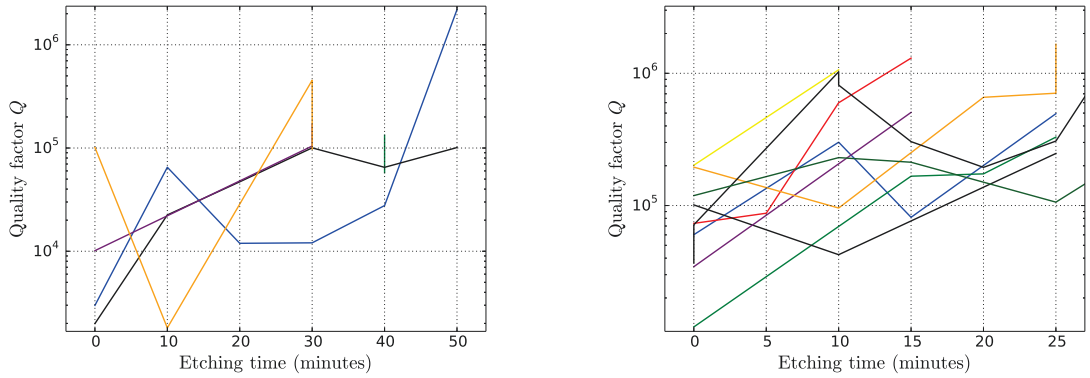


Figure 3.4: Quality factor of the compression mode of different micropillar samples versus etching time, where zero marks the etching time after which the wafer was broken into samples. Each line denotes a different sample. If a sample has multiple values at the same etching time, the later value shows the effect of photoresist removal or demetallization. Left: Samples 2012-10-02-C1 to -C9. Right: Samples 2012-10-09-C1 to -C9.

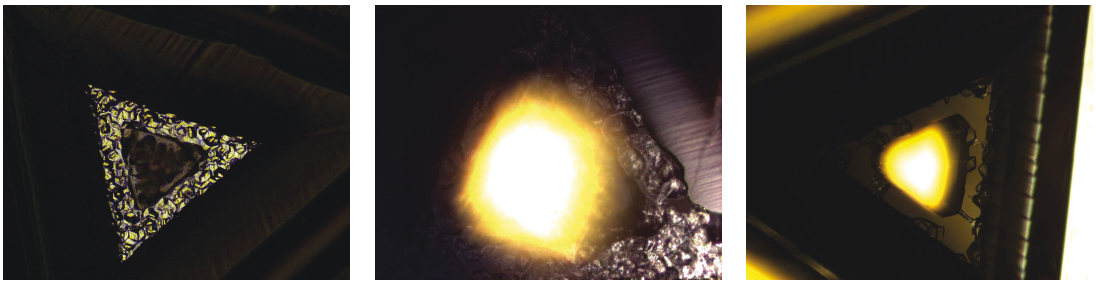


Figure 3.5: Microscope images of the membranes on which the micropillar is suspended for different samples. The micropillar is out-of-focus of the center of each image, while the decoupling frame can be seen at the border. The membrane consists of many interconnected microcrystals with a variable level of perforation that depends on membrane thickness. Left: Sample 20140113C6 before individual optimization ($t = 0$). No holes are discernible in the membrane. Middle: Sample 20121009C5 after 30 minutes of individual etching. The homogeneous gray zones in the membrane are holes, of which about 10 are discernible. Etching was stopped after this step as the sample would become too fragile. The quality factor at this point was $Q = 2.1 \times 10^6$ and increased after mirror coating in cryogenic operation to 65×10^6. Right: Sample 20100817C7 after optimization with almost no membrane left. Samples in this state are extremely fragile and do not support even 2 minutes of further etching.

removal step are necessary to fully remove all residues, which can be confirmed with the optical microscope. As both the frequency and quality factor change due to the reduced pillar mass, the pillar is characterized once more after this step.

Mirror coating

The fabrication process of the dielectric mirror on top of the micropillar was developed through a long series of trial and error [Kuh13, Zer17]. Initial attempts to use a wet photoresist failed as the photoresist became impossible to remove completely from the membrane after being hardened during the high-temperature mirror coating. Fabrication of the mirror before the mechanical structure failed since we were unable to find a way to protect the sides of the mirror during the 24 h of etching in hydrofluoric acid. A transfer procedure involving separately fabricated micromirrors and glue was investigated and deemed promising, but took longer to develop than the presented method due to the rare occasions where high-quality mirror coatings could be made. A series of different physical masks was tried, all with the goal to coat a high-quality mirror on the pillar surface but to prevent any coating on the pillar sides and the membrane, as the latter would reduce the mechanical quality factor. The final recipe transfers a dry-film photoresist onto the pillar with a circular hole centered on the pillar surface in the cleanrooms of Paris 7 university, after which the mirror coating is performed at the Laboratoire de Matériaux Avancées (LMA) in Villeurbanne, France. The detailed steps are as follows:

1. We purchase the dry film photoresist ALPHO NIT215 from Nichigo-Morton Co., Ltd. This film consists of a 15-μm thick negative photoresist, embedded between two thin polyester foils for protection and mechanical stability. A 1.5 cm $\times$ 1.5 cm square of this film is scotch-taped onto a 4" photomask plate with only a 100-μm diameter black circle in the center of the film. Before this, the upper protective foil of the resist is pulled off whereas the remaining lower protective foil prevents direct contact between photomask and resist. To avoid wrinkles and ensure a good adhesion of the film, it helps to pre-heat the photomask plate for 10 minutes at 120 °C before posing the film. We finish with another 4 minutes heating of the photomask plate at 120 °C to obtain a flat dry film surface without wrinkles or air enclosures.

2. The lithography mask is flood-exposed in a MJB4 photolithography machine with an exposure energy of 80 mJ/cm^2 of UV-light. Exposure must be done before approaching the quartz micropillar to the dry film, as the multitude of reflecting surfaces and UV-transparency of quartz would compromise the effect of the lithography mask. To avoid exposure of the photoresist in the 100 μm circle through reflections from the back, a black absorber is inserted into the substrate compartement of the lithography machine during exposure.

3. The lithography mask is left in the lithography machine and a clean micropillar sample (as in Fig. 3.6) is inserted in the substrate compartment of the machine. The unprotected side of the dry film points downwards towards the pillar in this

configuration. The alignment screws and microscope of the lithography machine are then used to slowly approach the circle on the dry film to the triangular top face of the micropillar. A lateral alignment error smaller than 5 μm is necessary to ensure that the circle lies inside the micropillar triangle. When dry film and sample touch each other well enough for the pillar to get lifted by and stick to the film, the mask plate is removed from the machine and flipped upside down to ensure that the pillar does not fall off. The contact between pillar surface and dry film is further consolidated by another bakeout of the ensemble at 120 °C during 4 minutes.

4. After cooling off, the pillar sample can be removed together with the dry film from the plate. The second protective foil can be pulled off the pillar with a tweezer, leaving the resist film attached to the pillar. Under an optical microscope, good alignment of the exposed motive of the dry film and good adhesion of the dry film to the pillar surface should be confirmed at this point (see Fig. 3.6).

5. The circular hole is developed during 45 s in a 1 % aqueous solution of Na_2CO_3. The sample is held in the developer without agitation in order not to break the fragile sample, before dipping it in deionized water for 15 s and gently blow-drying with nitrogen gas.

6. After the previous step, the dry film is usually not fully removed in the circular hole. We therefore finish the development in a reactive ion etching machine with an oxygen plasma, which etches about 0.1 μm of dry film resist per minute. This has also the advantage to thinning down the entire dry film as to reduce shadowing effects in the later mirror coating step. The plasma cleaning is performed with a pressure of 50 mTorr of oxygen and power set to 55 W during typically 10 minutes[1]. AFM and micromap measurements have been used to confirm that the surface roughness is not measurably increased for quartz or silica substrates by the plasma etching step.

7. The samples are now ready to be packaged and transported for the mirror coating. The dielectric mirror coating consists of alternating $\lambda/4$ layers of SiO_2 and Ta_2O_5 as low- and high-refractive index materials, where the first and last layer consists of SiO_2 and the last layer thickness is adjusted to obtain the desired residual transmission. The layers are deposited by ion beam sputtering (IBS) in a well-maintained "SPECTOR Ion Beam Deposition System for Optical Coating" from the company Veeco, which is kept in the class 1 cleanroom of the LMA in Lyon (see Fig. 3.7). For a residual transmission below 1 ppm, the total coating thickness amounts to about 7 μm.

8. After the mirror has been coated, the dry film resist is removed by posing each sample in a separate acetone bath for at least 10 minutes, after which the film can usually be pulled off with a tweezer. Care is taken to avoid residual shards of mirror coating from the dry film to end up on the round mirror on top of the

[1]The corresponding recipe of the reactive ion etching machine is called "Clean pure O2 30sccm 30W 50mtorr". Also for other mirror substrates, this cleaning recipe has often proved superior to wet cleaning techniques such as ultrasonic baths in acetone, isopropanol and water.

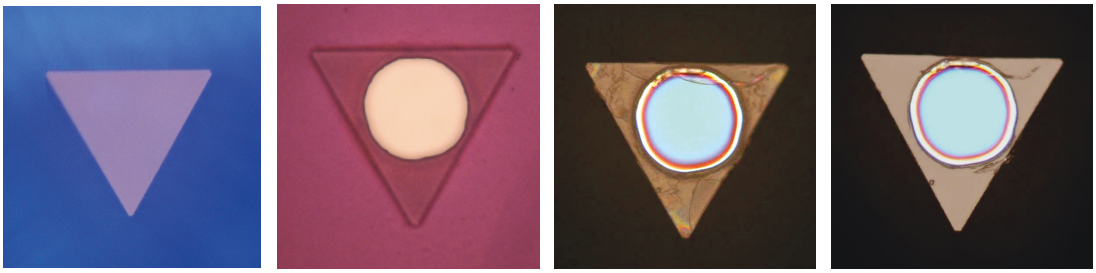


Figure 3.6: Steps of the mirror fabrication on top of the micropillar sample 20121009C5. From left to right: Clean pillar surface after demetallization and cleaning step, with dry film mask after plasma development step, with mirror coating, and after final plasma cleaning step. The sidelength of the triangular pillar surface is 240 μm.

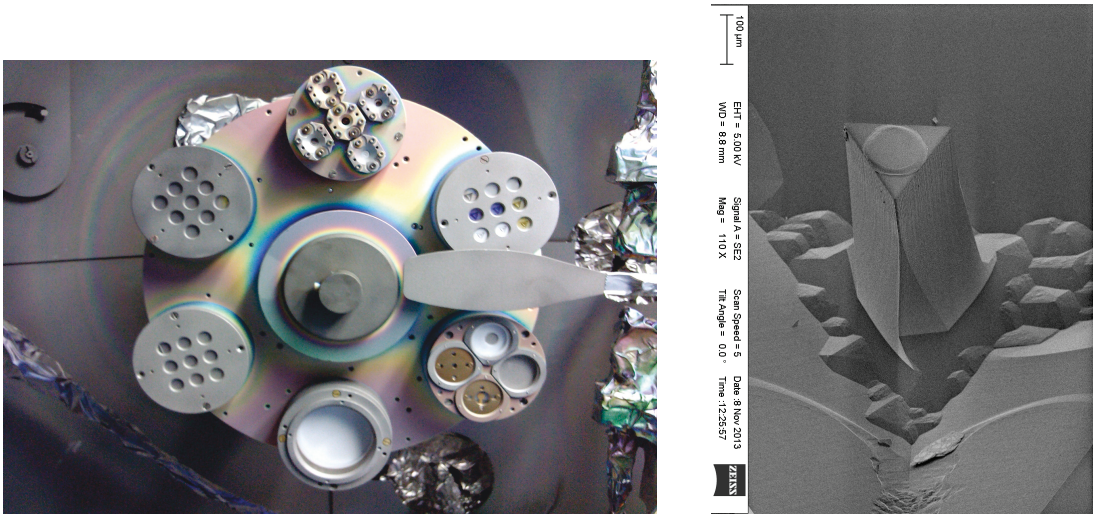


Figure 3.7: Left: Samples mounted inside the SPECTOR machine just before coating the dielectric mirrors. Out of the six circumferentially arranged satellites, the upper right one contains seven micropillar samples with different attempts of mask design. The two left satellites contain nine concave micromirror substrates each. Right: Scanning electron microscope image of a micropillar with a successfully coated mirror. Clearly, no mirror coating is seen on the side faces of the pillar.

pillar. The cleaning is finished by subsequently dipping the sample in another bath of acetone, isopropanol and deionized water.

9. Residues of the dry film or mirror shards compromise the cleanliness of many samples. While we have not found a recipe to completely remove these, another plasma cleaning step as in step 6 yields the best improvement and is frequently used to finalize the dry film removal (see Fig. 3.6).

10. To further lower the optical loss of the dielectric mirror coating, a 10-h bakeout at 450 °C is performed at the LMA. Both experience of the LMA and measurements on one sample suggest that this step reduces the optical loss by about 20 ppm.

Discussion

The effect of various fabrication steps on the frequency and quality factor of the samples are summarized in Table 3.1. The frequency change deviates little between different

Action	$\langle \delta f \rangle \pm \sigma_{\delta f}$ (kHz)	$\langle \delta Q \rangle \pm \sigma_{\delta Q}$ (10^3)	N
Etching (per minute)	-0.49 ± 0.06	$+40 \pm 74$	19
Photoresist removal	$+6.5 \pm 0$	$+72 \pm 0$	1
Demetallization	$+12.1 \pm 4.3$	$+48 \pm 928$	13
Mirror coating	-17.7 ± 1.1	-445 ± 1184	4
Cooldown below 4 K	$+50.2 \pm 0.2$	$+34000 \pm 29000$	2

Table 3.1: Effect of various fabrication steps on the micropillar's compression mode. The indicated values are the mean differences of the resonance frequency f and quality factor Q before and after the action, along with the standard deviations of the differences for the available N experimental values.

samples, except for the demetallization step, where the standard deviation of more than 30% can be explained by the fact that the frequency change for some samples includes the photoresist removal step as well. The etching step reduces the membrane thickness and thereby decreases the restoring force of the compression mode, thus resulting in a decrease of the frequency. Photoresist and metal removal decrease the pillar mass while the mirror coating increases it, thereby increasing or decreasing the frequency. Cooling the pillar leads to contraction and a higher material stiffness and thus increases the pillar frequency.

While the mean quality factor change is nonzero for all fabrication steps, the effect always varies by more than the mean change, indicating that the quality factor behaves rather randomly. The quality factor optimization through individual etching can be seen in more detail in Fig. 3.4, where we observe a general tendency towards higher quality factors. This is consistent with the fact that the given mean quality factor change per minute multiplied by the mean etching time for the 19 recorded samples shows that the quality factor has increased on average by about 900×10^3. From FEM simulations, a plateau of quality factors above 10^6 was expected for a nominal membrane thickness of 20 μm, or etching times of 25 minutes (left) and 10 minutes (right) in Fig. 3.4. While a local quality factor maximum can be identified for a few samples, this phenomenon is not universal and the maximum quality factor approaches

10^6 here only for 3 out of 19 samples. A number of observations can explain this difference:

- The membrane was simulated as a quartz slab of uniform thickness, while the real membrane is composed of tiny crystals and even has holes for longer etching times (see Fig. 3.5). Ongoing work at ONERA aims at a reduction of the membrane roughness to remove this problem.

- A strong dependence of the quality factor on temperature was observed: for one sample, quality factors ranging from 200×10^3 to 800×10^3 were measured on two consecutive days with the mounting in vacuum left unchanged. As fluctuations of the environment temperature was neither considered in the simulation nor in the characterizations, the optimization step in the fabrication may have missed some quality factor optima. Future simulations should ensure robustness of the design with respect to temperature change, especially in light of the envisioned cryogenic operation of the samples.

- The etching optimization for some samples was carried out with the photoresist layer left on top of the pillar. While this corresponds to measurements in a frequency region that has not been simulated and therefore may explain the non-observation of the quality factor optimum, it is rather desirable since the later mirror coating leads to a final compression mode frequency close to the one of a pillar with metal and photoresist layers (cf. Table 3.1).

- It was not clear in which parameter regime the robustness of the FEM design had been tested since the responsible engineer had already left ONERA.

The non-attainment of satisfactory quality factors for quantum experiments with the design membrane thickness led us to experimentally explore the regime of even thinner membranes, which eventually resulted in many samples with quality factors in the 10^6 region. We explain this improvement by the decreased mode density of the membrane and its increased elasticity, with the obvious limit of a fully isolated pillar for a zero membrane thickness. The downside of this regime is the increased fragility of the samples and the reduction in heat conductivity between the micropillar, which is expected to be heated by laser absorption in the experiment, and the sample chip, which can be well-thermalized at cryogenic temperatures.

FEM simulations performed with COMSOL after the conclusion of the fabrication [Zer17] have shown that the geometry of the pillar suspension is still useful for obtaining high quality factors: the sparse mechanical mode spectrum in the decoupling shield rarely overlaps with the comression mode frequency. The transmission of mechanical vibrations from pillar to the sample support is maximum when pillar and shield modes coincide and minimum for the opposite case. The non-overlapping case is of course more common for a sparse spectrum. The strong and seemingly random dependence of the quality factor on temperature, membrane thickness and added or removed weight is also explained by this simulation, as the observed shifts of the pillar frequency are typically larger than the frequency difference between two consecutive modes of the decoupling shield, thus leading to a random sampling of the vibration transmission spectrum of the shield. This mechanism is discussed in more detail in section 3.1.4.

3.1.3 Mechanical characterization

We characterize the mechanical properties of our samples by interferometric measurements of the motion of the pillar surface for different kinds of excitation. Several different Michelson interferometers have been used in this thesis, which differed mostly in the employed laser source and wavelength: the interferometer in the ONERA cleanroom uses a HeNe laser at 633 nm, while the interferometers at the LKB worked with different available lasers near 1064 nm. The schematic of a typical Michelson interferometer is presented in Fig. 3.8. The lens before the micropillar sample is chosen to produce a beam waist smaller than 100 μm near the pillar sample. To focus the optical beam on the micropillar surface, the sample is mounted on a mirror mount on top of a manual three-axis translation stage inside the vacuum chamber. While the chamber is open, the sample is first aligned normal to the beam with the mirror mount while the laser spot is reflected off a corner of the sample. The translation stage is then used to position the beam on the micropillar by mapping the sample surface as known from the photomask design (Fig. 3.3) with the reflected intensity measured by the photodiode. Finally, the interferometer alignment is optimized by balancing the power arriving at the photodiode from the signal and local oscillator path, by superposing the local oscillator beam on the signal beam and by aligning the beams on the center of the photodiodes. The calibration of the photodiode output in displacement units is described in 3.1.3.

Mechanically-driven response measurement

Once the Michelson interferometer is locked at a zero of the interference fringes (see Fig. 3.10), the electrical signal provided by the photodiode is maximally sensitive to the displacement of the micropillar surface and we can proceed with searching the resonance frequency of the compression mode of the sample under test. This is usually done by recording the displacement resulting from a mechanical drive of the sample support generated by a small piezo for the frequency interval around the expected resonance frequency with a network analyzer. A crude estimate of the resonance frequency of the micropillar compression mode can be given by dividing the sound velocity [Bec58] along the pillar axis $c = \sqrt{\frac{E_{33}}{\rho}} \approx 6.3\,\mathrm{km/s}$ by twice the thickness $2h = 2\,\mathrm{mm}$, which yields $\Omega_0/(2\pi) \approx 3.2\,\mathrm{MHz}$. If the displacement signal dominates the optical noise of the measurement, a typical result as shown in Fig. 3.9 (left) features a resonance peak on top of a complicated background displacement. The peak frequency is about 20 % above the estimated resonance frequency Ω_0 but much closer to values predicted from finite element simulations. Furthermore, since we expect from the pillar design the compression mode to have the highest quality factor of all modes, this peak likely corresponds to this mode. This assumption will later be confirmed by a measurement of the effective mode mass (section 3.1.3). In a small frequency interval around the resonance, the response is dominated by the compression mode and the observed mechanical susceptibility from Eq. 2.1.18 can be approximated as the sum of one mechanical mode with frequency Ω_0, mass m_0 and quality factor Q_0

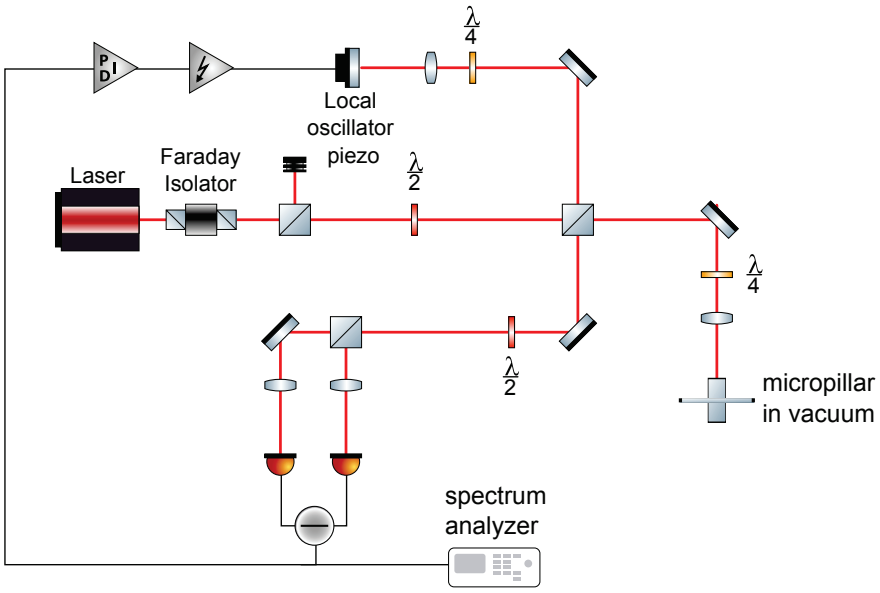


Figure 3.8: Schematic of the Michelson interferometer used for microresonator characterization at the LKB. The laser beam is split into the upper local oscillator (LO) and the right signal beam, which are recombined after reflection off a piezo-mounted LO mirror and the micropillar, respectively. A $\lambda/2$ waveplate and a polarizing beamsplitter cube cause the two beams of orthogonal polarizations to interfere, the corresponding signal being detected by a differential photodiode. The low-frequency photocurrent is amplified and used to stabilize the LO phase, while the higher frequencies are monitored by a spectrum analyzer or network analyzer. The micropillar is typically mounted in a vacuum chamber on a fast piezo that drives its motion for frequency response or ringdown measurements.

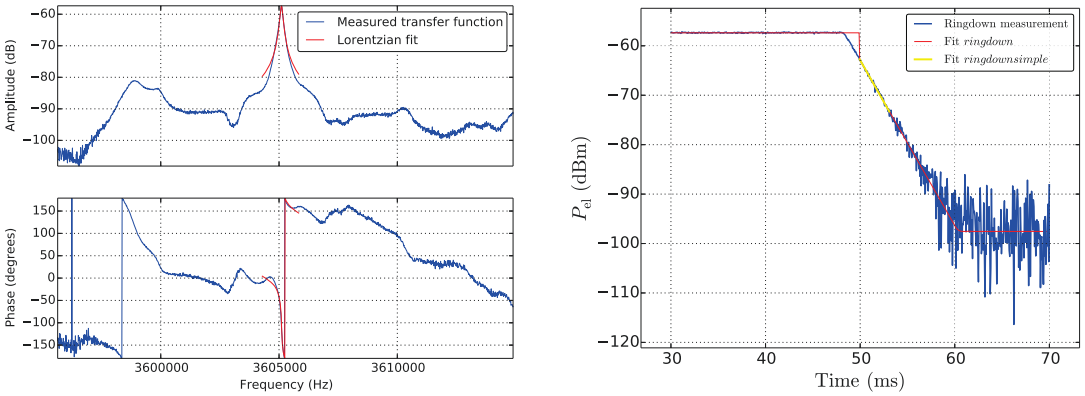


Figure 3.9: Left: Measurement of the response of sample 20121002C9 to a frequency sweep of the piezo excitation at room temperature and atmospheric pressure. The compression mode corresponds to the central peak. In a small region around the resonance, a Lorentzian fit can be used to extract the peak response, response phase and resonance width. The fit yields a quality factor of 32×10^3. Right: Ringdown measurement for sample 20121002C9 at room temperature and atmospheric pressure. The drive is visibly turned off at time 48 ms. The fit function "ringdown" yields a quality factor of 29.4×10^3 while the function "ringdownsimple" yields 29.1×10^3.

and a coherent background term $Re^{i\phi}$ with constant magnitude R and phase ϕ:

$$\begin{aligned} \chi_m(\Omega) &= \frac{1}{m_0\Omega_0^2(1 - i/Q_i - \Omega^2/\Omega_0^2)} + Re^{i\phi} \\ &\approx i\frac{Q_0}{m_0\Omega_0^2}\frac{1}{1 - i\frac{\Omega-\Omega_0}{\Omega_0/(2Q)}} + Re^{i\phi}\,. \end{aligned} \tag{3.1.1}$$

The Lorentzian approximation ($\Omega - \Omega_0 \ll \Omega_0$ and $Q \gg 1$) was made in the last step and yields an excellent approximation for the compression mode. With an additional global phase factor to account for electronic delay, this formula is used to locally fit the measured response to extract the resonance frequency and quality factor.

Three problems typically occur when the quality factor is estimated from a response measurement:

1. If the frequency is swept faster over the resonance than the typical mechanical response time $\tau_0 = \frac{Q_0}{\Omega_0}$, a mechanical ringing phenomenon analogous to the optical one presented in section 2.2.5 is observed. This makes the acquired response spectrum depend on the sweep rate. Since this effect becomes more pronounced for higher quality factors, it may even lead to miss a resonance altogether for fast broadband measurements. It thus imposes very long measurement times for high quality factors.

2. The quality factor extracted from a fit as in Fig. 3.9 where the background is not modeled in sufficient detail depends on the frequency interval over which the fit is performed. This effect becomes stronger for smaller quality factors which result in a lower relative resonance peak amplitude.

3. Drifts of the mechanical resonance frequency during the measurement may cause an artificial broadening of the resonance peak and thus an underestimation of the quality factor.

All three problems can be avoided when a ringdown measurement is performed to estimate the quality factor.

Ringdown measurement

For a ringdown measurement, the resonance to be characterized is driven until the resulting vibration reaches a steady state. The drive is then turned off abruptly and the decay of the vibrational energy is measured by a time-resolved detection of vibration signal power, where the raw signal is filtered by a sufficiently wide electronic bandpass filter centered around the mechanical resonance frequency. Any ringing phenomenon immediately disappears once the drive is turned off, since the ringing is a result of interference of the driven motion with the one stored in the resonator. If background modes within the filter bandwidth are excited, their lower displacement amplitude and quality factor (as can be inferred from the response measurement in Fig. 3.9) causes the vibrational energy in these modes to decay much faster than the one in the high-quality factor mode to be tested, and these modes are absent from the ringdown signal after a short transient. Finally, drifts of the resonance frequency that remain within the (flat) filter passband do not change the detected power significantly. From Eqs. 2.1.1 and 2.1.12, we find the equation of motion of a single harmonic oscillator with constant loss angle $\phi(\Omega) = 1/Q_0$ in absence of external forces

$$\ddot{x}(t) + (1 - i\phi_0)\Omega_0^2 x(t) = 0 \tag{3.1.2}$$

which can be solved by an exponential Ansatz for the initial conditions $x(t_0) = x_0$ and $\dot{x}(t_0) = 0$, where x_0 is the displacement when the drive is turned off and we assume without loss of generality that the velocity is zero at this moment. The solution is in the Lorenzian approximation

$$x(t) = x(t_0)e^{i\Omega_0 t}e^{-\gamma_0 t/2} \tag{3.1.3}$$

where the damping rate $\gamma_0 = \Omega_0/Q_0 = \Omega_0\phi_0$ was introduced. The energy stored in the harmonic oscillator, proportional to $|x^2(t)|$, thus decays with the damping rate γ_0.

The ringdown measurement can be implemented with a function generator in burst mode to generate an excitation signal of fixed duration and a spectrum analyzer in zero span mode to record a signal proportional to the squared displacement versus time. For an unambiguous measurement, the spectrum analyzer's resolution bandwidth must be much larger than the expected resonance width and much larger than the maximum expected frequency drift. A useful preliminary test of the measurement electronics is to perform a ringdown measurement with the function generator output directly connected to the spectrum analyzer input, as to characterize the lower limit of quality factor measurable with the electronic setup. The obtained value should be limited by the spectrum analyzer bandwidth setting. Fig. 3.9 shows a ringdown measurement with different fits to extract the quality factor. Both fits yield similar quality factors. While more parameters such as peak amplitude and background noise level can be

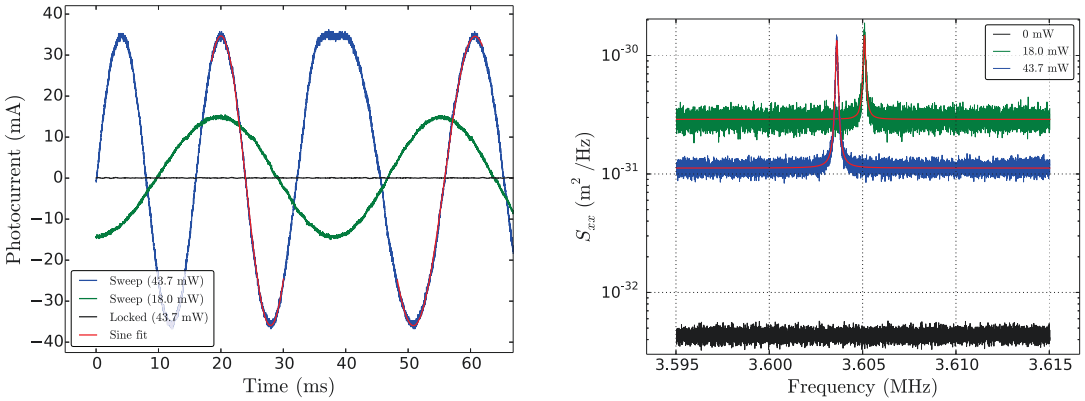


Figure 3.10: Left: DC-output signal from the photodiode for swept and locked Michelson interferometers. The interferometer's sensitivity is estimated with sinusoidal fits to obtain the fringe amplitude. These fits must be constrained to only 1 fringe per fit, as vibrations and inhomogeneous sweep rates make a multi-period fit impossible. Right: Thermal noise spectra of sample 20121002C9 at room temperature and atmospheric pressure. The black dark noise trace was converted to displacement units with the conversion factor of the blue trace. Assuming an effective mode temperature of 300 K, the red Lorentzian fits yield an effective mass of $(33.5 \pm 1)\,\mu$g for both measurements. See main text for calibration details.

extracted from the full fit "ringdown", the fit function "ringdownsimple" automatically detects the most relevant part of the ringdown curve and applies a linear fit to it, which leads to a more robust convergence for automated fits over a wide range of parameters. The latter fitting algoritm was fine-tuned by minimizing the variance of the quality factor extracted from 100 measurements with identical settings.

Measurement of Brownian motion

When no external force is applied to the micropillar, we expect its residual motion to be dominated by the drive from thermal noise as derived in section 2.1.7. From Eqs. 2.1.19, 2.1.22 and 3.1.1, we can derive the thermal displacement spectrum of a harmonic oscillator near resonance:

$$S_{\mathrm{xx,th}}(\Omega) = \frac{2k_B T Q_0}{m_0 \Omega_0^3} \frac{1}{1 + \left(\frac{\Omega - \Omega_0}{\Omega_0/(2Q_0)}\right)^2}\,. \tag{3.1.4}$$

To convert the measured photocurrent spectrum into displacement units, the interferometer can be calibrated by measuring the amplitude $\bar{I}$ of the interference fringes observed when local oscillator path length is slowly swept (see Fig. 3.10 left). Since the photocurrent depends sinusoidally on the difference length x between local oscillator and signal path $I(x) = \bar{I}\sin(2kx)$, we find the conversion factor from displacement to photocurrent units at the working point $\bar{x} = 0$ to be

$$\frac{\mathrm{d}I(x)}{\mathrm{d}x} = \frac{4\pi\bar{I}}{\lambda}\,. \tag{3.1.5}$$

Accounting for the frequency dependence of the photodiode gain, a measured electrical photocurrent spectrum is thus straightforwardly[2] converted into displacement units, as shown in Fig. 3.10. The figure shows that a flat noise background limits the interferometer's displacement sensitivity. The flat frequency dependence and the amplitude of the noise, together with its linear dependence on the mean photocurrent, suggest that the measurement is limited by the quantum fluctuations of the measured photocurrent (section 2.2.3) with spectral density $S_{II,\mathrm{shot}} = \bar{I}$, leading to an effective displacement noise floor of

$$S_{\mathrm{xx,shot}}(\Omega) = \frac{\lambda^2}{16\pi^2 \bar{I}}\,. \tag{3.1.6}$$

If the interferometer is misaligned, the fringe amplitude $\bar{I}$ underestimates the mean photon flux. In this case, the measured shot-noise exceeds the expected one by the ratio between real and estimated photon flux. Since the thermal noise driving the mechanical mode and the shot-noise are uncorrelated, their spectral densities may be added to yield a theoretical expression to fit the experimental spectrum in Fig. 3.10. By assuming the sample to be at room temperature, the effective mass m_0 of the compression mode can be extracted from the fit parameters with formula 3.1.4. The two measurements yield an estimated effective mass $m_0 = (33.5 \pm 1)\,\mu\mathrm{g}$ which is in good agreement with the theoretical prediction $m_{\mathrm{eff}} = \frac{1}{2} \times 2650\,\mathrm{kg/m^3} \times 1\,\mathrm{mm} \times \frac{\sqrt{3}}{4}(240\,\mu\mathrm{m})^2 \approx 33.0\,\mu\mathrm{g}$ from Eq. 2.1.24, where the presence of the membrane and coupling shield around the micropillar is neglected.

A measurement of Brownian motion in vacuum would yield an even better ratio between Brownian motion and shot-noise due to the higher quality factor of the compression mode. However, the measurement at atmospheric pressure was found more robust: in vacuum, the frequency drifts of the compression mode due to laser heating and absence of heat removal by the surrounding gas become comparable to the mechanical resonance linewidth within the measurement time and may lead to an underestimation of the mechanical quality factor. We confirm that this effect is negligible at atmospheric pressure by comparing the quality factor extracted from the thermal noise spectrum ($Q_{\mathrm{thermal}} = 28.2 \times 10^3$) with the value from a ringdown measurement ($Q_{\mathrm{ringdown}} = 29.3 \times 10^3$) and finding reasonable agreement.

To motivate the assumption that the mechanical mode thermalizes at room temperature and that the micropillar is not heated significantly by the incident laser, we can estimate the temperature from the frequency of the compression mode. For another sample, the compression mode frequency was measured to linearly depend on temperature near room temperature with a proportionality coefficient of $-248\,\mathrm{Hz/K}$ (section 3.1.4). From the difference frequency of 1.48 kHz between the observed thermal noise peaks for the presented measurements with two different powers, one can estimate a total frequency shift near 2.5 kHz with respect to the unheated frequency for the measurement with 43.7 mW, and thus expects the micropillar to be heated by 10 K by the laser, which causes an error of 3 % in the estimation of the effective mass.

[2] Care must be taken since the theoretically derived spectra in this thesis are symmetric double-sided spectra $S^{(\mathrm{ds})}(-\Omega) = S^{(\mathrm{ds})}(\Omega)$. Spectrum analyzers yield single-sided spectra $S^{(\mathrm{ss})}(\Omega) = S^{(\mathrm{ds})}(-\Omega) + S^{(\mathrm{ds})}(\Omega)$ that are only defined for positive frequencies. Plotted spectra are single-sided unless otherwise stated.

3.1.4 Limitations of the mechanical quality factor

Further experimental efforts have been made to understand the mechanisms of mechanical dissipation in our micropillar samples. In a first experiment, the influence of how the sample is mounted was investigated. The results are shown in Table 3.2 and demonstrate that mechanical strain on the sample support alters the observed quality factor. Since the expected quality factor limit from internal dissipation mechanisms of

Number of turns	0	1/4	2/4	3/4
Quality factor (10^3)	680 ± 30	600 ± 20	770 ± 10	930 ± 50

Table 3.2: Dependence of the quality factor of the compression mode of the micropillar on the clamping force exerted by the sample mounting. Sample 20121009C9 is pressed against a Teflon support by a thin (0.3 mm) sheet of brass that is screwed into the Teflon with two M1.6 brass screws. At 0 turn, the sample is held well enough such that it does not fall off the vertical support surface. The quality factor is measured by a ringdown measurement in vacuuum, and one screw is subsequently tightened in steps of 1/4 turn. An increase of the laser power by a factor of 10 did not lead to a change of the quality factor beyond the specified error.

4×10^6 at room temperature and 4 MHz [War60] is too far above the measured values to invoke a dependence of internal damping on strain as an explanation for the observed mounting-dependent damping, clamping loss[3] is likely the limit of the sample's quality factor. Clamping loss of other mechanical systems was shown to be strongly influenced by the modes of the surrounding support structure [JRK+11], which suggests to examine these modes for our samples.

A measurement of the mechanical response of different locations on the sample surface to a piezo excitation is shown in Fig. 3.11. The compression mode resonance corresponds to the pronounced peak in the response, which is also discernible in the response on other, rather far-away spots of the sample surface. This indicates that the compression mode is not confined exclusively to the pillar. Its nonzero displacement at zones of the sample that are subject to higher mechanical loss, e.g. due to radiation of ultrasound into the support structure, is strong evidence for clamping loss. Rigorously, however, one needs to compute the spatial overlap of the micropillar mode shape with all relevant modes of the support to exclude the possibility of destructive interference of the emitted waves [WRBV+11]. This would require a three-dimensionally resolved measurement of the displacement profile and would yield the spatial component of the coupling ratio of the compression mode to other, higher-loss modes. Experimentally unfeasible, this was done in FEM simulations with COMSOL [Zer17] and confirms qualitatively that the overlap is significant enough to lead to observable effects: The coupling ratio also has a temporal component which can be viewed as the degree of overlap of the compression and support structure modes in the spectrum. Changing the length of the central micropillar in the simulation, which is to first order a way

[3] We denote by clamping loss all loss mechanisms other than dissipation by the bulk material (structural damping), dissipation related to specific surface states and damping by surrounding gas or liquid. With this terminology, clamping loss is equivalent to the often used expressions "acoustic radiation loss", "phonon tunneling loss" or "suspension loss". If the text does not state clearly what part of the resonator is defined as the "suspension", it includes everything except for the central micropillar, i.e. the membrane, the decoupling shield, the quartz chip and the surrounding mechanical structure.

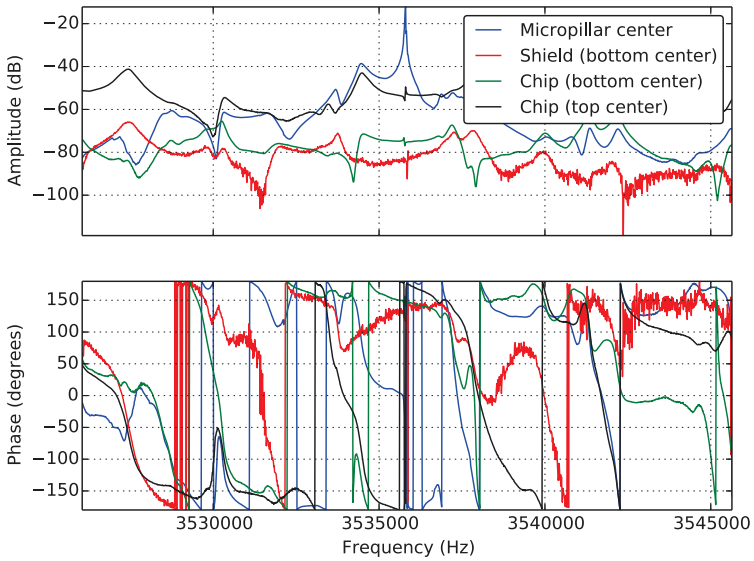


Figure 3.11: Response measurement from a piezo excitation to the vibration of different spots on the surface of sample 20121002C9 at room temperature in vacuum. The amplitude of the different measurements is not comparable since the reflection coefficient at the different measurement positions differed significantly. The compression mode corresponds to the highest peak at the center.

to tune its frequency while leaving the modes of the suspension unaltered, results in maxima of the mechanical dissipation when the micropillar frequency coincides with the frequency of a suspension mode.

The experimental presence of support structure modes in the proximity of the compression mode is also clearly visible in Fig. 3.11. Frequency shifts of the compression mode with respect to other, spatially overlapping modes can be caused experimentally by laser-heating of the pillar alone, by a change of micropillar mass or frequency during fabrication (section 3.1.2), or by a change of the sample temperature since different mode frequencies have different temperature-dependence. All three effects have been experimentally observed to influence the mechanical damping of the compression mode.

To study the temperature-dependence of the mechanical quality factor more systematically, the same sample that was used for the above measurements was mechanically characterized in the temperature range from 0.06 to 300 K. The sample, at the time still without mirror coating and with all metal layers removed, was mounted on a polished copper surface of the mixing chamber plate of the dilution refrigerator in a stack composed of a thin piezo, the micropillar sample and finally a thin sheet of brass that was pressed onto the sample by two lateral screws with a force comparable to 2/4 turns in Table 3.2. The cryostat was then cooled to dilution temperatures and heated back to room temperature over the course of five days in total. The results from continuously performed frequency response and ringdown measurements are shown in Fig. 3.12. The observed frequency-dependence of the compression mode is in agree-

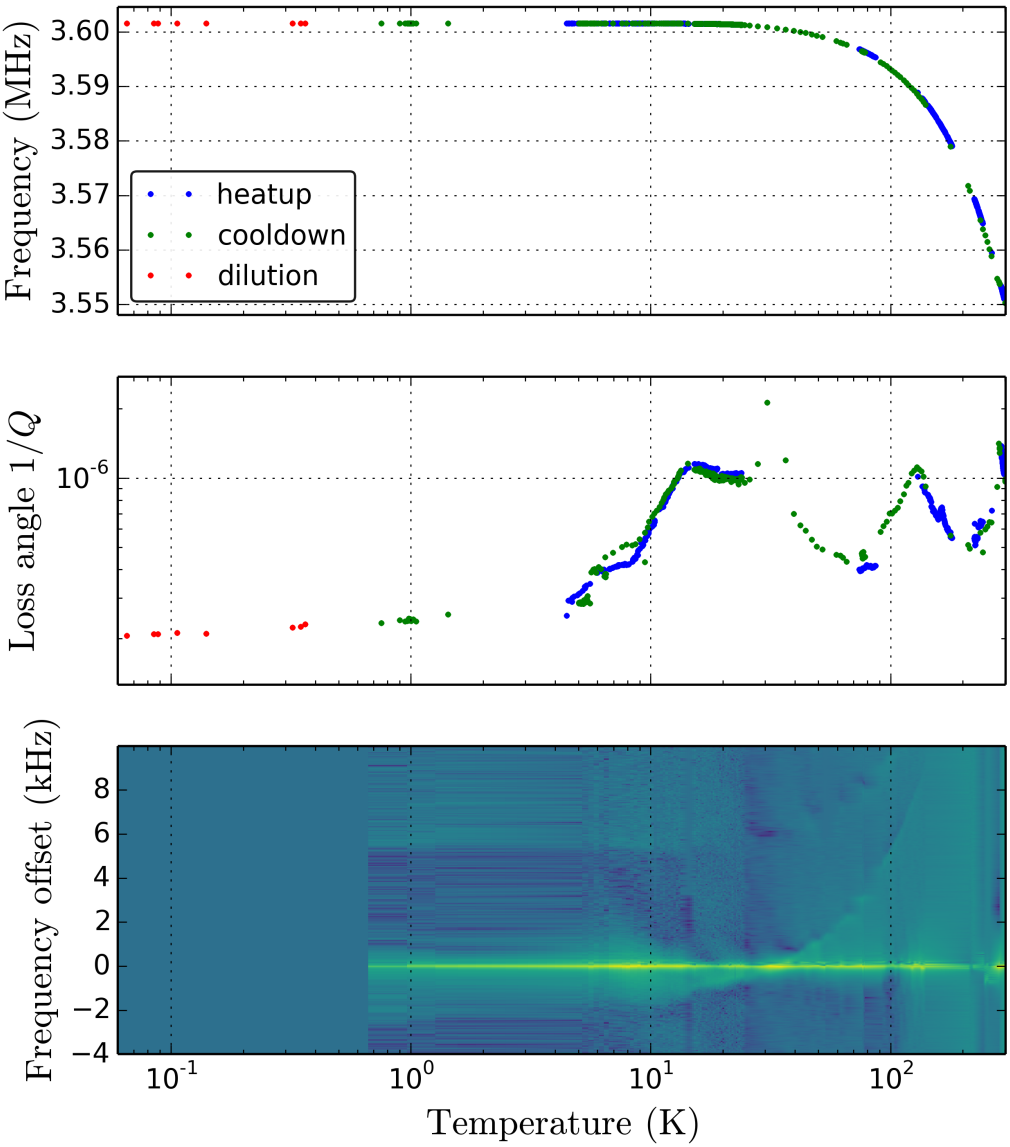


Figure 3.12: Mechanical characterization of sample 20121009C9 in vacuum during a cryostat run. The two upper plots show the frequency and loss angle of the compression mode measured by the ringdown method during cool-down, dilution operation and heat-up of the refrigerator. The lower plot shows the mechanical response spectra to a piezo excitation centered around the compression mode resonance frequency. The color scale ranges from dark blue (-90 dB) to bright yellow (-20 dB). Only three spectra were available below 4 K. The frequency of the curved spectral feature that crosses the mechanical resonance at about 30 K was confirmed to change with temperature in order to exclude the possibility of an electrical parasite. The feature is furthermore present both during cool-down and heat-up. All data were acquired automatically with a computer-controlled Michelson interferometer lock. Still at an early stage of the experiment, frequent computer crashes and divergence of the frequency search algorithm caused temperature zones with missing datapoints both during heat-up and cool-down. Care must be taken not to misinterpret the color plot due to missing data, which result in horizontal lines such as between 1 and 4 K.

ment with other measurements for compression modes in alpha-quartz [Sch08]. From linear fits of the frequency in the ranges between 0 and 10 K and between 250 and 300 K, we can deduce the useful temperature-coefficients of the resonance frequency of $\frac{\mathrm{d}f}{\mathrm{d}T} = -60\,\mathrm{mHz/K}$ and -247.8 Hz/K, respectively. The room-temperature value is reasonably close to the estimate of -278 Hz/K with literature values from [Kuh13]. The low-temperature value is prone to a large error since frequencies were only discernible with minimum steps of 1 Hz, and can be assumed to be zero for all practical purposes. The temperature-dependence of the loss angle is more complicated. We observe four maxima at 15 K, 33 K, 129 K and 283 K. The highest loss peak at 33 K can be seen to coincide with the crossing of another peak through the compression mode peak in the frequency response plot. Other loss peaks, in particular the one at 283 K, also show features resembling a resonance crossing in the response plot, which are however less clearly discernible, partially due to missing data in some temperature regions (see caption of Fig. 3.12). While we expect a maximum of intrinsic loss at 15 K and 33 K [Sch08], we attribute the other maxima to clamping loss resulting from a degeneracy of the compression mode and support structure modes. In the low-temperature region, this degeneracy becomes less likely due to extremely low intrinsic loss [GCI$^+$12] and the resulting lower ratio of the mean width per average support structure mode spacing. The lowest loss angle of 2×10^{-7} is reached at the lowest accessible temperature of 66 mK.

3.2 Coupling mirror

In order to build a high-finesse Fabry Perot cavity with the micropillar resonator as flat end mirror, a second concave input coupling mirror must be chosen carefully. In the subsequent subsections, we will outline the constraints on the choice of the coupler imposed by the micropillar end mirror, before describing the mirror fabrication and characterization. We will conclude with a description of the observed effects of optical coupling between the fundamental and higher-order Hermite-Gaussian modes in cavities constructed with the fabricated mirrors.

3.2.1 Experimental constraints

The aim of this section is to find a mirror geometry that allows for Fabry-Perot cavities with a finesse $\mathcal{F}$ of at least 10^5. This finesse is equivalent to a sum of optical loss and coupling mirror transmission of about $P+T = 2\pi/\mathcal{F} \approx 60$ ppm per roundtrip. Aiming for the loss to be smaller than the coupler transmission in order to be able to detect the intracavity light with a high quantum efficiency, we find at most 30 ppm of optical loss acceptable. The mirror coatings available from the LMA having very low absorption loss of the order of 1 ppm, we expect the dominant optical loss mechanisms to be scattering due to residual surface roughness of the mirrors and clipping due to finite mirror size.

Scattering loss

According to Eq. 2.2.2, the electric field of a fundamental Gaussian beam can be written in polar coordinates r, ϕ, z as

$$E(r,\phi,z) = E_0(z)\exp\left(-\frac{r^2}{w(z)^2}\right). \tag{3.2.1}$$

A rough surface with excess height $\xi(r,\phi)$ with respect to a surface perfectly matched to the beam's phase front yields a reflected beam with the excess phase profile $2k\xi(r,\phi)$. By approximating the intracavity field by a fundamental Gaussian beam, we find the fraction that remains in the cavity mode, i.e. the mirror's amplitude reflection coefficient r, by forming the overlap integral between the phase-distorted beam and a perfectly reflected one [Vir, Arc06]:

$$r = \frac{1}{\pi w(z)^2}\int_0^\infty\int_0^{2\pi} e^{-2ik\xi(r,\phi)}\exp\left(-\frac{2r^2}{w(z)^2}\right) r\mathrm{d}r\,\mathrm{d}\phi\,. \tag{3.2.2}$$

Computing the fraction of intracavity power that is lost through scattering $P_\mathrm{D} = 1-r^2$, we find through a second-order expansion of the last exponential

$$P_\mathrm{D} = 4k^2\bar{\xi}^2 \overset{1064\,\mathrm{nm}}{\approx} \left(\frac{\bar{\xi}}{1\,\text{Å}}\right)^2 \times 1.4\,\mathrm{ppm} \tag{3.2.3}$$

where the average square roughness $\bar{\xi}^2 = \int_0^\infty\int_0^{2\pi}\xi^2(r,\,\phi)\exp\left(-\frac{2r^2}{w(z)^2}\right)r\mathrm{d}r\,\mathrm{d}\phi$ was introduced as a measure of surface roughness over the transverse dimension characteristic of the optical beam. We see that an average roughness $\bar{\xi}$ below 3 Å is therefore required to keep the scattering loss of two mirrors below 20 ppm.

Clipping loss

Using the same reasoning as in the previous section and replacing the rough surface by a perfect circular aperture of diameter d, we find that clipping loss P_C amounts to

$$P_\mathrm{C} = e^{-\frac{1}{2}\left(\frac{d}{w(z)}\right)^2}. \tag{3.2.4}$$

The steepness of the exponential makes a sufficient ratio $\frac{d}{w(z)}$ with reasonable safety margin crucial for obtaining a high finesse cavity: while a ratio of 5 yields an acceptable clipping loss level lower than 4 ppm, a ratio of 4.5 already causes 40 ppm of loss per mirror.

From the fabricated micropillar mirrors (see Fig. 3.6) a useful mirror diameter of about 60 μm can be estimated. For a ratio of 5 between mirror size and beam waist, we must therefore design the optomechanical Fabry-Perot cavity with a waist of less than 12 μm. The waist w_0 on the flat mirror of a plano-concave cavity can be computed from the cavity length L and the mirror's radius of curvature R_c with the relations from Eq. 2.2.2:

$$w_0 = \sqrt{\frac{\lambda}{\pi}\sqrt{L(R_\mathrm{c}-L)}}\,. \tag{3.2.5}$$

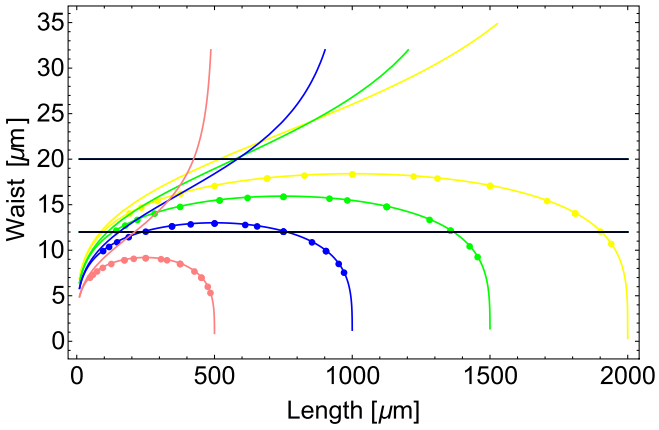


Figure 3.13: Beam waist size on the flat (lower curves with dots) and on the concave mirror (upper, diverging curves) of a plano-concave Fabry-Perot cavity as a function of the cavity length, for different radii of curvature (pink: 0.5 mm, blue: 1 mm, green: 1.5 mm, yellow: 2 mm). The dots indicate the cavity lengths where a higher-order transverse mode TEM_{nm} (up to $n + m = 10$) is expected to become degenerate with the fundamental TEM_{00} mode. Two horizontal lines show the upper design limit for the waist size on the flat (12 μm) and concave (20 μm) mirrors.

The waist $w(L)$ on the concave mirror is obtained through the normal evolution of a Gaussian beam as given in Eq. 2.2.2. Figure 3.13 plots the two beam waists in a plano-concave Fabry-Perot cavity for different radii of curvature as a function of cavity length. The plot shows that even for a very short radius of curvature (such as 1 mm), the requirement of a small waist size on the flat mirror constrains the cavity length to be either very short or very close to the radius of curvature. While for a short cavity length, the waist size on the second mirror remains comparable to the first waist, it diverges when the cavity length becomes comparable to the radius of curvature. To avoid stability problems in the latter regime, which were already investigated in other groups[GBP+06b], we choose to operate our optomechanical system in the limit of short cavity length. Building very short cavities with conventional 1/4" concave mirrors is impossible since the configuration where the two mirror samples are in contact imposes a minimum cavity length. Furthermore, a larger range of cavity lengths with low clipping loss is acceptable when unconventionally strongly curved mirrors are used: for example with a $100 - \mu$m diameter mirror with radius of curvature about $800\,\mu$m, Fig. 3.13 suggests that we can construct high-finesse cavities up to $500\,\mu$m long. If a shorter cavity length is acceptable, the concave mirror may even be smaller. With any given mirror, we expect the lowest optical loss for the shortest accessible cavity length.

3.2.2 Fabrication of coupling mirrors

Traditional polishing methods can yield spherically shaped mirror substrates of extremely low roughness (< 1 Å rms). Unfortunately, the lowest commercially available radii of curvature for such substrates are in the range of 5 to 10 mm. Low roughness (< 2 Å rms) and short radius of curvature (< 2 mm) mirror substrates are up to

now only accessible through CO_2 laser ablation of silica[HDB$^+$12]. In this process, a powerful (1 - 30 W) and focused (waist size 10 - 200 μm) pulse of CO_2 laser light is directed on a flat surface of fused silica. Since the wavelength 10.6 μm of CO_2 lasers lies deeply in the absorption band of silica, this procedure deposits energy onto the surface with a profile that reproduces the Gaussian shape of the laser beam, which results in a rapid heating of the material. The associated high surface temperature results in evaporation of the material [NBH15], thus creating a nearly Gaussian-shaped dip in the surface, which at its center can be used as an approximately spherical mirror substrate. A very low surface roughness of the order of 2 Å is routinely observed in the ablated region, which is most likely a result of the large surface tension of the molten silica during solidification after the laser pulse.

Single-shot fabrication

When a single laser pulse is used to create a concave structure, the beam waist primarily defines the diameter of the resulting feature, while the pulse duration and power control the depth, and can thus be used to tune the radius of curvature. Significant ablation is only obtained above a certain threshold power density, which is theoretically expected[NBH15] and experimentally confirmed to lie in the 0.1 MW/cm^2 range. Given a typically available laser power of 20 W, this constrains the diameter of obtainable structures to about 100 μm. Typical pulse durations of the order of 100 ms were used to obtain structures with a radius of curvature of about 1 mm. As the detailed dependence of the obtained geometries is very nonlinear in the laser power, pulse duration and waist size[HDB$^+$12], and further depends on the type of silica substrate used, we optimize the ablation pulse parameters experimentally. One realized fabrication protocol, which resulted in samples L1 to L21, is the following:

1. We begin by purchasing plates of polished Herasil of dimensions 75 mm × 25 mm × 1 mm from the company HTM in Voiron, France. Plates made of Suprasil were also used successfully.

2. A standard anti-reflection coating optimized for the laser wavelength of 1064 nm is applied to the back side of the glass plates by the company Fichou in Fresnes, France. Ignorant of the details of the later mirror fabrication steps at the moment of fabrication of this coating, it was later realized that this low-quality polymer coating was damaged by the final bakeout step at 500 °C (Fig. 3.14). Future samples should be AR coated with SiO_2 and Ta_2O_5 layers instead. While the damaged film significantly alters the geometric shape of a traversing laser beam, the existing samples may however be used to characterize the fabrication process or to fabricate phononic-crystal mirrors (see section 7.4.1).

3. The glass plates are cut into squares of 10 mm × 10 mm × 1 mm with a diamond saw in the glass blowing workshop at LKB. To avoid scratches on the surface during the cutting process, a layer of standard photoresist or nail polish is applied on both sides and removed with acetone after the cutting. Samples are then thoroughly cleaned by ultrasonic baths in acetone, isopropanol and water for about 30 s each, then dried after dipping in a fresh bath of isopropanol.

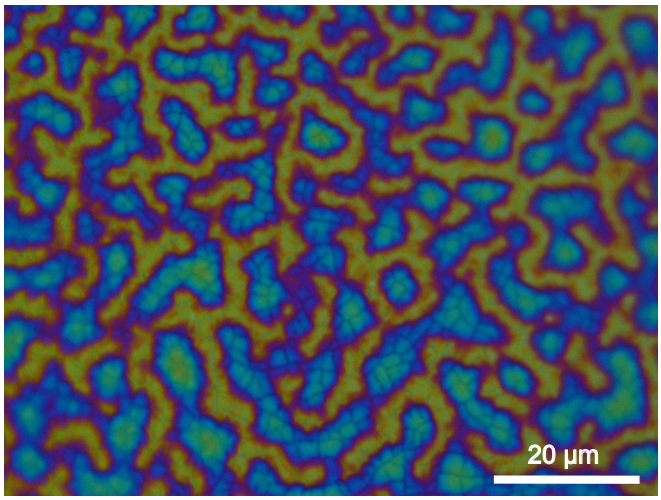


Figure 3.14: Burned AR coating on the back surface of sample L1. The colorful pattern is the result of the interference of light in the thin AR film of varying thickness after it was damaged in a bakeout step at 500 °C.

4. With one test sample, the ablation setup is calibrated. The setup consists of a Synrad CO_2 laser capable of delivering 27 W of light. The laser beam is superposed with an alignment beam from a red laser diode. The beam propagates in free space over a path length of roughly 2 m before it is focused with a plano-convex lens of 5 cm focal length. The sample is held near the focal plane with a tweezer that is mounted on a manual three-axis translation stage. Once the sample position is adjusted with the translation stage such that it is well in and normal to the beam path, a maximum-power laser pulse with a duration of 100 ms is generated. If the sample is close enough to the focal plane, a bright visible flash indicates that ablation has taken place. In this case, the sample is moved away from the focal plane, and another test shot is performed on a fresh zone on the sample. This process is repeated until the distance from the lens is found where no ablation takes place any more, that is where the available power density has decreased below the threshold value due to the beam's divergence. While increasing the pulse duration at this point allows to re-enter the ablation regime, and the displacement optimization along the z-axis may be repeated to further increase the beam diameter on the sample, this did not lead to significantly larger structures in our experiments. Once the distance from the lens is optimized, a series of structures are fabricated with varying pulse durations. The geometry of these samples is characterized with a stylus profiler and the relationship between radius of curvature and pulse duration established (Fig. 3.15). When the alignment laser is transmitted by a concave structure, interference fringes can be seen if a white screen is placed 10 to 50 cm behind the sample. The fringe pattern was found to be correlated with the obtained radius of curvature. A fast approximate in-situ characterization of fabricated samples is possible if the fringe spacing of the interference pattern is recorded along with

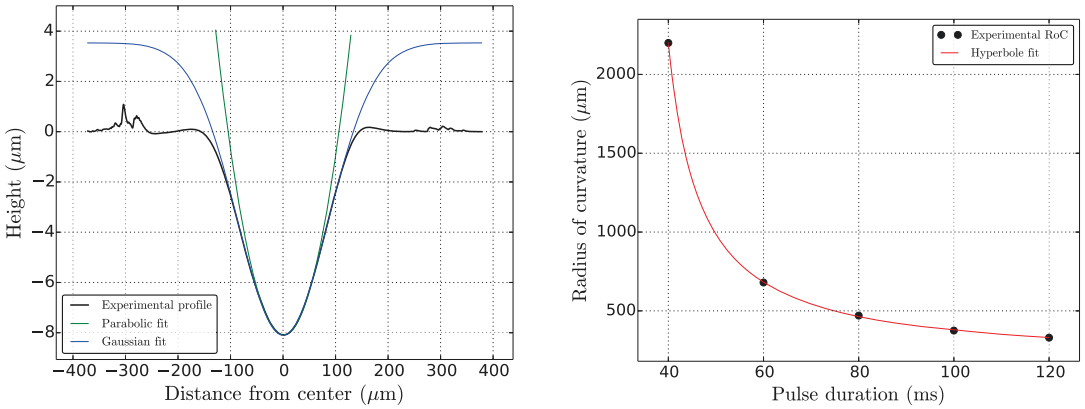


Figure 3.15: Left: Depth profile of the concave feature created by a calibration shot (duration of 120 ms) on a Herasil plate as measured with a stylus profiler. Gaussian and parabolic fits in the central region of the trace yield similar radii of curvature (here 700 μm). The distance between the inflection points of the Gaussian (here 140 μm) or the experimental profile's FWHM (here 138 μm) are practical measures of the usable mirror diameter and among the largest values obtained with the single-shot fabrication method. Right: Radius of curvature R_c for a series of calibration shots with the same laser power and beam diameter but different pulse duration. The presented hyperbolic fit $R_c \approx \frac{13500\,\mu\text{m ms}}{\tau - 33\,\text{ms}}$ is motivated by the fact that R_c is inversely proportional to the Gaussian's depth, which is roughly proportional to the pulse duration above a threshold duration (33 ms in this case).

the calibration shot parameters.

5. Samples are fabricated by inserting a new silica plate into the ablation setup at the right distance from the focusing lens and generating the pulse corresponding to the desired radius of curvature. Structures with larger ($> 500\,\mu$m) radii of curvature are very hard to discern under a microscope, which is why two position markers 3 mm left and right of the center are fabricated on the sample. To make the markers easily visible, the distance from the focal plane is reduced for these shots. To avoid ablated material from the marker shots to condense on the central mirror substrate, the markers are fabricated before the desired structure. Samples are kept a minimum of time in the "dirty" environment of the ablation setup to avoid contamination of the optical surfaces.

6. The samples are cleaned again after the ablation shots as in step 3. Finally, samples are cleaned once more in the RIE machine in the cleanroom of Paris-Diderot University with the recipe "Clean pure O2 30sccm 30W 50mtorr" for 10 minutes to remove any organic residues.

7. Samples are coated with dielectric mirrors of the desired transmission (typically in the range of 30 to 120 ppm) by the LMA (see details about the coating process in Fig. 3.7). The optical coating loss is further lowered by performing a standard bakeout at 500 °C for 10 hours, which reduces strain in the coating layers.

In order to protect our samples while gluing them on support plates, or to recover optical properties after a degradation by environmental dust, we routinely use "First Contact" solution from Photonic Cleaning Technologies, LLC [PZD10]. We have not observed any significant excess loss brought about by this cleaning procedure.

3.2.3 Optical characterization

The main optical design parameters are the attainable cavity finesse and the quantum efficiency for the collection of intracavity light. Once the finesse is known, the quantum efficiency can be deduced if the coupling mirror transmission is known. For ease of understanding, we will first discuss the measurement of the transmission and then explain the more complex finesse measurement. To conclude, we will try to understand the dominant optical loss mechanisms that limit both the finesse and quantum efficiency that can be obtained with our samples.

Transmission measurement

The mirrors on the micropillars are designed to transmit less than 1 ppm of the incident light. Since we do not use the resulting very small cavity transmission for anything but as an aid to align the cavity, we will count the residual transmission of the micropillar mirror among the optical loss mechanisms and do not require to carefully characterize it. On the contrary, the transmission T of the coupling mirror is important to know for a number of reasons:

1. As explained in the next section, the sum $T + P$ can be deduced from a measurement of the cavity finesse. In order to compare the optical loss P measured for cavities with different coupling mirrors, one therefore needs a precise figure of the mirror transmission T.

2. The ratio $\eta = T/(T + P)$ defines the fraction of light that leaves the cavity after interacting with the mechanical resonator, and is in our experiment the limiting factor of the overall quantum efficiency and thus important to characterize.

3. The ratio $R_0 = (T - P)^2/(T + P)^2$ gives the reflection coefficient of a cavity at resonance. If a larger resonant reflection is measured, this indicates a geometric mismatch of the incident and the intracavity beam shapes, which is another important factor entering the overall quantum efficiency whose knowledge allows an optimization of the optical setup.

Two independent measurements are used to estimate the power transmission of the coupling mirrors. The first measurement is performed by LMA in a designated setup directly after coating the mirrors on a silica witness sample from the same coating run. Since the coupling mirror substrates are made from a silica type with similar refractive index, it is a reasonable assumption that the power transmissions of witness and coupling mirrors are essentially identical.

The second measurement is performed at LKB with a setup schematized in Fig. 3.16. The main complication in measuring a transmission of the order of 10^{-5} is the associated large dynamic range of power measurements and to prevent reflected light

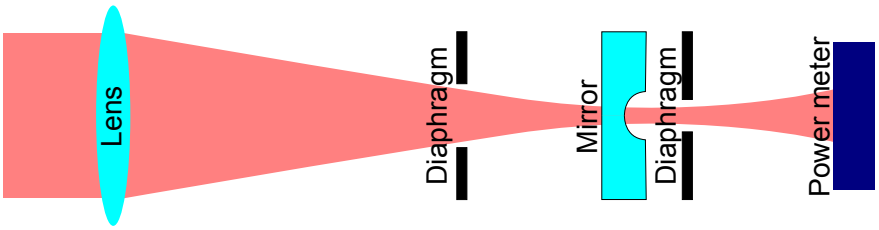


Figure 3.16: Schematic of the arrangement used to measure the transmission of the coupling mirrors. The mirror is mounted on a translation stage and the ensemble is shielded from stray light by connecting the components with lens tubes.

from being scattered onto the power meter. To tackle the first issue, a semiconductor power meter with large dynamic range and good linearity (Thorlabs S120C) is used for the measurement. The second issue is solved by careful alignment of diaphragms, and by redirecting reflected light on a beam block with a Faraday rotator. Before inserting the sample into the setup, the power of the incident laser is chosen close to the maximum rating of the power meter and recorded. With the coupling mirror sample inserted into the setup, a three-step alignment is performed by maximizing the power reflected by the sample: the laser beam is centered with respect to the first diaphragm, set to normal incidence on the sample and the beam waist is aligned to the surface of the sample by moving the lens along the beam axis. The beam waist is designed to be close to the value that will be defined by the later Fabry-Perot cavity geometry (10 to 20 μm). The first diaphragm aperture is then closed until it starts reducing the reflected power. The second diaphragm, which is concentric with the first one by design of the setup, is then closed to a reasonable aperture about 5 times the size of the expected beam diameter at the diaphragm. In this configuration, the ratio of measured transmitted power to the incident power is the power transmission of the coupling mirror. By scanning the lateral position of the sample, a transmission map can be created. For the mirror S6-7, which is characteristic for all samples S6-x, we obtain the values given in Table 3.3. We have observed transmission variations up to

Sample	T_{LMA}	T_{min}	T_{max}	T_{center}	T_{side}
S6-7	35	29	40	35	30

Table 3.3: Different measurements of the mirror transmission in units of ppm. The value T_{LMA} is measured on a witness at LMA, all other values were measured at LKB. T_{min} and T_{max} are minimum and maximum observed transmissions while scanning the sample, T_{center} is the value observed at the center of the concave structure while T_{side} is the value next to it.

30 % near the ablated structure. Since it is likely that these variations are caused by an etalon effect between the mirror coating and the uncoated but parallel back surface of the sample (4 % reflection off the backside of the sample and an incident waist of 16 μm can explain the observed variation), we will use the values T_{LMA} that were measured by the LMA in subsequent sections, which are reasonably close to our observations. The etalon hypothesis was confirmed later on another sample by scanning the laser wavelength while measuring the incident and transmitted power,

where a sinusoidal dependence of the transmission on wavelength was found, with a period corresponding to a cavity length of twice the sample thickness (about 3 mm by taking the refractive index of silica into account) and relative amplitude of the order of 30 %. An antireflection treatment on the backside of the coupling mirrors is certainly the best way to deal with this problem in order to obtain a maximum quantum efficiency for collecting the light injected and reflected from the measurement cavity. However, about half of the free spectral range of the etalon ($\nu_{\mathrm{fsr}} = 100$ GHz) is covered by the tuning range of our Nd:YAG laser, which opens the possibility of fine-tuning the input coupling mirror transmission after its fabrication.

As a third method to measure the coupling mirror transmission, one might use a measurement of the resonant cavity reflectivity R_0 to find T from $T+P$ with Eq. 2.2.50. The error in a measurement of R_0 due to residual geometric mismatch of incident and cavity fields, together with a minimum R_0 at the common condition $T = P$ render this measurement rather imprecise for typical situations. An estimation of the resonant reflectivity R_0 is however very useful to estimate the geometric mismatch factor and to test the consistency of the estimated transmission and loss figures.

Finesse measurement

The finesse of a cavity is defined as the ratio between the free spectral range (FSR), $\Omega_{\mathrm{fsr}} = 2\pi\frac{c}{2L}$, and the FWHM of an optical resonance peak $\kappa = (T + P)\Omega_{\mathrm{fsr}}/(2\pi)$ (see section 2.2.6). For low-finesse cavities, this ratio can be directly extracted from a multi-FSR spanning absorption or transmission spectrum, which can be acquired by recording the cavity reflection or transmission during a linear sweep of the cavity length or laser frequency. For cavities with $\mathcal{F} \gg 10^3$, this method is prone to yield imprecise results because of errors associated with a non-uniform sweep velocity. This error can be reduced when the resonance width and the free spectral range are measured separately.

To estimate the resonance width in units of frequency, we imprint a phase modulation of significant amplitude $\phi_0 \approx 1$ onto the single-mode laser beam used to measure the cavity reflection or transmission during a scan of the cavity length. According to Eq. 2.2.8, the resulting incident field can be viewed as three single-mode laser beams separated by the modulation frequency Ω_{mod}, each of which produces a resonance feature in the low-pass filtered reflection or transmission trace. From the known frequency separation Ω_{mod} between the three resonances, the width of each resonance can be calibrated in frequency units. Fig. 3.17 shows a trace with a fit of three Lorentzians with equal spacing. For this method to yield reliable results, the sweep velocity must only be constant during the passage of the three resonances, which is much easier to accomplish if we choose $\Omega_{\mathrm{mod}} \ll \Omega_{\mathrm{fsr}}$.

To measure the free spectral range, we record oscilloscope traces of the cavity reflection during a sweep of the cavity length over several free spectral ranges with enough datapoints (typically 5×10^6) to obtain good fits for the widths of each observed resonance with the method described in the previous paragraph. With the known resonance separation Ω_{mod}, each of the three-Lorentzian fits can also be used to extract a local sweep velocity in units of frequency per time. The ensemble of extracted sweep velocities with the associated times of resonance passage represents the sampled sweep

velocity function $\dot{\Omega}_{\mathrm{c}}(t)$, which can be fitted with an analytical model. To keep this model as simple as possible, a sinusoidal sweep voltage at frequency Ω_{sweep} is applied to the cavity piezo for this measurement[4]. To model the piezo nonlinearity, the voltage-dependent piezo displacement $x(V)$ is developed into a Taylor series that is truncated after the second order: $x(V) \approx x_0 + \alpha V + \beta V^2$. The model is easily extended to higher-order developments, which we find however unnecessary in our measurements. We find for the time-dependent displacement

$$x(t) = x\left(V_0 \sin(\Omega t)\right) \approx x_0 + \alpha V_0 \sin(\Omega t) + \frac{\beta V_0^2}{2}\left(1 - \cos(2\Omega t)\right) \tag{3.2.6}$$

We finally make a linear approximation of the cavity resonance frequency for the n-th longitudinal cavity resonance in the limit of small variations $x(t) \ll L_0$ of the default cavity length L_0:

$$\omega_{\mathrm{c}}(t) = \frac{2\pi n c}{2\left(L_0 + x(t)\right)} \approx \frac{\pi n c}{L_0}\left(1 - \frac{x(t)}{L_0}\right) \tag{3.2.7}$$

The resulting sweep velocity model $\dot{\omega}_{\mathrm{c}}(t)$ is the sum of two phase-shifted sines at Ω_{mod} and $2\Omega_{\mathrm{mod}}$ and can be fitted to the experimental values if enough values were sampled per period. For good fits, in practice at least three resonances must be passed per half-period of the drive voltage. From the fit parameters of $\dot{\omega}_{\mathrm{c}}(t)$, the frequency difference $\omega_{\mathrm{c}}(t_1) - \omega_{\mathrm{c}}(t_2) = \Omega_{\mathrm{fsr}}$ of pairs of adjacent resonances can be used to deduce the free spectral range and cavity length.

The described procedure to extract the cavity length and finesse from a measurement has been implemented in the python code "finesse.py". With this program, the finesse of a cavity can be measured within seconds such that cavity alignment by optimization of the finesse is possible in real time. Typical measurement data and the corresponding evaluation by the program are shown in Figs. 3.17 and 3.18. To estimate the error in the measurement of the free spectral range, Fig. 3.19 compares the cavity length estimated with a micro-positioner screw and the length deduced from ω_{fsr} measured and evaluated with the procedure described above. While repeated measurement results at the same length were observed to scatter by up to 5 %, the linear fit slope close to unity indicates a lower systematic error. Many fits of individual peaks such as in Fig. 3.17 yield a slightly different distance of the right and left sideband from the central peak, which suggests that the assumption of constant sweep rate $\dot{\omega}_{\mathrm{fsr}}$ during the passage of a resonance does not hold. Together with the low systematic error for the length measurement, this suggests that vibrations of the cavity length are the dominant source of error in the finesse measurement, and not for example an insufficient modeling of the piezo behaviour. We therefore resort to averaging when more precise estimates are desired.

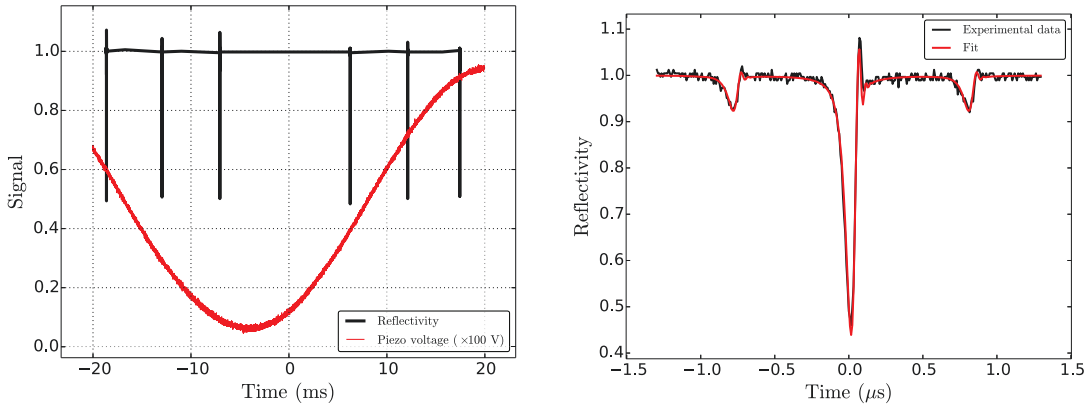


Figure 3.17: Left: Reflection spectrum of a Fabry Perot cavity constructed with the samples S6-9 (concave coupling mirror) and QT2 (flat mirror) with a length of roughly 80 μm. The piezo scan range includes three thin resonances, which are swept over in both directions. Reflectivity data in regions between peaks, where the reflectivity never lies below a threshold of 0.9, are thrown away at the moment of acquisition to save hard-drive space. The overshoot of the reflectivity is caused by the ringing due to the fast piezo sweep rate compared to the width of the resonances. Right: Zoom on a specific resonance dip of the left figure. Two sideband dips due to a 200-MHz phase modulation of the probe laser can be seen. All three resonances are fitted with Eq. 2.2.43. The sideband distance is used to convert the fit result to frequency units, which yields a linewidth of $\kappa/(2\pi) = 15.7\,\mathrm{MHz}$.

Optical loss versus cavity length

The automatization of finesse measurements has enabled us to measure the cavity length dependence of the optical loss for many combinations of mirrors, in order to characterize the different loss mechanisms that limit the performance of our optomechanical system. A sound understanding of the coupling mirror-induced loss has proved useful before considering the complications that arise from introducing the small micropillar mirror into the system. By coating a dielectric mirror of negligible transmission ($T < 1$ ppm) on a flat, superpolished quartz plate in the same coating run as the micropillar, we have fabricated a 'witness' sample for the upper limit of mirror quality that can be achieved with our mirror coatings. One such sample (QT2) has been used as back mirror to characterize many different concave coupling mirrors in a Fabry-Perot configuration. The measurement is performed in a cavity mount as seen in Fig. 4.13, which will later be used to construct cryogenic cavities. Table 3.4 shows the minimum loss obtained in these measurements, which is generally attained for rather short cavity lengths. Assuming all other types of loss negligible for an optimized cavity length, we find that the average observed loss $P = 11$ ppm can be explained by scattering loss (Eq. 3.2.3), for example with an average roughness of 1.5 Å rms on

[4]The commonly used voltage ramp can be developed into a series of sinusoidal components at odd multiples of the sweep frequency. It therefore probes the piezo response at many different frequencies, where it may differ both in amplitude and phase, which leads to a very complicated displacement and sweep velocity function. Experimentally, even regions with inverted sweep direction were observed for triangular voltage ramps.

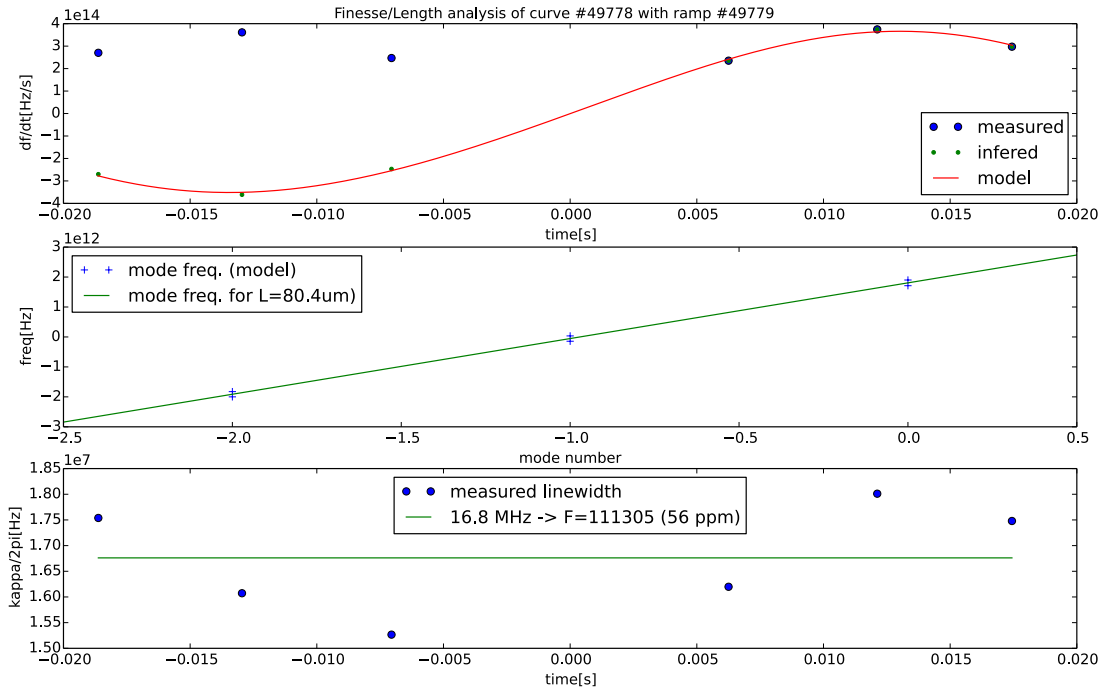


Figure 3.18: Output from the python code "finesse.py" which extracts the free spectral range and the finesse from the reflection of a Fabry-Perot cavity for a sinusoidal sweep. The upper plot shows the sweep rates inferred from the distance of the two sidebands for each resonance. Since only the absolute value of these rates is deduced from the fit, half of the values are multiplied by -1 based on the piezo sweep direction. The red fit corresponds to the model for $\dot{\omega}_c(t)$ explained in the main text. The middle plot shows the integral of the red fit of the upper plot as a function of the resonance number by grouping identical resonances together. The slope of the linear fit corresponds to the extracted free spectral range, used to deduce the cavity length. The lower plot shows the measured linewidth for each resonance, computing the finesse from the mean linewidth and free spectral range.

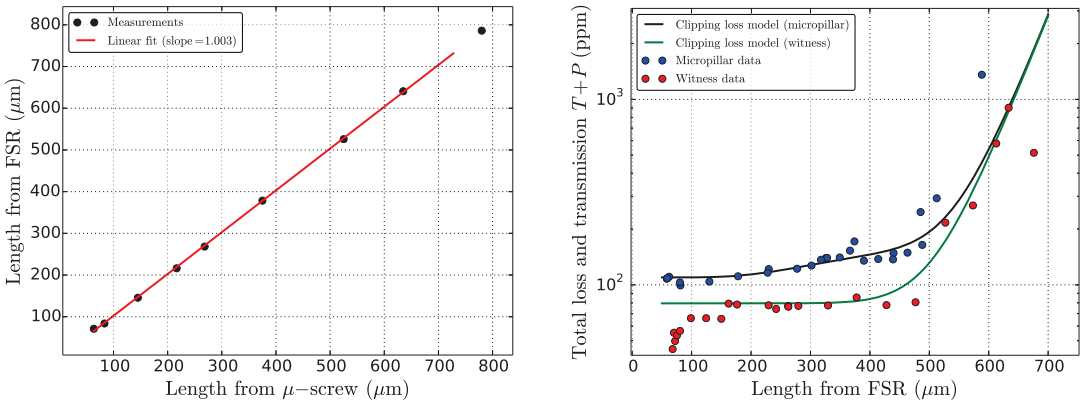


Figure 3.19: Left: Length estimated by the program "finesse.py" versus length estimated by turning a micropositioner screw. The abscissa has an arbitrary offset. Right: Estimated $T + P$ versus cavity length for two cavities, constructed with the same concave coupling mirror (S6-9, $T = 35$ ppm) and either the quartz witness sample QT2 (large, homogeneous flat mirror, baked out after mirror coating) or the micropillar sample 20121009C4 (without bakeout). The clipping loss model consists in the sum of a loss offset and the clipping loss from a circular aperture at each mirror surface. Both curves assume an equal aperture at the coupler (82 μm) and the independently measured radius of curvature $R_c = 900\,\mu$m.

the coupler and 2.4 Å rms on the quartz sample. These values are consistent with the expectations from the fabrication process.

A measurement of the cavity length dependence of optical loss is shown in Fig. 3.19 (red datapoints). We see that the minimum loss only occurs for the shortest cavity lengths (50 to 100 μm). While clipping loss is expected to depend on the cavity length because of the length dependence on the optical beam diameter on the mirrors, we do not expect a significant change in the waist size in this region. We rather assume that physical contact between the mirror substrates led to a change in the angle between the two mirrors for the shortest cavity lengths, which accidentally improved the finesse. This hypothesis is supported by the observation that the incident beam had to be readjusted in this length region much more than usual to match the cavity mode when the cavity length was changed. For intermediate cavity lengths (100 to 500 μm), the angle between the mirrors could not be re-adjusted in the measurement setup, which

Sample	L_{opt} (μm)	$\mathcal{F}_{\mathrm{max}}$ (10^3)	$(T+P)_{\mathrm{min}}$ (ppm)	T (ppm)	P_{min} (ppm)
S6-7 top	66	122	52	35	17
S6-7 right	79	144	44	35	9
S6-7 center	178	128	49	35	14
S6-9	72	126	50	35	15
L9	150	128	49	44	5
L14	189	108	58	44	14

Table 3.4: Minimum loss P_{min} observed with various coupling mirrors in a Fabry-Perot cavity of length L_{opt} with the quartz witness sample QT2.

suggests to explain the observed plateau $P \approx 45$ ppm by increased scattering loss due to the fact that the coupling mirror was normal to the flat mirror on a suboptimal spot. The final tendency of rapidly increasing loss for $L > 500\,\mu$m can be explained by modeling the coupling mirror as a perfect spherical mirror covered by a $82\,\mu$m aperture clipping larger and larger fractions of the widening optical beam as the cavity length increases (green trace in Fig. 3.19).

For the blue datapoints in Fig. 3.19, the same coupling mirror was used and the witness sample exchanged for a micropillar with dielectric mirror coating (sample 20121009C4). Due to the small size of the micropillar mirror, at each length the finesse was optimized by tuning the lateral position of the coupler due to a slight angle between the cavity axis and the axis whose length is tuned. One immediately notes an increase in the loss plateau for short cavities by about 20 ppm with respect to the previous section, i.e. a total loss of $P \approx 70$ ppm. This increase can partially be explained by the micropillar sample lacking the bakeout step, which is known to lead to 15 - 20 ppm of excess loss. Indeed, another sample (20121009C5) was found to enable cavities with total loss as low as 30 ppm after a bakeout treatment. Apart from this loss plateau, we observe the loss to slightly increase to 120 ppm for 300 - 400 μm length. By using a clipping loss model with identical parameters as for the witness measurement, we can explain this excess loss by adding a second clipping aperture of $56\,\mu$m at the micropillar surface. We therefore conclude that despite a nominal mirror diameter of 100 μm defined by the photolithography mask, barely more than half this diameter contributes to the useful mirror surface. This fact is expectable from the optical microscope image of the micropillar mirror (Fig. 3.6), which suggests that at least one third of the mirror diameter is unusable due to border effects from the mirror coating process.

To conclude, we note that our best micropillar mirror causes at least 10 ppm of excess loss compared to a flat mirror for cavity lengths at which clipping loss is negligible. The useful coating diameter is below 60 μm and restricts the cavity length to below 200 μm if no excessive clipping loss is to be tolerated. With these values ($T + P = 60$ ppm, $L = 200\,\mu$m), the cavity linewidth $\kappa/(2\pi) \approx 7$ MHz is about twice the mechanical frequency of the micropillar. This excludes the resolved sideband regime of optomechanical coupling if a high quantum efficiency readout, i.e. lowest-possible optical loss, is desired.

3.2.4 Loss due to higher-order transverse mode degeneracy

While the simple clipping loss model of section 3.2.3 explains the general tendency towards higher loss in Fig. 3.19, some observations are in stark contrast with the model:

1. The red datapoint at length $680\,\mu$m in Fig. 3.19 has lower loss than the data for shorter cavities, while the clipping loss model predicts a monotonous increase of loss with length in this case.

2. The same clipping loss parameters for the coupling mirror are unable to simultaneously explain the two measurements in Fig. 3.19.

3. Excess loss restricted to specific cavity lengths has been observed.

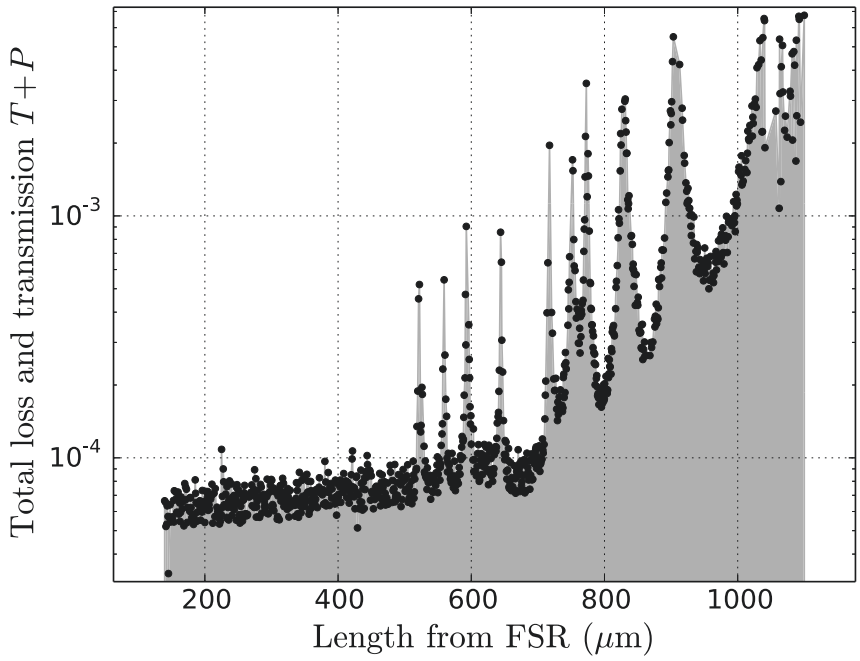


Figure 3.20: Measurement of $T + P$ for the fundamental TEM_{00} mode of a Fabry-Perot cavity versus cavity length with at least one datapoint per free spectral range. Apart from a loss plateau around 60 ppm for lengths below 400 μm and a clipping-related loss increase for higher lengths, many loss peaks are observed for higher cavity lengths. These peaks are related to degeneracies of the fundamental and higher-order cavity modes.

4. For specific lengths, an extra resonance peak near the original TEM_{00} resonance is observed with similar properties (polarization, spatial profile). This peak cannot be removed by aligning the incident beam or polarization without losing alignment to the original resonance.

In order to understand these effects, we have performed a finesse measurement of every accessible resonance of a cavity constructed with samples S6-4 and QT2 in a dedicated setup in the context of the Master project of Remi Metzdorff. The setup is similar to the one presented above, with the differences that the angle between two cavity mirrors can be aligned and that the cavity length can be tuned in an automated way with a stepper motor. One measurement result is presented in Fig. 3.20. One observes the same tendency towards higher loss for longer cavities as before, but additionally a number of loss peaks of variable height and width. We interpret these peaks as a result of the degeneracy of the TEM_{00} and a higher-order TEM_{lm} cavity mode at specific cavity lengths [BHM+15], which provides a resonant loss channel for the fundamental cavity mode since higher-order modes often experience higher clipping loss due to their larger lateral size (Eq. 2.2.3). These degeneracies occur when twice the excess Gouy phase of the higher-order beam (see Eq. 2.2.3) after progagation over the cavity length equals an integer multiple p of 2π:

$$L = R_\mathrm{c} \sin^2\left(\frac{p}{l+m}\pi\right) . \tag{3.2.8}$$

A few predicted degeneracies are marked as dots in Fig. 3.13. While the degeneracies in our measurement roughly occur at these predicted lengths, the astigmatism of up to 20 % of our mirrors leads to a different degeneracy length for pairs of different transverse modes with $l_1 + m_1 = l_2 + m_2$, therefore many more peaks than expected from a simple model. Furthermore, since the real mirror shape is not spherical but rather Gaussian, the effective radius of curvature depends on the spot size on the mirror, with an increase of the radius of curvature for longer cavities. Thus, the exact degeneracy length depends on the detailed shape of the mirror and the mode number. Because of these complications, we have resorted to a numerical simulation of the cavity loss as a function of cavity length with the measured mirror depth profile with the program OSCAR [Deg08, BDFF14]. We have obtained good agreement with the measurement and the simulated mode profiles confirm the hypothesis of transverse mode degeneracy [MND+17].

Unfortunately, the degeneracy test setup does not permit to study the degeneracy effect for a cavity involving a micropillar, as a lateral alignment degree of freedom is currently not implemented. Nevertheless, the observed presence of loss peaks and the dependence of the effective radius of curvature on the cavity length can also explain the inability of a simple clipping loss model to explain the coarsely sampled data in Fig. 3.19 in detail. While we have observed no pronounced loss peaks at the short cavity lengths where we intend to operate our optomechanical system, awareness of the existence of loss peaks leads us to perform small cavity length scans before assembling an optomechanical cavity in order to potentially avoid a cavity length where the TEM_{00} and a higher-order mode are degenerate and excess loss is to be expected.

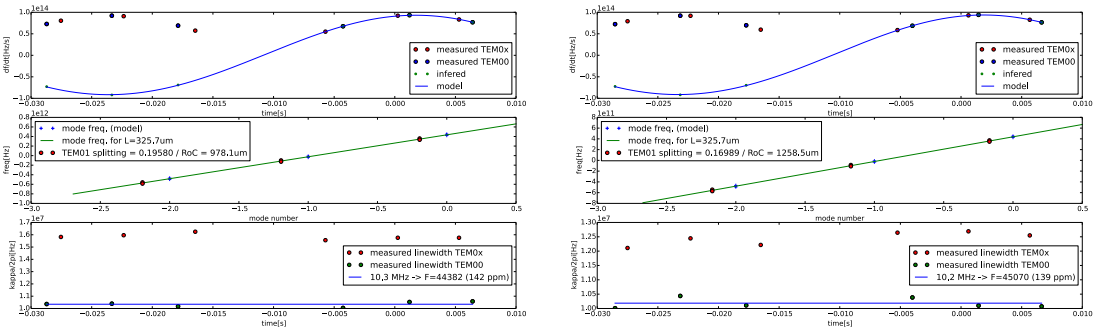


Figure 3.21: Measurement of the radius of curvature with the program "finesse.py" in a cavity constructed with pillar 20121009C4 and coupling mirror S6-9. First, the incident beam is misaligned in order to obtain a reflection trace with a similar resonance depth for the TEM_{00} and TEM_{01} (left) or TEM_{10} (right) modes. The program automatically labels the different modes and estimates the excess Gouy phase from the frequency distance of the two mode families. From the Gouy phase and cavity length, the effective radius of curvate in the corresponding vertical or horizontal direction can be extracted. Both bottom plots show that the first higher-order modes already experience about 50 % more optical loss than the fundamental mode.

3.2.5 Measurement of radius of curvature

When the incident quasi-TEM_{00} mode is perfectly aligned to the cavity axis, its overlap integral with higher-order transverse modes of the cavity is nearly zero and these resonances therefore do not appear in the reflection spectrum. If the incident beam is laterally displaced, these higher-order modes start to appear in the spectrum while the fundamental resonance dip decreases in depth due to its reduced overlap integral. By aligning the incident beam to a configuration where the TEM_{00} and the TEM_{01}[5] resonance dips have similar depth, their resonance frequency difference ν_{01} modulo the free spectral range ν_{fsr} can be calibrated in a similar way as the free spectral range as described in section 3.2.3. This procedure was implemented in the program "finesse" to allow for a rapid and reproducible measurement of this kind and is shown in Fig. 3.21. Similarly to Eq. 3.2.8, the radius of curvature R_{C} of the mirror along the x-(y-)direction can be deduced from the Gouy phase $\psi(L)$ of the higher-order beam, which is extracted from the ratio $\frac{\nu_{10}}{\nu_{\mathrm{fsr}}} = \frac{\psi(L)}{\pi}$ (equivalent expression for ν_{01}) with the relations from Eqs. 2.2.2 and 2.2.3:

$$R_{\mathrm{C}} = L \sin^{-2}\left(\pi \frac{\nu_{10}}{\nu_{\mathrm{fsr}}}\right) . \tag{3.2.9}$$

For the pairing of micropillar 20121009C4 and coupling mirror S6-9, Fig. 3.21 shows that we find radii of curvature of 980 μm and 1260 μm in the two perpendicular directions at a cavity length of 325 μm. The observed radius of curvature has a weak dependence on the cavity length (typical variation of 10% over the full range of accessible lengths). This dependence can be explained by the Gaussian profile of

[5] Even without a camera monitoring the cavity transmission, the different modes are easily identified up to the third transverse order by scanning the incident beam position with the closest lens over the coupling mirror surface and counting the maxima of the resonance depth for a particular resonance.

the mirror which leads to larger effective radius of curvature for larger spot sizes. We therefore also expect larger radius of curvature for subsequently higher-order transverse cavity modes such as the TEM_{02}, although we have not investigated this effect in detail. From this argumentation, it becomes clear that since the spot size of the TEM_{10}-mode is different and generally larger than the fundamental mode shape, the obtained radius of curvature is only an approximation of the effective radius of curvature of the fundamental mode.

The observed astigmatism is typical of our samples and can be attributed to an astigmatism in the mode shape of the CO_2 laser used for photo-ablation. The astigmatism of the mirror is expected to result in a fundamental cavity mode shape with slightly different waist sizes along the two transverse directions. To perfectly couple an incident laser beam with Gaussian shape to such astigmatic cavity modes, one should best use cylindrical lenses to adapt the incident mode shape to the cavity. In light of the increasingly complex alignment procedure with a cylindrical lens, we first estimate the effect of the observed astigmatism on the coupling efficiency by computing the overlap integral of two Gaussian beams with different waists. From Eq. 2.2.2, we find[6] for the ratio η of the power of a beam with Rayleigh length z_{R1} contained in a concentric beam with Rayleigh length z_{R2} and waist at a distance z_0 from the waist of the first beam the useful expression

$$\eta = \frac{4 z_{R1} z_{R2}}{(z_{R1} + z_{R2})^2 + z_0^2} \,. \tag{3.2.10}$$

In the limit of short cavities, we find $z_R^2 \propto R_C$ and compute for the coupling ratio between two modes adapted to the two above observed radii of curvature $\eta \approx 99.6\%$. This value is only a lower bound of the coupling ratio since our estimation assumed that a beam adapted to the larger radius of curvature is coupled to a cavity having the lower radius of curvature in both directions, while this is only true in one direction. We therefore deem the mirror astigmatism totally negligible for the coupling efficiency, while it has the welcome effect of introducing a birefringence to the cavity (see section 3.2.5).

Cavity birefringence

Birefringence occurs when light of different polarization experiences different effective path lengths when crossing a birefringent medium. Birefringence of the medium inside a Fabry-Perot cavity manifests itself as a non-degeneracy of optical modes of different polarizations. We routinely observe this effect with our measurement Fabry-Perot cavities as shown in Fig. 3.22. The only non-vacuum medium interacting with the cavity field are the two mirror coatings. Mirror birefringence is most naturally expressed as a phase-shift $\Delta\phi$ between two beams of orthogonal polarizations after reflection by the mirror. This phase-shift can be connected to the experimentally more relevant frequency splitting $\Delta\nu$ between the cavity eigenmodes through $\Delta\nu = \nu_{fsr} \frac{\Delta\phi}{2\pi}$. While the amorphous mirror coating materials are not intrinsically birefringent, mechanical stress causes different refractive indices of the material along and orthogonal to

[6] The calculation becomes much simpler if expressions involving the complex beam parameter $q(z) = z + i z_R$ are used instead of the real-valued mode shape from Eq. 2.2.2.

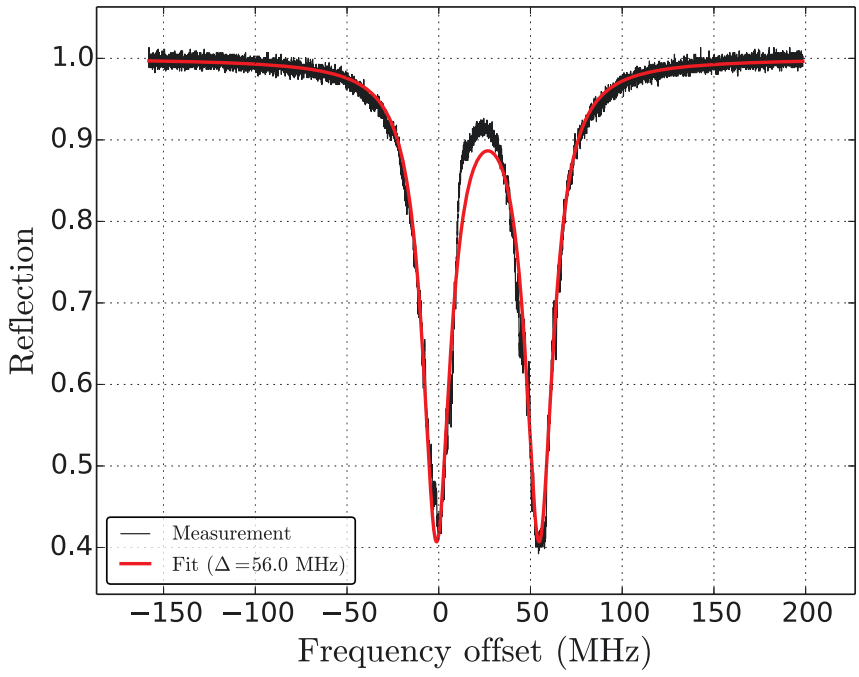


Figure 3.22: Cavity reflection for incident light with linear polarization at an angle of 45° with respect to the eigenmode polarization. Two well-separated optical resonances are visible. If the polarization of the incident light is rotated, either the right or the left peak disappears, depending on the direction of rotation. The frequency separation of both peaks is calibrated with phase-modulation sidebands that are not visible in the plot.

Flat mirror	Concave mirror	T	Date	R_1	R_2	L	$\Delta\nu_{\mathrm{m}}$	$\Delta\nu_{\mathrm{th}}$
QT2	L14	300	04/9/14	267	319	78	34.1	31.5
QT2	S6-4	300	15/4/16	1700	1260	204	8.2	4.1
20140113C5	S6-9	300	16/9/14	1256	991	120	44.0	7.2
20140113C5	S6-9	0.1	18/9/14	1256	991	102	56.0	8.4
20121009C5	S6-9	0.1	01/3/16	1256	991	71	221.0	12.1
20121009C5	S6-9	0.1	25/6/16	1256	991	46	436.0	18.7

Table 3.5: Birefringence splitting $\Delta\nu$ observed for various Fabry-Perot cavities. Radii of curvature R_1 and R_2 and length L are given in microns, measured and theoretical frequency splittings $\Delta\nu_{\mathrm{m}}$ and $\Delta\nu_{\mathrm{th}}$ in MHz, and temperatures T in K.

the direction of the stress. Furthermore, birefringence may be caused by rotationally asymmetric mirror shapes. For an astigmatic mirror with the radii of curvature R_1 and R_2, reference [UBRR15] both predicts theoretically and confirms experimentally the cavity eigenmode splitting frequency

$$\Delta\nu = \frac{\nu_{\mathrm{fsr}}\lambda}{4\pi^2}\frac{R_1 - R_2}{R_1 R_2}\,. \tag{3.2.11}$$

In order to elucidate the cause of birefringence in our cavities, a compilation of observed splittings is listed in Table 3.5. To compare the measured values with the theoretical expression, the two radii of curvature are measured with the procedure described in section 3.2.5.

With the flat witness sample QT2, we observe frequency splittings in close proximity to the predictions from Eq. 3.2.11. Since we do not expect any significant birefringence from the flat witness, we conclude that stress birefringence is negligible in the coupling mirrors, in agreement with the measuremennts in [UBRR15]. With micropillar samples instead (all data from the third row on), Eq. 3.2.11 can only explain small fractions $(5 - 15\%)$ of the observed splittings. While aligning a cavity that contains a micropillar sample, the incident laser polarization must be frequently adjusted, implying that the birefringence of the micropillar mirror also depends on the position of the beam on the sample. The observed frequency splitting furthermore depends on temperature, even if its dependence on the cavity-length dependent free spectral range is taken into account. All evidence suggests that the mirror on top of the micropillar is significantly birefringent.

While the micropillar mirror is mostly flat (see Fig. 3.7), the steep height step at its edge might lead to a geometric contribution to the birefringence. At the same time, thermal stress is expected in the coating since the mirror is coated at about 200 °C and the thermal expansion coefficient of quartz is different from the one of amorphous silica and tantala. The observed weak temperature dependence of the birefringence may be explained by both effects, since we expect further thermal stress to build up during cool-down, but cannot exclude the possibility of small misalignments of the cavity during cool-down either, which would cause the intracavity field to probe a different region of the micropillar mirror with potentially different geometric or stress-induced birefringence.

We conclude that we cannot precisely predict the birefringence splitting of a new

cavity, but must experimentally characterize this parameter for each newly aligned cavity. If the alignment and temperature of a cavity are not changed, the birefringence splitting remains remarkably constant. The observed birefringence is of the order of $(10^{-5} - 10^{-3}) \times \nu_{\mathrm{fsr}}$. For cavities with $\mathcal{F} \approx 10^5$ in the 100-μm length range, this results in two well-resolved polarization eigenmodes with typical splittings of the order of 100 MHz. This splitting is very useful in order to inject two independent orthogonally polarized laser beams into the cavity, which can be derived from the same laser source with an acousto-optic modulator.

Chapter 4

Experiment - Optomechanical cavity inside a cryostat

Door meten tot weten

— Kamerlingh Onnes

This chapter deals with the immediate environment of the samples. The first section discusses the dilution refrigerator that is employed to reduce the temperature of this environment to temperatures in the 50-mK range. The second section discusses the mounting structure of the samples, i.e. the mechanical, electric and gas environment of the samples and concludes with a characterization of the obtained performance.

4.1 Cryogenics

Thermal force noise is the main obstacle in attempts to observe quantum effects with mechanical systems in the MHz frequency range. Equation 2.1.22 shows that the thermal force noise spectrum is proportional to the temperature of the sample environment. Quantum optical experiments therefore benefit linearly from a temperature reduction of the sample with classical cryogenic techniques. With optical powers in the 100 μW to 1 mW range absorbed by our measurement Fabry-Perot cavity, commercial ^{3}He or dilution refrigerators are the most straightforward solution to reach temperatures in the 100 mK range. To go to even lower temperatures, techniques such as adiabatic demagnetization refrigeration with paramagnetic salts or nuclear demagnetization are available, but unrealistic to be efficient in light of the high optical power absorbed by the sample. After a quick reminder of the fundamentals of the dilution refrigeration technique, we will discuss the cryostat used for our experiments, along with some technical details necessary to successfully carry out the experiment.

4.1.1 Dilution refrigeration

The transformation from liquid to gaseous helium is a first-order phase transition with an associated latent heat L_{evap}. This latent heat makes evaporation of liquid helium one of the simplest methods to achieve low temperatures. The cooling power $\dot{Q}$

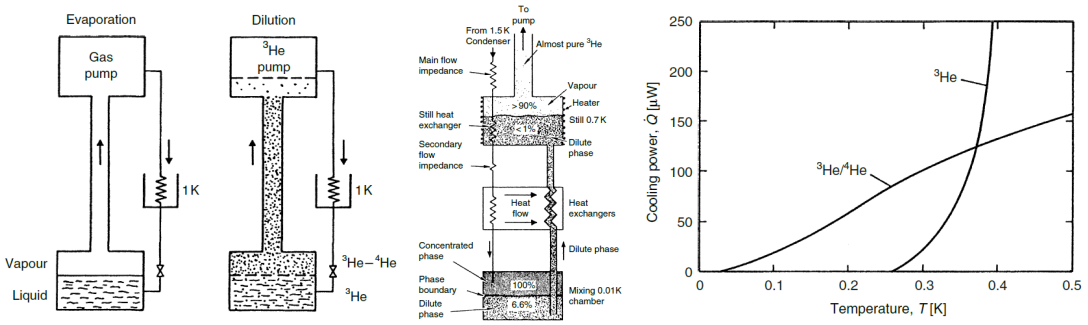


Figure 4.1: Left: Simplified schematics of a ^{3}He and a dilution refrigerator. In reality, the concentrated ^{3}He phase is lighter and therefore on top of the dilute phase. Middle: More detailed illustration of a dilution refrigerator, showing how a still solves the challenge of pumping on the heavier diluted phase and how good heat exchanger design can lead to reinjection of ^{3}He into the mixture chamber at nearly the mixing chamber temperature. Right: Cooling powers of a ^{3}He and dilution refrigerator if the same helium pump is used. Figures are taken from [Pob07] (left and middle) and [Lou74] (right).

available when $\dot{n}$ moles of helium per second evaporate is therefore $\dot{Q} = \dot{n}L_{\text{evap}}$. When pumping on liquid helium, the evaporation rate is proportional to the vapor pressure $P_{\text{vap}}(T) = P_0 \exp\left(-\frac{L_{\text{evap}}}{RT}\right)$. We therefore find that the cooling power of an evaporation refrigerator scales with the temperature as $e^{-1/T}$. While using ^{4}He as refrigerant enables reaching temperatures down to 1.3 K before the vapor pressure drops below a reasonable pumping pressure, using a ^{3}He-refrigerator makes temperatures down to 0.3 K accessible. Since no substance with a lower boiling point than the one of ^{3}He is known, methods different from evaporation must be used to reach even lower temperatures.

In the 1960s, it was realized that, contrary to classical expectations from the third law of thermodynamics, liquid ^{4}He can dissolve finite amounts of ^{3}He even at zero temperature due to the quantum properties of the two isotopes. A mixture of both isotopes forms a dilute phase[1] consisting of all available ^{4}He with up to 6.6 % of ^{3}He. If any excess ^{3}He is available in the mixture, a second phase of concentrated ^{3}He forms and remains on top of the dilute phase due to its lower density. For each ^{3}He atom that we artificially remove from the dilute phase, another such atom will cross the phase boundary from the concentrated phase in order to re-establish the equilibrium concentration. This phase transition is accompanied by a latent heat of mixing $L_{\text{c}\to\text{d}}$ due to the higher binding energy of ^{3}He in the dilute phase, which is a consequence of a lower inter-atomic distance due to the lower zero-point motion of ^{4}He compared to ^{3}He. As in the case of evaporative cooling, a molar flow $\dot{n}_3$ of ^{3}He across the phase separation boundary brings about a cooling power $\dot{Q} = \dot{n}_3 L_{\text{c}\to\text{d}}$. The flow of ^{3}He is driven by pumping on the surface of the dilute phase, where more ^{3}He than ^{4}He evaporates at a given temperature due to the higher vapor pressure of ^{3}He. A specially designed still is necessary for this as the dilute phase surface would otherwise be covered with the concentrated phase (see Fig. 4.1). Since the cooling occurs at the

[1]The terms "dilute" and "concentrated" are used with respect to ^{3}He.

phase separation boundary in the so-called mixing chamber at T_{mc} and the evaporation at the spatially separated still surface at T_{s}, the rate $\dot{n}_3$ can be much higher than the vapor pressure associated with the mixing chamber temperature if mixing chamber and still are thermally well-isolated from each other. This trick removes the temperature-dependence of $\dot{n}_3$ that limits the temperature range accessible with evaporative cooling and has enabled dilution refrigeration to temperatures as low as 2 mK. In continuous operation, the ^{3}He that has been pumped from the still is re-liquefied and injected into the concentrated phase. In the limit where the re-injected helium is at the mixing chamber temperature, the cooling power $\dot{Q}$ in presence of a heat leak P_{leak} can be approximated at $T < 50\,\mathrm{mK}$ to be

$$\dot{Q} + P_{\mathrm{leak}} = \dot{n}_3 T_{\mathrm{mc}}^2 \frac{84\,\mathrm{W.s}}{\mathrm{mol.K}^2} = 210\,\mu\mathrm{W} \frac{\dot{n}_3}{250\,\mu\mathrm{mol}/s} \left(\frac{T_{\mathrm{mc}}}{100\,\mathrm{mK}}\right)^2 . \tag{4.1.1}$$

The quadratic temperature dependence of the cooling power with dilution refrigeration arises from integrating the linear specific heat of ^{3}He over temperature to obtain the temperature-dependend latent heat of mixing. This dependence, compared to the exponential evaporative cooling power dependence on temperature, causes that any pair of reasonably-designed dilution and ^{3}He refrigerators reach the same temperature for a high enough threshold cooling power $\dot{Q}_{\mathrm{th}}$ (see Fig. 4.1). While dilution refrigeration enables lower temperatures than ^{3}He evaporation for heating powers $P < \dot{Q}_{\mathrm{th}}$, the situation is reversed when higher powers are requested. This makes the choice of preferred cryostat unambiguous if the required cooling power is known.

The presented description of refrigeration techniques can be found in much more detail in the references [Pob07, Lou74].

4.1.2 Horizontal dilution refrigerator from CryoConcept

The dilution refrigerator model HDR-JT-S200-30 from the company CryoConcept from Courtabœuf, France, has been purchased by our research group before the beginning of this thesis. Its design has been optimized to suit the needs of our optomechanical system and is described in detail in another thesis from our group [Kuh13]. While detailed instructions for the operation of the cryostat can be found in Appendix E, we only summarize the most important cryostat parameters here:

- The dilution refrigerator has a specified **cooling power** of 200 μW at 100 mK. The nominal molar flow rate is $\dot{n}_3 = 300\,\mu\mathrm{mol/s}$ (Eq. 4.1.1). The specified base temperature is 30 mK. With optimized operation settings, powers up to 10 mW can be dissipated on the mixing chamber plate for short amounts of time ($< 1\,\mathrm{min}$) without interrupting the dilution operation.

- Besides the normal still, a second "**double still**" is available which functions in two configurations: during cooldown, it is used in the inverse direction to increase the throughput of mixture and thereby the thermalization rate of the dilution circuit due to the double still's lower gas flow impedance. In dilution operation, it is used in the normal direction in parallel with the actual still to increase the evaporation rate of ^{3}He (see dilution circuit schematic in Fig. 4.2).

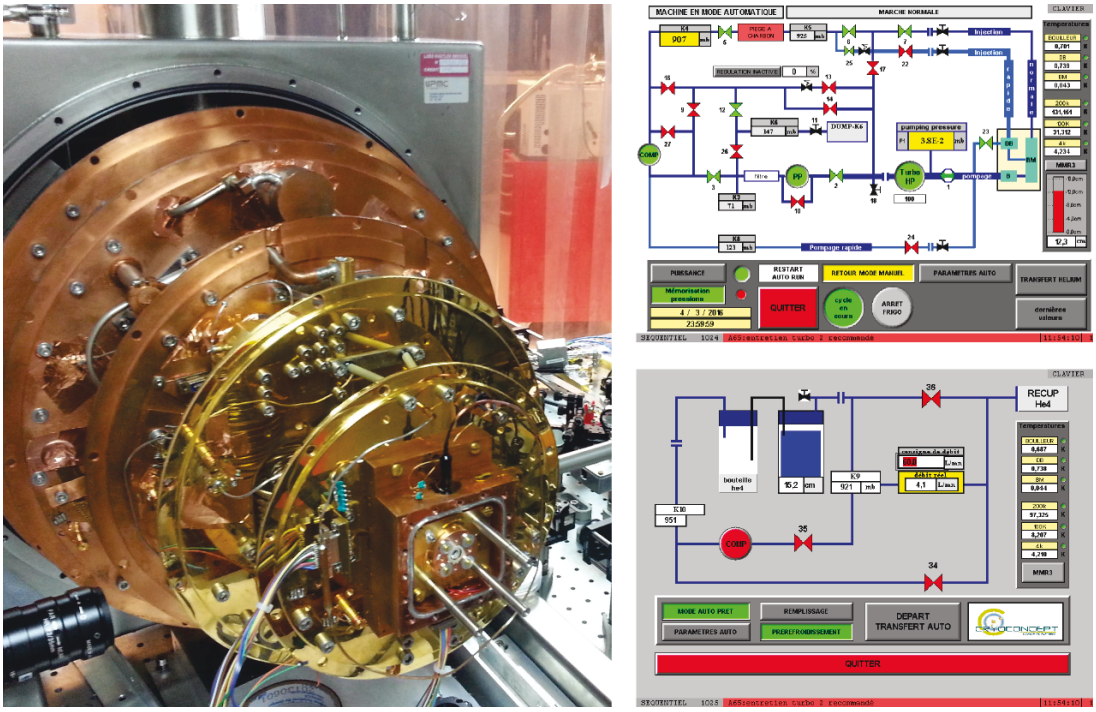


Figure 4.2: Left: Photograph of the dilution refrigerator with all heat shields removed. Mounted on the square steel frame in the back, one can see the various plates that are named after their nominal temperature: "200 K", "100 K", "4 K", "1 K" and "BM" (mixing chamber). At the center of the mixing chamber plate, the measurement Fabry-Perot cavity is mounted inside a separate copper vacuum chamber (IVC). Right: Screenshot of the touchscreen interface of the cryostat controller. The upper screen shows the controls of the ^{3}He - ^{4}He mixture circuit, the lower screen the controls of the ^{4}He circuit.

- In order to fit the cryostat on an optical table, the various temperature stages are arranged **horizontally** instead of vertically as for most dilution refrigerators (see Fig. 4.2). Consequently, the dilution unit is not lowered into the heat shield, but instead five heat shields and the outer vacuum chamber (OVC) are mounted individually when the cryostat is closed.

- The cryostat is pre-cooled to 4 K by circulating ^{4}He gas, and possesses a **^{4}He reservoir** capable of holding about 10 l of liquid helium, which allows operation of the cryostat without the vibrations introduced by the ^{4}He filling tube.

- With a minimum consumption of about 0.35 l/h of liquid ^{4}He in dilution operation, the cryostat has an **autonomy** of 25 h once the reservoir is completely filled.

- Closing the cryostat takes about two hours, pumping down the OVC to reach a vacuum below 10^{-5} mbar takes one day, and a **cool-down** to dilution operation from room temperature takes at least 1.5 working days and 100 l of ^{4}He. Heating up the cryostat takes a similar timescale if gas is circulated in the cryostat. Without circulation, a heat-up cycle takes up to 5 days.

- Two windows in the OVC, three more windows in some heat shields and holes in the remaining heat shields and cryostat plates allow a **free-space optical beam** to be injected along the axis of the cryostat from the room temperature environment to a sample mounted on the mixing chamber plate, and its transmission to be monitored through the second OVC window. The transmission aperture in the 4 K plate can be closed with a mechanical feedthrough during refrigerator operation to reduce the radiated heat load on the mixing chamber. All windows have been replaced by anti-reflection coated windows such that a 1064 nm laser beam only experiences about 5 % of optical loss per trip through the cryostat. Windows are made from SF57 silica to keep stress birefringence changes during thermal cycles low. Since the optical holes have been kept small to reduce the thermal radiation heat leak on the mixing chamber plate, clipping of the optical beam is the main concern in achieving a high transmission efficiency.

- The cryostat can be operated in **continuous dilution operation** for about 10 consecutive days, after which we observe a degradation of the cavity mode alignment with respect to the incident beam, which cannot be recovered without inducing excessive clipping loss (95 %). We believe that ice originating from vacuum leaks in the OVC forms on the cold cryostat surfaces and induces either a deformation of the mechanical structure, or a roughness increase of the optical surfaces due to the formation of crystals. The alignment degradation can usually be predicted one to two days in advance by a significant increase in cracking sounds originating from the cryostat during the refilling of the ^{4}He reservoir. By heating the cryostat to 90 K and pumping on the OVC for a few hours, original alignment can be recovered (see Fig. 4.3).

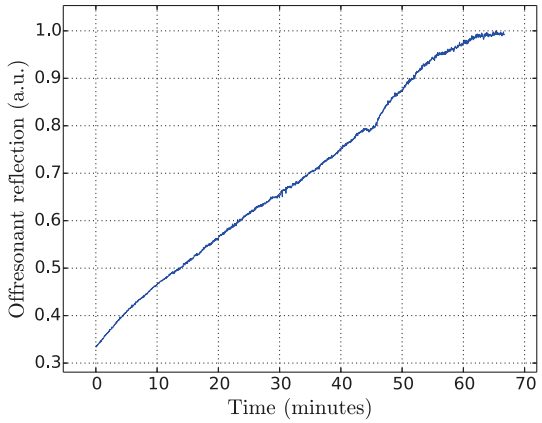


Figure 4.3: Alignment recovery of the measurement cavity during a thermal cycle of the cryostat around 90 K. The plot shows the off-resonant reflection from the cavity while the incident beam remains optimally aligned to the fundamental cavity mode and the incident power is kept constant. When the refrigerator is in continuous dilution operation, the cavity alignment abruptly degrades every 9-10 days and can be recovered from about 5% to near 100% of the level of room temperature alignment by heating up the cryostat to about 90 K and pumping on the OVC with a turbo pump for a few hours.

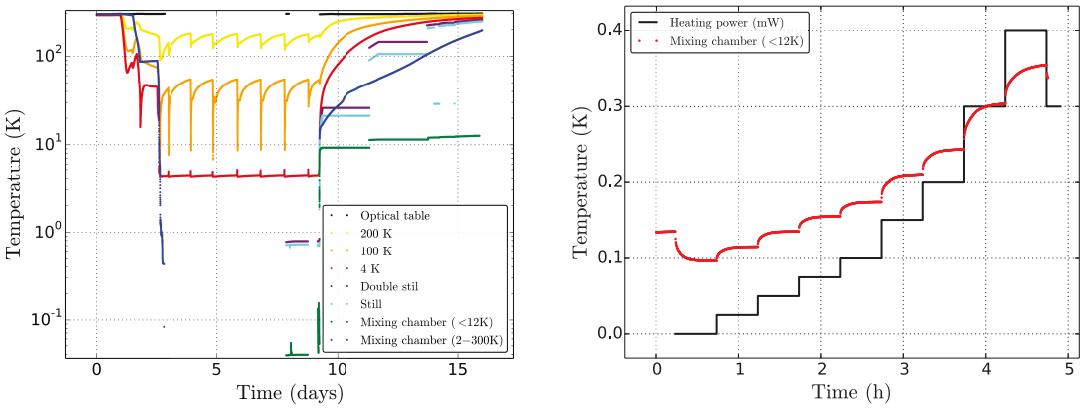


Figure 4.4: Left: Temperatures of the various plates during a full cryostat run. During the experiments, the thermometers on the coldest parts were turned off in order to decrease the electronic noise induced by the temperature measurement. The refilling of the ^{4}He reservoir is clearly visible as temperature oscillations with a period of 24h. Right: Mixing chamber temperature T_{BM} versus power P dissipated by the heating resistor. A model with $T_{\mathrm{BM}} \propto P^{2/3}$ reasonably fits the asymptotic temperatures. We explain the discrepancy from the theoretical $P^{1/2}$-dependence (Eq. 4.1.1) with the fact that the equation is only an approximation above 50 mK, the fact that the still temperature was not stabilized by a heater during the measurement and with a temperature-dependent heat leak induced by a CuNi capillary that was installed to fill the IVC. Furthermore, electrical pickup led to a wrong base temperature reading of the thermometer during this measurement. With proper shielding of the electrical cables leading to the cryostat, the observed base temperature of 90 mK drops to 40 mK without heating the still.

4.1.3 Temperature measurement

The cryostat has been equipped with six thermometers by the manufacturer that are installed at the following cryostat parts: 200 K plate, 100 K plate, 4 K plate, still, double still, mixing chamber. All but the last thermometers yield reasonable readings over the full range of accessible temperatures. The mixing chamber thermometer made from RuO_2 saturates at 12 K and therefore prohibits knowledge of the sample temperature during cool-down or heat-up. In order to extend the available range of temperature readings of the mixing chamber plate, we have installed an additional carbon ceramic resistance thermometer[2] whose resistance-versus-temperature relation has been calibrated by the manufacturer from 1.5 to 375 K. A Keithley 2000 6.5-digit multimeter is used to perform four-point measurements of the resistance of this and other sensors (optical table temperature, liquid nitrogen level, IVC fill tube pressure). All thermometers are interfaced with the experiment computer and their values are logged in 20 s intervals during measurements and thermal cycles, unless the resistance measurement electrically perturbs the experiment. Typical records are shown in Fig. 4.4.

Furthermore, the manufacturer has equipped the still with a heating resistor. By feeding an appropriate current to this resistor, the evaporation rate of ^{3}He can be increased by heating the still to an optimal temperature of about 850 mK. However, for the measurements reported in this thesis, the resulting improvement was estimated negligible compared to the added complexity and risk of introducing electrical noise, so that the heating resistor was left disconnected. Nevertheless, this option should be considered for future, optimized versions of the experiment.

Another 1 kΩ heating resistor has been installed on the mixing chamber plate. By passing a known current through this resistor and measuring the resulting voltage with two separate leads, the dissipated electrical power can be precisely estimated. This power can be used to characterize the temperature-versus-cooling power relation of the cryostat, which is shown in Fig. 4.4. More importantly, by finding the electrical power P_{el} leading to the same asymptotic mixing chamber temperature as an optical beam of known power P_{in}, the absorption coefficient of the measurement Fabry-Perot cavity can be estimated as $P_{\mathrm{el}}/P_{\mathrm{in}}$. We typically find an off-resonant absorption at the level of 3 % and a resonant absorption within 5 % of the theoretically expected level of $1 - R_0$ for a one-sided Fabry-Perot cavity (see Eq. 2.2.50).

An almost immediate change of the reading of the RuO_2 thermometer has been observed when a cable is connected or disconnected to the external electrical feedthrough connectors of the cryostat, or when the transmission aperture in the 4 K plate is opened. The timescale of this apparent temperature change (faster than 5 s) is very different from the one of temperature changes induced by the heating resistor on the mixing chamber plate (several minutes, see Fig. 4.4). We therefore do not believe that the thermometer reading reflects the mixing chamber temperature in this scenario, but instead suppose that the RuO_2 thermometer is prone to pickup of electrical parasites which either heat the thermometer through its cable leads above the mixing chamber temperature, or otherwise electrically perturb the temperature measurement. During experiments, we usually search the electrical configuration leading to the lowest read-

[2] Calibrated CCS Type A2 N653 from ICE Oxford, Oxon, United Kingdom.

ing of the mixing chamber temperature, which we assume to be an upper bound of the real temperature. This configuration corresponds to a closed ^{4}He plate transmission aperture, with as few cables connected to the cryostat as possible and unpopulated connectors shorted to the cryostat ground.

4.1.4 Electrical connections and noise

Due to the way the cryostat is constructed, all metallic cryostat parts are electrically connected to the optical table ground and to the ground potential of the pump stand. To minimize the negative effect of cryostat ground potential fluctuations, efforts are made to realize all electrical devices such as piezos and thermometers as electrically floating with respect to the cryostat ground. The cryostat has been delivered with 12 twisted pairs of thin manganin wire connecting two 12-pin Jaeger connectors outside the cryostat with a 25-pin micro Sub-D connector on the mixture chamber plate. Each pair of wires has a series resistance of $2 \times 25\,\Omega$ with less than 10 % change during cooldown and is rated for a maximum voltage of 100 V. We have removed an earlier installed custom cable capable of delivering 1000 V to the mixing chamber due to heat leak concerns. Out of the 24 wires, we use

- 4 for a four-point measurement of the heating resistor on the mixing chamber,
- 4 for the custom carbon ceramic thermometer on the mixing chamber,
- 2 for a quartz tuning fork inside the IVC used as a pressure sensor,
- 6 for three coarse piezos (max. voltage 100 V) to tune the length of the measurement cavity,
- 2 for a fast piezo to improve the locking performance of the measurement cavity,
- 2 for an antenna near the quartz micropillar for direct excitation of mechanical displacement through the piezoelectric effect.

The details of the wiring are shown in Appendix D. The most important cables are the ones leading to the actuators of the measurement cavity on the mixing chamber plate, since these actuators are most susceptible to transduce electrical noise into displacement noise of the cavity. The full circuitry leading to the actuators is shown in Fig. 4.5. The introduction of low-pass filters formed by cryogenic capacitors and room-temperature resistors near the cryostat case successfully reduces the displacement noise floor of the cavity. For the capacitors to work at low temperature, chemical capacitors should be strictly avoided. By testing a number of capacitors in liquid nitrogen, it was found that the capacity of larger capacitors with higher voltage rating changes less when cooled to low temperatures. The capacity of the used capacitors is reduced by about 10 % when the mixing chamber plate reaches mK temperatures.

Initial measurements of the Brownian motion of the micropillar compression mode led to estimated environment temperatures up to 600 K for a mixing chamber temperature below 100 mK, while room-temperature measurements resulted in estimates near 300 K. This paradox can be explained by considering the proportionality of the thermal force noise spectral density in Eq. 2.1.22 both on the absolute temperature

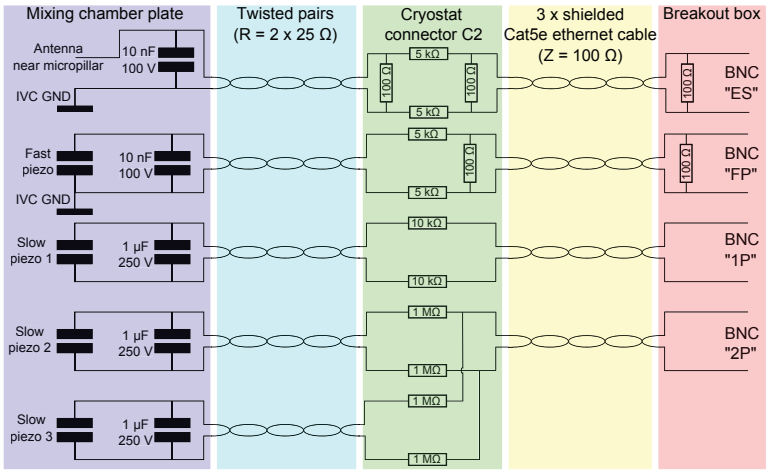


Figure 4.5: Electrical connections of the cryogenic actuators from the mixing chamber plate to the breakout box in the electronics rack. For all connections, low-pass filters and attenuators are formed as close to the actuators as possible by the capacitors on the mixing chamber plate and resistors inside the Jaeger connector screwed to the cryostat case. For high-frequency actuators, impedance matching is provided by additional resistors in the breakout box in the rack. The outer shields of the ethernet cable are connected to the cryostat ground.

as well as on the inverse mechanical quality factor: The thermal force noise spectral density changed from 300 K, where $Q \approx 3 \times 10^6$ was observed, to 100 mK, where $Q \approx 60 \times 10^6$ was measured, by a factor of 6×10^4. Thus, a temperature-independent electrical force at only 10 % of the room-temperature thermal force can explain the observed 6000-fold overestimation of the thermal bath at cryogenic temperatures. It was found through trial and error that proper shielding of the cables leading to the cryogenic actuators removes this problem. In a minimalistic experiment where all actuators but one slow piezo were disconnected, using the interior connectors of two BNC cables to feed the two terminals of the slow piezo while connecting the exterior conductors of both cables to the cryostat case resulted in the observation of compression mode environment temperatures of the order of the mixing chamber temperature down to 100 mK. Disconnecting the exterior conductors from the cryostat and connecting them either to the ground of the function generator driving the piezo or leaving them floating resulted in the previously observed drastic excess force noise with deduced temperatures varying with the position of the cables. The finally implemented configuration presented in Fig. 4.5 uses twisted pairs of doubly shielded ethernet cables to connect the cryostat connector to BNC connectors in a breakout box in the electronics rack. The connections from the output ports of our electronic devices to this breakout box are kept as short as possible. The ethernet cable is wrapped in a braided copper sleeve for additional shielding. All shields are connected to the cryostat housing through the Jaeger connector screwed to the cryostat feedthrough. No excess force noise was observed in this configuration near the compression mode frequency of the micropillar. However, crosstalk from electronics on the optical table to the measurement cavity length at lower frequencies ($< 200\,\mathrm{kHz}$) is still observed. We suppose that

this effect is a consequence of the ground connection between the two fast actuators in Fig. 4.5 and the optical table, which is unavoidable unless the measurement cavity support is electrically disconnected from the mixing chamber plate, which would lead to thermalization problems. We therefore plan to remove the antenna actuator in future experiments despite its usefulness and to connect the fast piezo in a fully floating configuration.

4.1.5 Inner vacuum chamber (IVC)

As discussed in chapter 7, one of the biggest experimental limitations to observing quantum effects with our optomechanical system is heating of the mechanical resonator due to laser absorption, which may lead to sample temperatures of the order of 10 K for the highest injected optical powers, while the mixing chamber plate remains well below 1 K. Finite-element simulations of heat propagation have shown that the largest thermal resistance arises from the membrane that holds the micropillar to its surrounding quartz structure. The thermal resistance of the membrane is expected to be inversely proportional to its thickness, while we have observed the highest mechanical quality factors with the thinnest fabricated membranes (see 3.1.2). One straightforward way to lower the thermal resistance without compromising the mechanical quality factor is to use heat removal channels other than conduction through the quartz membrane. While heat removal by conduction through a surrounding gas can be realized with low-pressure ^{3}He[3], another interesting solution is to use the heat removal provided by the evaporation of the superfluid ^{4}He film [RS39, Wil67] that covers the micropillar surface if a sufficient amount of ^{4}He is available for condensation. Both solutions require a vacuum chamber around the measurement cavity that can be filled with helium gas of either isotope. Since the cryostat's OVC pressure must remain below 10^{-5} mbar for proper heat isolation, an additional inner vacuum chamber (IVC) must be constructed.

Mechanical design

The schematic and a photograph of the implemented IVC can be seen in Fig. 4.6. A number of requirements must be met by the IVC, which translate straightforwardly into implementation details:

1. The IVC should not put a significant thermal resistance between the mixing chamber plate and the sample or gas inside the IVC. It is therefore made from oxygen-free high-conductivity copper (OFHC copper), with all contact surfaces gold-coated to prevent any significant Kaptiza resistance from oxidized copper.

2. Up to 12 electrical wires must arrive inside the IVC at the measurement cavity actuators. We have therefore implemented a 15-pin micro-sub-D vacuum electrical feedthrough. Unfortunately, the feedthrough could not be soldered to the IVC without risking its destruction. It was therefore glued with Stycast 2850FT as suggested by the literature [RS98].

[3] All gases other than ^{3}He have negligible vapor pressure below 1 K and are therefore not available at sufficient pressure to provide a significant heat exchange mechanism.

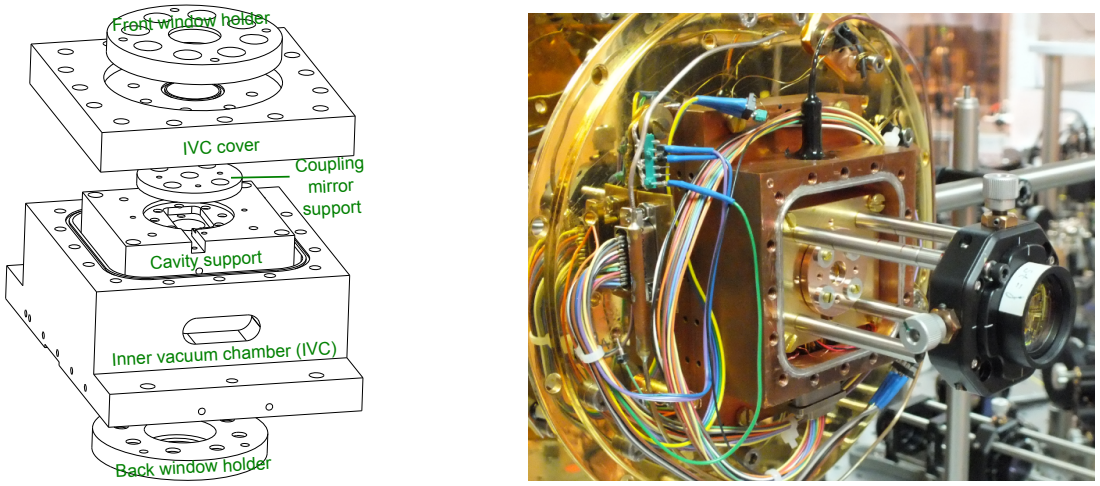


Figure 4.6: Left: Schematic of the IVC design, including the cover, front and back window holders and cavity assembly. The window covers are made from aluminum, all other pieces from OFHC copper. Brass screws are used preferredly because of their large thermal contraction. Right: Photograph of the opened IVC mounted on the mixing chamber screen of the dilution refrigerator. The measurement cavity is mounted inside the IVC and four metal rods hold an alignment lens, which is mounted on the IVC cover once the chamber is closed. The silver residue around the circumference of the IVC walls is a residual of the indium joint. The fill capillary arrives at the top of the IVC. The solder joint leak tightness is double-secured with a coating of excess Stycast 2850FT. The electrical feedthrough with arriving cables can be seen at the IVC bottom. All screw holes are designed such that no air is trapped by tightened screws.

3. The IVC must be leak-tight at low temperature, in particular when superfluid ^{4}He is to be used. For example, a circular hole with a diameter of 1 μm is expected to cause a superleak rate of the order of 2×10^{-5} mbar.l / s. The number of cryogenic joints must therefore be kept minimum, with the only required removable joint between the IVC body and cover implemented by strongly compressing indium wire between two copper surfaces [RS98]. The large number of screws required to exert the necessary pressure can be seen by the number of holes in the IVC cover in Fig. 4.6.

4. The optical beam must arrive inside the IVC at the measurement cavity. Therefore, two (for injection and transmission) AR-coated windows are implemented along the axis of the IVC. The same type of indium wire as for the cover joint is pressed between a copper surface and the glass window to form leak-tight joints for the windows.

5. The amount of helium gas inside the IVC should be variable without re-heating the cryostat to room temperature in order to allow for the proper characterization of the effects induced by the presence of helium around the sample. By using long (2.5 m) and thin (outer diameter 1.5 mm, inner diameter 0.6 mm) tubing of very low heat conductivity (CuNi), a fill tube from the outside of the cryostat to the mixing chamber has been implemented. The capillary was soldered to an appropriate aperture of the IVC, thermally anchored at the 1 K, 4 K and 100 K screens and fed outside the cryostat through the OVC pumping joint to a custom-made KF flange. The introduction of the fill CuNi capillary raised the base temperature from 30 mK to 39 mK. Future versions should therefore use an even longer capillary between the 1 K and mixing chamber plates, and potentially a cryogenic valve [SS98] to prevent superfluid flow along the fill tube to cause heat leaks when the IVC contains ^{4}He.

6. Since considerable amounts of helium may be present inside the IVC, the risk of failure-induced rapid vaporization of large gas volumes must be mitigated. By far the best solution is a middle vacuum chamber (MVC) containing all cryostat parts at or below liquid helium temperature[4], which avoids that gas leaks from the IVC thermally link cold cryostat parts with the room-temperature OVC walls and potentially produce a self-amplifying heating mechanism. The absence of an MVC in our cryostat implies serious constraints: No cryogenic overpressure valves (e.g. blow-off-valves, see [LS14]) can be implemented as this would have the undesired effect of breaking the OVC vacuum. The high flow-impedance fill capillary is therefore the only way to evacuate gas from the IVC. Due to the small inner diameter (0.6 mm) of the capillary, it may become clogged by frozen air and render evacuation impossible up to liquid nitrogen temperatures, where helium has already expanded 200-fold with respect to its liquid volume. Therefore, we restrict the amount of injected helium to at most the quantity required to fill the IVC volume at room-temperature and atmospheric pressure. Since this implies a maximum liquid ^{4}He-volume of 1/866 times the IVC volume, the sample cannot be totally immersed in liquid helium. We further reduce the risk of failure in our

[4]This scheme is implemented for example in the group of S. Balibar at the ENS Paris.

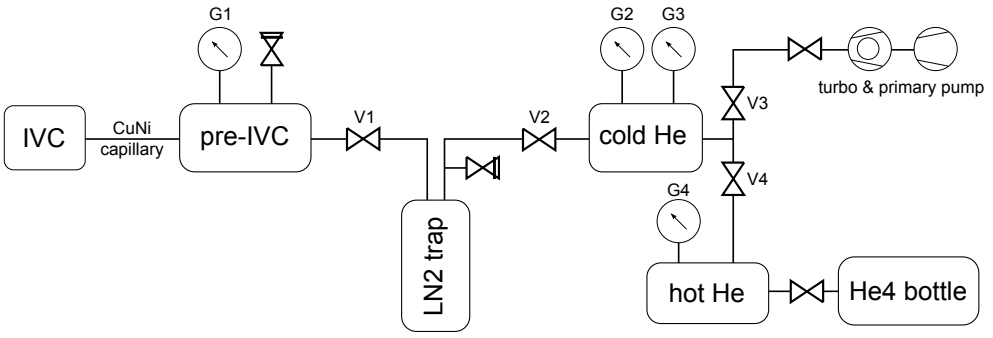


Figure 4.7: Vacuum components before the IVC. Vacuum tubes with a volume comparable to the IVC are shown as boxes. Overpressure valves are installed at the room-temperature tubes connected directly to the IVC (pre-IVC) and at the liquid nitrogen trap (LN2 trap). All valves V1-V4 are manually controlled, gauge G1 is a membrane gauge ($10^{-2} - 130$ mbar), G2 and G4 are Bourdon gauges ($10 - 1000$ mbar), and G3 is a combined Pirani and cold cathode gauge ($10^{-4} - 10$ mbar).

implementation by placing over-pressure valves at the room-temperature parts connected to the IVC capillary, and by passing all injected gas through a liquid nitrogen trap to filter away leaked air and other impurities (see Fig. 4.7).

Helium injection

To test the proper operation of the fill capillary and the external vacuum components for helium injection (see Fig. 4.7), the linearity of the pressure-dependence of mechanical damping was confirmed at room-temperature. For this measurement, the IVC is first pumped down through the fill tube and a leak-test is performed[5]. After closing V1 and V3, the cold LN_2 trap is filled with 10 to 30 mbar of ^{4}He from the bottle through V4 and V2, after which V4 is closed and V2 left open. After about one minute, the readings of G2 and G3 stabilize and V1 is opened for roughly 1 s to inject a fixed amount of helium into the pre-IVC space while continuously recording the pressure of G1 as shown in Fig. 4.8. The pressure discontinuities occuring when V1 is opened are clearly visible. After closing V1, the large impedance of the capillary between IVC and pre-IVC leads to a slow exponential decay of the pressure towards the equilibrium value (yellow fits in Fig. 4.8) with decay time $\tau = p_{\text{avg}}^{-1} \times 660\,\text{s.mbar}$. To theoretically predict the decay time, we first estimate the mean free path of helium at room temperature and a pressure of 10 mbar to be $15\,\mu$m, which places the flow inside the fill capillary in the viscous regime for an inner diameter of about 30 times the mean free path. In this regime, the Poiseuille equation [RS98] leads us to predict an exponential

[5]We perform various leak tests: for a quick estimation of leaks, the IVC is evacuated through the fill tube with a turbo pump in series with a leak detector, and helium is sprayed around the IVC (ultimately by filling the OVC to a considerable helium pressure). Due to the large flow-impedance of the fill tube, a better leak test is done in the inverse direction: the leak detector pumps on the OVC while the IVC is filled with ^{4}He through the fill tube. This test is continuously performed when the cryostat is cooled down to detect the emergence of temperature-dependent leaks.

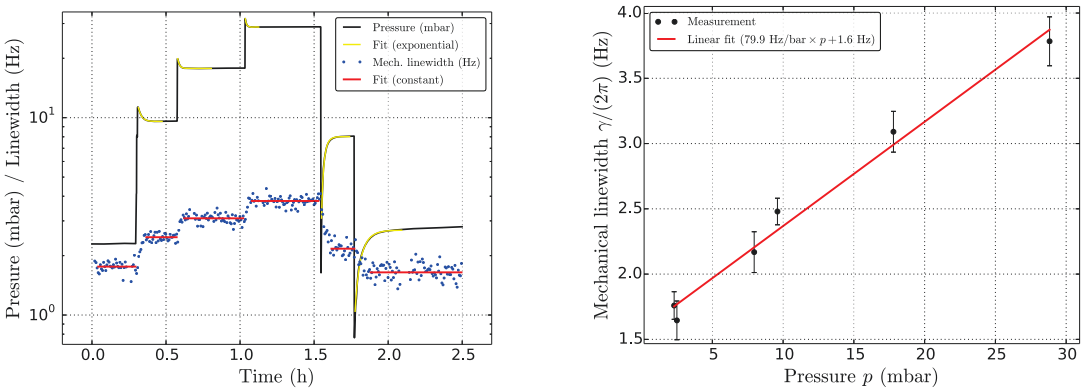


Figure 4.8: Left: Pressure of G1 (pre-IVC) and simultaneously measured mechanical damping rate $\gamma/(2\pi)$. At the discontinuities of the pressure, ^{4}He is injected to or removed from the IVC by opening valve V1 for about 1 s to the nitrogen trap holding a different helium pressure. The pressure is fitted with exponential decays towards an asymptotic value. The decay-time τ for the establishment of pressure equilibrium through the CuNi capillary between IVC and pre-IVC is inversely proportional to the average pressure p_{avg} and of the order of $\tau = p_{\mathrm{avg}}^{-1} \times 660\,\mathrm{s.mbar}$. Right: Extracted damping rate dependence on the helium pressure. The indicated measurement values are the averages of the measured damping rates over the indicated red lines in the left plots, while the indicated error corresponds to their variance. The offset of the linear fit indicates an intrisic quality factor of 2.2×10^6 while the slope yields a pressure damping rate of 79.9 Hz/bar.

decay[6] of the pre-IVC pressure $p(t) = \bar{p} + \delta p \exp(-t/\tau)$ towards its asymptotic value $\bar{p}$ with a decay time τ for small initial pressure amplitudes δp. We find the numerical decay time to be

$$\tau = \frac{1}{p_{\mathrm{avg}}} \frac{128\eta L}{\pi d^4} V\alpha(1-\alpha) \approx \frac{375\,\mathrm{s.mbar}}{p_{\mathrm{avg}}} \tag{4.1.2}$$

where $L = 2.5\,\mathrm{m}$ denotes the capillary length, $d = 0.6\,\mathrm{mm}$ its inner diameter, $V = 0.12\,\mathrm{l}$ the total volume of IVC and pre-IVC, $\alpha = 0.7$ the ratio of the pre-IVC volume to V, $\eta = 2.0 \times 10^{-5}\,\mathrm{kg.m^{-1}.s^{-1}}$ the viscosity of ^{4}He at room temperature and p_{avg} the average pressure between IVC and pre-IVC vessels, which we approximate by the asymptotic pressure $\bar{p}$ in the fits. The theoretically predicted value coincides only to within a factor of 2 with the observed one. This significant discrepancy is not surprising because a number of poorly justified approximations were made in Eq. 4.1.2 (e.g. assuming a purely viscous flow and equating p_{avg} and $\bar{p}$) and many experimental parameters such as the volumes or actual inner diameters are not precisely known.

The damping rates deduced from simultaneously performed ringdown measurements are also shown in Fig. 4.8, along with the extracted linear dependence of the damping rate on the surrounding gas pressure with a slope of 79.9 Hz/bar. The observed gas damping of mechanical motion will be further discussed in section 4.1.6.

In a recent experiment, the IVC was filled with 70 mbar of ^{4}He at room temperature

[6] If the assumptions of constant p_{avg} or of small pressure amplitude $\delta p \ll p_{\mathrm{avg}}$ are dropped, different functional dependencies hold. For example for constant pressure in one of the vacuum chamber such as during pumping, a t^{-2} time-dependence is found.

after which the cryostat was cooled down with a leak detector constantly pumping on the OVC. While no leaks appeared down to about 2 K, the leak rate increased to a level above 10^{-2} mbar.l/s during the condensation of the ^{3}He-^{4}He mixture. The associated heat leak prevented to reach dilution operation with the cryostat. While no leaks were found before the cryogenics run, a leak of the order of 10^{-7} mbar.l/s was detected at room temperature after the cryogenics run and could be located to be related the joint of the electrical feedthrough of the IVC. The leak increased further after subsequent cryostat runs, despite attempts to fix the leak by applying small amounts of Stycast 2850FT to suspicious parts of the joint. We therefore plan a re-design of the electrical IVC feedthrough to prevent superleaks in the future.

Quartz tuning fork pressure sensor

A commercial 32 kHz quartz tuning fork [FC07] has been mounted inside the IVC as a pressure sensor. In order to couple the tuning fork to the surrounding gas, its cylindrical metal housing has been opened at the top. For the 32 kHz mode, we have measured quality factors of the order of 10^5 in vacuum and 10^4 at ambient pressure by a simple measurement of the electrical reflection with a network analyzer. The gas-dependent damping has been characterized at room temperature and was found to enable pressure measurements of the order of 10 mbar, limited by a low signal-to-noise ratio when gas damping becomes comparable to the intrinsic damping. Quality factors above 10^6 have been measured at low temperature in agreement with reported values from the literature [MMM08], resulting in even better pressure sensitivity, especially in light of the $1/\sqrt{T}$-increase of gas damping rates for constant pressure. Driving the tuning fork for longer (> 2 s) amounts of time near its resonance frequency with an unnecessarily high electrical power of 10 dBm resulted in the destruction of the tuning fork during the cryogenic characterization: one prong of the fork was later found at the bottom of the IVC. Nevertheless, the obtained data show that a pressure of 1 mbar of ^{4}He should lead to a quality factor reduction by a factor of 2 for a quartz tuning fork at 4 K, which therefore constitutes a cheap (< 1 Euro), compact and rapidly implemented pressure sensor for our purposes that is useful in the range from 0.1 to above 1000 mbar. Similar conclusions with more extensive characterization measurements were reported by another group [HVVT14].

4.1.6 Gas damping

A theoretical expression for the damping rate of the micropillar's compression mode in the viscous regime can be derived in analogy to [Bri04]. From Eqs. 2.1.13 and 3.1.3, we find that the energy $E(t)$ stored in the vibration of the compression mode for an initial displacement amplitude x_0 of the pillar surface at $t = 0$ decays at the rate

$$\frac{\mathrm{d}E(t)}{\mathrm{d}t}\Big|_{t=0} = -\frac{1}{2}\gamma_0 m_0 \Omega_0^2 x_0^2 \tag{4.1.3}$$

where the effective mode mass according to Eq. 2.1.24 relates to the density $\rho_{\mathrm{quartz}} = 2650\,\mathrm{kg/m^3}$, height h and surface area A of the pillar through $m_0 = \frac{1}{2}\rho_{\mathrm{quartz}} A h$. By neglecting the gas damping that occurs at the transversally moving micropillar sides,

Gas	Gauge	Pressure (mbar)	$\frac{\gamma_{\text{measured}}}{2\pi\,\text{Hz/bar}}$	$\frac{\gamma_{\text{theory}}}{2\pi\,\text{Hz/bar}}$	Reference	Sample
air	atmospheric	1000	104.3	114.9	Fig. 3.9	20121002C9
^{4}He	Bourdon	1500	42.0	35.1	28/11/13	20121009C9
^{4}He	atmospheric	1033	44.0	35.1	27/11/13	20121009C9
^{4}He	membrane	$10-100$	62.2	35.1	Fig. 4.9	20121009C9
^{4}He	membrane	$2-30$	79.9	35.1	Fig. 4.8	20121009C5

Table 4.1: Different measurements of gas damping rates γ_{measured} normalized to 1 bar by assuming a linear pressure dependence. The effect of intrinsic damping has been removed from the data if necessary. Theoretical values γ_{theory} are obtained from Eq. 4.1.5. If no figure in this thesis shows the corresponding measurement, a date as reference indicates the date of observation (see lab notebook for further information).

we can estimate the energy radiated by the top- and bottom surface of the micropillar per time $\mathrm{d}t$ into the surrounding gas with density ρ_{g} and sound velocity c_{g} by integrating twice the energy density of a planar acoustic wave over the volume $Ac\,\mathrm{d}t$. Since the gas velocity must coincide with the surface velocity at the pillar surface $\dot{x}(t) = \Omega_0 x_0 \sin(\Omega_0 t)$, we eventually find

$$\frac{\mathrm{d}E(t)}{\mathrm{d}t}\Big|_{t=0} = \rho_{\text{g}} A c_{\text{g}} \Omega_0^2 x_0^2 \,. \tag{4.1.4}$$

Equating the power loss of the compression mode with the radiated energy allows to find the gas damping rate

$$\gamma_0 = \frac{4 c_{\text{g}}}{h} \frac{\rho_{\text{g}}}{\rho_{\text{quartz}}} \,. \tag{4.1.5}$$

For an ideal gas at temperature T with atomic mass m_{g}, we find that the sound velocity [LL80]

$$c_{\text{g}} = \sqrt{\frac{\gamma k_B T}{m_{\text{g}}}} \tag{4.1.6}$$

only depends on the temperature. The adiabatic index γ is 5/3 for monoatomic and 7/5 for diatomic gases. At constant temperature, the damping rate is therefore proportional to the pressure p through the gas density, which for an ideal gas reads $\rho_{\text{g}} = p m_{\text{g}}/(k_B T)$.

Various experiments have been performed to study the effect of gas damping on the compression mode of the micropillar. The measurement in Fig. 4.9 confirms a linear dependence of damping rate on pressure for pressures above 10 mbar. A more exhaustive list of the available data is shown in Table 4.1. We find that the measurement with air at atmospheric pressure agrees to better than 10 % with the theoretical prediction. The measurements with ^{4}He at or above atmospheric pressure also agree reasonably well with Eq. 4.1.5. The remaining damping rates for pressures measured with the membrane gauge G1 are significantly too high, and furthermore do not agree perfectly with each other. A number of possibilities come to mind in order to explain this disagreement:

- Both measurements were performed with the same membrane gauge. The simplest explanation of our observation is that the gain of the gauge is off by a factor of two and that it has furthermore drifted over time.

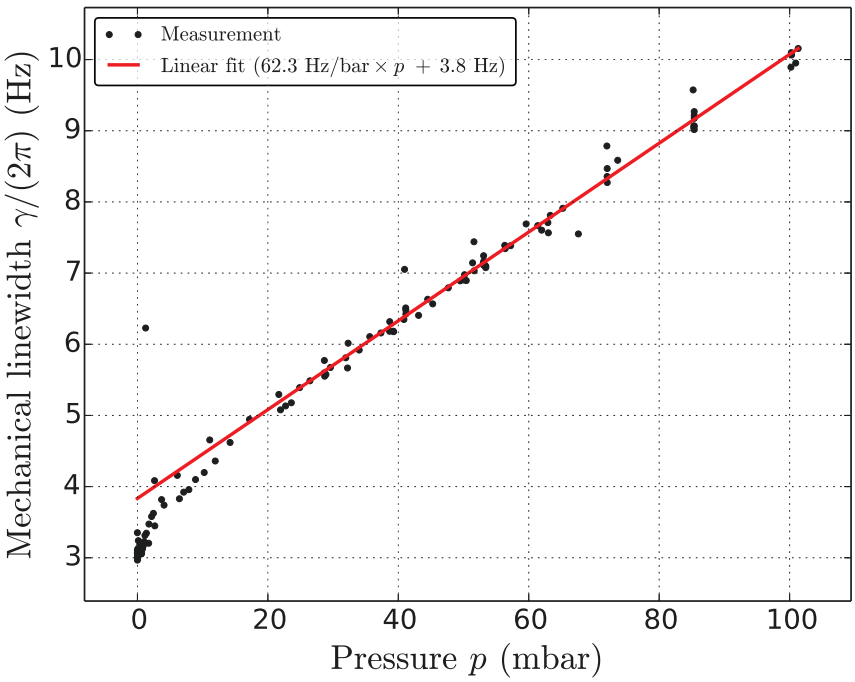


Figure 4.9: ^{4}He gas damping rate or the micropillar's compression mode measured by variation of the helium pressure inside an older version of the IVC. Pressures are measured with the membrane gauge G1, damping rates extracted from ringdown measurements. From $10 - 100$ mbar the data are well-modeled with a linear fit. For lower pressures, the damping is lower than expected, which is a signature of the passage from the viscous to the molecular regime that calls for a different model. The outlier at 1.3 mbar is caused by a short delay between pressure and ringdown measurements, during which the helium pressure changed significantly.

- Even though Fig. 4.9 suggests a linear dependence and we do not expect a change of regime above the viscous regime, the ensemble of measurements suggest that the linear gas damping coefficient in units of Hz/bar decreases at higher pressures, which is equivalent to a sublinear dependence on pressure.

- The samples were mounted differently for the two low-pressure measurements: one was used in an interferometer configuration, the other one in a very short Fabry-Perot cavity. The presence of other surfaces less than 150 μm away from the vibrating pillar surface is likely to affect the gas damping rate, such that our model is expected to become imprecise.

Despite a possible error of the pressure damping rate by a factor of up to two, we can conclude that gas damping is low enough to allow operation of resonators with quality factors of 10^6 in ^{4}He pressures of the order of 50 mbar. While nano-optomechanical systems typically require operating pressures in the high-vacuum range, the macroscopic size $h = 1\,\mathrm{mm}$ of the micropillar is directly responsible for this robustness against gas pressure, which can be seen from its appearance in Eq. 4.1.5.

4.1.7 Superfluid ^{4}He damping

As an extension to the measurement presented in Fig. 3.12, the effect of ^{4}He on the mechanical properties of a micropillar was studied as a function of temperature. For this measurement, an older version of the IVC was used. In a first cryogenic run, the sample was mounted on the older IVC support as described in section 3.1.4, the vacuum chamber was closed and its fill capillary left connected to the OVC, which resulted in a vacuum inside the IVC once the OVC had been pumped down. The inferred mechanical loss angle is shown in Fig. 4.10. After heating the cryostat back to room temperature, the vacuum of the IVC was naturally broken when the OVC was filled with nitrogen gas in order to open the cryostat. With the cryostat open, the IVC fill tube was connected to a pump and pumped down, purged a few times with helium and finally filled with a pressure of 1.5 bar of ^{4}He through the fill capillary. The copper fill capillary was then closed by mechanical squeezing, cut with a wire cutter to separate it from the pump tubing, and quickly sealed with a drop of standard electronics solder. By observing the strongly gas-damping dominated mechanical loss over a few hours during which the cryostat OVC was pumped to vacuum and leak-tested, it was confirmed that no leaks were present in the system. During the whole procedure, the IVC was not opened in order not to modify the mounting of the sample as to avoid a change in its mechanical quality factor. The mechanical loss angle was measured during two subsequent cryogenic runs with the same procedure as for loss angles in vacuum. The inferred loss angle is also shown in Fig. 4.10.

From Eqs. 4.1.5 and 4.1.6, we expect the gas damping rate to be proportional to the square root of the temperature. Indeed, a square root function yields a good fit of the data down to a few Kelvin. Excess damping is observed when the loss measured in vacuum becomes comparable to the gas damping, in particular for the peaks at 15 K and 33 K. Below 2 K, the gas damping rate decreases abruptly, indicating the start of condensation of the helium in agreement with tabulated values of the saturated vapor pressure [Pob07]. We expect gas damping to rapidly become negligible for

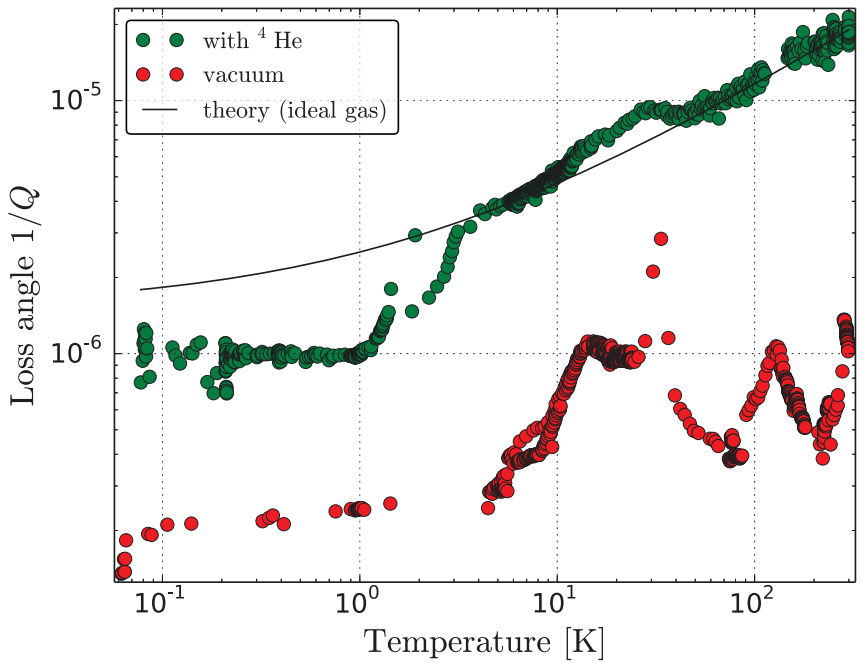


Figure 4.10: Loss angle of the micropillar versus cryostat temperature with the same sample 20121009C9 inside an older version of the IVC. The environment of the micropillar is either vacuum or a fixed amount of ^{4}He at constant volume.

temperatures below 1 K due to the exponential decay of the vapor pressure. However, we observe a saturation of the mechanical loss angle at a constant level in a range from 0.1 to 0.8 K, corresponding to a damping rate of the order of $\gamma/(2\pi) = 4.0$ Hz.

We have attempted to explain the observed damping by assuming the following simple model. As the micropillar vibrates with an amplitude of x_0, its surface area oscillates by the amount $6ax_0$, where $a = 240\,\mu$m is the sidelength of the triangular micropillar surface and we multiply by two to account for both sides of the micropillar. The excess surface must be periodically covered by the saturated Rollin film of thickness $d \approx 30$ nm [Wil67], which requires a periodic flow of the film towards and away from the micropillar. In the simplest model, we can assume that the periodic flow propagates through the film in the form of a surface wave, also known as third sound [Atk59]. The dispersion relation of third sound is linear with a sound velocity that can be assumed to be of the order of $c_3 = 1\,$m/s [EAD64] for our first estimate. For an upper estimate of the damping resulting from the wave emission, we assume that all radiated energy in the created third sound wave is lost and neglect the fact that the micropillar dimension is much larger than the wavelength of the emitted wave of roughly 0.3 μm. With an analog calculation as the one presented in section 4.1.6, we find for the third-sound induced damping rate γ_3 the expression

$$\gamma_3 = \frac{48 c_3 d}{\sqrt{3} h a} \frac{\rho_{^4\mathrm{He}}}{\rho_{\mathrm{quartz}}} \tag{4.1.7}$$

which yields the estimate $\gamma_3/(2\pi) \approx 0.027$ Hz for our parameters. The observed value is about two orders of magnitude above the theoretical upper estimate, thereby excluding third sound emission as a significant mechanical loss channel. A possible alternative explanation for the observed damping is the evaporation of the superfluid film on the laser-heated micropillar surface: random momentum kicks of emitted Helium atoms are expected to create a damping mechanism similar to gas damping in the ballistic regime. A measurement of the power-dependence of the damping rates associated to the superfluid film could clarify this proposition. Another valid explanation for the observed damping is a change in the clamping of the mechanical sample due to thermal cycling, thereby resulting in a different low-temperature limit of the quality factor. A change of the mechanical quality factor after four thermal cycles has been observed for another sample (see section 7.2.5) and remains the most likely explanation. It is therefore desirable to repeat this experiment with variable amounts of helium in the IVC without thermal cycling, which will be possible with the more recently implemented fill capillary once the superleaks of the new IVC have been removed.

4.2 Optomechanical cavity

While the cryostat provides a low-temperature environment for the samples, the heart of our experiment is the optomechanical Fabry-Perot cavity (measurement Fabry-Perot cavity, short FPM) that is mounted on the cryostat's mixing chamber plate. Having already presented the properties of the samples in chapter 3, we will discuss in more detail here how to combine the samples into a Fabry-Perot cavity that is compatible with the constraints imposed by cryogenics.

4.2.1 Requirements for a cryogenic optomechanical cavity

The performance of our optomechanical system will not only depend on the sample properties alone, but on the detailed manner how these samples are mounted into a cavity assembly. Most of these details are consequences of various performance requirements, which we list in the following:

1. As concluded in section 3.2.3, the **cavity length** should be shorter than 200 μm to avoid excessive optical clipping loss by the micropillar mirror. The resulting free spectral range $\nu_{\mathrm{fsr}} > 750\,\mathrm{GHz}$ does not lie in the tuning range of our laser. The cavity length must therefore be adjustable over at least $\lambda/2$ to match the cavity resonance frequency with the laser frequency. Inserting a thin piezo disc between the two cavity mirrors for cavity length tuning would require the piezo to be maximally strained by 0.25 % of its thickness, while typical piezoelectric materials are short of this number by a factor of about 20 at cryogenic temperatures. The piezo must therefore be mounted behind either one of the two mirrors, which makes the cavity design less compact.

2. **Vibrations** caused by boiling helium or various pumps are easily a problem for sensitive experiments inside cryostats. Preliminary experiments [Kuh13] in our dilution refrigerator suggested that the length of cavities constructed from composite copper structures holding the micropillar on top of a piezo on one side and the coupling mirror on the other side with vacuum in between the two samples vibrated with mean amplitudes of more than ten times the optical linewidth $\lambda/(2\mathcal{F}) \approx 10\,\mathrm{pm}$. This level of vibration prevents stabilization of the cavity resonance frequency to the laser frequency due to the unavailability of sufficiently broadband and large-range frequency modulators. Unable to compactify this design because of the piezo requirement in point 1, attempts were made to mitigate the level of vibration by inserting a hard spring between the two sample surfaces. The requirements that the spring should

   - allow to be strained by more than 0.25 % by the piezo at low temperatures,
   - be sufficiently hard to rigidify the two sample surfaces so that the level of relative surface vibration is reduced by at least a factor of 100,
   - have a large damping coefficient to avoid the introduction of another mechanical resonance peak,
   - and not introduce excess cavity length noise near the mechanical measurement frequency

   can all be realized by placing thin rings of Kapton between the two samples to define the cavity length (see section 4.2.1).

3. The two mirror samples should be **parallel** to about 1° in order to place the intracavity beam at the center of the coupling mirror where optical loss is minimum. All critical mechanical parts are therefore fabricated with this parallelism specification. Furthermore, the Kapton spacer is cut from a sheet of uniform thickness and the mirror angle is fine-tuned with the tension of the screws that compress the spacer by optimizing the observed finesse.

4. For alignment and later diagnostics, the **transmitted light** of the cavity should be accessible. Therefore, all piezos and other mechanical parts have a central hole at the cavity axis.

5. The **thermal contact** between each sample and the mixing chamber plate of the cryostat should be good enough that the sample temperature coincides with the mixing chamber. Therefore, all OFHC copper parts are gold-coated to avoid the Kapitza resistance of the otherwise present oxide layer. Samples are glued with Stycast 2850FT to the polished and clean surfaces of the copper holder. A copper slice is inserted between the piezo and the micropillar sample to thermally shortcut the piezo.

6. **Thermal contraction** upon cooling down the cavity must be taken into account for two reasons. First, thermal stress of the micropillar sample induced by different thermal expansion coefficients of the quartz wafer and the supporting structure may lead to a degradation of the micropillar's mechanical properties [Wil12]. To mitigate this effect, the micropillar sample is glued on top of a quartz support which contracts at the same rate as the sample. Thermal stress will therefore principally appear at the interface between the quartz support and the metal structure underneath where it cannot affect the mechanical oscillator. Secondly, the cavity may become misaligned due to thermal contraction. This effect is kept negligible by making the design as cylindrically symmetric as possible. Furthermore, the different cavity parts are held together with brass screws on top of stainless steel wafers, which increases the screw pressure and results in even better stability at lower temperatures because the screws contract more than the copper and steel parts.

7. Since the cavity assembly can be considered as a part of the optomechanical sample used in the experiment, its design should be **modular** such that different assembled and aligned cavities can be easily exchanged in the experiment without affecting the properties of an aligned cavity. Furthermore, micropillar and coupling mirror samples should be exchangeable without risking the destruction of either sample, i.e. samples should only be glued to removable parts of the cavity assembly.

8. Finally, a **quick and easy alignment** of the cavity allows testing more samples and thus a better optimization of the optomechanical system. To this end, both the coupling mirror alignment parts and the lens needed to focus the incident laser beam into the cavity are held concentrically by metal rods screwed to the cavity support (Fig. 4.13). Lateral alignment of the beam and of the coupling mirror is done by using commercial micropositioners that can slide along the rods in the axial direction. While these positioners still have hystereses at the 10 μm level, alignment is much easier than with the previous home-built positioners.

Kapton spacer

The first measurement cavities were constructed as stacks of a micropillar sample, four $1\,\mathrm{mm} \times 1\,\mathrm{mm}$ squares of 125 μm at the sample corners, and a coupling mirror (see the

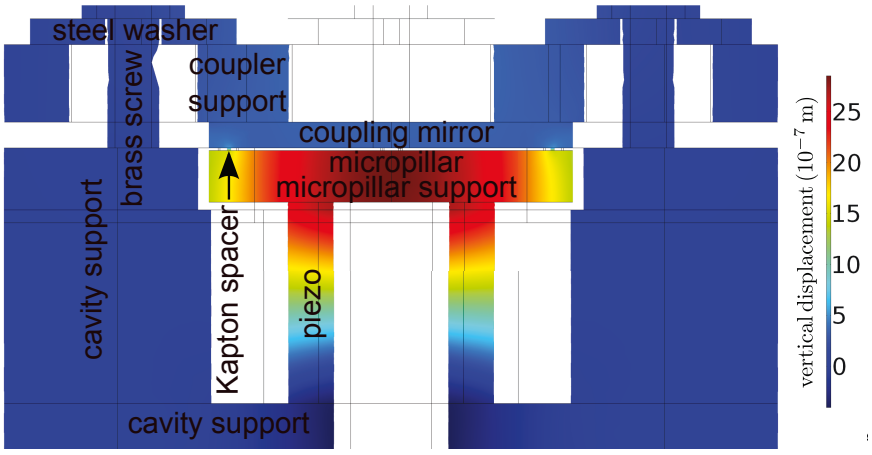


Figure 4.11: Simulated displacement along the measurement cavity axis for the maximum piezo voltage in a simplified geometry. Between the center of the micropillar sample and the coupling mirror, about 2 μm of displacement is achieved which is sufficient for tuning the cavity length over several free spectral ranges. The majority of the strain occurs in the four 0.1 mm $\times$ 1 mm $\times$ 125 μm Kapton spacers.

middle of Fig. 4.12). In this design, four screws act on the coupling mirror support to compress the stack situated on the quartz sample support on top of a piezo stack. If the piezo support and the coupling mirror can be assumed rigidly clamped and the samples are assumed much harder to compress than the Kapton film, the expansion of the piezo compresses the Kapton film and leads to a cavity length change. Kapton as the material of choice because of its low elastic modulus of $E < 8$ GPa even at cryogenic temperatures [DBM92] in comparison with the one of quartz or copper in the 100 GPa range. The three[7] employed piezos PD080.30 from Physik Instrumente in Lederhose, Germany, have a blocking force specification above 1 kN, which should allow to compress 4 mm^2 of 125-μm thick Kapton by more than 3 μm, which is sufficient. The remaining design challenge of the Kapton spacer cavity is thus to build a sufficiently rigid piezo and coupling mirror support structure whose deformation is small in comparison with the Kapton spacer. To this end, a finite-element simulation of the cavity assembly was performed with COMSOL with a result shown in Fig. 4.11. The presented geometry achieves a cavity displacement of about 1.5 μm with the maximum piezo voltage of 100 V. This value must be divided by a factor of 7/2 for low-temperature operation, since the piezoelectric constant is expected to decrease about 7-fold, which is partially compensated by the possibility to use both voltage polarities in this temperature regime. With this configuration, it was experimentally observed that about three subsequent cavity resonances could be tuned to the laser frequency, while barely one resonance was accessible in cryogenic operation. The very small error margin of the piezo displacement is a result of the fact that the measurement cavity was designed and constructed with back-of-the envelope estimations, long before the

[7] At the time of the experiment, only piezo stacks with a nominal displacement of 2 μm were available. Since the displacement at cryogenic temperatures is known to be about $6-7$ times lower than the room-temperature value, three such piezos are stacked on top of each other.

above presented Kapton spacer simulations were conducted.

Only after the successful implementation of this design, it was realized that the previously observed vibrations (see point 2 in section 4.2.1) were the consequence of an optomechanical instability imposed by the geometry of the micropillar sample and not caused by the cavity assembly design. A low-vibration operation of the cavity without Kapton spacer was therefore possible and enabled a simpler alignment procedure and resonance tuning at low temperatures. In this cavity design, the micropillar sample is glued to its quartz support and the coupling mirror support disc is pressed against metal spacers on top of the cavity support block. To achieve the desired cavity length, a large amount of metal spacers of different thickness is available. In this design, at least two optical resonances are accessible by using only two of the three piezos.

4.2.2 Final cavity design

The design considerations from the previous section have led to the Fabry-Perot cavity assembly shown in Fig. 4.12. The coupling mirror sample and the system used for alignment can be seen in Fig. 4.13. Due to its convenience for alignment, the measurement cavity setup was also used to characterize the coupling mirror samples as presented in 3.2.3.

4.2.3 Cavity alignment protocol

While the alignment of our Fabry-Perot cavities is straightforward if a well-designed aligment protocol is followed, much time can be lost if it is not. The following steps lead to an optimal alignment with our setup.

Cavity assembly

The cavity assembly is performed in a cleanroom environment to prevent exposure of the optical surfaces to dust.

1. Since it takes about one to two days from the assembly of a new measurement cavity to the mechanical characterization in vacuum, one should first confirm in a dedicated characterization interferometer that the mechanical resonator sample has the desired properties.

2. For a new cavity, we start by inserting the piezo alignment cylinder from the back of the cavity support through the three piezos and place the micropillar support plate on that cylinder from the top. This should make the holes in cavity support, piezos and micropillar support plate concentric. The support plate is then clamped by screwing the two foreseen screws finger-tight.

3. A new quartz support can be easily fabricated from an unusable or broken micropillar sample by removing the decoupling frame and micropillar and keeping only the chip with the triangular hole. This support is glued onto a new micropillar support plate with Stycast 2850FT. Next, the micropillar sample is glued with three tiny dots of the same glue by using a syringe to apply the glue and gently spreading the glue by pressing on top of the micropillar sample with a tweezer.

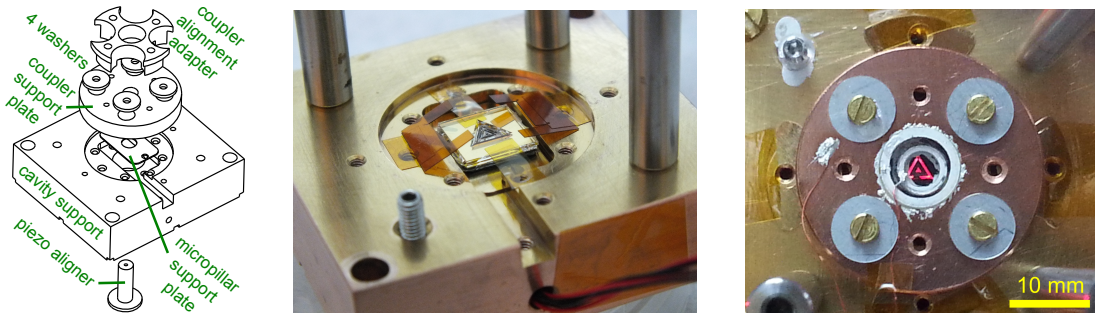


Figure 4.12: Left: Schematic of the designed cavity assembly. Middle: Photograph of a micropillar sample mounted with its support plate on the cavity support structure. Three piezos are pressed against the back of the micropillar support plate by the two screws at the top (not visible here) and bottom of the micropillar. The micropillar is only posed on a quartz support of the same dimension that is glued to the micropillar support plate, and prevented from falling down during alignment by Kapton scotch tape at the top and bottom of the sample. From the left and right of the sample, "U"-shaped pieces of Kapton provide four spacer squares at the corners of the micropillar sample with dimensions of roughly $1\,\mathrm{mm} \times 1\,\mathrm{mm}$. These spacers will be compressed on the micropillar sample by the coupling mirror substrate and thereby hold the micropillar precisely in place. Right: Photograph of the measurement cavity after completion of alignment, ready for a cryogenic experiment. The micropillar sample is illuminated with a red flashlight from the back to illustrate its macroscopic size: the central triangular dot visible by the naked eye corresponds to the mechanical oscillator that will be cooled close to its quantum ground state. The silver-white ring around the central triangle is a fast piezo glued to the backside of the coupling mirror.

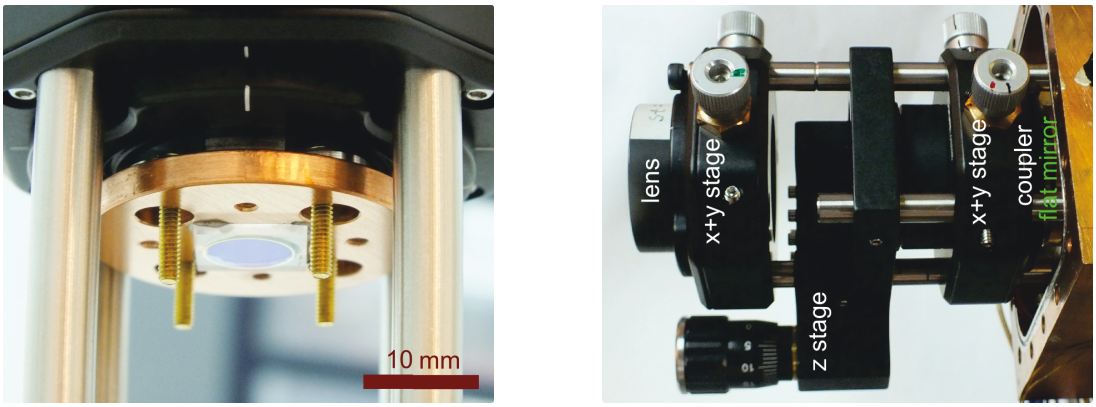


Figure 4.13: Left: Photograph of a coupling mirror glued to a copper disc. The disc is screwed from the back to the black 2-axis translation stage, which can slide along the third direction on the four steel rods. The four brass screws sticking through the disc can be used to mount the copper coupling mirror on top of another mirror sample to construct a stable Fabry-Perot cavity assembly from which the translation stage can be removed. Right: Photograph of the mechanical parts used to align a Fabry-Perot cavity. The copper support on the right edge holds the back mirror of the cavity on top of a piezo stack which can be used for fast scans of the cavity length. Four cylindrical steel rods are screwed into the support, which hold the coupling mirror mounted on a lateral translation stage (Thorlabs CXY1, also visible on the left photo) that is needed to align the cavity. The lateral translation stage can be moved along the rods to tune the cavity length by the axial translation stage (Thorlabs SM1Z) behind it, which is screwed to the rods. The rods are also used to hold a lens on a lateral translation stage needed to focus the incident laser beam to a waist size of the order of $10\,\mu$m.

Attention is paid to center the micropillar sample to within 1 mm from the center of the support. During the 24 h of curing, the sample surface is covered with a clean lid in order to prevent dust deposition on the micropillar mirror.

4. In the meantime, a suitable coupling mirror is chosen and glued to a new coupling mirror support plate. To this end, the coupling mirror support plate is screwed to the lateral positioner (Thorlabs CXY1) equipped with the coupler alignment adapter. The lateral positioner is rigidified with an axial positioner (Thorlabs SM1Z). The positioners are shown in Fig. 4.13. The positioning screws are brought to a central position and the positioners are slid onto a cage system with a clean endplate with a concentric hole. The coupling mirror is placed at the center of the endplate with the mirror surface facing down and manually centered on the hole. Glue is applied to the coupling mirror support plate and this plate is slowly lowered until it is pressed against the sample. In earlier stages of the experiment, a conductive silver glue was used for the possibility of later removal of the glue. For permanent gluing, Stycast 2850FT is expected to yield better results. The coupling mirror surface may be protected with a layer of FirstContact during this step.

5. Metal spacers needed to achieve the desired cavity length are now scotched on the cavity support block such that the scotch tape can be removed once the spacers are compressed between cavity and coupling mirror support.

6. After curing of the glue, the positioners holding the coupling mirror sample are removed from the gluing cage system and instead slid onto the four rods screwed to the cavity support block. The coupling mirror is carefully lowered until it is about 0.5 mm above the surface of the micropillar (visible with the bare eye). At this point, the axial positioner can be screwed to the sliding rods as the remaining cavity length tuning will be done with the microscrew of the positioner.

7. Using the axial positioning screw, the cavity length is made as short as possible without risking the physical contact between the two samples. The free movement of the lateral stage must not be compromised by touching sample surfaces.

8. The lateral positioners are adjusted such that the concave structure of the coupling mirror is visibly placed on top of the micropillar. This can be achieved by careful inspection with the bare eye by drawing a straight line on the plexiglas window between the experimenter and the cavity in the dust-free environment and moving the coupling mirror such that both the concave feature and the micrpillar appear on the line with the cavity support block placed on a flat surface and rotated in steps of 90 °.

9. Four steel washers are placed on the the corresponding holes of the coupling mirror support and four M2 brass screws are inserted and screwed to the point that the washers can only barely move along the screw axis.

The ensemble is now ready to be mounted on the cryostat for further alignment.

Cavity alignment

1. Before mounting the cavity block on the cryostat, a **preliminary waist adjustment** is performed. From the desired cavity length and estimated coupling mirror radius of curvature, a cavity waist estimate is computed. By removing the periscope for injecting the incident laser beam to the cryostat and allowing the beam to freely propagate over the same distance as to the coupling mirror surface, the waist produced by a convex lens with 60 mm focal length can be conveniently monitored with a scanning slit beam profiler (Thorlabs BP series). The injection telescope lenses (see section 5.4) are adjusted in order to produce the desired waist (typically in the 10 μm range) with the lens.

2. After the IVC is cleaned and dust blown away from it, the cavity assembly can be **mounted** on the IVC on the mixing chamber plate. Before placing the cavity, its various electrical cables must be connected to the feedthrough inside the IVC as shown in Fig. D.1. The assembly is then fixed tightly with four M4 brass screws.

3. With the periscope in place, the injected beam is first brought to **normal incidence** on the coupling mirror by tuning the closest mirror mount before the cavity (on the periscope). Once normal incidence is coarsely achieved, the alignment can be further improved by closing a diaphragm at the end of the injection telescope and optimizing the reflection monitored on a photodiode.

4. With the diaphragm opened again and a circular disc with central hole or a diaphragm mounted concentrically on the cage system of the cavity support block, the second-nearest mirror to the cavity (mirror between injection telescope and periscope) is used to **center the incident beam** with respect to the cavity support block. About five iterations of the previous and this step are performed until the incident beam hits the coupling mirror both at its center and normal to the surface.

5. Now the 60 mm-**lens** mounted on a lateral positioner (Thorlabs CXY1) is slid on the cavity support cage system and coarsely centered with respect to the incident beam. The distance of the lens from the coupler is adjusted such that the reflected intensity is maximum. This means that the waist of the incident beam is placed on the coupling mirror surface.

6. The incident laser is now aligned on the **concave feature** of the coupling mirror by laterally scanning the lens position along parallel lines about 50 μm distant from each other, while observing the reflected intensity with the photodiode. Depending on the precise distance of the lens from the coupling substrate and the radius of curvature of the sample, the concave feature produces a dip or peak in the reflected intensity with can be relatively unambiguously identified by its lateral size, which should be close to the size of the concave structure as characterized during fabrication.

7. Once the concave feature is found, the lens position is adjusted to match the incident beam's wavefront to the radius of curvature of the mirror. This is done

by slowly (100-μm steps) advancing the lens towards the cavity while laterally scanning the incident spot over the concave mirror. Progressively, the "V"-like dip should become transformed into a **"W"-like** scan trace on the photodetector, indicating that the reflectivity on the central region of the coupling mirror becomes comparable or larger than on the flat regions of the coupling mirror. It might be necessary to scan the spot position in both lateral directions to find the central spot of the coupling mirror. Once the "W"-feature becomes apparent, the reflectivity of the central peak is maximized with the lens position.

8. We now install an **imaging setup** capable of delivering a magnified image of the back surface of the micropillar. Since the camera does not fit between the 1K and mixing chamber plate of the cryostat, a mirror must be mounted between those plates with the camera and lenses at the side of the cryostat. The imaging system is focused on the back surface of the micropillar. Flashlights are helpful to achieve this alignment.

9. The power of the incident laser is increased to **a few mW**. When the 60-mm lens position is scanned laterally, the small transmitted fraction of the laser should occasionally be visible with the imaging setup when it passes through a hole in the membrane around the micropillar. The position of the transmitted laser on the image is not a trustworthy indication of the actual location of the laser spot on the micropillar sample since the imaging system does not transform the Gaussian laser beam in the same manner as the optical image.

10. A 1 Hz sinusoidal voltage with sufficient amplitude to scan the cavity length by a free spectral range (typically 50 V) is applied to the cavity piezo. The lens position is scanned laterally again until a position is found where the laser spot on the image **flashes** at the frequency of the piezo scan voltage. This indicates a low-finesse resonance of a Fabry-Perot cavity formed between a flat part of the coupling mirror and a flat surface of the micropillar sample, i.e. a spot on the quartz chip, the decoupling shield or on the pillar. With the positions on the image where the flashing appears, a map of the pillar surface can be created and the position of the pillar estimated. It is then straightforward to locate the pillar on the image and position the incident laser beam on it, which should produce a more pronounced flashing feature due to the higher-finesse resonance.

11. By re-locating the concave structure on the coupling mirror with the lens, it becomes clear how the coupling mirror support must be laterally displaced to **superimpose** the micropillar and the concave coupling mirror. While this condition is approached, the flashing feature in the image becomes increasingly brighter and more pulse-like as the finesse of this quasi-mode increases.

12. When both mirrors are superimposed, one starts to see **absorption dips** in the reflected intensity measured by the photodiode. Once this happens, the depth of the dips is optimized by laterally displacing the coupling mirror and lens. At his point, the imaging system is not needed any more and the piezo scan frequency can be increased to $20 - 50$ Hz in order to monitor the reflection spectrum with a faster update rate.

13. Using the axial positioner of the coupling mirror, the cavity length is now **shortened** while the reflection spectrum is monitored and continuously re-optimized with the lateral positioners. With the program "finesse" as described in section 3.2.3, the cavity length can be estimated and the lateral position of the coupling mirror adjusted to optimize the cavity finesse.

14. When the coupling mirror support **touches** the cavity spacer, the axial positioner screw becomes less effective and harder to turn, and residual angles might lead to misalignment of the cavity that is reversible when the cavity length is increased again. At this point, the lateral positioners become ineffective due to the friction between the mechanical parts. A measurement of the cavity length should indicate a length close to the design length. In order to obtain the configuration where the coupling mirror support is pressed against the spacers and the lateral alignment is optimum, several iterations of the approach of the parts may be necessary due to residual angles and hysteresis of the positioners.

15. Once a good configuration is achieved, the four brass **screws** that press the coupling mirror support against the cavity support block can be tightened. Care must be taken in order not to misalign the cavity by asymmetrically tightening the screws. Once the screws are tight, the two screws holding the coupling mirror support to the coupler alignment adapter can be removed. When the axial positioner is turned as to increase the cavity length, no length change should be measured any more since the parts are not held together any more. The positioners can thus be removed from the cage system.

16. While a 60-mm lens was used for the alignment because of constraints imposed by the positioners, an optimal lens can now be placed on the cage system. Shorter focal lengths allow for smaller diameters of the beam crossing the cryostat windows, which results in lower clipping loss. As a compromise between lens aberrations and clipping loss, we generally install a 40 mm lens. With the cavity completely aligned, the incident beam can now be **optimally aligned**:

    - The waist position is optimized by the distance of the coupling lens.
    - The waist size is optimized by moving the injection telescope lenses.
    - The angle of the incident beam is optimized with the last mirror before the cryostat.
    - The position of the incident beam is optimized with the second-last mirror before the cryostat.
    - The incident and reflected beam are rendered parallel to the cryostat axis with the lateral position of the coupling lens, in order to keep clipping loss from the cryostat windows at a minimum. An alignment diaphragm is used for this purpose as discussed in Appendix E.

4.2.4 Electrode actuator

Since crystalline quartz is piezoelectric, it is desirable to understand whether the application of an oscillating electric field near the micropillar results in a significant

displacement. This consideration is interesting because

- it makes using an auxiliary laser to drive mechanical motion unnecessary, and therefore results in less absorption heating and much less experimental complexity,
- it makes actuation with a piezo unnecessary, resulting in less electrical heating of the mixing chamber plate,
- it sheds light on a potential noise source in our system that might limit the experiment's performance.

We start with a crude estimate of piezoelectrically-induced displacement of the compression mode of the micropillar by modeling the pillar as a block of quartz between two laterally placed electrodes. The piezoelectric strain constants $d_{ij} = \epsilon/E$ relate an electric field E applied in one of the three directions $i \in 1, 2\,3$ to the resulting mechanical strain $\epsilon = \delta L/L$ in one of the six "directions" $j \in 1, \ldots, 6$, where the last three values indicate shear strain [LL58]. In the following, the optical axis of the quartz points along the z-direction ($i = 3$, pillar axis), while the piezoelectrically active axis of quartz points in the x-direction ($i = 1$, parallel to the pillar surface). For α-quartz, we find [Bec58] $d_{11} = 2.3\times10^{-12}\,\mathrm{m/V}$, which indicates that a voltage of $U = 1\,\mathrm{V}$ applied on a quartz block along the electrical axis with perfect electrodes results in a deformation of the order of $10^{-12}\,\mathrm{m}$ along that direction. All other piezoelectric strain constants are much smaller and are neglected here. The deformation along the electrical axis leads to a roughly tenfold smaller symmetric deformation along the z-direction according to the Poisson ratio $\sigma = 0.12$ [War92]. Half of this deformation $\delta x = \frac{1}{2}\nu d_{11} U$ is the displacement of the micropillar surface. To convert this deformation to an equivalent radiation-pressure force F, multiplication by the spring constant of the compression mode $k = m_0\Omega_0^2$ leads to

$$\frac{F}{U} = \frac{1}{2}\sigma d_{11} m_0 \Omega_0^2 \approx 2.5\,\mu\mathrm{N/V}\,. \tag{4.2.1}$$

For comparison, integrating the thermal force noise spectrum over one mechanical linewidth for a micropillar compression mode with quality factor $Q = 10^3$ at room temperature yields an rms-force of $F_{\mathrm{th}} \approx 200\,\mathrm{pN}$, which is about four orders of magnitude below the estimated force for 1 V applied to the electrodes. Actuation with an electrode can thus easily dominate thermal noise, especially at low temperature and for higher quality factors. Even the Johnson noise of a 50 Ω resistor at 300 K connected to the modeled electrode results in a force spectral density that equals the thermal noise of a compression mode of $Q = 10^6$ at 0.1 K. We conclude that care must be taken to implement a rather ineffective electrode geometry in order to avoid the introduction of excess displacement noise into the optomechanical system.

In a preliminary experiment, a thin wire ($200\,\mu$m in diameter) tip connected to the output of a network analyzer was approached to the back surface of a micropillar while the displacement of the micropillar's front surface was monitored interferometrically with a network analyzer. The cable ground was attached to the optical table. Care was taken to avoid any piezoelectric substance other than the micropillar sample nearby. A calibration of the interferometer's sensitivity as described in section

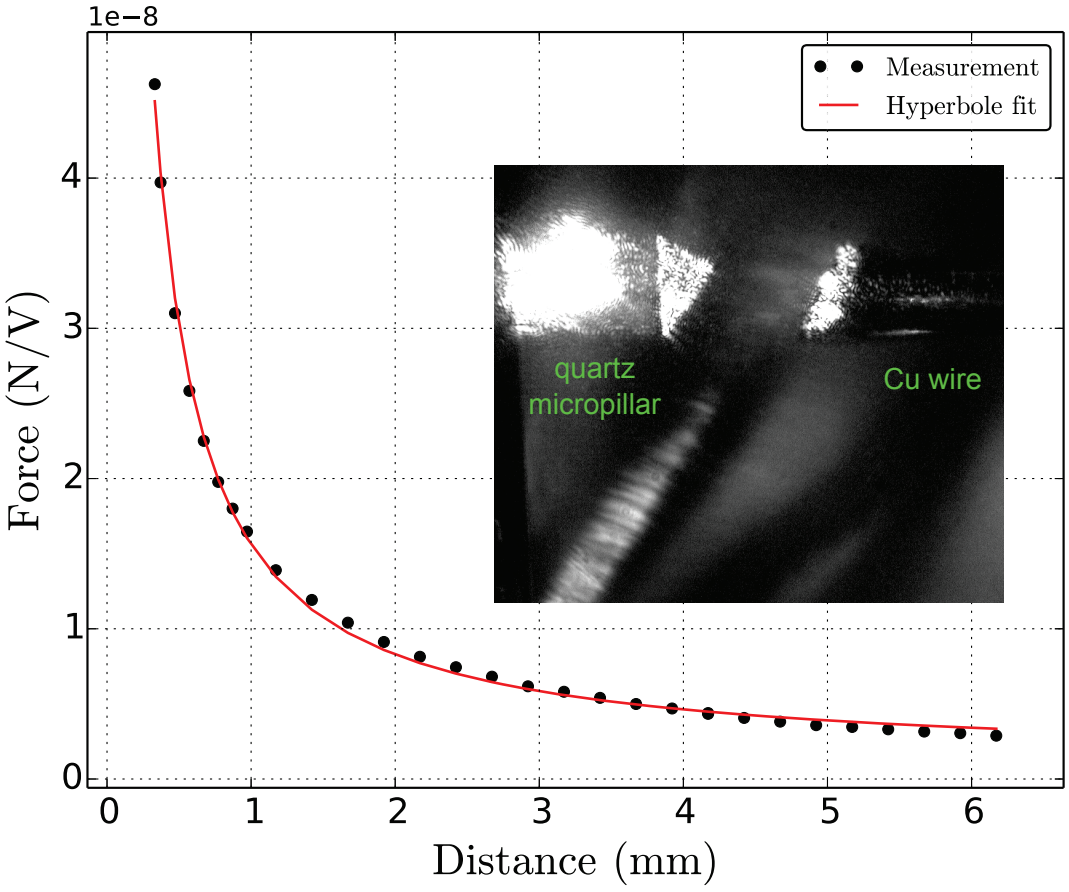


Figure 4.14: Characterization of electric actuation of the micropillar sample 20121002C3. The inset shows a microscope picture taken at an angle with respect to the bare quartz micropillar surface while a laser beam is incident on the other micropillar surface. The transmitted light hits the copper electrode nearby and is scattered into the microscope. The main plot shows the deduced equivalent radiation-pressure force per electrode voltage exerted by the electrode at the pillar's resonance frequency as a function of electrode distance. A hyperbolic fit reproduces well the observed behaviour. The zero of the horizontal axis corresponds to the one of the hyperbolic fit. The measurement at the shortest distance nearly touched the pillar surface. If the wire is approached even further, the resonance frequency increases over the distance of a few microns and the quality factor starts to degrade.

3.1.3 enabled conversion of the network analyzer traces into units of displacement per voltage, and subsequent division by the compression mode's mechanical susceptibility (known from its quality factor, frequency and effective mass) resulted in measurements of the equivalent radiation-pressure force provided by the electrode in units of Newton per Volt. The electrode was mounted on a manual 3-axis micropositioning block and positioned along the z-axis by monitoring the configuration at an angle with an optical microscope.

The maximum response is achieved when the wire is displaced by about 1 mm away from the pillar along the x-axis, which is expected from the model since a perfectly centered wire cannot create an electric field component in the x-direction in the micropillar volume. The measured force as a function of the distance d between a nearly centered electrode tip and the micropillar surface is shown in Fig. 4.14. The obtained force is about two orders of magnitude lower than the theoretical estimate for a perfect electrode geometry and decays as $1/d$. To verify that electrostatic forces, which scale with the square of the electric field, play no role in the actuation mechanism, a bias of up to 200 V was added to the voltage modulation from the network analyzer, and the response recorded as a function of bias voltage. Since no measurable dependence on bias voltage was observed, we conclude that the actuation mechanism is dominated by the piezoelectric effect and thereby a probe of the electric field over the volume of the micropillar. The observed $1/d$-scaling of the electric field has been observed with similar geometries [Arc06].

With a satisfactory preliminary electrode characterization, the same 200-μm thick wire was fed through a copper tube of an outer diameter 1 mm, clamped at one end to the copper tube with a drop of Stycast 2850FT epoxy, and cut about 1 mm behind the end of the tube. This construct was bent and screwed to the back of the cavity support block such that the end of the thin wire was about 1 mm behind and 1 mm to the side of a well-positioned micropillar sample. This way, efficient electrostatic actuation can be carried out while not blocking laser light transmitted by the micropillar. Testing the electrode with the same sample as used for Fig. 4.14, mounted on the cavity support block, yielded an actuation force of 7 nN/V, which would correspond to a distance of about 2.5 mm in the characterization experiment, to be compared to the real distance of about $1-1.5$ mm. Such overestimation of the electrode distance in the measurement cavity setup can be explained by the different ground potential geometry and could be qualitatively reproduced in the test setup by inserting grounded metal rings between the micropillar and electrode.

Chapter 5

Experiment - Optical setup

Nature and nature's laws lay hid in night;
God said "Let Newton be" and all was light.

— Alexander Pope

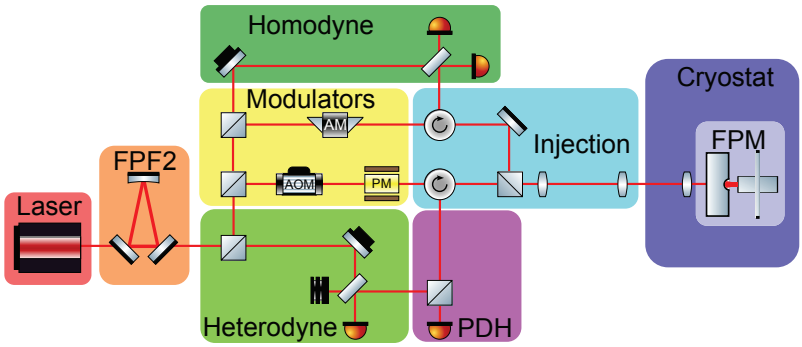


Figure 5.1: Simplied schematic of the optical setup.

This chapter presents the implemented optoelectronics for preparing and measuring the light field that interacts with the optomechanical system. A simplified schematic of the optical setup is shown in Fig. 5.1, while a more detailed schematic can be found in Fig. 5.2. The simplified schematic can be divided into functional groups which provide a natural structure for this chapter: A preliminary characterization of the laser noise (section 5.1) has shown an excess of classical amplitude and phase noise at the mechanical frequency. The laser output is therefore directed through the filter cavity FPF2 (section 5.2) and is subsequently split into two local oscillator beams and two signal beams. The homodyne signal beam passes through the amplitude-modulator AM (section 5.3.1) in order to stabilize its power, to allow for automated measurement series with variable signal beam power, and for driving the mechanical system with radiation-pressure modulations. The heterodyne signal beam is frequency-shifted by one or two acousto-optic modulators AOM (section 5.3.2) such that the homodyne and signal beams can be made simultaneouly resonant with the two polarization eigen-

modes of the measurement cavity. The heterodyne signal beam then passes through a phase-modulator PM (section 5.3.3) that is used to produce sidebands for the later generation of a PDH error signal and to imprint a phase modulation calibration marker on the beam. Each signal beam then passes through a circulator and the two beams are combined with orthogonal polarizations and shaped with a telescope for the injection into the cryostat (section 5.4). The reflected homodyne beam is directed by the circulator to the homodyne detector (section 5.6) while the reflected heterodyne beam is divided into a part that is directly detected to generate the PDH error signal (section 5.8) and another one that goes to the heterodyne detector (section 5.7). We conclude the chapter by discussing how the ensemble of optical instruments is interfaced by a computer in section 5.9.

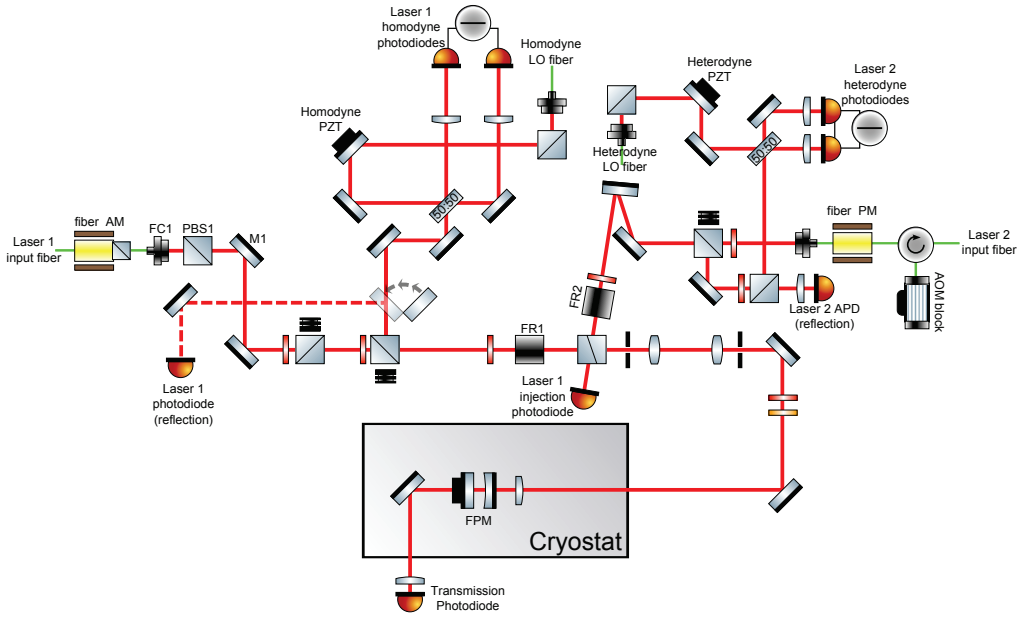


Figure 5.2: Detailed schematics of the optical setup. The four green input fibers come from the filter cavity FPF2 setup. Laser beams are shown in red. The beam-path denoted with "Laser 1" is often termed "homodyne laser", while the beam-path "Laser 2" is also called "heterodyne laser". The "AOM block" retroreflects the incoming light after a double-pass through one or two AOMs. The last mirror before the cryostat represents two vertically superimposed mirrors (injection periscope). Red (yellow) slices crossing the laser beam represent $\lambda/2$ ($\lambda/4$) waveplates.

5.1 Laser source

Due to its superior noise performance and stability, we use a Mephisto Nd:YAG laser from Innolight/Coherent for most experiments presented in this thesis. The output beam has been characterized to have a waist of $w_0 = 197\,\mu$m situated inside the laser head. Directly behind the laser, we place a Faraday isolator which has been measured to achieve an isolation of 39.5 dB and a transmission of 88%. At the maximum pump

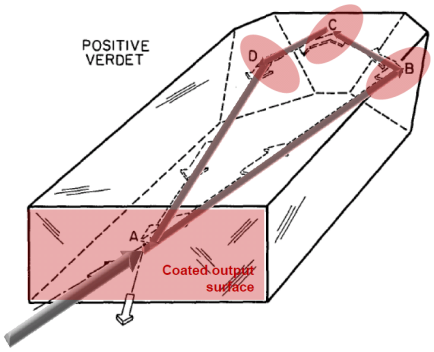


Figure 5.3: Schematic of the nonplanar optical ring resonator inside the Nd:YAG laser. Faces B, C, D reflect the intracavity light by total internal reflection, while face A has a reflective coating to transmit a fraction of the generated laser emission towards the experiment, and to serve as a coupler for the incoming pump light. By geometry, face A has a slightly different transmission for the two possible linear polarizations. Due to the nonplanar arrangement of the beam path, linearly polarized light experiences a slight rotation upon each roundtrip. The Faraday effect is used to partially cancel this rotation for one propagation direction, thereby lifting the degeneracy of the two modes with different transmissions. The lower-transmission mode thereby acquires the highest lasing gain, such that lasing only occurs into this travelling-wave mode. This trick avoids spatial hole burning effects, which is part of the reason for the obtained narrow laser linewidth. Figure from [Coh14].

current of 2.6 A, we obtain 540 mW of optical power behind the isolator at the laser wavelength of 1064 nm. The Nd:YAG crystal is operated at 24 °C. A representation of the laser cavity is shown in Fig. 5.3.

The laser wavelength can be tuned in two ways: through a change of the crystal temperature, which shifts the cavity length and peak gain frequency on a timescale of the order of 1 Hz, and through a piezo that rapidly modulates the cavity length. Figure 5.4 shows the achieved frequency-shifts versus crystal temperature and piezo modulation frequency. The total temperature tuning range of about 30 GHz corresponds to the free spectral range of a Fabry-Perot cavity with a length of 5 mm. The impracticability to construct such long cavities with our optomechanical system explains why we must use piezo actuators to tune the measurement cavity resonance frequency to the laser frequency, instead of the other way around. The transfer function of the piezo actuator in the laser is particularly well-suited for laser stabilization: with a simple RC low-pass filter before the piezo, the phase delay remains $\pi/2$ up to 100 kHz and no major resonance peaks appear in that frequency region. With a few notch filters between 100 and 200 kHz, we can expect to achieve unity-gain bandwidths of the order of 100 kHz with this actuator.

The intensity noise of the laser is characterized in Fig. 5.5. Classical noise in the radio-frequency range is mainly caused by the relaxation peak of the laser at 1 MHz. By comparing the noise spectra of light reflected by a Fabry-Perot cavity for different detunings, it was confirmed that classical noise in the phase quadrature of the laser is of comparable magnitude as in the amplitude quadrature. The Mephisto laser has an integrated noise eater to suppress the effect of the relaxation peak. However, we

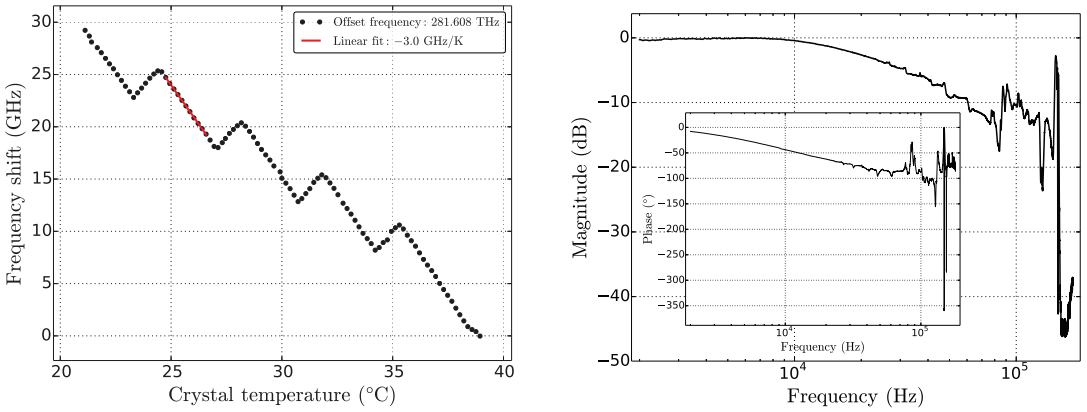


Figure 5.4: Left: Laser frequency tuning through the crystal temperature. The data were measured with a wavemeter. The points on upwards slopes indicate mode-hops over one free spectral range of the laser cavity, where the wave-meter interpolates between two simultaneously emitted laser frequencies. The timescale of frequency modulation through the temperature is of the order of a few seconds. The temperature was tuned through a voltage input with a gain of the order of 1 K/V. Right: Magnitude and phase (inset) of the measured Nd:YAG piezo actuator for frequency tuning. The magnitude was normalized to the DC-response. The response amplitude is about 5 MHz/V with a maximum voltage of 100 V. The effect of an analog low-pass filter before the piezo with a cut-off frequency of about 20 kHz is clearly visible in the plot.

are mainly interested in the noise level at the resonance frequency of our mechanical system. While the noise eater reduces the peak level of the relaxation oscillation at 1 MHz, it also broadens the relaxation peak such that noise at the mechanical frequency of 3.5 MHz becomes larger with the noise eater than without. We therefore operate the laser without the noise eater. From the measurement, we expect classical intensity noise of the laser at the mechanical frequency of 3.5 MHz to equal the shot-noise at a power of the order of 30 μW. Since up to 100 times higher powers will be needed for the measurement, classical laser noise must be suppressed by at least 20 dB with a filter cavity.

5.2 Filter cavity

The noise characterization of the available laser source clearly shows that classical laser noise is about 20 dB above shot-noise at the mechanical frequency and planned maximum power. To reduce the laser noise, we may choose between two common methods:

1. A **noise eater** uses a beam-splitter to reflect part of a laser beam on a photodetector. The measured intensity signal drives an amplitude-modulator situated before the beam-splitter in a negative feedback configuration to stabilize the power measured by the photodetector, thereby reducing the classical laser noise after the modulator. A good description of this method can be found in [Arc06]. The scheme has three potential issues:

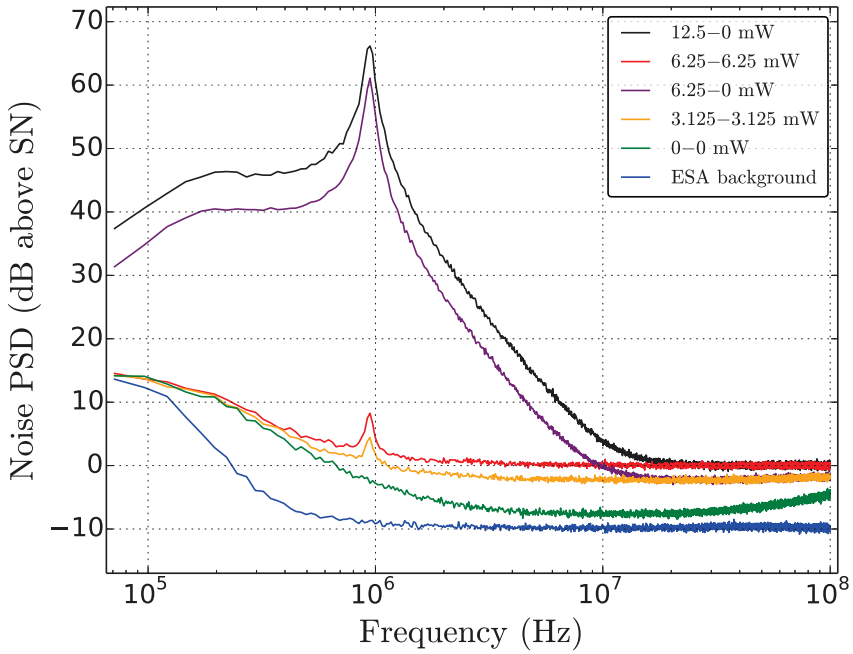


Figure 5.5: Intensity noise of the Mephisto Nd:YAG laser. The level of 0 dB corresponds to the measured shot-noise for 12.5 mW of light (equivalent to a relative intensity noise of -165 dB). For the measurement, the balanced photodetector from section 5.5.2 was used. The labels indicate the power incident on either one of the two photodiodes. When the power is balanced, we observe that classical noise is suppressed by about 60 dB, which allows to observe the shot-noise from $2-100$ MHz. It has been verified that the shot-noise scales linearly with the power. If only one photodiode is illuminated, we observe excess classical noise, which mainly originates from the pronounced relaxation peak of the laser at 1 MHz. This noise scales quadratically with the optical power.

- **High-frequency limit** If the noise-eater should yield a noise reduction by a factor f at frequency Ω, the feedback system composed of detector, servo-amplifier and modulator must have a unity-gain-frequency of $f\Omega$. This makes the scheme challenging to implement above 1 MHz.
- **Power limit** In order to reduce classical noise of the transmitted beam to a fraction f of its shot-noise, the detected power must be f times the one of the transmitted beam. This constraint is both problematic by limiting the available power and by the fact that larger photodiodes are required, potentially limiting the unity-gain-frequency due to their larger capacity.
- **Phase quadrature** Since the phase quadrature of the reflected beam is not measured, phase noise cannot be corrected for. Of course, a second noise eater for the phase quadrature can be implemented, further increasing complexity and decreasing the power yield by a factor of two.

2. The **filter cavity** scheme resonantly transmits the noisy laser beam through a symmetric Fabry-Perot cavity. The transmitted beam is a sample of the intra-cavity field, whose noise spectrum is proportional to the incident field's spectrum multiplied by a Lorentzian with the linewidth of the filter cavity. Despite the benefit of also filtering the phase noise and the mode shape of the transmitted light, filter cavities also have drawbacks:

   - **Low-frequency limit** Noise at frequencies below the cavity linewidth is practically not suppressed. Dimensions of table-top experiments and reasonable values of the optical finesse place a lower limit on filter cavity linewidths at about 10 kHz.
   - **Power limit** As we will see in section 5.2.3, the onset of optical nonlinearity imposes an upper limit of the power handling capability of a filter cavity with a given geometry. With our second design, 300 mW can be transmitted.
   - **Phase noise** Since the frequency reference that the transmitted light is stabilized to is defined by the resonance frequency of the filter cavity, fluctuations of the effective length of the filter cavity are imprinted on the transmitted beam. A stable length and low-noise mirrors must therefore be used.
   - **Lock bandwidth limit** If the incident laser frequency should be modulated at low modulation frequency Ω_{mod} with a large amplitude $\delta\omega$, the filter cavity resonance frequency must follow this modulation. If no feedforward scheme is used, the unity-gain frequency Ω_{ug} of the filter cavity lock must be large compared to Ω_{mod} to yield a small deviation between laser frequency and cavity resonance. For a filter cavity linewidth κ, this means:

     $$\Omega_{\mathrm{ug}} \gg \Omega_{\mathrm{mod}} \frac{\delta\omega}{\kappa} . \tag{5.2.1}$$

     Amplitude modulations of the incident laser are simply low-pass filtered by the cavity. Practically achievable unity-gain frequencies below 100 kHz thus require to place all frequency and amplitude modulators after the filter cavity, with the possibility to introduce additional noise.

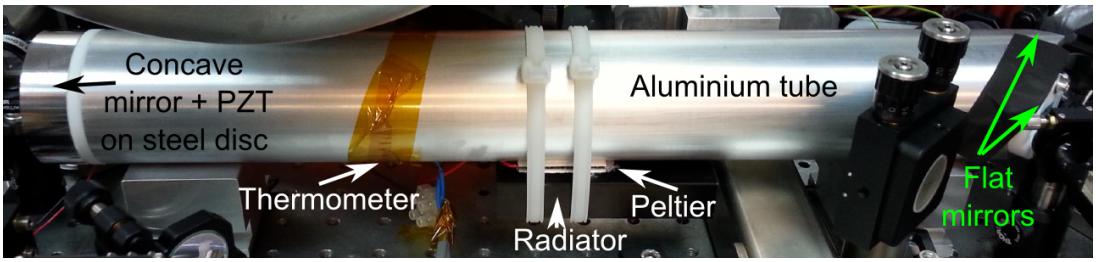


Figure 5.6: Photograph of the filter cavity on the optical table including the surrounding optical elements. The large radiator held with cable binders presses a Peltier element against the aluminum filter cavity tube to thermally tune its length. Left from it, the thermistance used as a sensor for temperature stabilization of the filter cavity tube is scotched to the tube with Kapton tape. Thermally conducting paste is used to improve the thermal link of both the Peltier element and the thermometer. The curved mirror is glued on a piezo mounted on the left end-plate of the tube, while the two planar 45° mirrors are situated at the right end of the cavity tube.

Since our mechanical system resonates at 3.6 MHz, a filter cavity is the only practical solution to reduce the classical laser noise.

5.2.1 Implementation

The filter cavity is constructed as a ring resonator with three mirrors arranged in an isosceles triangle. A concave, highly reflective mirror with a radius of curvature of 1 m is mounted on a piezo at the apex. Along the base of the triangle, two partially transmitting ($T = 475$ ppm each) planar mirrors serve as input- and output couplers at incident angles of nearly 45°. Fig. 5.6 shows how the three mirrors are mounted on a 45-cm long aluminum tube. Between the aluminum tube and the steel disc supporting the concave back mirror, we insert a ring of Teflon which can be slightly compressed with the screws that fix the steel disc on the tube, thereby permitting to perform fine adjustments of the cavity alignment. With a round-trip length $L = 2 \times 404\,\mathrm{mm} + 32\,\mathrm{mm} = 84\,\mathrm{cm}$, the cavity has a free spectral range of $\Omega_{\mathrm{fsr}}/(2\pi) = 357$ MHz. A design finesse of the order of 10^4 thus yields a linewidth of the order of 36 kHz, providing 40 dB of classical noise suppression at the mechanical frequency of 3.6 MHz. The optical and electronic setups around the filter cavity are shown and described in Fig. 5.7.

Fig. 5.8 shows the PDH error signal that is digitally computed in the RedPitaya's FPGA. The digital implementation of the controller allows automated locking sequences and greatly increases the ability of our setup to perform autonomous measurements (see Fig. 5.11). At the moment, we use both the piezo of the filter cavity and the piezo of the Mephisto laser to stabilize the laser-filter cavity frequency difference. The filter cavity is mounted on the optical table near a number of pumping tubes, which results in vibrations of the unstabilized cavity length of the order of its linewidth. The required high locking bandwidth was not achievable with the cavity piezo alone. This can be understood by looking at the measured transfer function of this actuator, as shown in Fig. 5.9. In particular, the two piezo-support resonances

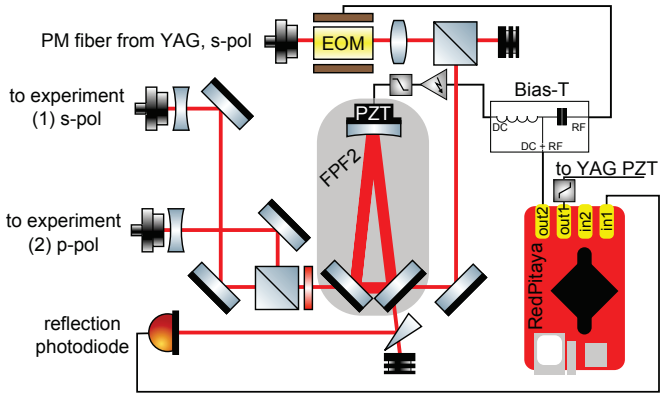


Figure 5.7: Optical and electronic setups around the filter cavity. The laser source is provided through a polarization-maintaining optical fiber in order to keep the setup modular. A resonant 50-MHz electro-optic modulator (EOM) is used to imprint two sidebands one the beam which will be used to create a PDH error signal later on. A subsequent polarizing beamsplitter provides additional polarization-filtering after the EOM. Light is injected into the cavity at an angle near 45 °and a fraction of the reflected light is reflected by a wedge to a photodetector. The photocurrent is digitally demodulated in the RedPitaya with a phase-shifted version of the 50-MHz signal sent to the EOM. After low-pass filtering of the digital signal, the resulting PDH error signal is sent with different gains and cut-off frequencies to both the cavity piezo (low frequencies) and to the faster piezo situated inside the Nd:YAG lase source (not shown here). Passive RC-lowpass filters with cutoffs of 15 Hz before the filter cavity piezo and 1 kHz before the Nd:YAG piezo reduce the electronic noise fed to the piezos and thus drastically improve the cavity frequency stability. The beam transmitted by the cavity is split and injected with orthogonal polarizations into two polarization-maintaining fibers.

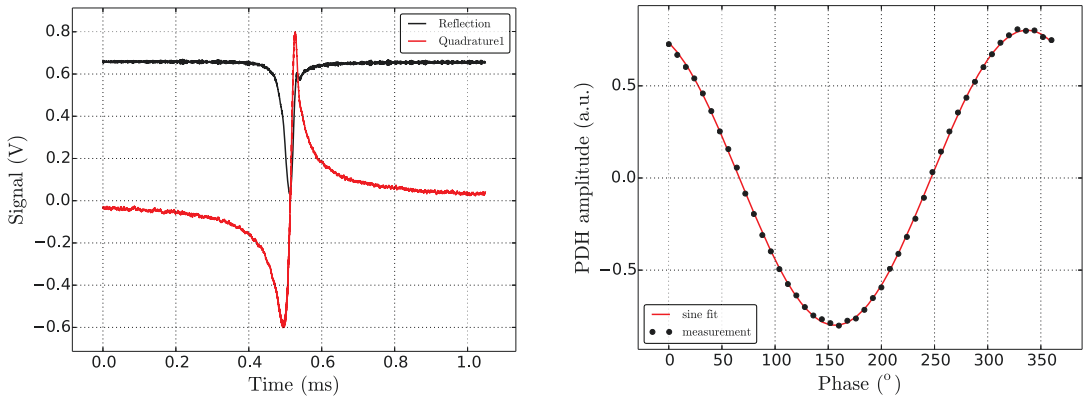


Figure 5.8: Left: Reflection and digital PDH error signal used to lock the filter cavity. The phase-modulation sidebands are too far away to be visible in this plot (the modulation frequency is about one third of the free spectral range). Due to the fast scan velocity compared to the optical linewidth, the onset of optical ringing is visible in the reflection trace, which leads to an asymmetric PDH error signal. The stationary PDH error signal is however symmetric. Right: PDH error signal amplitude as a function of the demodulation phase. Initial versions of PyRPL performed an automatic scan over a range of demodulation phases to find the optimal value. More recent versions find the optimum phase by searching for phase of the maximum (in magnitude) of the complex PDH error signal composed of the two demodulated quadratures, which only requires a single PDH trace acquisition. With the newer version, a PDH lock can be achieved at detunings where a conventional PDH error signal has a slope of zero.

between 20 and 30 kHz clearly visible in the magnitude response plot start to perform self-sustained oscillations if the unity-gain frequency is set above 1 kHz. While this problem was solved by combining the large range of the filter cavity piezo and the high bandwidth of the Nd:YAG piezo, a future update could use the more recently implemented IIR filter that is presented in section 6.2.5 to overcome the lock challenge with the cavity piezo alone. Fig. 5.10 shows the locking sequence that was implemented to achieve autonomous locking of the filter cavity.

5.2.2 Performance

The two planar mirrors have been measured to have transmissions of $T = 475 \pm 5$ ppm, yielding an upper finesse limit of 6,600. From the measured linewidth of $\kappa = 61.2 \pm 3$ kHz (see the fit in Fig. 2.3), we deduce an actual finesse of 5,500[1]. From the finesse, we can deduce a cavity round-trip loss of $P = 190$ ppm, while only 70 ppm were expected from the mirror specifications. The discrepancy can be explained by clipping loss at the edge of one 45°mirror, which is visible with an infrared viewer and can be drastically increased with small changes of the angle of the apex mirror. While a small correction to the angles of the cavity spacer surfaces should be able to get rid of this clipping effect, it was not found necessary for the planned experiments. We

[1]Due to the angular incidence, the p- and s-polarization eigenmodes have different finesses. While we only mention the high-finesse mode (s-pol.) in the main text, the p-polarization mode has been measured to have a finesse of 286 ± 30 and is often practical for alignment or debugging purposes.

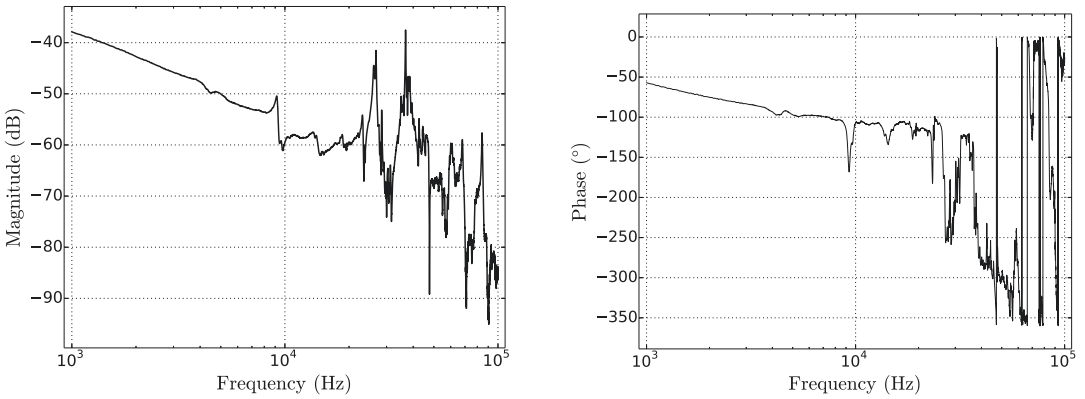


Figure 5.9: Magnitude (left) and phase (right) of the FPF2 piezo transfer function with an analog 15 Hz low-pass filter inserted before the piezo.

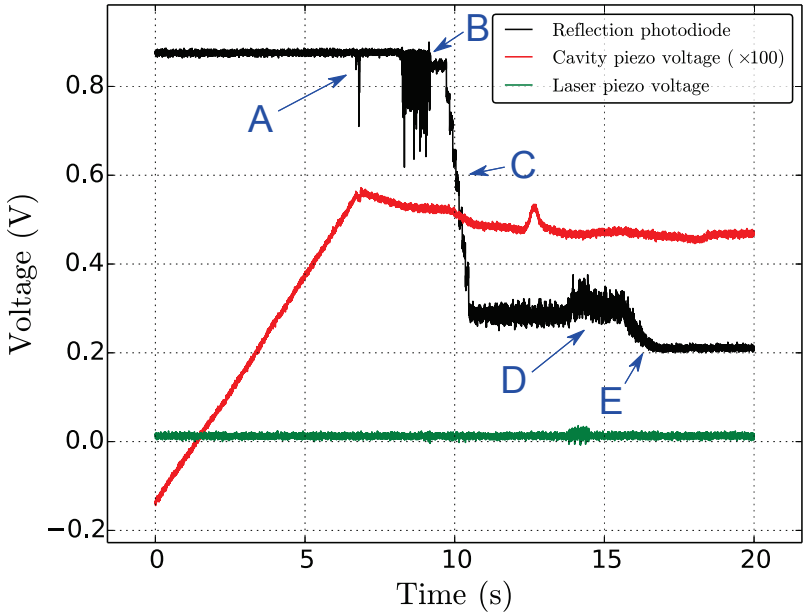


Figure 5.10: Automatic locking sequence of the filter cavity. At the beginning, the cavity piezo voltage is ramped up while an oscilloscope is triggered on the reflection signal with a threshold slightly below the off-resonant reflection. This way, the passage of a resonance (A) is detected and the voltage ramp is stopped. The low-pass filter before the slow piezo ensures that the piezo voltage is above the resonant voltage setting shortly after the passage. An integrator feeding the reflection signal to the cavity piezo is engaged with a setpoint such that the cavity length slowly drifts towards resonance and stays beside it (B). Once this is detected with the reflection signal, the integrator gain is optimized and the setpoint ramped towards the resonance (C). The error signal for the lock is then switched from the reflection to the PDH signal (D), and the setpoint is again ramped to the final setting at resonance (E). Due to the narrow linewidth of the filter cavity resonance, the effective gain of the laser piezo, which provides a frequency-modulation, is extremely high, which explains the low amplitude of the laser piezo voltage modulations.

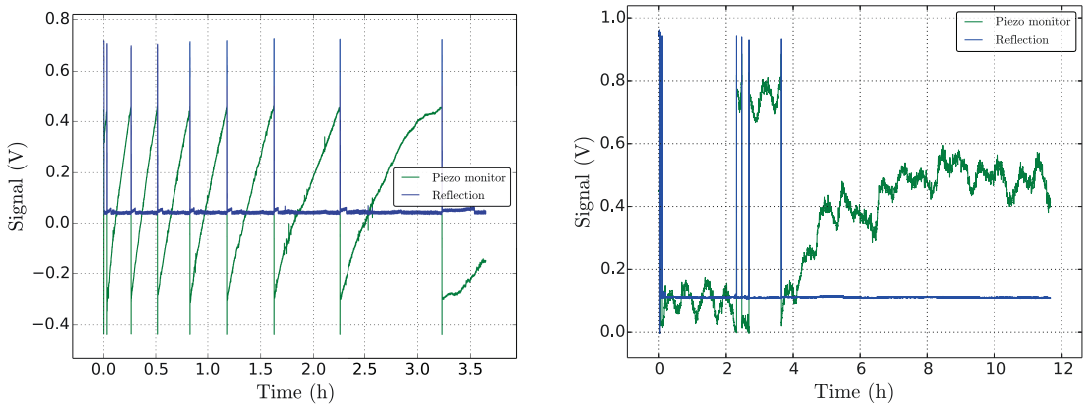


Figure 5.11: Long-term stability of the filter cavity lock. The left plot shows the behavior without, the right plot with temperature-stabilization of the cavity spacer. The time scales of the two plots are different. A reflection near unity corresponds to an unlocked cavity, while a reflection near 0.1 corresponds to a locked cavity. The nearly vertical lines in the piezo voltage correspond to the automatic resonance search. In both plots, the cavity automatically re-locks within seconds, and only loses lock due to the saturation of the actuator range.

measure a minimum resonant reflectivity of $R_0 \approx 4\%$ which is very close to the value of 2.8% expected from the above stated transmission and loss, thus indicating good alignment and proper input mode shape. The clean input mode shape is a result of the spatial filtering provided by the fiber that delivers the input beam from the laser. When locked at resonance, we find that 46.5% of the incident power is transmitted by the cavity. The theoretical expectation of $4T^2/(2T + P)^2 \approx 69\%$ is not attained, probably again due to clipping loss. We have injected up to 300 mW into the cavity without observing any nonlinearities.

In order to make optimal use of the 14-bit resolution of the digital-to-analog converters of the Red Pitaya, the gain of the high-voltage amplifier that drives the cavity piezo was adjusted such that the maximum output voltage swing of the Red Pitaya only covers about 1.5 free spectral ranges of the cavity. While this provides excellent short-term performance, temperature drifts of the laboratory by several °C over a few hours caused length drifts of the long aluminum spacer and required several times per hour to re-lock the cavity to another resonance. Consequently, the spacer temperature was stabilized by a temperature controller connected to a thermometer and Peltier element at the center of the cavity spacer (see Fig. 5.6), which increased the typical duration of re-lock-free filter cavity operation to several hours. The obtained performance is shown in Fig. 5.11.

Using the homodyne detection (section 5.6), it has been confirmed that classical amplitude and phase noise at powers of 1 mW are below 10% of the shot-noise at the frequency of 3.5 MHz.

5.2.3 Raman lasing in FPF1

At an earlier stage of the experiment, a shorter, linear filter cavity had been constructed with two mirrors of transmission $T = 100$ ppm. The parameters of the old filter cavity (FPF1) are compared with the ones for the finally implemented cavity from the previous section (FPF2) in Table 5.1. Due to the higher finesse and larger beam diameter in FPF1, the maximum intensity on the cavity mirrors was almost three times as high as in FPF2. In this power range, we have observed that the filter cavity transmitted light at an optical frequency that was absent from the injected light. This was observed in two ways: by sending the transmitted light of the filter cavity to another cavity (analysis cavity) with a much larger free spectral range, additional absorption dips were present in the absorption spectrum. We have checked that these absorption dips were not due to higher-order transverse modes of the analysis cavity and that their spatial mode closely matched the fundamental Gaussian mode. By varying the pump current of the laser source, it was furthermore confirmed that no optical feedback to the laser could explain the observed behavior. Fig. 5.12 shows that the intensity of these absorption dips relative to the carrier was observed to follow a threshold-like behavior as expected from the onset of lasing action. Measurements at higher power were impossible due to difficulties in maintaining the filter cavity locked in this regime. Independently, the right plot of Fig. 5.12 shows the transmission spectrum of the filter cavity measured with an optical spectrum analyzer. An additional peak shifted by about 8.4 nm towards longer wavelength appears in the spectrum if the filter cavity is operated above threshold. Due to the low resolution of the optical spectrum analyzer, this measurement is rather imprecise. The corresponding frequency-shift of 2.2 THz can be measured more precisely by combining it with a different measurement method. The analysis cavity reflection spectrum enables to measure the difference frequency modulo its free spectral range (1.5 THz in this case) by the method described in section 3.2.5 with a relative error around 5%. From the sequence of possible difference frequencies obtained by this measurement, the value closest to 2.2 THz is selected. By this method, we find a difference frequency of $\Delta\omega/(2\pi) = 1.9 \pm 0.13$ THz.

The observed frequency lies at the lower end of typical frequencies of molecular vibrations. In collaboration with Johan Biscaras, we have obtained Raman spectra of similar dielectric mirror coatings as the ones in the FPF1 with the equipment of the IMPMC, Paris, to independently confirm our observation. Raman spectra are usually expressed in units of wavenumber $\Delta\lambda^{-1} = 1/\lambda_1 - 1/\lambda_2 = \Delta\omega/c$. The above deduced difference frequency is equivalent to a wavenumber of $\Delta\lambda^{-1} = 62.7 \pm 4.3\,\mathrm{cm}^{-1}$. The

Cavity	Length (cm)	Finesse	Waist	Injected power	Max. intensity
FPF1	36.0 cm	29000	320 μm	51 mW	212 kW/cm^2
FPF2	84.0 cm	5500	535 μm	250 mW	75 kW/cm^2

Table 5.1: Comparison of the parameters of the two constructed filter cavities FPF1 and FPF2. The length indicates the round-trip length of the cavity. FPF1 starts Raman lasing at the stated injected power, while no such effects were observed with FPF2. The lower maximum intensity on the mirror surface, primarily due to the larger waist size in FPF2, can explain this difference.

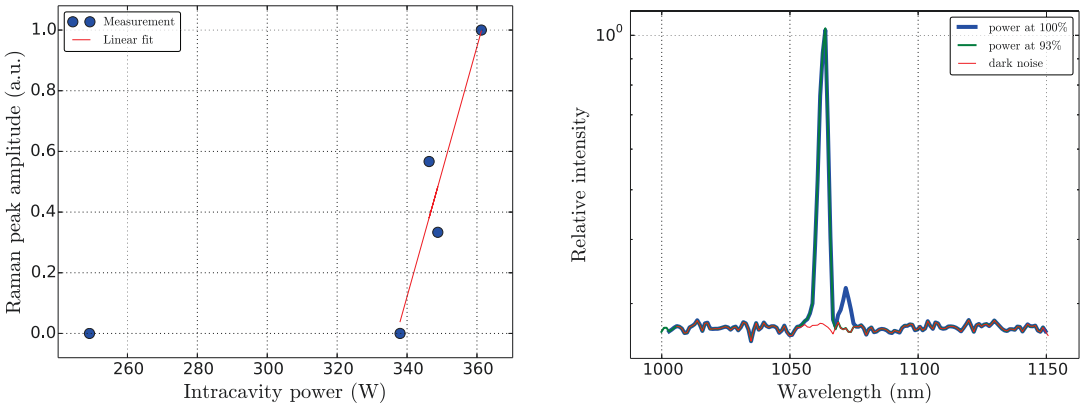


Figure 5.12: Left: Amplitude of the Raman sideband transmitted by the filter cavity versus estimated intracavity power. A strongly nonlinear power dependence is observed with a threshold for the onset of Raman lasing at about 340 W. Right: Transmission spectrum of the filter cavity slightly below (93%) and slightly above (100%) Raman lasing threshold. A sideband at lower wavelength is observed above the thresold. The poor resolution of the employed spectrometer only allows a crude estimate of the wavelength separation of $8.4 \pm 2.0\,\mathrm{nm}$.

spectra in Fig. 5.13 feature a pronounced peak at $58 \pm 1\,\mathrm{cm}^{-1}$, which lies within the error range of the frequency estimated with the analysis cavity.

We therefore interpret the observed nonlinearity as Raman lasing in the filter cavity. With a free spectral range of 830 MHz $\ll$ 1.9 THz, the lasing wavelength is essentially defined by the Raman gain of the mirror coatings. From Table 5.1, we find that Raman lasing occurs for intracavity intensities of the order of 200 kW/cm^2. Since we expect the lasing threshold also to be proportional to the final density of states, the product of peak intracavity intensity and cavity finesse should be a valid constant to compare lasing thresholds among different setups. This value amounts to 6.1 MW/cm^2 for FPF1. While the parameters of FPF2 are far away from this threshold, we find for our measurement cavity ($\mathcal{F} = 10^5$, $w_0 = 10\,\mu$m) a threshold power of the order of 30 μW. While instabilities were observed in this power region (section 6.3), these could be explained as consequences of low-frequency ($<$ 100 kHz) mechanical modes. Due to a cavity length of the order of 100 μm, the free spectral range of the measurement cavity of 1.5 THz leads to a strong suppression of Raman gain at 1.9 THz. Furthermore, higher-quality mirror coatings, which might have a different Raman spectrum than the coatings of the mirrors of FPF1, are employed for this cavity.

5.3 Modulators

The filter cavity setup in Fig. 5.7 has two fiber outputs. Fig. 5.2 shows that the remaining optical setup has four fiber inputs. While the the first fiber from the filter cavity setup directly feeds about 10 mW to the amplitude modulator described in section 5.3.1, the second fiber is connected to two 50:50 fiber splitters to generate the remaining three fiber outputs. With a typical setting of the waveplates in the setup, 46 mW are directed to the homodyne local oscillator beam, 20 mW to the heterodyne

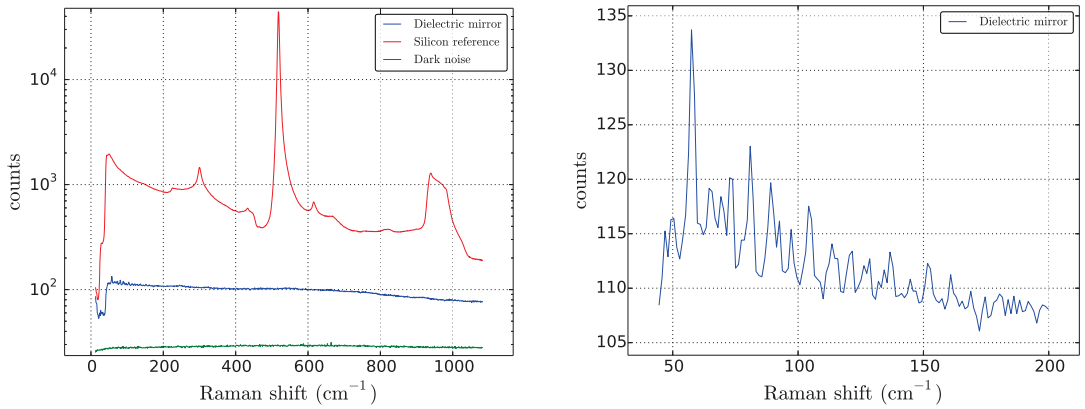


Figure 5.13: Raman spectra of a dielectric mirror of similar type as the one used in the filter cavity. The left plot also shows spectra for a silicon wafer and the spectrometer's dark noise as a reference. The right plot is a zoom of the left spectrum for small wavenumbers. A peak at $\Delta\lambda = 58\,\mathrm{cm}^{-1}$ is clearly visible.

local oscillator, and an equal power of 20 mW is directed towards a circulator that passes the laser through two acousto-optical modulators (section 5.3.2) before it is fed to the fiber phase modulator described in section 5.3.3.

5.3.1 Non-resonant fiber amplitude-modulator (AM)

One goal of the experiment presented in this thesis is to study the trade-off between the noise in the estimation of a mechanical displacement and the back-action force that arises from such a measurement. In a quantum-limited measurement, the ratio between the spectral densities of these two noise sources is proportional to the square of the optical power that is employed for the measurement. Since it is expected that many measurement series with variable power will be performed, an automatized tuning of the power of the measurement beam is therefore desirable and easiest accomplished by passing the measurement beam through an amplitude-modulator before it interacts with the optomechanical system. The availability of an amplitude-modulator has further advantages: it can be used to stabilize the power of the measurement beam in order to suppress the power fluctuations of the laser source which are of the order of 2%, and it may be used to actuate the mechanical system by a radiation-pressure modulation. In the presented setup, the modulator model NIR-MX-LN-10-P-P-FA-FA-30dB from Photline Technologies in Besançon, France, is employed. The fiber modulator is specified to be able to handle more than 100 mW of incident power, has a 3-dB bandwidth of at least 10 GHz and both a high-impedance voltage input for low-frequency modulations and a 50-Ω input for RF signals. Since the modulator is implemented as a Mach-Zehnder interferometer with electro-optic modulators in both arms, the transmission depends sinusoidally on the applied voltage difference. The period of these sinusoids is of the order of 8 V. If the modulator is operated at the maximum slope of the sinusoid with the given parameters, Johnson noise at the 50-Ω input results in a classical intensity noise that equals the shot-noise of a 120-mW laser

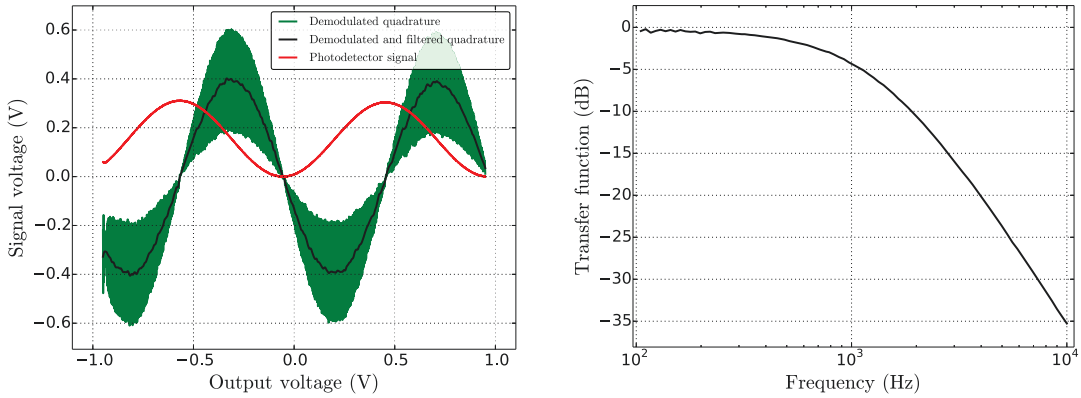


Figure 5.14: Left: Signals used to stabilize the injected power with an amplitude modulator. The red trace shows a voltage proportional to the power measured by the injection photodiode for a sweep of the modulator bias voltage. Since the modulator is constructed as a Mach-Zehnder interferometer, the power has a sinusoidal dependence on the voltage. In order to stabilize the power at a value of zero, a modulation at 10 kHz is fed to the bias voltage port of the modulator and digitally demodulated to generate the derivative of the power-versus-voltage curve. A low-passed version of this signal is fed to a PID which can stabilize the power at a value of zero. Right: Transfer function of the injection photodiode and DC-port or the amplitude modulator. A second-order low-pass filter is implemented to minimize effects of the power stabilization feedback loop at frequencies in the measurement band.

beam. Equivalently, for a 1.2-mW laser beam, shot-noise is equivalent to a white electronic noise source 20 dB above Johnson noise. Only electronic devices with extremely low noise should therefore be connected to the RF input, and drive signals should be attenuated as possible.

Fig. 5.14 shows the transmission of the modulator versus the output voltage of a RedPitaya that drives the DC port of the modulator through an amplification stage and an analog low-pass filter. The right plot shows the realized transfer function. A cut-off frequency of the order of 1 kHz was implemented in order to reduce the potential introduction of intensity noise through the DC-port. A RedPitaya is easily implemented to lock the transmission to the desired setpoint. However, this scheme fails if the setpoint is close to a minimum or maximum of the power transmission. To solve this problem, we imprint a weak amplitude-modulation (in our case at 10 kHz) on the optical beam and demodulate the detected signal. The result is proportional to the derivative of the transmission curve and thereby the modulator is easily locked to a transmission close to zero (see Fig. 5.14). In order to inject powers of the order of 100 μW into our cryogenic measurement cavity, a careful locking sequence is required to avoid instabilities: First, the amplitude modulator is locked to zero transmission. Next, the measurement cavity is locked with another laser beam. Once the cavity lock is stable, the setpoint of the amplitude modulator transmission is slowly ramped up to the desired value, while possibly adjusting feedback gains as a function of the injected power. As soon as it is possible to use the DC-transmission of the modulator as the error signal for the power stabilization feedback loop, the 10 kHz-modulation is turned off in order to avoid to drive the optomechanical system with a radiation-pressure force

at this frequency.

5.3.2 Acousto-optic modulator (AOM)

Acousto-optic modulators are optomechanical devices that have been commercially available for many years. Light passing through a crystal is diffracted by a grating that results from a spatial refractive-index modulation accomplished by driving standing or travelling ultrasonic waves in the crystal with a piezo-electric actuator. Since the ultrasonic wave must propagate nearly perpendicularly to the propagation direction of the laser beam to satisfy the Bragg condition, the wavevector of the diffracted light picks up a component in the transverse direction. Energy conservation requires that its frequency ω_{L} is Doppler-shifted by the frequency of the ultrasonic wave Ω_{AOM}. Thereby, an acousto-optic modulator simultaneously achieves three tasks:

- it frequency-shifts the light to $\omega_{\mathrm{L}} \pm \Omega_{\mathrm{AOM}}$, where the sign depends on whether the diffraction order +1 or -1 is used,
- it modulates the power of the diffracted beam in proportion to the RF-power of the drive tone,
- it deflects the propagation direction of diffracted light by an angle proportional to the drive frequency.

While the first two effects are desired in our setup to generate a frequency-shifted laser beam, the third effect would lead to a modulation of the mode-matching of the laser beam with respect to the measurement cavity mode. We therefore use the AOM in a double-pass configuration: the transmitted beam is retro-reflected by a mirror and passes a second time through the modulator. While this trick doubles the frequency-shift, it nullifies the angular deflection and squares the diffraction efficiency[2].

The frequency-shift accomplished by the AOM should correspond to the birefringence splitting of the measurement cavity. This way, the frequency-shifted laser can be locked to the cavity simultaneously with an unshifted beam of orthogonal polarization. From Table 3.5, we see that birefringence splittings of our measurement cavities range rom 40 to 450 MHz. We therefore seek to implement AOMs with drive frequencies from 20 to 225 MHz. Because commercial AOMs only achieve frequency modulation ranges of at most one third of the center frequency shift, we use two AOM's with center frequencies of 200 MHz and diffraction orders of +1 and -1 in series to subtract their frequency-shifts, which enables to cover AOM frequency differences from 0 to 100 MHz. For higher frequency shifts, only one AOM is used. Figure 5.15 shows the implemented optical setup with two modulators of model MT200-A0,4-1064 purchased from AA OPTO-ELECTROMIC from Orsay, France, that are optimized for drive frequencies from 150 to 250 MHz. Our setup employs a fiber-based optical circulator such that the beam reflected from the AOM setup can be re-injected into the free-space-fiber coupler in the opposite direction. This implementation is very helpful to quickly achieve proper alignment of the double-pass scheme since an optimization of

[2]Since the diffraction efficiency $\eta \in]0, 1[$, the double-pass efficiency is smaller the single-pass value: $\eta^2 < \eta$.

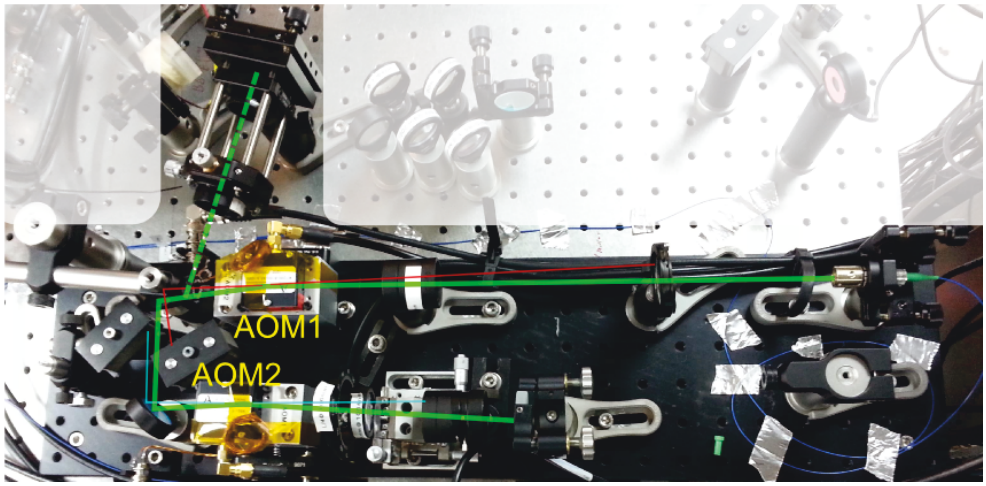


Figure 5.15: Optical setup of the acousto-optical modulator (AOM) block. Laser light is injected from the fiber coupler on the left. Green lines indicate the path of the light through both AOMs (full line) or optionally only through the first (dashed line) AOM. At the end of either path, the laser beam is retro-reflected to the fiber coupler and separated from the incident beam by a fiber-coupled optical circulator. While AOM1 uses the diffraction order +1 which results in a laser frequency equal to the sum of incident optical and AOM drive frequency, AOM2 uses order -1 to generate the difference frequency. The non-diffracted beams are indicated by thin red (lower frequency relative to the diffracted beam) or blue (higher frequency) lines, and are directed to beam blocks. Proper alignment of the double-pass configuration is achieved by mounting the lens behind AOM2 on a three-axis manual translation stage and optimizing its position while monitoring the AOM block transmission for an AOM drive frequency sweep.

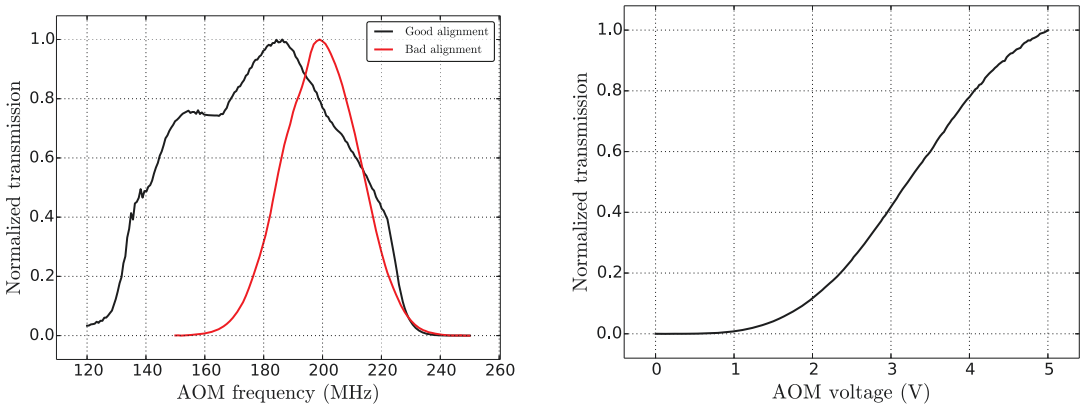


Figure 5.16: Left: Transmission of the AOM block versus drive frequency and voltage. The frequency transmission curve can be tuned with the alignment of the setup. The working point for the feedforward stabilization is typically set to 0.2. Right: AOM block transmission versus intensity modulation input voltage of one AOM driver. The function is clearly non-linear.

the transmission of the combination of the AOM block and the circulator generally results in a very good overlap between the beams crossing the modulators in both directions. Furthermore, it allows the light to pass through the modulators in both directions with the same polarization, which was found to enable a 10% increase in diffraction efficiency in our setup. One modulator is driven by the amplified output from a voltage-controlled oscillator (VCO) that allows to analogically modulate the frequency-shift, while the second one is driven by the amplified signal from direct digital synthesis (DDS) with 31 bits of frequency resolution. Fig. 5.16 shows the achieved transmission of the AOM block for variable DDS drive frequencies. The two plots show that the useful range of modulation frequencies depends strongly on the alignment of the setup. Even for an optimal alignment, the frequency dependence is not flat because resonant analog electronic circuits are used to enhance the drive signal in the AOMs. In order to achieve a flat frequency response with our setup, we record the dependence of AOM block transmission versus drive voltage amplitude and frequency for the final alignment of the setup and adjust the drive voltage at each frequency setting as to result in a constant transmission. We usually set the working point of this feed-forward scheme to a relative transmission of 0.2 and achieve stability of the order of 5% over a frequency span of 100 MHz. Of course, this scheme is ineffective for fast modulations of the AOM frequency. Since we only change the AOM frequency during locking sequences or in measurement series that involve a scan of the difference frequency between two laser beams, this poses no problems.

The drive frequency generated by the VCO was found to perform slow temperature-drifts, which is reasonable for an analogically driven device. Since our heterodyne detection (see section 5.7) requires a stable beatnote frequency, we electronically mix the drive voltages of the two AOM's and use a differential counter of zero-crossings in the RedPitaya to derive an error signal proportional to the integral of the difference frequency between a set frequency reference and the difference frequency of the AOM drivers. Fig. 5.17 shows how feeding this error signal with an appropriate gain to the

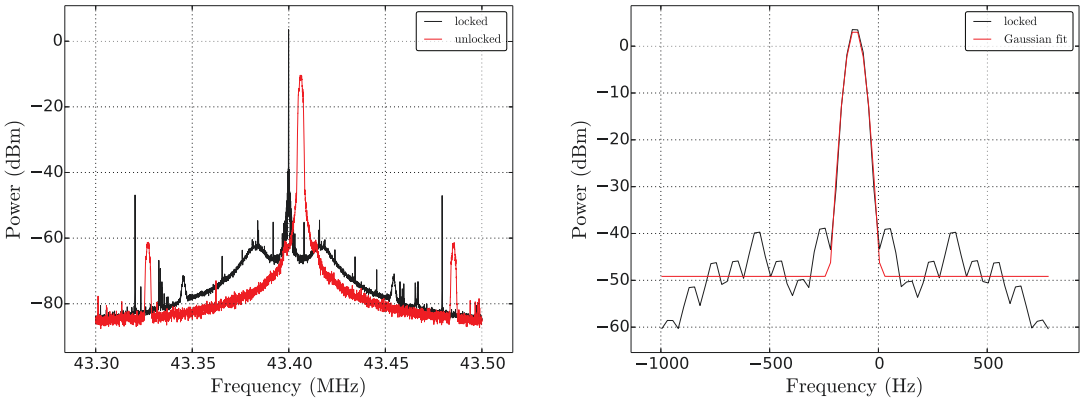


Figure 5.17: AOM difference frequency lock performance with the simplistic frequency detector implemented in the FPGA. Right: Zoom on the peak of the black trace in the left plot. The abscicssa shows the difference frequency from the setpoint of 43.4 MHz. The deviation of 107.5 Hz corresponds to a relative frequency difference of 2.5 ppm, which is to be expected for the quartz in the RedPitaya. A feedforward scheme to correct the frequency difference has been implemented. Work towards synchronizing the FPGA clock with the spectrum analyzer one is in progress.

VCO input results in a good difference frequency lock between the two AOM drivers.

We conclude by noting one problem with our implementation of the AOM block. The employed fiber circulator was found to cause a cross-talk of about 1% of the incident to the transmitted light. This cross-talk was found not to arise from the fiber coupler or end faces of any fiber connectors and is therefore attributed to the circulator implementation. Since passing twice through two AOM's and a fiber coupler leads to an overall transmission efficiency of at best 10%, the residual cross-talk is comparable to the signal beam and causes a beatnote at the difference frequency of the two laser lines. Since the polarization of the parasitic beam is orthogonal to the resonant mode of the filter cavity, the only problem caused by this beatnote is a parasitic modulation of the PDH error signal if it lies close to a harmonic of the PDH modulation frequency.

5.3.3 Non-resonant fiber phase-modulator (PM)

The principal electro-optic phase-modulator used to generate PDH sidebands and to calibrate spectra in this thesis is the model NIR-MPX-LN-10-P-P-FA-FA from Photline Technologies, Besançon, France. The modulator comes with two FC/APC polarization-maintaining fiber connectors and has an integrated polarizer such that only the light polarization that experiences optimal phase-modulation is transmitted. With properly adjusted incident polarization, we have measured a power transmission of 57 % for our modulator. The specified minimum 3 dB bandwidth is 10 GHz and the modulator can handle at least 100 mW of injected optical power.

A phase-modulator imprints a phase-shift $\phi(t)$ on the transmitted light that is proportional to the applied voltage $V(t)$:

$$\phi(t) = \frac{\pi}{V_\pi} V(t)\,. \tag{5.3.1}$$

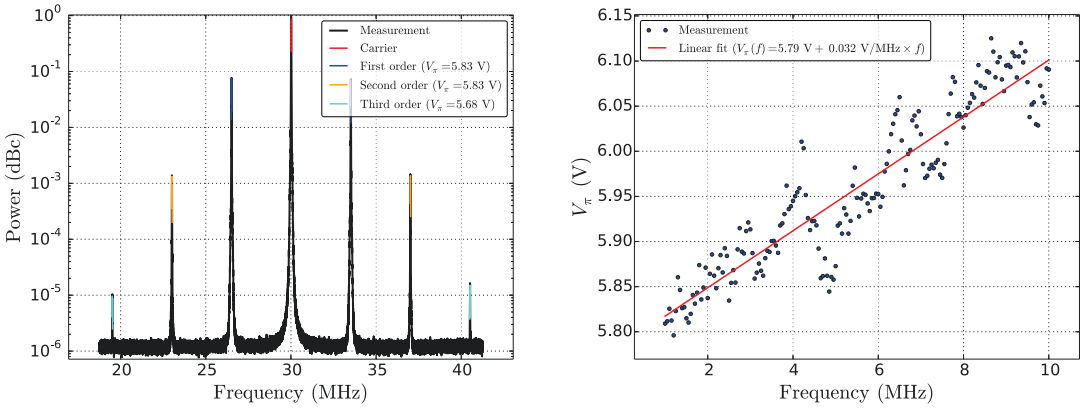


Figure 5.18: Left: Heterodyne spectrum of a phase-modulated beam. The central carrier at 30 MHz and symmetric modulation sidebands up to the third order are visible. By comparing the sideband and carrier peak power, the phase-modulator's V_π can be calibrated (see main text). Right: Measured frequency-dependence of V_π for an electrical modulation power of 0 dBm. The values are extracted from a spectrum similar as the one in the left plot which is acquired in the max-hold mode of the spectrum analyzer while the modulation frequency is swept over the full frequency range. The obtained data are reasonably-well approximated by a linear relation.

While the proportionality constant V_π does not depend on frequency or voltage amplitude for an ideal modulator, such linear and frequency-independent modulator performance can only be an approximation in practice, which must be confirmed experimentally.

We have calibrated our modulator by performing heterodyne spectroscopy of a phase-modulated beam. To this end, light transmitted by the phase-modulator is brought to interference with a frequency-shifted reference beam. The reference beam is created from the same laser source by passing it through two acousto-optic modulators in double-pass configuration (see 5.3.2). The interfering beams are detected with the avalanche photodiode (see 5.5.1) which has a bandwidth of 50 MHz. For a difference frequency of roughly 30 MHz, a sinusoidal voltage of amplitude 1.0 V and frequency 3.5 MHz, the measured spectrum is shown in Fig. 5.18. We see the central peak that corresponds to the beatnote of the two lasers at the difference frequency of 30 MHz and three subsequently smaller symmetric sidebands at the right and left of the carrier. As explained in section 2.2.2, the ratio of the amplitudes of the n-th sideband and the carrier is given by the ratio of the Bessel functions $J_n(\phi)/J_0(\phi)$. All peaks are locally fitted by Gaussians in linear scale to extract their amplitude, and sideband amplitudes are extracted as the mean value of upper and lower peak power in order to correct for any linear slope in the detector response[3]. From the deduced modulation phase and known modulation voltage we extract the calibration constant V_π. We see from the fit that the extracted values $V_\pi = 5.83\,\mathrm{V}$ for the first two sideband orders agree to better than 0.2%, while the third order leads to an underestimation by 3% that is likely caused by the lower signal-to-noise ratio in the peak estimation. The good agreement

[3]We observe asymmetries of the order of 3% in this frequency range.

between the first- and second-order values suggests that the detector response is not a major source of error up to modulation frequencies of 7 MHz. By varying the optical power of the signal laser beam by one order of magnitude, the nonlinearity of the photodetector was confirmed to lead to less than 1% variation in the estimated V_π. Changing the electrical modulation power from -30 dBm to +10 dBm resulted in about 2% variation. By repeating the experiment for different frequencies, up to 10% lower V_π at lower frequencies were found (e.g. 5.2 V at 500 kHz). This tendency is in agreement with the manufacturer's specification of a V_π of $4-5$ V at 50 kHz. As shown in Fig. 5.18, the frequency-dependence of V_π was measured between 1 and 10 MHz and is best expressed with the linear dependence

$$V_\pi(f) = 5.79\,\mathrm{V} + f \times 32\,\mathrm{mV/MHz} \tag{5.3.2}$$

and a relative error of 2% to account for the systematic variation of V_π. The measured V_π was cross-checked and found to agree to within 20% with an estimation through transduction of a phase modulation on the slope of a PDH error signal. We expect however larger uncertainties from this method due to the large amount of parameters that must be measured in order to calibrate the phase modulator (cavity linewidth, error signal amplitude, amplifier transfer functions). We will therefore employ the above presented calibration in the rest of this thesis.

5.4 Cryostat injection optics

To combine the two orthogonally-polarized signal beams in Fig. 5.1, a Glan-Taylor prism is used for optimal polarization separation. After this, the beams cross a 50-cm long cage system with two lenses which can be slid along the optical axis. This telescope configuration allows to prepare a vast range of optical waist sizes and positions at the mixing chamber, and is used to re-optimize the optical alignment when the cryostat is cooled down. To obtain good alignment of the laser beams to the cage system axis such that sliding the lenses along the axis does not lead to misalignment, a diaphragm is mounted at either end. Behind the telescope, a pair of $\lambda/2$ and $\lambda/4$ waveplates are placed, which allow to adjust the polarization of the incident beams to one of the cavity eigenmodes. After the waveplates, a periscope lifts the laser beams by about 25 cm to the axis of the cryostat and allows alignment to the measurement cavity with a last mirror mount. Due to a lack of space available between the periscope and the cryostat window, the two waveplates could not be placed after the periscope.

5.5 Photodetectors

While a number of standard photodetectors are used in the setup as shown in Fig. 5.2 for the purpose of diagnostics, two classes of photodetectors in our setup fulfill the important role of performing the linearized quantum measurement described by Eq. 2.2.17. While we will only briefly discuss the commercially available avalanche photodetectors used for measurements with weak power in section 5.5.1, we perform the more important measurements at higher power with the diffential photodiode design described in section 5.5.2.

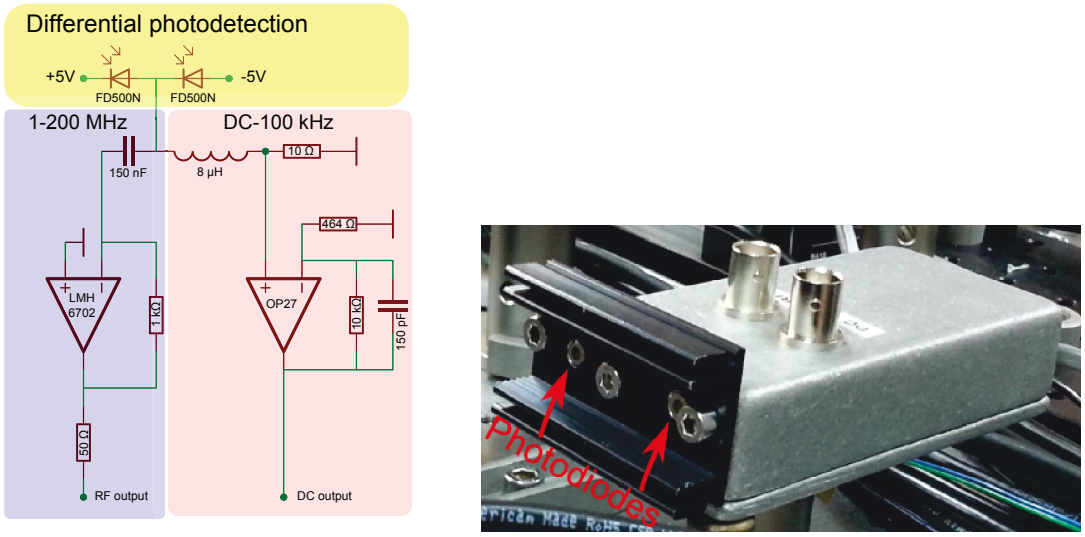


Figure 5.19: Left: Simplified schematics of the differential photodetector. A more detailed schematic can be found in Appendix C. See main text for explanation. Right: Photograph of the photodiode box with radiator.

5.5.1 Avalanche photodetector

Any photodiode can be used more or less successfully in avalanche mode: A large bias voltage (typ. 100 V) is applied to the photodiode in reverse direction. An electron-hole pair created by an incident photon rapidly acquires enough energy in the semiconductor to generate other electron-hole pairs and thereby results in a current of several photoelectrons per incident photon. The mean number of electrons per photon is referred to as the avalanche gain M. Operating a photodiode in avalanche mode has the advantage that a given current and voltage noise of the transimpedance amplifier of the photodiode can be overwhelmed by the photocurrent corresponding to weaker optical signal than for normal photodiodes. For the detection of beams with low optical power in our setup, we use the Thorlabs APD110C InGaAs avalanche photodetector which is specified to have a noise-equivalent power (NEP) of 0.5 pW/$\sqrt{\text{Hz}}$. The NEP can be equivalently expressed as a threshold incident power P_{th} at which the shot-noise signal is equal to the dark noise. For the APD110C, we find $P_{\text{th}} = \text{NEP}/\sqrt{2\hbar\omega_{\text{L}}} = 1.1\,\mu\text{W}$, which is orders of magnitude below standard photodiodes and about a factor of four below the saturation power of the detector. With a 3-dB bandwidth of 50 MHz and flat response up to 10 MHz, the photodiode is ideally suited to perform measurements of displacement at optical powers low enough not to disturb significantly the mechanical sytem in the measurement cavity. One draw-back of avalanche photodetectors is a strong temperature- and bias voltage-dependence of the avalanche gain, which requires to calibrate the photodiode gain simultaneously with a measurement.

5.5.2 Differential photodetector

As discussed in section 2.2.7, a homodyne detector requires to measure the difference of the photocurrents in the two outputs of an interferometer. Earlier attempts during this thesis work used two identically-built amplified photodiodes and a separate box to generate the difference of the two signals. The modular design is very practical in order to simultaneously create an output with the sum of both photodiode signals, which corresponds to the total optical power from local oscillator P_{LO} and signal P_{s} beams and is very helpful for calibration. However, problems in precisely balancing the gains of the two signal paths both at DC and at higher frequencies made the difference signal in practice not a reliable estimate of the broadband difference current. We therefore adapt a scheme where two photodiodes in the same device are operated in series as shown in Fig. 5.19. For an incident photon flux $I_{1,2}$ on either photodiode, an electrical current of $i_{1,2} = \eta_{\mathrm{PD}} e I_{1\,2}$ can pass through the diode, where $e = 1.6022 \times 10^{-19}\,\mathrm{C}$ is the fundamental charge and η_{PD} is the quantum efficiency of the photodiode, i.e. the number of electrons created on average per incident photon. The excess photocurrent $i_{-} = i_1 - i_2$ of the two photodiodes cannot pass from the $+5\,\mathrm{V}$ to the $-5\,\mathrm{V}$ bias terminal and must go through the amplifier section and therefore constitutes the measured signal. If both photodiodes have similar characteristics, this difference current is a good estimate of the difference of the photon flux on both diodes.

Photodiodes

We use the InGaAs photodiodes FD600N-1064 from Fermionics in Simi Valley, California for their reported high quantum efficiencies. The models we purchase come without protective window to avoid reflection loss at the air-glass interface and are optimized for our laser wavelength by an anti-reflection coating on the photosensitive surface. The photosensitive area has an effective diameter of 500 μm which is necessary for the desired capability to handle high incident powers. The large area causes the photodiodes to have a significant capacity of the order of 10 pA, which ultimately limits the RF bandwidth. While the photodiode is only rated for maximum reverse currents of 5 mA (6 mW), we do not observe any undesired effects for incident powers up to 30 mW. While we have not tested Fermionics photodiodes with even higher incident powers in fear of permanent damage to the photodiode, tests with older InGaAs photodiodes with identical case and diameter, but of slightly lower quantum efficiency (Epitaxx ETX-500) have shown linear behaviour up to 120 mW, where the output voltage of the photodiode was limited by the supply voltage (see Fig. 5.20). Incident powers up to 330 mW did not induce any observable damage to the ETX-500 photodiode. The high power handling capability certainly is a result of the rather low (5 V) bias voltage used, since the power dissipated in the photodiode is the product of bias voltage and photocurrent. Furthermore, it might be a consequence of the heat removal provided by a radiator pressed against the photodiode case in our implementation (see Fig. 5.19).

DC and RF frequency divider

The amplifier section in Fig. 5.19 first divides the signal into a DC and a radio-frequency (RF) component by passing it through an inductor in series with a resistor

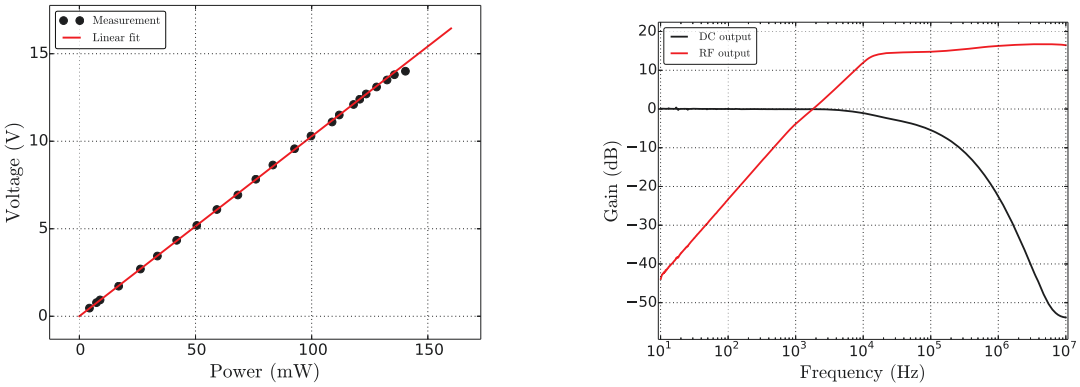


Figure 5.20: Left: Measured DC voltage from a photodetector constructed with a Epitaxx-500 photodiode for different incident powers. The detector voltage is linear up to incident powers of 120 mW, above which it saturates near 14 V, in agreement with the maximum output voltage specification of the amplifier OP27 for a supply voltage of 15 V. Power was further increased up to 330 mW without damaging the photodiode. From the slope of the linear fit, we can deduce a photodiode quantum efficiency $\eta_{\mathrm{PD}} = 0.77 \pm 0.04$. Right: Measured frequency-dependent gain of the DC and RF photodiode outputs. All gains are normalized to the DC gain at low frequency. The gain was measured by recording the photodiode response to an amplitude-modulation created with an acousto-optic modulator. The RF-output is flat to within 0.5 dB from 1 to 10 MHz. Since the transfer function of the modulator is not necessarily flat, these gain curves can only give an approximate idea of the real transfer functions.

or through a capacitor, respectively. The crossover frequency given by the capacitor $C = 150$ nF and resistor $R = 10\,\Omega$ alone is $f_{\mathrm{DC/RF}} = 1/(2\pi RC) \approx 100$ kHz. At this frequency, the inductor $L = 8\,\mu$H has an impedance of $|Z| = 2\pi f_{\mathrm{DC/RF}} L \approx 5\,\Omega$, which causes a correction of the crossover frequency towards even lower values. Fig 5.20 compares the measured transfer functions of the DC and RF outputs.

DC section

In the low-frequency amplification section, the DC voltage $R\bar{i}_-$ drops at the $R = 10\,\Omega$ resistor. The resistor value is deliberately chosen low such that the effective mean bias voltage $\pm 5\,\mathrm{V} - R\bar{i}_-$ changes by less than 10% even for photocurrents as high as 50 mA. The relatively low potential at the resistor terminal is then amplified with a gain such that the interference fringes can be resolved without saturation of the amplifiers. The gain is typically adjusted with a potentiometer to yield an output voltage of 1V/10mW to facilitate calibration (see detailed schematics in Appendix C) and allow the measurement of more than 100 mW of light if desired. For further convenience, we install a potentiometer to tune the offset of the DC amplification stage to zero (see Appendix C.1).

For each photodetector we construct, the dependence of the DC-voltage on the incident power $P_{1,2}$ on the two photodiodes is measured as shown in Fig. 5.20. For the measurement, the waist size of the laser beam on the photodiode is adjusted to about 25% of the photodiode diameter such that less than 0.1% of light is lost by

clipping (see Eq. 3.2.4). By dividing the measured slope by the one expected from the DC transimpedance gain for a perfectly converted photocurrent $i_{1,2} = \frac{e}{\hbar\omega_l} P_{1,2}$, we obtain an estimate of the quantum efficiency of the photodiode. Since the gain is a function of the employed resistors, we perform four-point measurements of the values of the used components with a digital multimeter (Keithley 2000) before assembling the photodiode. The resulting gain error is estimated to be well below 1%. For the photodiodes from Fermionics, we have measured quantum efficiencies in the range $\eta_{\mathrm{PD}} = [0.93, 0.97] \pm 0.07$. The observed efficiency fluctuations of a few percent are expected since different photodiodes are fabricated at different wafer locations and might suffer from defects or dust. The large indicated error of 7% is systematic and comes from the specified error of the power-meter that was used for the power measurement. Indeed, having compared our power meter (Thorlabs S120C) with the same model from another laboratory, it was found that our power meter overestimates powers by 6.7% with respect to the other one, indicating that the actual quantum-efficiencies of our photodiodes might be as low as 87%.

RF section

The RF component of the photocurrent is treated in a transimpedance amplifier configuration with a gain of typically $R = 1\,\mathrm{k\Omega} = 1\,\mathrm{kV/A}$, such that weak signals result in output voltages capable of dominating other electronic noise sources. For example, the shot-noise of incident light with single-sided spectral density $S_{ii,\,\mathrm{shot}}(\Omega) = 2e(\bar{i}_1 + \bar{i}_2)$ causes an output voltage spectrum

$$S_{VV,\,\mathrm{shot}}(\Omega) = 2eR^2(\bar{i}_1 + \bar{i}_2) \approx \frac{\bar{i}_1 + \bar{i}_2}{1\,\mathrm{mA}} \times \left(1.8\,\mu\mathrm{V}/\sqrt{\mathrm{Hz}}\right)^2, \tag{5.5.1}$$

which is about 60 dB above Johnson noise (50 Ω) for a typical incident optical power of 1 mW. The dark noise of the RF-output is dominated by the current noise $\sqrt{S_{ii}} \approx 18.5\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$ of the operational amplifier LMH6702, which results in a shot-noise-to-dark-noise ratio of unity for an incident power of 1.4 mW. The bandwidth of conventional transimpedance photodiode amplifiers is given by the lowpass filter formed by the gain resistor (here $R = 1\,\mathrm{k\Omega}$) and the photodiode capacitance (here $C = 2 \times 10.8\,\mathrm{pF}$), which would yield a value of only 7 MHz for our design. The employed amplifier[Tex02] was primarily chosen because of its current feedback architecture[Tex92], which results in the replacement of the gain resistance in the calculation of the transimpedance lowpass filter cutoff frequency by the input impedance of the inverting amplifier input, which is of the order of $70\,\Omega$ and yields an RF bandwidth of the order of 100 MHz. In this frequency range, implementing a well-connected ground plane on the back of the PCB board has drastically improved the flatness of the RF stage. The low distortion of the amplifier makes it in practice possible to measure RF signals with amplitudes up to 1/10th of the supply voltage (5 V) before nonlinearity or saturation effects become significant.

Fig. 5.5 at the beginning of this chapter shows a measurement of laser noise with the photodiode. From the plot, one can see that

- the achieved response to shot-noise is very flat up to 100 MHz,

- the dark-noise is about at the level of the shot-noise at 1 mW,
- the dark-noise increases slightly for frequencies above 30 MHz,
- no distortion occurs around the peak at 1 MHz, i.e. the dynamical range of the detector is larger than 70 dB,
- our design achieves common mode suppression of at least 58 dB (see the difference between the relaxation peak and the residual peak for balanced illumination),
- the photodiode output should be amplified further to decrease the background noise from spectrum analyzer for measurements that require the highest-possible detection-efficiency.

We conclude that the overall performance of the photodetector is satisfactory.

5.6 Homodyne detection

The homodyne detector (see Fig. 5.2) is composed of the balanced photodetector presented in section 5.5.2 and a mirror mounted on a piezo-electric actuator in the beam path of the local oscillator beam that is able to adjust the relative phase between the local oscillator and the signal beam. The piezo-mounted mirror is shown in Fig. 5.21. The mirror and piezo design is optimized low mechanical cross-talk to the measurement cavity in the nearby cryostat and a high locking bandwidth. In comparison with a navie approach, our solution resulted in at least 10 dB lower crosstalk to the measurement cavity[4] and a ten-fold improvement of the locking bandwidth from 2 to 20 kHz. Fig. 5.22 shows the measured piezo transfer function. The first mechanical resonance at about 15 kHz is related to the slow piezo and can in principle be attenuated by adjusting the relative gains of the fast and slow piezo. More significant resonances occur only around 150 kHz with a resonance enhancement of about 25 dB. The elevated peak response on those resonances in association with a rapid phase shift usually leads to a significant locking bandwidth penalty by a factor of about 10 with respect to actuators where these resonances are overdamped. By implementing an IIR filter (see Fig. 5.22 and section 6.2.5) in the servo controller with a transfer function given by the inverse actuator response within the locking bandwidth and a combination of strongly damped poles and lightly damped zeros near the piezo resonance frequencies around 150 kHz, we have been able to achieve a in-series transfer function of IIR filter and piezo that does not feature these resonance enhancement and that completely cancels the 15-kHz resonance. The resulting transfer function is additionally passed through a first-order low-pass filter to achieve a smooth roll-off of the loop gain. The final locking bandwidth is not limited by the piezo resonances any more, but rather by the delay imposed by the employed high-voltage amplifier.

When the homodyne detection is employed the measure the phase fluctuations of a laser beam reflected off the measurement cavity, a problem related to the near-perfect optical impedance matching $T \approx P$ of the cavity arises: since the resonant cavity

[4]We believe that the residual cross-talk is predominantly electrical, which makes it impossible to assess the residual mechanical crosstalk.

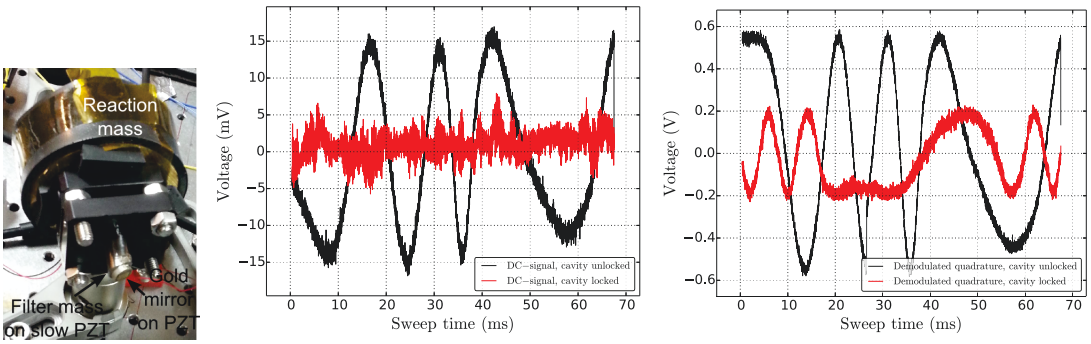


Figure 5.21: Left: Piezo-mounted mirror for controlling the phase of the local oscillator beam. Mechanical crosstalk from the piezo to the measurement cavity led to the presented design: a gold-coated wafer shard of about 4 mm × 4 mm is glued on a thin (250 μm) piezo disc which itself is glued to a steel cylinder with angular front surface. The steel cylinder is glued on a larger piezo for coarse actuation that is glued to a heavy (0.5 kg) reaction mass. The reaction mass is mechanically decoupled from the optical table by wrapping all screws in rubber. Middle: Homodyne DC output voltage during a scan of the local oscillator piezo voltage, while the signal beam is off-resonant or locked to the resonance of the measurement cavity. The error signal is drastically reduced when the cavity is locked due to the low reflectivity at resonance. Furthermore, the amplitude of the error signal is modulated by the power modulation of the signal beam that results from the vibrations of the cavity. Right: Homodyne error signal generated by demodulation of an amplitude modulation at 12.121 MHz. In contrast to the middle plot, this error signal decreases only by about a factor of three when the cavity is locked and remains a good estimator for the local oscillator phase despite cavity vibrations.

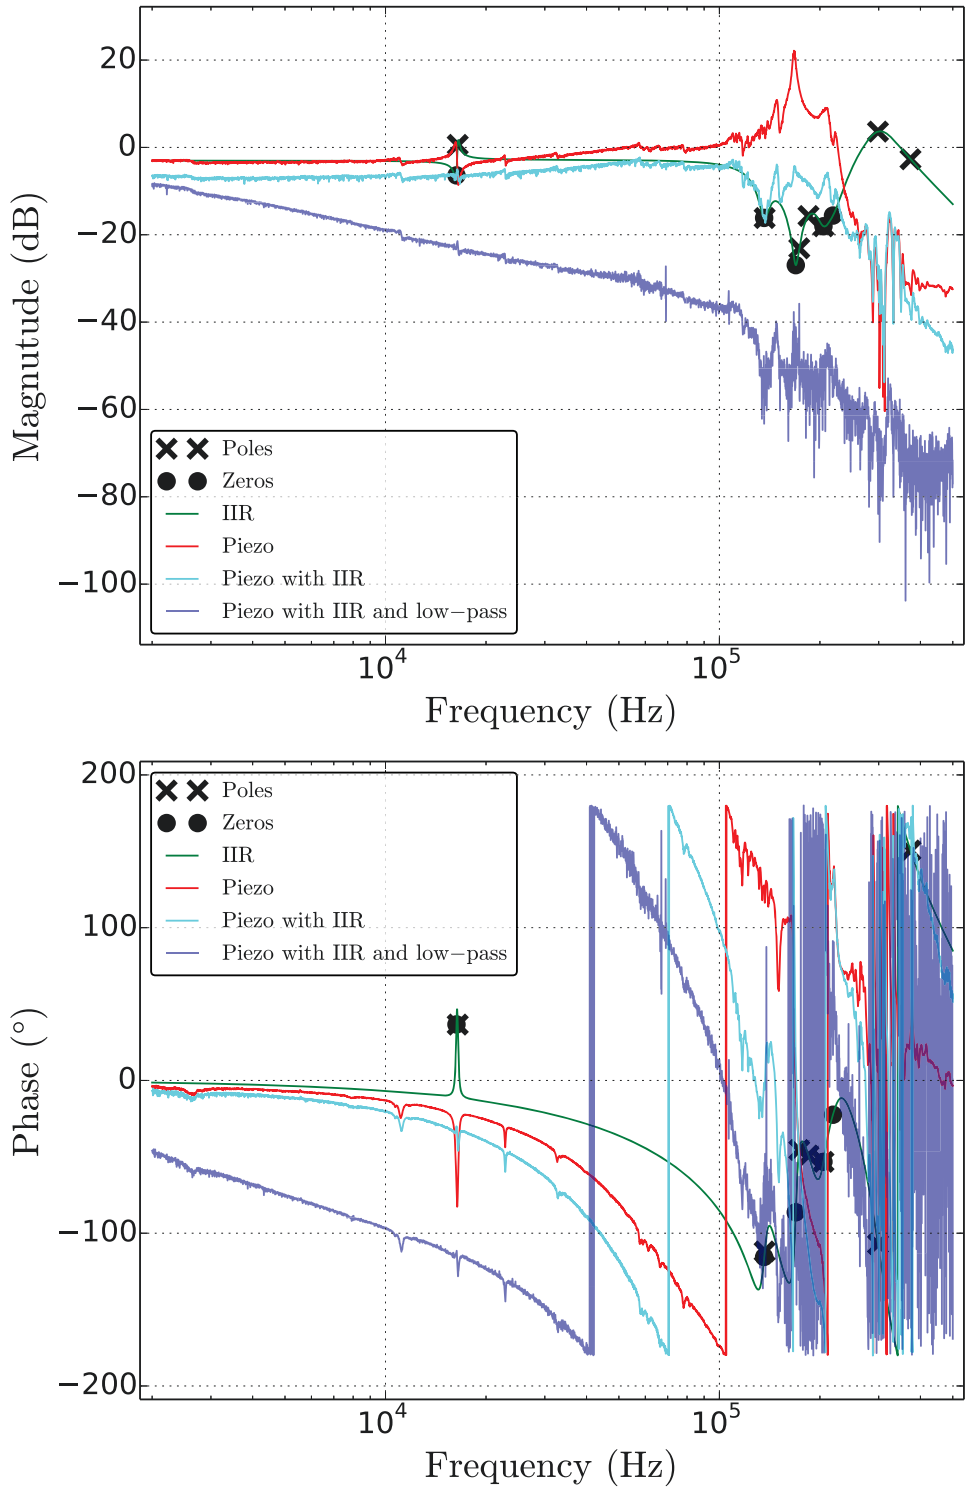


Figure 5.22: Magnitude and phase of the open-loop transfer function from the error signal input of the servo controller to the displacement of the piezo-mounted mirror in the local oscillator beam path. The bare piezo response is nearly flat up to 100 kHz but suffers from a significant delay that is caused by the employed high-voltage amplifier. The resonance peaks in the piezo response are cancelled by applying an IIR filter to the error signal before it is fed to the piezo. While we are able to achieve a near-perfect cancellation for the resonance peak at 15 kHz, the resonances above 100 kHz are more efficiently suppressed by the combinations of highly damped poles with lightly damped zeros in the IIR transfer function. The magnitude of the resulting low-pass filtered transfer function suggests that locking bandwidths up to 100 kHz are possible if the excess delay of the high-voltage amplifier can be reduced.

reflectivity R_0 is nearly zero, almost no light is available to yield an error signal for the homodyne phase at the DC output of the detector (see Fig. 5.21, middle). Furthermore, vibrations of the measurement cavity of the order of its bandwidth result in a strong modulation of the DC-fringe visibility, which effectively results in a the homodyne detector DC error signal that is proportional to the third power of the relative phase between signal and local oscillator beams. By imprinting a weak amplitude modulation at a frequency (here 11.2 MHz) comparable to the cavity bandwidth (here 16.7 MHz) on the homodyne signal beam, we can derive a much more robust error signal for the local oscillator phase by demodulation of the measured homodyne signal at the same frequency (see Fig. 5.21, right). The improvement can be understood by the fact that only a small fraction of the amplitude-modulation sidebands interacts with the cavity due to their off-resonant frequency, while a large part of the sidebands is directly reflected off the input coupling mirror of the measurement cavity. The resulting reduction in absorption is associated with a reduction of information about the length of the measurement cavity. The implemented scheme therefore stabilizes the local oscillator phase with respect to the surface of the input coupling mirror of the measurement cavity and decouples the stabilization loop of the measurement cavity and the of the homodyne phase. Future work might use the DC-part of the homodyne detector or a similarly demodulated signal at a frequency within the cavity bandwidth as an error signal to lock the measurement cavity.

5.7 Heterodyne detection

The optical and piezo-mounted mirror design of the heterodyne detection is similar to the one of the homodyne detection presented in section 5.6. Since the local oscillator of the heterodyne provided by the light transmitted by the filter cavity while the signal beam is frequency-shifted by the AOM block (section 5.3.2) to a difference frequency close to the cavity birefringence splitting frequency, the heterodyne beatnote frequency lies in the vicinity of the birefringence splitting of the measurement cavity. Birefringence splitting frequencies up to 450 MHz require to use a photodetector with an equivalently high bandwidth. Since our home-built differential photodetector only has a bandwidth of 100 MHz, we employ a commercial Newport 1611 freespace photodetector with a bandwidth of 1 GHz. While a balanced heterodyne detector would allow a 3-dB reduction in the detection noise floor (see the term p_{LO} in section 5.7), we have not found balanced photodetector at the required bandwidth. To achieve maximum detection efficiency, an asymmetric 90:10 beamsplitter is used to combine the signal with a strongly attenuated local oscillator. The local oscillator power is set close to the saturation threshold of the photodetector (1 mW).

To lock the phase of the heterodyne local oscillator beam with respect to the signal beam, the RF signal from the photodetector is demodulated twice with an analog signal at the difference frequency of the two AOM drivers, since the beatnote frequency is twice the AOM difference frequency due to the double-pass configuration (section 5.3.2). The resulting signal resembles the error signal of the homodyne detection and is used for a similar feedback loop to stabilize the relative phase. The resulting improvement in frequency resolution of the homodyne detection is easily seen in the

spectrum around of the beatnote frequency in Fig. 5.23, and is mandatory if signals from narrow mechanical resonances are to be monitored with the heterodyne detection. The heterodyne detection suffers from similar issues as the homodyne detection: when the measurement cavity is locked at resonance, the above desribed error signal cannot be used to to stabilize the LO phase. We therefore plan to implement a similar lock-in method for the heterodyne lock in the near future by imprinting an amplitude modulation on the heterodyne signal beam with the acousto-optic modulator.

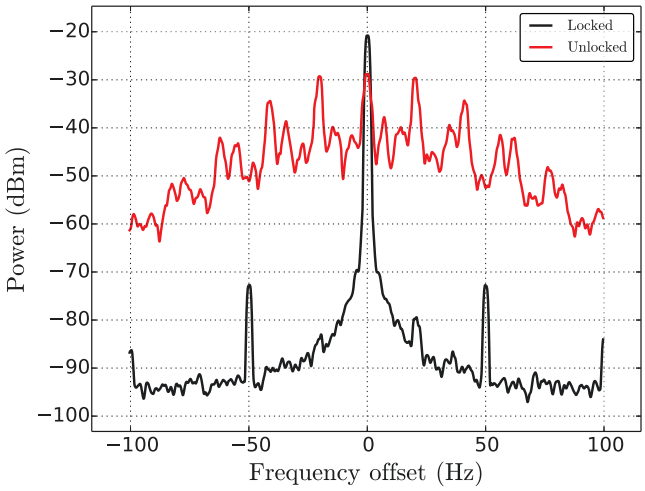


Figure 5.23: Heterodyne local oscillator lock performance. When the lock is engaged, the width of the carrier peak becomes limited by the resolution bandwidth of the spectrum analyzer (1.9 Hz here).

5.8 PDH detection

Over the course of development of the experiment, a variety of slightly different PDH error signals were generated to serve both as an error signal to lock the measurement cavity resonance frequency to the laser (frequencies below 200 kHz), and as a detector for the displacement fluctuations of the mechanical oscillator (above 1 MHz). For all generated PDH signals, phase modulation sidebands are imprinted on the homodyne signal beam with the fiber phase modulator and the reflected beam is directed to the avalanche photodetector. Due to constraints imposed by the photodiode bandwidth, modulation frequencies are chosen in the range from 20 to 80 MHz. Figure 5.24 shows a typical PDH error signal along with the corresponding reflection traces measured with the DC part of the APD and the laser 1 photodiode (see Fig. reffulloptics). While PDH error signals were generated digitally with the RedPitaya, we create an analog PDH error signal for the measurement cavity in order to perform calibrated displacement measurements on the micropillar compression mode by imprinting an additional phase modulation peak at 3.6 MHz on the laser beam and acquiring the PDH error signal spectrum with a spectrum analyzer. In order to avoid a dependence of the obtained spectra on other locking electronics connected to the PDH signal, we have built two

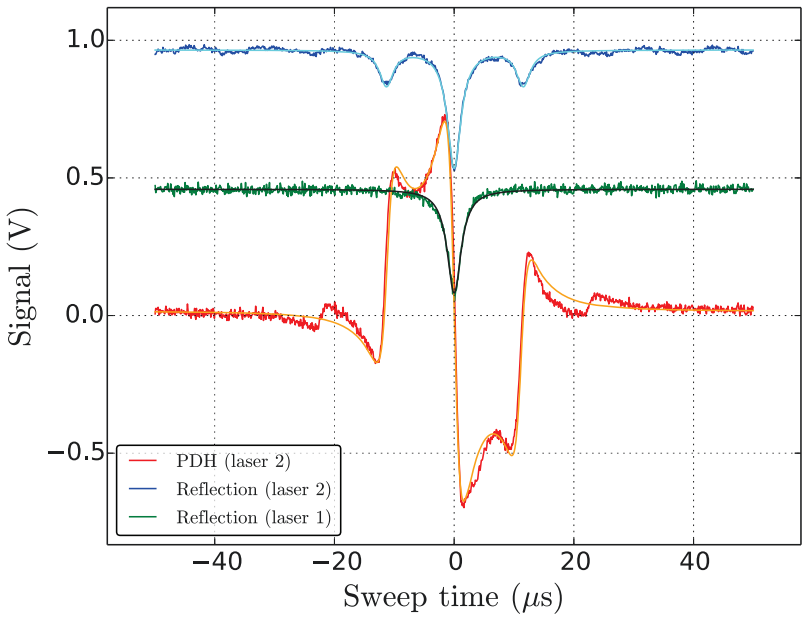


Figure 5.24: Reflection and PDH error signal traces. All traces were centered and both reflection traces were multiplied by a factor of 10 for better visibility. The PDH sidebands imprinted on laser 2 result in a lower resonant reflection with respect to laser 1.

identical unity-gain amplifiers, the output of one is used for cavity locking and the other for spectrum acquisition.

At a later stage of the experiment, it was necessary to modulate the intensity of the heterodyne signal beam with a signal proportional to the cavity-laser detuning to prevent instabilities of mechanical modes below 110 kHz. Using the same laser both for driving and detecting the motion was not straightforwardly possible with a standard PDH signal because of cross-talk from the intensity modulation to the error signal. While a PDH error signal only depends on the laser intensity to second order for perfectly resonant detunings, a linear power dependence arises in our experiment from an imperfect lock and small residual detunings. In order to remove this residual dependence, we use the low-frequency part of the avalanche photodetector signal as an estimate of the laser power. A normalization module implemented within the PyRPL framework multiplies the broad-band photodetector signal with the inverse of the power estimate in order to yield a normalized signal which does not depend on the laser power[5]. The normalized signal is subsequently fed through a standard PDH

[5]Since division is not a standard operation in the FPGA logic, we use a feedback-loop inside the FPGA to implement this operation: The input signal is multiplied by a number g. The difference between the product and a a previously defined set-point is computed, and a bit-shifted version of the difference is subtracted from the previous value of the number g. With a reasonable value for the set-point and bit-shift, the transfer function from the input signal to the product resembles a low-pass filter with typical cut-off frequencies of 1-5 MHz. However, since the suppression of low-frequency components is accomplished in a multiplicative way, all frequency components are affected by a low-frequency modulation, which is not apparent from the a single-tone characterization with a network analyzer.

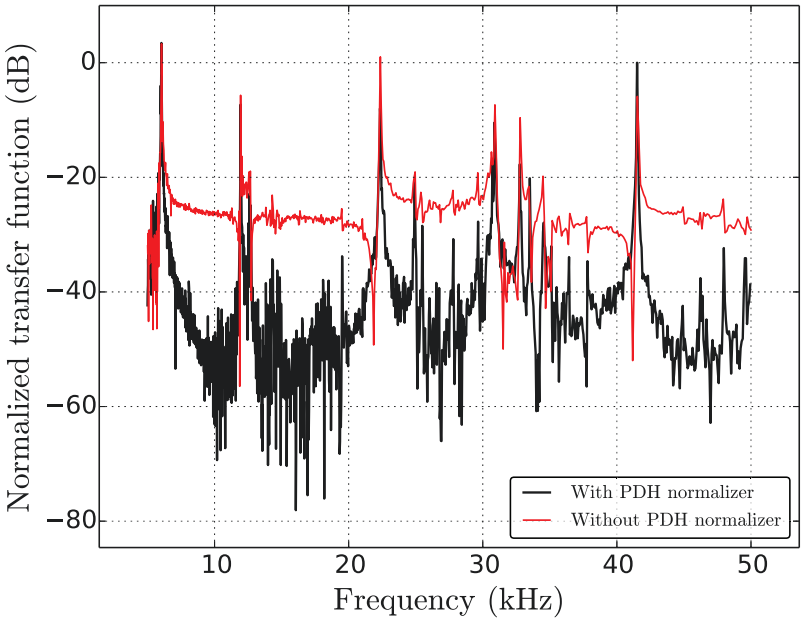


Figure 5.25: Transfer function from an amplitude modulator to the PDH error signal generated with and without normalization of the RF photocurrent before demodulation. Without normalization, the residual sensitivity of the PDH error signal to amplitude modulations results in a cross-talk that severely limits the measurement of the response of the cavity to an amplitude modulation. This cross-talk is suppressed by more than 25 dB when the normalizer is engaged. The mean detuning during the acquisition was 1.5% of the optical linewidth.

signal chain involving demodulation and low-pass filtering. Fig. 5.25 shows that this procedure is effective in removing the power dependence from the PDH signal. While this procedure also results in a steeper slope of the error signal because the absorption dip at resonance is effectively removed, the signal-to-noise ratio of the PDH error signal is of course not improved by the transformation.

5.9 Computer interface of the combined setup

All opto-electronic devices presented in the previous parts of this chapter are interfaced with the a computer that is employed to control the data acquisition with a programm written in Python. For this purpose, the following class structure was defined:

- At the highest-level, a class **Acquisition** contains objects that represent logical parts of the experiment. These are at the moment
  - The **cryostat** object is used to interface the resistance heater and temperature measurement. Its method "params()" returns a dictionary with all availble temperatures and heater settings at the moment of execution.
  - The **laser** object provides a method "set(p1, d1, d2, amfb, esfb)" that prepares the laser setup in a state where the homodyne laser (laser 1) is locked to the measurement cavity with an incident power of p1 and detuning d1, laser2 is locked at detuning d2 (its power is currently fixed by the setting of a waveplate), the amplitude modulator feedback gain amfb and the feedback gain through the antenna actuator esfb. The locking sequence performed by the "set"-method is optimized to prevent instabilities of the measurement cavity. The laser class provides another method "params()" that collects a number of diagnostics parameters provided by an identically-named method of all its members. The laser object contains the following objects to implement these functions:
    * The **FPF** object represents the filter cavity and ensures that it remains locked throughout the experiment by controlling the RedPitaya connected to the filter cavity piezo.
    * The **AOM** object represents the AOM block. It sets the frequency shift of the AOM block to provide the requested detuning by interfacing the DDS and adjusting the AOM drive voltage as described in section 5.3.2.
    * The **AM** object controls the amplitude modulator. It implements the locking sequence of the homodyne signal beam power with a RedPitaya and is used for feedback with the amplitude modulator.
    * The **FPM** object controls the RedPitaya that is used to lock the measurement cavity at the requested detuning. In particular, it scans the voltage of the two coarse piezos used to bring the measurement cavity in the locking range of a third, faster piezo. Since one of its inputs is the PDH error signal, it is also used to implement the feedback with the antenna actuator. The FPM object further controls the phase-modulation calibration peak of the PDH detection with a network analyzer.

    - The **homodyne** object controls the homodyne detector. It is able to lock the detector to a given relative phase either with the DC error signal or with the lock-in method presented in section 5.6. The homodyne RedPitaya also generates a calibration peak for the homodyne detection.
    - The **heterodyne** object controls the AOM difference frequency and the heterodyne detector in a similar manner as the homodyne detector.
  - The interfaces of all measurement instruments, i.e. two oscilloscopes, the network analyzer and the spectrum analyzer are represented in the Acquisition-class by associated acquisition methods.
  - To conveniently store the performed measurements, the database of our data management system "pyinstruments" is interfaced by the acquisition class. All diagnostics parameters are saved in the database along with references to the files that contain the measured traces. The traces are organized hierarchically in the database.

- The Acquisition-class is sub-classed by particular classes implementing specific measurement sequences.

Chapter 6
Feedback control

Why did the multithreaded chicken cross the road?
to To other side. get the

— unknown

This chapter will deal with feedback and stability of linear systems. The first section discusses basic concepts and terminology of the theory of linear and time-invariant systems in presence of feedback. The second section introduces in detail the universal feedback-controller PyRPL that was developed and extensively employed for this thesis. The last section discusses the effects arising from the presence of both photo-thermal and radiation-pressure forces in our optomechanical system. The described suppression of the photo-thermal instabilities of our optomechanical system at higher incident optical powers is a particularly interesting problem in controller design: the stabilization of a system with unstable poles and unstable zeros through external feedback.

6.1 Brief introduction to feedback

Feedback is omnipresent in quantum optics experiments to minimize the effect of fluctuating environmental conditions on the experiment. The principle of feedback control is shown in Fig. 6.1. In the following, we will discuss a few concepts from the theory of feedback control that will be helpful in the following sections of this chapter. A large amount of literature is available for more detailed explanations, for example [Bec05].

Plant and controller

The system to control is traditionally called the plant and can be modeled by its impulse-response $p(t)$ to a Dirac-$\delta(t)$-like input. By assuming that the impulse-response of the plant can be modeled as linear and time-shift invariant with respect to the input signal, it follows that the output signal $x(t)$ of the plant is given by the convolution of the input signal $f(t)$ with the impulse-response of the plant:

$$x(t) = p(t) * f(t)\,. \tag{6.1.1}$$

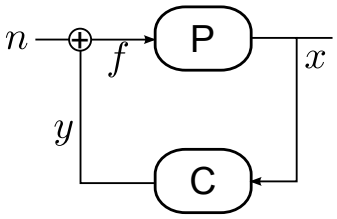


Figure 6.1: Schematic of a plant P and a feedback controller C. The sum $f = n + y$ of noise n and controller output y drives the plant, which yields an output signal $x = Pf$. The plant output is transformed by the controller C into the signal $y = Cx$. By solving these coupled equations, we find the closed-loop transfer function for noise through the plant to be $x/n = P/(1 - PC)$, which suggests that a properly designed controller may diminish the effect of noise on the plant output. A more detailed and rigorous treatment is given in the main text.

According to the convolution theorem, the response can be expressed as a product by Laplace-transforming[1] the equation:

$$X(s) = P(s)f(s)\,. \tag{6.1.2}$$

Since the transformation accomplished by the plant is usually well-modeled as an ordinary linear differential equation involving various derivatives of the input, its Laplace-transform can be written as a rational function

$$P(s) = K\frac{\prod_{i=1}^{M}(s - z_i)}{\prod_{j=1}^{N}(s - p_j)}\,, \tag{6.1.3}$$

where K is a constant scaling factor, p_j are the complex poles, and z_i the zeros of the transfer function. The controller is similarly modeled by the Laplace-transform $C(s)$ of its impulse response, and will also be assumed to be a rational function here. Commonly, one additionally defines

- the loop funcion of the combined system $L(s) = C(s)P(s)$,
- the sensitivity function $S(s) = \frac{f(s)}{n(s)} = \frac{1}{1-C(s)P(s)}$, and
- the closed-loop transfer function $T(s) = \frac{x(s)}{n(s)} = \frac{P(s)}{1-C(s)P(s)}$.

Stability

We define a system to be stable if and only if bounded inputs result in bounded outputs. Since any bounded input signal can only be non-zero at Laplace frequencies s with $\mathrm{Re}(s) > 0$, stability is equivalent to the absence of poles with positive real part in the transfer function. Poles or zeros with negative (positive) real part are therefore referred to as stable (unstable) poles or zeros.

[1] We define the Laplace transform of a function $f(t)$ as $F(s) = \int_0^\infty e^{-st} f(t)\mathrm{d}t$. We use the Laplace transform instead of the Fourier transform because transient instead of stationary signals must be considered in feedback design. The Laplace-transform formalism $s = i\Omega + \Gamma$ simplifies to the Fourier treatment for $\Gamma = 0$. Experimental measurements of transfer functions are usually plotted against the Fourier frequency Ω.

Delay

We define all-pass filter A_{z_0} at Laplace frequency z_0 as a linear, time-invariant system with $\mathrm{Im}(z_0) = 0$, $\mathrm{Re}(z_0) < 0$ and

$$A_{z_0}(s) = -\frac{s + z_0}{s - z_0}\,. \tag{6.1.4}$$

We see that the magnitude response of the all-pass is unity for all Fourier frequencies $s = i\Omega$, while the phase response changes from 0 at $s = 0$ through $-\pi/2$ to $-\pi$ for $|s| \to \infty$. We will use all-pass filters in order to model any kind of delay in a real-world system. The real part[2] of z_0 must be negative because the filter would be unstable for an unstable pole z_0 (see the denominator). Furthermore, $\mathrm{Re}(z_0) > 0$ would result in a phase-advance towards higher frequencies with a flat magnitude response, i.e. an acausal response. The delay of any system with a stable zero z_i can be increased without modifying the magnitude response by multiplying the system with an all-pass filter A_{z_i}. A stable system which only contains stable zeros is therefore referred to as a minimum-phase system. Inversely, a stable non-minimum-phase system cannot be converted to a stable minimum-phase system by multiplying with the corresponding all-pass filters, since these all-pass filters would be unstable due to the unstable pole that is needed to realize the filter.

Feedback

If a feedback controller $C(s) \neq 0$ is employed, the system response $x(s)/n(s)$ is changed from the open-loop response $P(s)$ to the closed-loop response $T(s) = P(s)/(1 - C(s)P(s))$. For frequencies at which $|C(s)P(s)| > 1$, the effect of noise n is significantly reduced in the output x, and the feedback is effective. The degree to which feedback can be used is limited by the stability of the closed-loop system. If we assume $P(s)$ to be a stable system, the poles of the closed-loop system are given by the zeros of $1 - C(s)P(s)$. For large gain $|C(s)| \gg 1$, the zeros of $1 - C(s)P(s)$ are approximately given by the zeros of the plant, unless the poles of C cancel the zeros of P. For a simple controller, e.g. $C(s) = -k$, we therefore conclude that the closed-loop system is stable with arbitrary gain $k > 0$ if and only if all zeros of the plant are stable. Any kind of delay, which can be modeled by introducing unstable zeros into the plant or controller, therefore jeopardizes the degree to which feedback can be effective. Attempts to cancel unstable zeros of the plant by introducing unstable poles into the controller usually do not work, since the output of an unstable controller tends to saturate for arbitrarily small imperfections of the pole-zero cancellation. In practice, the challenge of feedback is the design of a controller with a large gain at frequencies where delay is insignificant in order to yield an effective noise reduction, and a small enough gain to avoid instabilities at (typically higher) frequencies where delay is significant, without introducing excessive delay with the controller.

[2] While we have required the imaginary part of z_0 to be zero for a simplified discussion, complex values can be modeled if a second-order all-pass filter $A_{z_0}A_{z_0^*}$ composed of a complex-conjugate pair of filter frequencies z_0, z_0^* is considered. The requirement of stable zeros for minimum delay also holds for complex-valued zeros.

6.2 Python RedPitaya Lockbox (PyRPL)

The Red Pitaya [red16] is an affordable (< 250 Euros) field-programmable gate array (FPGA) board with fast analog inputs and outputs. This makes it possibly useful for quantum optics experiments, in particular as a digital feedback controller for analog systems. Based on the open source software provided by the manufacturer of the board, we have created the software package PyRPL [Neu16b] (Python RedPitaya Lockbox) which implements many devices that are needed for optics experiments with the Red Pitaya. The user-interface and all high-level functionality are written in Python, but an essential part of the software is hidden in a custom FPGA design written in the hardware description language Verilog. In this section, we will briefly present the hardware platform Red Pitaya, then discuss the Python and Verilog infrastructure of PyRPL in section 6.2.2, and end this section with a presentation of the various digital signal processing (DSP) modules that are implemented in PyRPL: a few simpler DSP modules will be discussed in section 6.2.3, the more complex IQ and IIR modules in sections 6.2.4 and 6.2.5, and higher levels of abstraction in section 6.2.6.

6.2.1 Red Pitaya

The Red Pitaya board (recently renamed "STEMlab 125-14") was first released in 2013 as a Kickstarter project and has since gained great popularity in the scientific and electronics communities. The board combines a Xilinx Zynq-7010 System-on-chip (SoC) [Xil16] with two-channel analog-to-digital (ADC) and digital-to-analog (DAC) converters with 14 bit resolution and 125 MHz sampling rate. The term SoC generally indicates that various systems are implemented on the same chip. In the case of the Zynq-7010, an FPGA allowing programmable logic (PL) is combined with a dual-core processing system (PS), composed of an ARM Cortex-A9 processor with standard peripherals such as memory and input-output interfaces, and a high-bandwidth connection to the FPGA. The PS allows to run a Linux operating system (OS) on the board and thereby to use, interface, and program it like a standard personal computer. The PL is defined by a hardware description language, for example Verilog, whose compiler generates a so-called bitfile for the FPGA that can be loaded via the PS with a Linux command. The manufacturer of the Red Pitaya provides an example Verilog source code that realizes the interface to the ADC and DAC, a two-channel oscilloscope, a two-channel arbitrary function generator, and four PID controllers. Furthermore, the example code implements a bus between PS and PL that allows a Linux user connected to the board through an ethernet connection to read and write registers of the FPGA, for example to retrieve the data measured by the oscilloscope or to set the gains of the PID controller. The example code allows even beginners to modify the Verilog code for customized PL functionality.

Analog performance

While the PS clock rate is 1 GHz, the FPGA clock rate is 125 MHz. The FPGA clock is by default derived from a quartz oscillator on the board, but can be replaced by an external clock source in order to synchronize the clock with other laboratory devices. The DAC and ADC are clocked synchronous with the FPGA. However, an

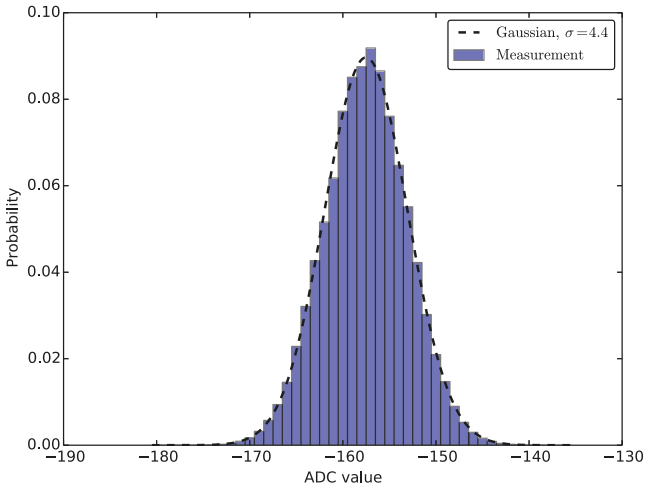


Figure 6.2: Histogram of 1.6×10^6 samples of unprocessed ADC values for a Red Pitaya input terminated with 50 Ω. The distribution is Gaussian with a standard deviation of $\sigma = 4.4$. The effective number of bits is therefore ENOB $= 14 - \log_2 4.4 \approx 12$. The distribution does not broaden measurably when a strong signal (1 V ramp at 1 MHz) is generated with the DAC while the ADC data is recorded.

analog signal passing through the ADC, PL, and DAC experiences a significant delay of about 20 clock cycles, which is mainly caused by the employed ADC and DAC converter architecture. Therefore, the RedPitaya is only interesting for feedback loops below 2 MHz.

The resolution of the ADC and DAC is 14 bits and both the differential nonlinearity and the transition noise of both components are specified in the datasheets [Sem12, Lin16] to be of the order of the least significant bit (LSB). However, the implementation of these components on the Red Pitaya board results in a slightly worse performance. In agreement with [Sei15], we have measured effective resolutions of 12 bits for the ADC and 11 bits for the DAC (see Fig. 6.2). Much worse than this, we have observed various cross-talk issues from the digital supply voltage to the analog outputs of the board, which lead to a DAC output voltage that is not a monotonic function of the digital signal value. This problem could be resolved by a small modification of the analog circuit of the Red Pitaya board [Neu16c]. The modification involves the removal of 4 resistors from the board and modifies the output voltage range from [-1, +1] V to [0, +2] V. After the modification, the analog transfer function of the board is flat, monotonic in the gain, and the effective output resolution is better than 13 bits [Neu16c]. Furthermore, the modified output voltage range [0, +2] V is more practical in most of our applications. Subsequent discussions with the manufacturer have triggered developments within the company of Red Pitaya to produce boards with improved noise performance, which are now sold under the name "STEMlab 125-14". These boards use a different design modification (similar to the one described in [Neu16a]), which leads to an improved performance compared to the original boards, but slighty worse noise performance than what can be obtained with by removing the 4 resistors from the board.

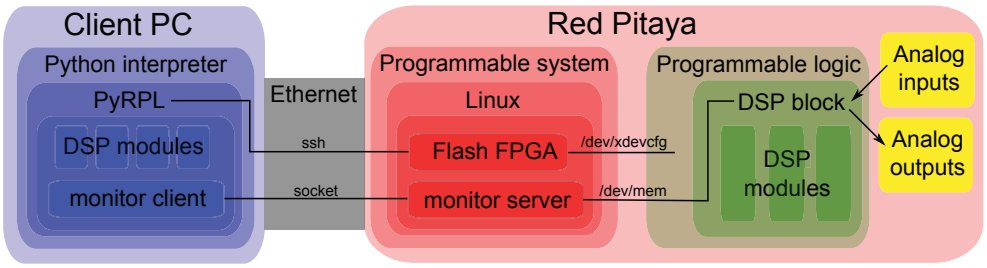


Figure 6.3: Overview of the PyRPL architecture. Blue blocks indicate software written in Python, red blocks software in C, and green blocks indicate Verilog code.

The noise in Fig. 6.2 clearly follows a Gaussian distribution. Therefore, averaging may be employed to reduce the effective noise at the cost of a reduced bandwidth. At a typical locking bandwidth of 100 kHz, about 1250 samples from the analog-to-digital converter may be averaged, in order to yield a resolution improvement of about $\log_2 \sqrt{1250} \approx 5$ bits. In a similar way, the analog output resolution may be enhanced through averaging by placing an analog low-pass filter behind the output. If care is taken to adapt the typical signal range to the full input and output ranges of the ADC and DAC, noise performance is rarely an issue in standard feedback applications.

6.2.2 PyRPL overview

The software controlling the Red Pitaya board can be divided into various parts and is schematized in Fig. 6.3. We will first discuss the communication layer of the software that runs on the client computer, the PS and the PL. Then we will discuss the Verilog implementation of the DSP multiplexer, which serves as a router between analog signals and various DSP modules.

Communication with the client computer

When a PyRPL instance is started, the client computer connects to the Red Pitaya PS through a secure shell (ssh) connection. The client then uploads a bitfile which contains the compiled PL design and loads this file into the FPGA through a Linux command. Next, the client uploads and executes the small program "monitor_server" that is written in C. The server program allows the client to connect to it through the TCP/IP protocol and listens for commands from the client. The possible commands are 'read' and 'write', followed by an address and the number of 32-bit packets of data to read or write from or to the FPGA, starting from the given address. If the address lies within the address space that the OS of the RedPitaya has assigned to the FPGA bus, the read or write request becomes available to the PL on the various bus signals. Here, the request is interpreted by the logic implemented in the Verilog code, which essentially maps a number of the registers of each DSP module to specific addresses within an address subspace reserved for the module. By knowing the names and addresses of all registers of a DSP module, a Python object corresponding to the DSP module can be created on the client computer, such that the properties of the object correspond to the FPGA registers of the DSP module. When these properties

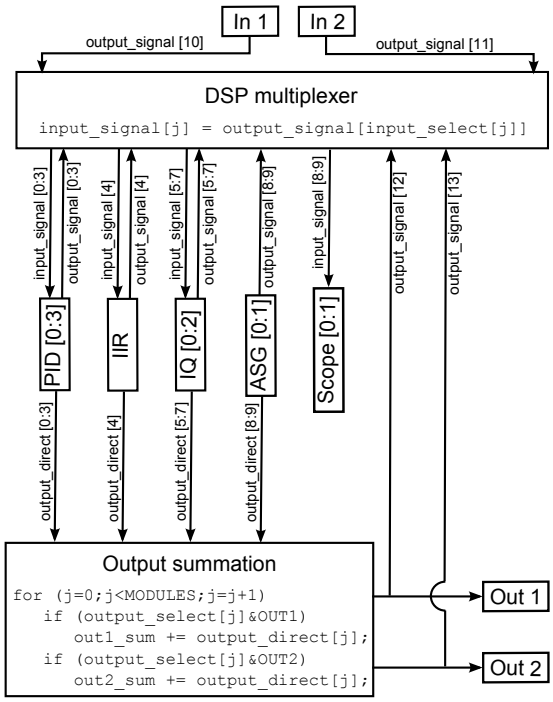


Figure 6.4: Schematic of the DSP multiplexer that acts as a router for the 14-bit signals between the different DSP modules (PID 0-3, IIR, IQ 0-2, ASG 0-1, Scope 0-1) and the analog inputs (In1, In2) and outputs (Out1, Out2).

are read or written, the Python value is automatically synchronized with the content of the corresponding registers through the described communication chain. The end result is that a Python programmer can assign or read the FPGA registers in the same way as local Python variables.

Communication speed was a major concern in our implementation. Without excessive code optimization, we achieve typical read and write delays for a single 32-bit register of the order of 300 μs. When a block of registers with adjacent addresses are read or written, for example the memory storing an oscilloscope trace, the read/write delay scales sublinearly. For reading or writing 2^{14} 32-bit values, a time of the order of 7-10 ms is required. The delay is limited by the speed of the ethernet connection, and could be further improved by compressing the data that is send over the network or by using another communication protocol such as UDP.

DSP multiplexer

The DSP multiplexer serves as a router between the different DSP modules of PyRPL. Its Verilog implementation is schematized in Fig. 6.4. The module definition fits into few lines of code, but is nevertheless very resource-intensive. The multiplexer can accommodate up to 16 DSP modules that provide each a port for a 14-bit input signal ("input_signal"), a 14-bit output signal ("output_signal"), and an additional 14-bit signal called "output_direct" which can be routed directly towards the analog outputs. The value of one register for each module called "input_select" determines

which other module's output signal is connected to the former module's input port. Another 2-bit register "output_select" per DSP module is used to select whether the module's signal "output_direct" is sent to the analog output 1, output 2, to both outputs or to none. As shown at the bottom of Fig. 6.4, the sum of the selected "signal_direct" of all DSP modules is computed for each of the two analog outputs and then sent to the DACs. The analog input signals from the ADCs and the output signals that are sent to the DACs are also represented in the form of DSP modules which only provide an "output_signal" port in the DSP multiplexer in order to keep the interface simple. Some DSP modules, such as the oscilloscope (Scope) and the arbitrary signal generator (ASG) have by definition only inputs or only outputs. For the PID and IIR modules, the signals "output_direct" and "output_signal" are identical. By combining a large number of DSP modules in series, or by rapidly (< 1 ms) modifying the connections between modules in real time, the multiplexer module allows highly complex functionalities to be implemented. Furthermore, it allows the oscilloscope, network analyzer and spectrum analyzer to be connected to any DSP module without the need of signal conversion to the analog domain, which has proven a highly practical tool for debugging, setup and in keeping analog cross-talk low.

6.2.3 Simple DSP modules

This section discusses the various DSP modules. While emphasis is placed on explaining the theory behind the modules here, a more hands-on introduction can be found in the IPython notebook tutorial in reference [Neu16b].

Arbitrary signal generator module

The ASG module is an adapted version of the two-channel ASG provided by the Red Pitaya manufacturer. Our Python interface made it possible to slightly reduce the necessary FPGA resources for the ASG in order to reserve more FPGA space for other modules. Waveforms defined by 2^{14} values are easily loaded through the Python interface. The ASG supports frequencies from 0.1 Hz to 62.5 MHz and various burst and pulse modes. We have added extra triggering functionality to the ASG in order to allow arbitrary delays for the turn-on or turn-off of the ASG with respect to the arrival time of an external trigger signal, which has been useful for the modified ringdown measurement shown in Fig. 7.8. Furthermore, we have implemented a pseudo-random noise generator (PRNG) based on a Lehmer pseudo-random number generator [PRB69]. While the current PRNG has a period of the order of 2 s, true random number generation is possible with FPGA's [Bae08] and might be implemented in the future.

Oscilloscope module

The two-channel oscilloscope module is based on the example oscilloscope provided by Red Pitaya. Each channel allows to save a time-trace of any DSP signal with 2^{14} points and adjustable duration. We have extended the triggering functionality and added various time-stamps which allow to superpose traces at different times. The fast communication interface presented above allows the display of both oscilloscope

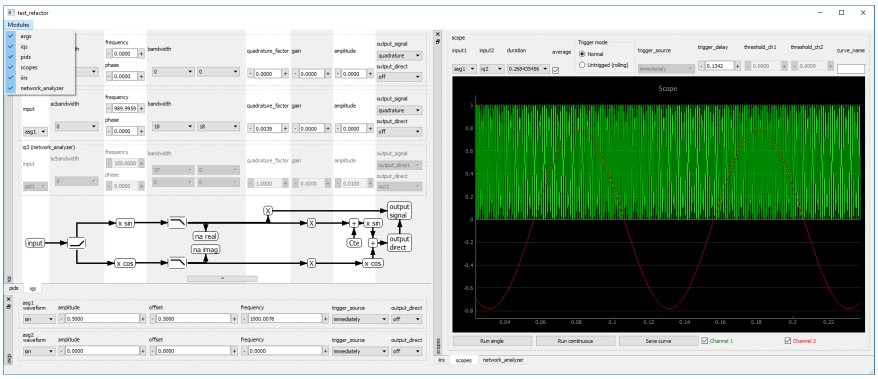


Figure 6.5: Graphical user interface (GUI) of the oscilloscope, network analyzer, spectrum analyzer and a number of additional controls.

channels in real time on the client computer (see the graphical user interface in Fig. 6.5). By computing the windowed Fourier transform of an oscilloscope trace, a simple spectrum analyzer has been implemented. By connecting the two oscilloscope channels to the two quadratures of an input signal demodulated with an IQ module (see section 6.2.4), an improved spectrum analyzer has been implemented that allows to acquire the spectrum of an input signal in a narrow window around the demodulation frequency, thereby allowing for excellent resolution bandwidths below 1 Hz despite the relatively small size of the oscilloscope memory (2^{14} values per channel) with the only drawback of a maximum span of about 10^3 resolution bandwidths. With a relatively small improvement of the Python code, a span above 0.5 MHz with arbitrarily small resolution bandwidth could be realized in the future by continuously streaming oscilloscope data to the client computer. A major update involving the implementation of the fastest streaming rates achieved by other Red Pitaya users [red16] could result in spans up to 30 MHz with arbitrarily small resolution bandwidth. For this, a direct interface from the FPGA to the RAM of the PS and an improved communication program with the client computer based on the UDP protocol would have to be implemented.

PID module

The PID modules are also extended versions of a Red Pitaya example. The output signal s_{out} computed by a standard PID module with integral gain g_{i}, proportional gain g_{p} and derivative gain g_{d} can be written as

$$s_{\mathrm{out}}(t) = g_{\mathrm{i}} \int_{-\infty}^{t} e(t)\mathrm{d}t + g_{\mathrm{p}} e(t) + g_{\mathrm{d}} \frac{\mathrm{d}}{\mathrm{d}t} e(t)\,, \tag{6.2.1}$$

where $e = s_{\mathrm{in}}(t) - s_0$ is the difference between the input signal s_{in} and the setpoint s_0. A very practical modification of the PID module is the possibility to modify the current value of the integral in the above equation through a register termed "ival" that is accessible in the Python interface. This way, the PID integrator can be re-set to arbitrary output voltages, or voltage ramps can be easily implemented as for-loops in Python that periodically change the value of the integral.

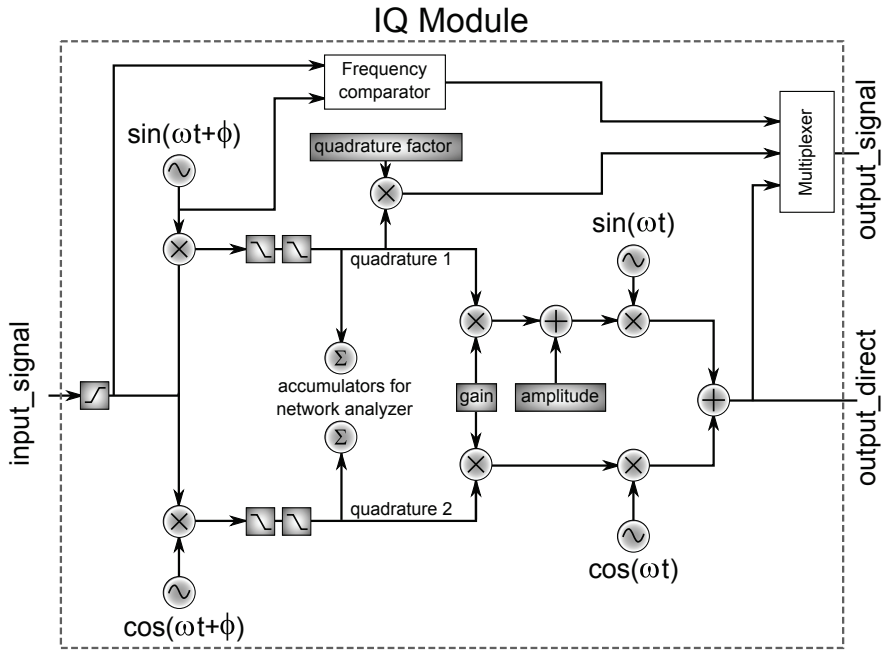


Figure 6.6: IQ module layout. Dark rectangles denote registers that can be used to select the functionality implemented by the IQ module.

In addition to this, we have added 4 first-order filters with selectable low-pass or high-pass cut-off frequencies in series to pre-filter the input signal of the PID module. An additional saturation stage allows to define arbitrary maximum and minimum voltages for the PID output signal. We note that, while the saturation stage is often practical to protect sensitive components connected to a Red Pitaya output from undesired voltages, it is always preferable to build an analog attenuator circuit instead of using digital saturation when possible, in order to benefit from a maximum dynamic range and a better analog noise performance. Similarly, it is always preferable to replace digital low-pass filters acting on the output signal of a Red Pitaya by analog ones in order to improve the averaging of the DAC output noise.

6.2.4 IQ module

The three IQ modules available in PyRPL can be used as a lock-in detection for the generation of PDH error signals, as a network analyzer, as a tunable band-pass filter, and for detecting and locking the frequency of a VCO to a frequency reference provided by the IQ module. The design of the module is schematized in Fig. 6.6. Each four phase-shifted sine functions required for each IQ-module are generated from a look-up table (LUT) of 2^{11} 17-bit values stored in read-only memory (ROM) of the FPGA. For minimum resource requirements per IQ module, only a quarter period of the sine function is stored in the LUT. The phase of the sines is given by the most significant bits of a 32-bit register that is incremented each clock cycle by the value of the register that defines the frequency.

PDH-like error signal generation

In order to use the IQ module for the generation of a PDH error signal, the register "gain" is set to zero, the registers "amplitude" and "quadrature_factor" to non-zero values and the multiplied value of "quadrature1" is chosen as "output_signal" with the IQ multiplexer. With this setting, the signal "output_direct" can be sent through one of the analog outputs to a modulator such as an EOM or a piezo. The detected analog signal is digitized and connected to "input_signal". The signal is optionally high-pass filtered at the beginning of the IQ module and then demodulated at the frequency of the modulation signal at "output_direct". One or two subsequent low-pass filters remove the demodulation component at twice the modulation frequency before an amplified version of one demodulated quadrature is available as "output_signal" of the IQ module. The resulting PDH error signal is typically optimally amplified by an additional PID module before it is fed to some actuator for stabilization of the error signal value. An example trace of a PID error signal generated in this way can be found in Fig. 5.8. The figure also shows how the error signal amplitude is optimized by varying the demodulation phase.

Network analyzer

In a similar way, the IQ module may be employed as a network analyzer. Settings are identical as for the PDH error signal, with the difference that a for-loop in Python iterates over a list of frequencies to record. Updating the frequency register of the module[3] triggers a first counter of a pre-defined number of cycles specified by a register "sleep_cycles". This is done to allow for the settling of the response of the system under test before acquisition of its transfer function begins. Next, a number of cycles defined by a register "na_cycles" are counted down, during which the values of both demodulated quadratures are additively accumulated by the two accumulators shown in Fig. 6.6. Once the second counter reaches zero, accumulation stops and the Python client reads the values of the two accumulators as 62-bit numbers and computes the averaged complex transfer function from the output of the IQ module to its input at the current frequency. While the analog performance of the network analyzer is worse than commercial network analyzers, a number of custom adjustments often allow better functionality. For example, by iterating over a list of frequencies in reversed order, frequency sweeps in both frequency directions have been performed. This feature is commonly inaccessible in commercial network analyzers and has enabled the observation of an optomechanical Duffing effect as shown in Fig. 6.7. Another practical customization involves the stabilization of the magnitude of the input signal by adjusting the amplitude of the output tone, which was found to yield better measurement results in systems with response of high dynamic range and nonlinear response, such as high-Q mechanical resonances detected with a Fabry-Perot cavity.

[3] In zero-span mode, the frequency register is updated with its current value, which also triggers the acquisition of a new network analyzer value.

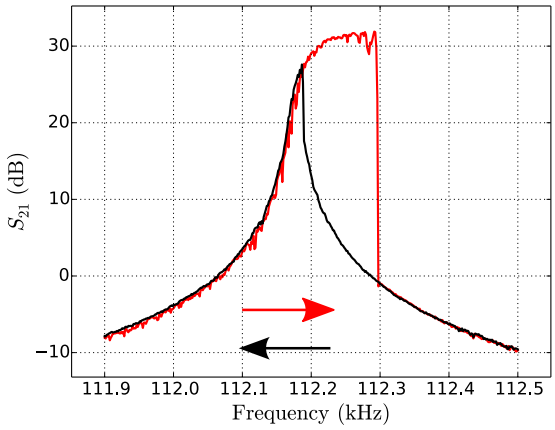


Figure 6.7: Optomechanical Duffing effect observed on a mechanical mode of sample 20121009C5 near 112 kHz with the PyRPL network analyzer. The arrows indicate the direction of the frequency sweep. For the measurement, a laser was locked to the red side of an optical resonance such that a strong optical spring effect significantly shifted the frequency of the mode towards lower frequencies. When the mode is strongly driven, the large vibration amplitude results in a reduced time-averaged optical spring due to the nonlinear optical lineshape, and thereby results in an effective increase of the mechanical frequency. While locking the laser on the blue side of the optical resonance was not possible due to instabilities, we expect the sign of the Duffing nonlinearity to be reversed in this case. Probing the system's response in the direction towards lower frequencies could demonstrate this effect.

Band-pass filter with tunable phase

For the band-pass filter operation mode, the value of the "amplitude" register is set to zero and instead the "gain" register is set to the desired filter gain. An incoming signal is then demodulated with a phase-shifted sin/cos-pair, low-pass filtered, and modulated with the un-shifted sin/cos-pair. The center frequency of the filter is thus determined by the frequency register of the IQ module, while the phase lag can be tuned by the relative phase-shift between the demodulation and modulation oscillators [GVK12]. The filter bandwidth is defined by the cut-off frequencies of the low-pass filters acting on the demodulated quadratures. If only one low-pass filter is used, the shape of the band-pass filter is Lorentzian. By combining up to seven band-pass filters of the described type in parallel, complex transfer functions can be realized. The digital quadratures of narrow band-pass filters can benefit from significant reduction of the ADC noise through averaging and are therefore internally represented with 24 instead of 14 bits. This advantage is easily lost through the noise added by the DACs. For optimal performance, it is therefore crucial to adjust the filter gain such that the output signal nearly saturates the available voltage range, and to use analog attenuation afterwards, which attenuates both the signal and the DAC noise.

Offset frequency lock

For implementing an offset frequency lock, for example to stabilize the difference frequency between two lasers, a simple frequency comparator was implemented in the IQ module. This comparator detects the changes of sign of the (optionally high-pass

filtered) input signal and of an internally generated sine wave at a specified frequency. The comparator output register value is incremented by "1" for each change of sign of the input signal, and decremented for each change of sign of the internally generated sine. This way, a robust signal proportional to the integral of the frequency difference between the input signal and internal reference oscillator is generated, which can be fed through a PID module with appropriate gain and sign to the piezo of a laser or an external VCO. Fig. 5.17 shows the performance of this sub-module in the stabilization of the frequency difference between two VCOs used as drivers for acousto-optic modulators. Lock acquisition is achieved even with initial frequency differences of 10-20 MHz. By adding the output signal of the frequency comparator with a user-defined gain to the value of the frequency register of the IQ module, a phase-locked loop (PLL) could be constructed instead, which tracks the frequency of an external oscillator fed to the input port of the IQ module.

6.2.5 IIR module

While relatively complex filters can be realized with combinations of several PID and IQ modules, the infinite impulse response (IIR) filter module is usually the best choice when the complicated transfer function of a piezo must be partially compensated for an optimal feedback loopshape. We will briefly discuss the theory of digital filters in order to present our implementation of an IIR filter in PyRPL. The theoretical discussion closely follows Reference [OS75].

Digital signals can be represented as sequences of numbers $x(n)$ at the discrete sampling times $t = nT$, where T denotes the inverse sampling rate and n an integer numbering the digital clock cycle. A linear and shift-invariant digital filter h is characterized by its impulse response $h(n)$, such that the action of the filter h on the signal x produces an output signal y given by

$$y(n) = \sum_{i=-\infty}^{\infty} x(i)h(n-i) = \sum_{i=-\infty}^{\infty} x(n-i)h(i) =: x(n) * h(n)\,, \tag{6.2.2}$$

where the last part of the equation defines the convolution sum of x and h. A filter capable of real-time signal processing must be causal, i.e. its output at cycle $n = n_0$ may only depend on the input signal for $n < n_0$, which requires its impulse response to be a causal sequence, i.e. $h(n) = 0$ for $n < 0$. We furthermore define a stable filter as one for which bounded input signals produce bounded outputs, which is equivalent to the condition

$$\sum_{i=-\infty}^{\infty} |h(i)| < \infty\,. \tag{6.2.3}$$

FIR filter theory

A simple way to implement a digital filter is to numerically compute the convolution sum of the input signal with a filter sequence h of finite length n_{FIR} (also called the number of filter taps) by storing the last n_{FIR} samples of the input signal in memory. This type of filter is known as finite impulse response (FIR) filter and has been implemented on the Red Pitaya by other groups [RS16] with $n_{\mathrm{FIR}} = 25{,}600$. Advantages of

FIR filters are the ease of implementation and stability by design, since a finite sum of finite filter coefficients is always below infinity (see Eq. 6.2.3). FIR filters were considered during early stages of PyRPL development, but eventually abandoned due to a central limitation of FIR filters: since the filter transfer function is essentially given by the discrete Fourier transform of the impulse response h, the frequency resolution δf in the definition of the transfer function is given by the sampling rate $f_s = 1/T$ divided by the number of taps n_{FIR}. The limited availability of FPGA resources puts constraints on the maximum number of taps and thereby usually requires a significantly reduced sampling rate T with respect to the hardware limit to obtain the frequency resolution that is necessary to cancel narrow piezo resonances. By design, a discrete-time filter has a minimum delay of $T/2$, which may jeopardize the loopshape for feedback applications and therefore requires a compromise between delay and frequency resolution. In [RS16], this compromise limited the achievable unity-gain frequency of the feedback loop to 2.8 kHz. We will show that the IIR filter of PyRPL can realize similar filter complexity with sampling frequencies of several MHz.

IIR filter theory

The output y of an IIR filter of length (M, N) is given by the equation

$$y(n) = \sum_{i=0}^{M} b_i x(n-i) - \sum_{j=1}^{N} a_j y(n-j)\,. \tag{6.2.4}$$

By recursively inserting the values of $y(n-j)$ into the right-hand side of the equation in order to write the equation in the form of Eq. 6.2.2, we realize that the equivalent impulse response of this type of filter has an infinite length for nonzero coefficients a_j. By feeding past output values $y(n-j)$ for $j > 1$ to the filter input, this infinite impulse length is realized with a finite amount of computations per cycle, which suggests that IIR filters provide a very resource-efficient filter implementation. However, we will see that the infinitely long impulse response sequence requires care in the filter design and implementation in order to guarantee the stability of the IIR filter.

In order to derive the frequency-response of the filter, we first introduce the z-transform as the discrete-time equivalent of the Laplace transform in continuous time. We define the z-transform of a sequence $x(n)$ as

$$\zeta[x(n)] \equiv X(z) := \sum_{n=-\infty}^{\infty} x(n) z^{-n}\,, \tag{6.2.5}$$

where z is the complex-valued discrete frequency. From the definition, it follows for the z-transform of a shifted sequence that

$$\zeta[x(n+n_0)] = z^{n_0} X(z)\,. \tag{6.2.6}$$

With these two relations, we can now return to Eq. 6.2.4, bring the sum involving $y(n-j)$ to the left-hand side of the equation sign, and define $a_0 = 1$ in order to incorporate $y(n)$ into the sum over y. We then apply the z-transform to both sides and re-order terms to obtain

$$H(z) := \frac{Y(z)}{X(z)} = \frac{\sum_{i=0}^{M} b_i z^{-i}}{\sum_{j=0}^{N} a_j z^{-j}}\,. \tag{6.2.7}$$

H can be seen as the impulse response to an input signal $x(n) = \delta_{n,0}$ in z-domain by inserting the z-transform $X(z) = 1$ into the above equation. We will later discuss how discrete frequencies z are mapped to continuous frequencies s.

According to the fundamental theorem of algebra, the numerator and denominator of Eq. 6.2.7 can be written in the form of products with a constant pre-factor C, a number of M zeros z_i, and N poles p_j:

$$H(z) = C\frac{\prod_{i=1}^{M}(z^{-1} - z_i^{-1})}{\prod_{j=1}^{N}(z^{-1} - p_j^{-1})}\,. \tag{6.2.8}$$

While these zeros and poles determine the stability of the filter, Eq. 6.2.4 suggests that the filter can be conveniently implemented as the convolution sum of a finite number of previous input and output signals with the coefficients a_j and b_i from Eq. 6.2.7. For a given set of coefficients computed from the poles and zeros, round-off errors due to finite numerical precision in the filter implementation may lead to arbitrarily large errors for the actually implemented poles and zeros as the orders M and N of the polynomials grow for more complex filters. This may lead to imprecisely implemented or even unstable filters. For this reason, IIR filters are generally separated into second-order sections (also known as "biquads"), i.e. rational functions of at most second degree, for which the round-off error is manageable. Several second-order sections can then be implemented in series or in parallel. While a few drawbacks arise from the parallel form, notably the impossibility to realize poles or zeros of multiple order, we have chosen this option in order to achieve a slightly better noise performance with arbitrary filter shapes.

IIR filter implementation

When designing an IIR filter, we specify the desired transfer function as a list of N complex poles p_j and M complex zeros z_i at Laplace frequencies $s = i\omega + \gamma$. A filter is termed proper if $N \geq M$, such that the filter transfer function for $|s| \to \infty$ remains bounded. We restrict our discussion to proper IIR filters since improper filters are expected to result in unstable closed-loop behavior at frequencies with insufficient phase margin due to the delay of the ADCs and DACs. We furthermore require no identical poles in order to be able to perform the partial fraction expansion below. With these constraints, the analog filter transfer function can be written in Laplace domain[4] in the form

$$H_\mathrm{a}(s) = C\frac{\prod_{i=1}^{M}(s - z_i)}{\prod_{j=1}^{N}(s - p_j)} = D + \sum_{j=1}^{N}\frac{r_j}{s - p_j}\,, \tag{6.2.9}$$

where C and D are constants and the residues r_j in the partial fraction expansion after the second equation sign can be computed with low numerical error from the poles and zeros with the Heaviside cover-up method [TF88, Wik16c]. D is only non-zero for strictly proper filters, i.e. if $N = M$. In order to find the corresponding filter

[4] We define the Laplace transform of a function $f(t)$ as $F(s) = \int_0^\infty e^{-st} f(t)\mathrm{d}t$.

in discrete time, we first compute the expression for the impulse response in continuous time domain through the inverse Laplace transform:

$$h_{\mathrm{a}}(t) = D\delta(t) + \sum_{j=1}^{N} r_j e^{p_j t}\Theta(t)\,, \tag{6.2.10}$$

where $\Theta(t)$ denotes the Heaviside step function. The impulse response in discretized time is

$$h(n) = h_{\mathrm{a}}(nT) = D\delta(nT) + \sum_{j=1}^{N} r_j (e^{p_j T})^n \Theta(n)\,, \tag{6.2.11}$$

which allows to compute its z-transform to be

$$H(z) = D + \sum_{j=1}^{N} \frac{r_j}{1 - e^{p_j T} z^{-1}}\,. \tag{6.2.12}$$

For the implementation of the filter, we combine pairs of terms after the summation symbol into second-order sections. Simple algebra yields an expression of the form of Eq. 6.2.7 for each second-order section, from which the coefficients $a_0 \equiv 1$, a_1, a_2, b_0 and b_1 for a second-order section implementation after Eq. 6.2.4 can be read-off. For these coefficients to be real and thereby implementable, poles and zeros from the filter specification of Eq. 6.2.9 with non-zero imaginary parts must be accompanied by their complex conjugates. The constant feedthrough term D is implemented as a separate second-order section with $b_0 = D$, $a_0 = 1$, and all other coefficients equal to zero.

In the Verilog code, we only implement a single biquad filter module, which computes the output for the coefficients of each second-order section in subsequent FPGA clock cycles. The cumulative sum of all second-order section outputs is the filter output. The presented implementation leads to a reduced effective sampling rate equal to the FPGA clock rate divided by the number of second-order sections, i.e. by half the filter order. Filter coefficients are represented as fixed-point numbers with 3 (29) bits before (after) the radix point. As a compromise between filter complexity and FPGA resources devoted to the IIR filter, we allow for a maximum filter order of 28, i.e. 14 second-order sections. The filter is preceded by a first-order low-pass filter to avoid aliasing when the sampling rate is significantly lower than the FPGA clock frequency.

IIR filter design

In order to design the filter loop shape, we usually measure the open-loop transfer function of the actuator that is to be improved. Various manual approaches to filter design are facilitated with the graphical interface shown in Fig. 6.8. Results accomplished with the IIR filter can be found in Fig. 5.22 and in section 6.3.

6.2.6 Higher abstraction layers

While the various DSP modules that were presented in previous sections implement the basic functionality of PyRPL, a limiting factor in early stages of development was the efficient organization of these basic building blocks into logical units. To this

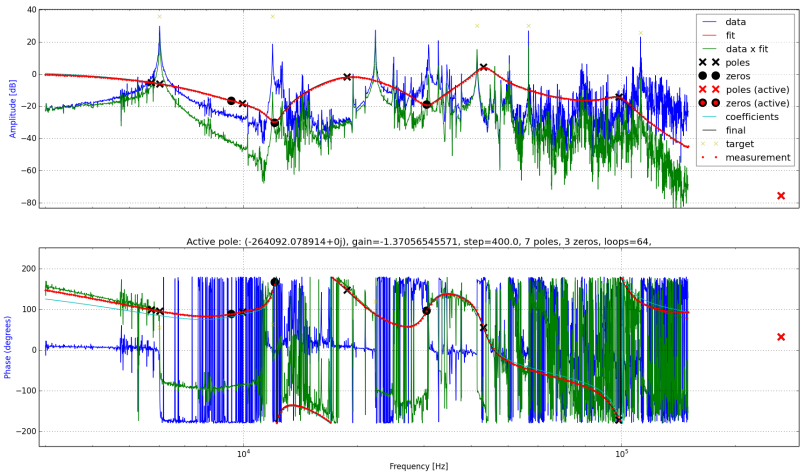


Figure 6.8: Graphical user interface (GUI) for the IIR filter design. The real and imaginary parts of the filter poles and zeros are adjusted with keyboard shortcuts while the resulting transfer function is simulated in real-time. Another keyboard shortcut performs a network-analyzer measurement of the actual filter transfer function and plots the result. Simulation and measurement agree with each other as long as no saturation occurs and the filter specification corresponds to a stable filter. The GUI can also be employed for manually fitting a measured transfer function with a rational function and subsequent inversion-based filter design.

end, an additional layer of abstraction was added to the software by defining higher-level objects that automatically acquire and manage the DSP modules necessary to accomplish the requested behavior. While a PyRPL user can always access the lower-level DSP modules directly, testing our software with internship students demonstrated that the implemented higher-level objects yield a sufficient description of a servo-controller such that the software user does not require knowledge about the DSP modules in order to successfully implement automated feedback loops with PyRPL [Win16].

MemoryTree

Storing the calibration and configuration data of a servo-controller in a dedicated file allows for the clean separation between data and code. Choosing a human-readable format such as YAML, together with the possibility to add comments to the configuration file, allows for convenient manual configuration with a text-editor. A hierarchical organization of the configuration file in a dictionary-like structure further improves the human-readability of the file. The class "MemoryTree" implements the interface between such configuration files and PyRPL in a way that allows the user to modify the configuration data with a text-editor while PyRPL is running. At the same time, PyRPL can modify its configuration file as well, for example to store new calibration values. The configuration file content can be conveniently addressed from the Python code through the implemented mapping of the configuration file hierarchy onto an apparent class hierarchy of the MemoryTree object, which also allows auto-completion

functionality with standard Python interpreters such as IPython or Jupyter Notebook.

Signal

The class "Signal" represents any signal that can be measured with the PyRPL oscilloscope. Particularly important signals with assigned custom names (e.g. "reflection", "transmission", "pdh") are the analog input signals and possibly derived lock-in signals that serve as error signals. The configuration file entry defining a derived signal such as "pdh" specifies parameters such as the demodulation frequency and phase in a dedicated section named "setup". All signals account for possible offsets[5], provide scaling in the units chosen by the user (if sufficient information is available), and have methods such as "sample", "mean", "rms", "max", "min", and "curve" that allow for a convenient analysis of the signal properties from the PyRPL API.

OutputSignal

The class "OutputSignal" inherits all properties from "Signal". The configuration file entry corresponding to an output signal contains additional information about the device connected to the analog output that is described. For example, a piezo actuator can be specified by its maximum voltage rating, its DC-gain in units of m/V, the cutoff frequencies of analog low-pass filters between the Red Pitaya and the actuator, and a reasonable unity-gain frequency for which no oscillation in a feedback application occurs. Furthermore, an IIR filter that is applied to the output signal can be specified, either directly through the IIR GUI, or by providing a list of poles and zeros. By internally associating a PID controller with each output signal, the information about the output transfer function can be used to add a simple method "lock" to the OutputSignal class, which enables a feedback loop with optimal gain for the output and which takes as arguments the Signal object to use as error signal for the lock, the setpoint to lock at and the expected sensitivity of the error signal.

Model

By providing a generic base class to inherit from, the "Model" class yields a way to adapt the lockbox functionality to the required type of servo loop. Standard models for interferometers and Fabry-Perot cavities are available and may be extended for customized behavior. The "Model" class provides a function "calibrate" to implement an automatic calibration of error signals. The default calibration function enables a sweep for all lockbox outputs and records the obtained traces for all input signals, saving the maximum and minimum of the traces in the signal section of the configuration file. The calibration method can be improved by fitting the obtained traces, which however makes it less robust. The calibration information allows to add model functions identical names as the input signals to the "Model" class which predict the expected input signal value as a function of a model parameter. This model parameter is for example the detuning in Fabry-Perot models and the relative phase of the

[5]Offsets are accounted for by adding an offset correction from the configuration file to the measured signals. The offset is typically obtained by the offset calibration function "get_offset" that is called once the analog signals are connected to the Red Pitaya board but the laser is off.

two arms in interferometer models. The model functions are used in conjunction with the function "lock" from the "OutputSignal" class in a way that the setpoint can be specified in the form of the model parameter, while the expected sensitivity is computed automatically from the derivative of the model function at the setpoint. A few additional global parameters from the configuration file, such as the laser wavelength and finesse of a Fabry-Perot cavity, are necessary for this calculation. The resulting method "_lock" for the model class only requires to specify the model parameter value to lock at, in addition to the input and output signals that should be used for the lock. By default, all defined outputs are used for locking. Additional functions may be provided to the Model class, for example the function "find_coarse" in the Fabry-Perot model for the efficient search of cavity resonances with very slow actuators.

Locking sequences

Locking sequences are realized by a function "lock" of the class "Model", that requires a sequence definition from the configuration file. Defining a locking sequence in the configuration file has several advantages: the sequence can be rapidly set up by several iterations of manual modifications of the sequence with a text editor and immediate testing, the history of configured sequences can be conveniently archived by programs such as "git", and the optimal settings for one cavity type can be efficiently exchanged between colleagues. For automatic measurement sequences that require non-default final locking parameters, for example for detuning scans, the sequence parameters from the configuration file can of course be overridden at the call of the "lock"-function that starts a locking sequence.

The method "lock"[6] of the class "Model" iterates over a list of locking stages of the configuration file. Each stage specifies a list of parameters that are passed by default to the "_lock"-function of the "Model" class, or to another custom function if this is specified by a stage name that starts with "call_". This way, locking sequences can be conveniently defined in the configuration file. For example, a typical sequence to enable the PDH lock of a filter cavity reads as follows in the configuration file:

```
stages:
  call_calibrate: {}
  drift:
    input: reflection
    detuning: 0.5
    time: 1.0
    offset: 1.0
  pdh:
    input: pdh
    detuning: 0.0
```

In this example, a calibration of the error signal amplitudes is performed at the beginning by calling the function "calibrate" of the "Model" class without parameters. The other locking stage names ("drift", "pdh") are arbitrary. In the "drift"-stage, the cavity

[6] Note that the higher-level function "lock" for enabling a locking sequence differs from the low-level function "_lock" that is called by the higher-level function during the locking sequence.

is locked to the reflection signal at a detuning of 0.5 cavity bandwidths. The ramp required to deterministically approach the resonance is automatically accomplished with the integrator setting for the lock and by setting the initial value of the integrator memory to the maximum voltage with the parameter "offset". After a time of 1.0 s (see the line "time: 1.0"), we expect the lock on the reflection signal to be established, such that the next stage "pdh" can switch to the PDH error signal at zero detuning. Various other functions are available, for example "scan_to_resonance", which performs a piezo ramp without integrator until a resonance is detected and then allows to pass to the next locking stage. A similar locking sequence as the one described here is shown in Fig. 5.10.

Another function "islocked" of the "Model" class employs the measured values of input signals and the history of function calls to PyRPL to determine whether a cavity is still locked at the desired setpoint. This enables efficient diagnostics and automatic re-locking.

6.3 Photo-thermal instabilities

This section presents the methods used for the optimization of the feedback loop that stabilizes the detuning between the laser and the measurement cavity. First, we present the large modulation synchronous detection (LMSD) scheme that allows to overcome the detrimental effects of vibrations on the characterization of the actuator responses in our system. Next, we present the response of our optomechanical system to modulations of the intracavity power at frequencies from 5 to 120 kHz. The measurements suggest that photo-thermal effects are present at the lower frequencies and modify the stability of the optomechanical system at higher laser powers. The last section presents various attempts to enable the stable operation of our optomechanical system in the unstable regime of high incident power and non-zero detuning.

6.3.1 Large modulation synchronous detection

A common problem during the first characterization measurements at cryogenic temperatures were acoustic vibrations of the cryostat, causing excessive detuning fluctuations of the order of the cavity linewidth. Since common error signals such as side-of-fringe, PDH or homodyne detection are linear in the detuning only over regions of a fraction of the cavity linewidth, the vibrations caused a nonlinear transduction of the displacement fluctuations. Problems arising from such nonlinear transduction mechanisms are

- effective gain fluctuations of the feedback loop, thereby an inefficient reduction of the fluctuations by feedback,
- a large uncertainty in measurements of the transfer functions of cavity length actuators,
- difficulties in measuring the actual amplitude of laser-cavity detuning fluctuations, and

- the appearance of peaks at sum and difference frequencies of actual peaks in the measured displacement spectrum, which are caused by the nonlinearity of the detection, and which render the identification of the original peaks difficult.

The listed measurements were needed in order to design an optimized feedback loop, capable of reducing the large vibration amplitudes to a fraction of the optical linewidth to allow standard detection mechanisms to be used for further optimization and measurements.

LMSD method

We have conceived a detection technique called large modulation synchronous detection (LMSD) that is simple to implement and that can provide a linear error signal over several tens of optical linewidths. The implementation of the technique is identical to PDH or synchronous detection, but requires an uncommon parameter regime. In detail, we require the following:

- As for PDH or synchronous detection, the modulation of the laser-cavity detuning may be accomplished either by acting on the laser frequency or on the cavity length. The modulation frequency Ω_{mod} sets the upper limit for the detectable frequencies in the error signal.
- The modulation depth $\delta\omega_{\mathrm{L}}$ should be significantly larger than the cavity linewidth κ, and larger than the expected vibration amplitude. This ensures that the effective sweep over the optical resonance accomplished by the large modulation covers both an entire resonance and the full vibration amplitude.
- For best performance, the maximum sweep velocity should be small enough to avoid optical ringing (see section 2.2.5): $\delta\omega_{\mathrm{L}}\Omega_{\mathrm{mod}} < \kappa^2$.

In order to derive a theoretical expression for the expected error signal, we start by writing the instantaneous laser frequency as $\omega(t) = \omega_{\mathrm{L}} + \delta\omega_{\mathrm{L}} \sin\Omega_{\mathrm{mod}}t$. In an idealized LMSD scheme, we can approximate the resonance dip at frequency $\omega_{\mathrm{L}} + \Delta$ in the high-pass filtered signal of the photocurrent of the reflected beam as a delta-peak: $r(\omega) = \delta(\omega_{\mathrm{L}} + \Delta - \omega)$. Demodulation at Ω_{mod} produces a signal proportional to the product $r(\omega)\sin(\Omega_{\mathrm{mod}}t)$, and subsequent low-pass filtering can be modeled as the integral over one half-period:

$$\begin{aligned}
\mathrm{LMSD}(\Delta) &\propto \int_{-\pi/2\Omega_{\mathrm{mod}}}^{\pi/2\Omega_{\mathrm{mod}}} \sin(\Omega_{\mathrm{mod}}t)\delta\left(\Delta - \delta\omega_{\mathrm{L}}\sin(\Omega_{\mathrm{mod}}t)\right)\mathrm{d}t \\
&\propto \int_{\Delta-\delta\omega_{\mathrm{L}}}^{\Delta+\delta\omega_{\mathrm{L}}} \tan\left(\arcsin\frac{\Delta - z}{\delta\omega_{\mathrm{L}}}\right)\delta(z)\mathrm{d}z \\
&\propto \frac{\Delta/\delta\omega_{\mathrm{L}}}{\sqrt{1-(\Delta/\delta\omega_{\mathrm{L}})^2}}
\end{aligned} \tag{6.3.1}$$

In the second line, the substitution $z = \Delta - \delta\omega_{\mathrm{L}}\sin(\Omega_{\mathrm{mod}}t)$ was performed. The linear range of the resulting error signal is of the order of $\delta\omega_{\mathrm{L}}$ instead of κ[7]. We note

[7] κ does not appear in the equation due to the δ-peak approximation. While the derived result qualitatively holds even for finite κ, the divergence of the error signal at $\Delta = \pm\delta\omega_{\mathrm{L}}$ is prevented for finite κ.

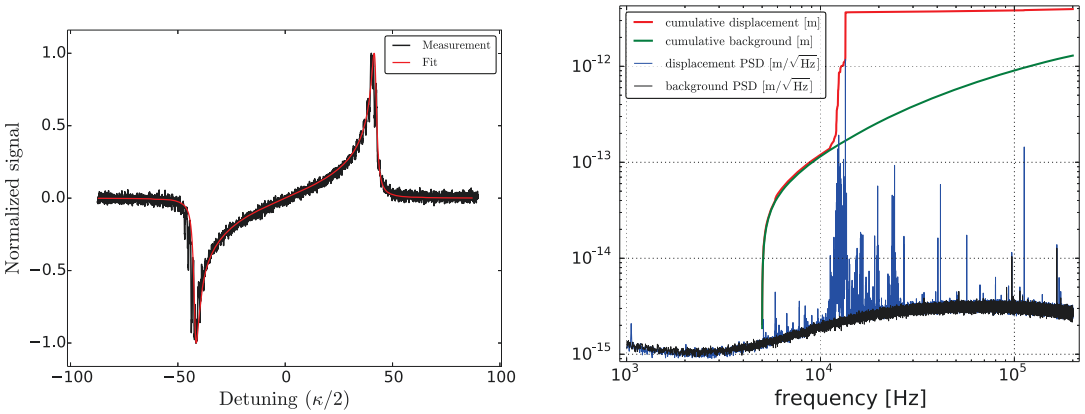


Figure 6.9: Left: Large modulation synchronous detection (LMSD) error signal. The fit was used to calibrate the detuning axis, and reveals that the error signal is linear over several tens of optical bandwidths. Right: Displacement noise of the cryogenic cavity without any feedback above 1 kHz. Since the cumulative noise is of the order of the cavity linewidth, the spectrum was measured with the large modulation synchronous detection. The spectrum was calibrated with a phase-modulation calibration peak.

that the idealized derivation may be realized by replacing the experimental reflection signal by the output of a well-designed trigger on the reflection signal. The presented error signal enables a trade-off between linearity and sensitivity: in the limit of small drive amplitude, the error signal is identical with standard synchronous detection, while in the limit of infinite drive amplitude, laser and cavity only interact during an infinitesimal amount of time per period.

Implementation

In our experiment, we realize the large modulation depth by strongly driving the compression mode of the micropillar at $\Omega_{\mathrm{mod}} \approx 2\pi \times 3.6$ MHz with the antenna actuator, and perform the demodulation with the IQ module of PyRPL. Drifts of the micropillar resonance frequency, which would modify the resulting modulation depth and thereby the error signal slope, are corrected for with four software-implemented feedback loops with an update rate of the order of 1 s, which stabilize the demodulation phase, demodulation gain, drive amplitude and drive frequency[8].

Measurements

Fig. 6.9 shows the experimentally realized error signal, which is reasonably linear over about 20 cavity linewidths. The fit was realized by a numerical calculation of the expected error signal. The right plot of the figure shows the measured noise spectrum and the cumulative integral over the spectrum. While the rms-displacement noise

[8] In principle, the stabilization of drive amplitude and drive frequency should suffice. However, due to the small mechanical linewidth of the order of 50 mHz and the associated long settling time of the mechanical oscillator upon changes of the drive, the performance of the stabilization scheme was significantly improved by adding the demodulation phase and gain, which have an immediate response.

of our measurement cavity is of the order of its linewidth, almost all of this noise originates from a group of peaks between 10 and 15 kHz. The amplitude of these vibrations was found to be correlated with the amount of ^{4}He in the cryostat reservoir. We therefore assume that the dominant source of these vibrations are of mechanical origin. Along with the spectrum, we have used the LMSD error signal to measure the transfer functions of the available cavity length actuators for the laser-cavity detuning stabilization. The measurement has enabled to design a feedback controller for the slow cavity piezo which allowed to maintain a PDH-lock with rms-fluctuations of the error signal of the order of 15% of the peak amplitude.

6.3.2 Photo-thermal instability

After having optimized the PDH-lock as described in the previous section, we achieved a satisfactory stability of the laser-cavity detuning. However, for powers above 10 μW injected to the measurement cavity, we observe the onset of self-sustaining oscillations of the cavity length at frequencies much lower than the compression mode of the micropillar. We have verified that these oscillations are not caused by the feedback loop that stabilizes the laser-cavity detuning in three ways: by a variation of the feedback gain, a variation of the amplitude of the PDH error signal, and by observing the oscillations even during a sweep of the cavity length with no feedback enabled. If the sweep measurement is performed at powers of the order of 100 μW, oscillation frequencies differing by as much as a factor of 10 occur in different regions of the red-detuned side of the optical resonance, while the dynamics on blue side is dominated by the self-oscillation of the compression-mode of the micropillar.

Measurement of the amplitude-modulation response

For a better understanding of the dynamics of our optomechanical system, we have measured the cavity response to an amplitude-modulation of the injected laser beam. For this measurement, two laser beams resonant with the two polarization eigenmodes of the measurement cavity were generated with the AOM setup described in 5.3.2. The first laser beam (heterodyne signal beam) was locked at resonance with the PDH error signal while the detuning of the second laser beam (homodyne laser) was tunable with the AOM frequency and its power modulated with the amplitude-modulator (section 5.3.1). The results of one detuning series of this pump-probe measurement are shown in Fig. 6.10.

Radiation-pressure effects

The five narrow peaks in the response could be approximately identified as mechanical modes of the suspension system of the micropillar by performing finite-element simulations of the mechanical structure with COMSOL and comparing the mode frequencies and masses. Furthermore, these peaks can be excited by any of the other available actuators (piezos, antenna). The damping of these peaks is significantly lower at cryogenic temperature than at room temperature. Zooming on the mechanical peaks, as shown in Fig. 6.11, reveals an optical spring-effect that is linear in the optical power and that shifts the resonance frequencies towards lower values for a red-detuned pump

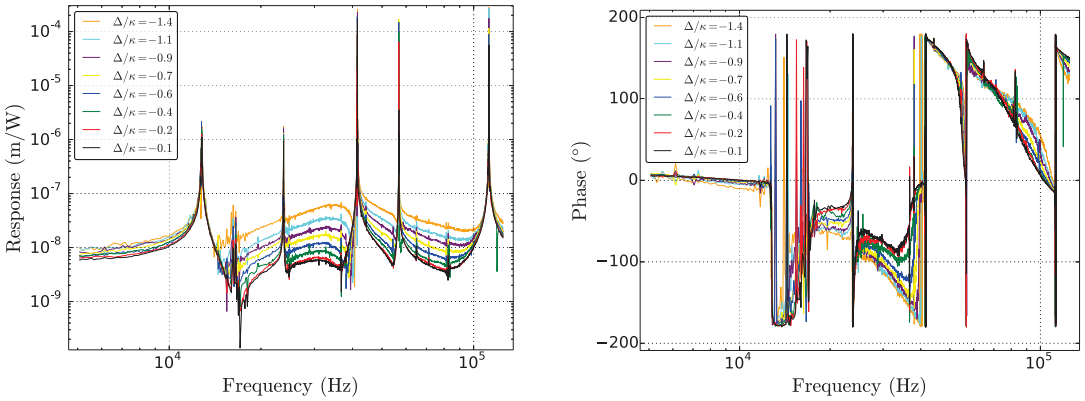


Figure 6.10: Left: Response of the cavity length to a modulation of the radiation-pressure at cryogenic temperature for an injected power of 10 μW. The measurements were acquired with the PDH detection of the heterodyne laser while the motion was driven by a second amplitude-modulated laser into a separate cavity mode. Above the shown frequency range, the response is essentially flat up to the compression mode frequency at 3.6 MHz. Five narrow peaks are discernible between 10 and 110 kHz, which correspond to mechanical modes of the micropillar with effective masses in the 0.1 to 10-mg range. The response between the peaks was limited by electrical cross-talk in this measurement. Right: Phase-response for the curves on the left. At the passage of each mechanical resonance, a phase-shift of π is observed. At the lowest frequencies, the phase envelope features a small phase-advance, which turns into a small phase-delay at higher frequencies.

beam. This sign is expected for an optical spring from radiation-pressure. However, for a pure radiation-pressure response, one would also expect small optomechanical damping for a red-detuned laser beam because the response of the intracavity field to a length-modulation from the mechanical oscillation is slightly delayed due to the finite cavity bandwidth, which results in a radiation-pressure force not in phase with the mechanical motion. While such damping is observed for the 110-kHz mode, we find anti-damping for the lowest-frequency modes at 13, 25 and 56 kHz.

Photo-thermal effects

To explain this effect, we first note that optical delay is negligible at these low frequencies due to the small ratio between the mode frequencies and the cavity bandwidth (14 MHz in this measurement). A mechanism that can explain the phase-advance of the optical force with respect to the intracavity field is the photo-thermal effect: absorption of intracavity light by the micropillar leads to an expansion of the mirror coating and remaining mechanical structure and shortens the cavity length for higher powers, while radiation-pressure tends to elongate the cavity. At very low frequencies, we therefore expect the radiation-pressure and photo-thermal forces to have opposite phase. Towards higher frequencies, the photo-thermal force is expected to experience a low-pass response with a cut-off frequency associated to the characteristic thermal response time of the structure, which leads to an additional delay of $\pi/2$ above the cut-off. The high-frequency photo-thermal force therefore has a $\pi/2$ phase-advance with respect to radiation-pressure, and the sum of the forces can therefore yield any

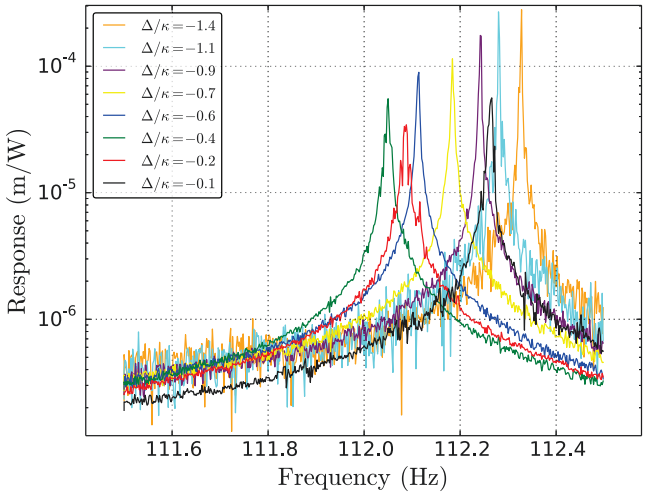


Figure 6.11: Zoom on the response of the drum mode of the micropillar at 112 kHz. Both a strong optical spring and a weak optical damping are observed, both maximum near $\Delta = \kappa/2$.

phase-advance with respect to a pure radiation-pressure response between $\pi/2$ and 0, depending on the ratio of the two force magnitudes. The phase-response plot in Fig. 6.10 reveals that for frequencies around 5 kHz, a phase-advance of about 10° is observed. After the phase-jump of π of the first mechanical mode at 13 kHz, the overall phase response remains still below $-\pi$, which indicates that the response of this mode certainly experiences phase-advance. While the lighter mechanical mode at 112 kHz can be laser-cooled simultaneously with the compression mode of the micropillar and thereby does not jeopardize the stability of our optomechanical system, the lower-frequency modes with significant phase-advance due to photo-thermal effects become unstable at higher powers. Therefore, we have to use external feedback to ensure the stabiliy of our system in this high-power regime.

6.3.3 Feedback-suppression of photo-thermal instabilities

For the stabilization feedback loop, we may choose from all available actuation mechanisms in our system: the antenna that drives the micropillar structure piezo-electrically, the fast or slow piezos that directly change the cavity length, or a modulation of the injected power. We will see that a modulation of the optical power is the only viable option with the current measurement cavity design.

Antenna and fast piezo actuation

Fig. 6.12 shows the measured transfer functions for the antenna and fast piezo actuators. A successful feedback controller was implemented with the antenna actuator at early stages of the experiment. However, the phase response of this actuator is very steep in nearly every frequency region. When the optical power was increased, the resulting optical spring effect resulted in a large phase-shift of the feedback drive at the effective mechanical frequency, which rendered the feedback-controlled system unstable. Attempts to adjust the feedback controller loopshape during the power ramp

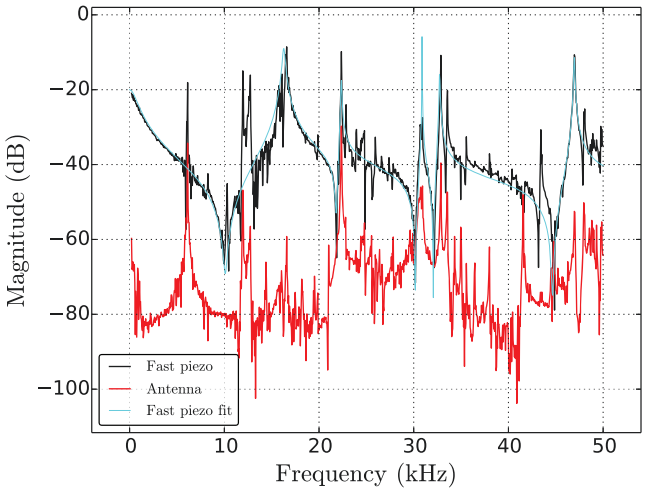


Figure 6.12: Open-loop electro-mechanical transfer function of the fast piezo and antenna actuators, measured by PDH detection with a cavity temperature below 1 K. The coarse features of the fast piezo response can be modeled by a rational function with alternating poles and zeros as shown by the fit. The antenna response has many sharp features and the phase-plot (not shown) reveals that it has non-minimum phase. The voltage signal sent to the antenna actuator is attenuated by about 50 dB at the entry port of the cryostat, such that the transfer function in units of m/V would be above the fast piezo response. This suggests to explain the sharp deviations between the fast piezo response and the fit as direct piezo-electric actuation of the micropillar structure.

were made, but are both impractical and do not resolve the problem that the overall feedback gain is limited through the complicated phase response of the antenna actuation.

As an alternative, we attempted to employ a carefully designed fast piezo which can be seen in Fig. 4.12. The geometry and mounting of the actuator were optimized with finite-element simulations of the response by anticipating the possibility to flatten the effective actuator response with an IIR filter from section 6.2.5. The possibility to use an IIR filter relaxes the design constraints with respect to classical design techniques [BYC+10] such that no flat piezo response, but only a collocated response, i.e. a response without unstable zeros, is required. This is realized in practice by choosing a small piezo size and mounting the piezo as close as possible to the spot that is to be displaced. The realized piezo response is shown in Fig. 6.12 as well. While the response features an alternating sequence of poles and zeros that is typical for collocated actuator-sensor pairs, a number of features characteristic of the antenna actuator appear in the piezo response as well. These features are almost invisible at room-temperature, where the mechanical modes of the quartz structure surrounding the micropillar experience stronger damping. We therefore conclude that most likely, the piezo electrodes act as antenna actuator as well, which is a reasonable assumption for the electrode geometry when the characterization of the piezo-electric response of quartz in Fig. 4.14 is considered. The phase of the antenna features is random with respect to the piezo response envelope, which results in unstable zeros in the piezo transfer function. Consequently, the IIR filter technique cannot be employed to make

the fast piezo actuator useful for feedback at these frequencies. A future cavity design will employ careful electrical shielding between the fast piezo and the quartz structure in order to obtain a collocated piezo transfer function.

6.3.4 Feedback through amplitude-modulated light

The only remaining actuation mechanism for feedback is the amplitude modulation of a laser-beam that is injected into our optomechanical system. While the response to an amplitude-modulation certainly has the best of all presented transfer functions for feedback (see Fig. 6.10), it has the drawbacks that an extra laser beam must be available for actuation, and that this extra laser beam also heats the cryogenic optomechanical system through absorption. Whenever possible, we therefore use the PDH error signal of the heterodyne signal beam to provide the measurement for feedback and the stronger homodyne signal laser-beam as actuator. As long as the relative amplitude-modulation depth of the homodyne signal beam is small, we do not expect the amplitude-modulation feedback, which is limited to frequencies below 200 kHz, to negatively affect the homodyne measurement at frequencies near the compression mode resonance of the micropillar. At later stages of the experiment, measurement and actuation with the same laser were implemented with a modified PDH error signal that is described in section 5.8.

Simple amplitude-modulation controller

Fig. 6.13 shows the transfer function of a simple controller for amplitude-modulation feedback, and the resulting improvement in the measured displacement noise spectrum as a function of feedback gain. The controller transfer function was designed manually with the IIR GUI with the goal to center the response phase of the plant and the controller in series at the peak frequency of each mechanical resonance at $\pi/2$, which is equivalent to the Bode stability criterion[9]. Implementing a simple controller despite the possibility to achieve more complex behavior has the main advantage that the phase and amplitude slopes are generally small. Small phase slopes result in similarly small changes of the effective feedback gain and phase when the optical spring effect shifts the resonances frequencies of low-frequency modes, and thereby render the feedback controller more robust. The achieved improvement is satisfactory: the appearance of squashing around the 112-kHz mode shown in Fig. 6.13 for gains above 1 indicates that the feedback performance is limited by detection noise rather than the stability of the controller. The strong feedback gain generates a displacement noise that is anti-correlated with the detection noise: an out-of-loop measurement with another detector would show excess noise when squashing occurs because of the missing correlation between detector noise and displacement noise. The optimal feedback gain for the 112-kHz mode is therefore $g = 1$. We see that the more massive lower-frequency modes are also damped for higher feedback gains. However, this visibly occurs only at gains above 1, and therefore a compromise between the noise of low-frequency modes and the noise at the 112-kHz mode had to be found with the controller. With gains between 1

[9]The Bode stability criterion states that a closed-loop system is stable if its open-loop transfer function at all frequencies with gains greater or equal to 1 has a phase delay less than π.

and 2, up to 800 μW of light could be injected into the cryogenic measurement cavity while maintaining the system stable.

Complex amplitude-modulation controller

The responses to an amplitude-modulation shown in Figs. 6.14 and 6.10 are visibly different, despite the fact that both curves were measured with the same measurement cavity in different cryogenic runs: most notably, a new peak around 6 kHz is present in the later measurement of Fig. 6.14. We believe that the most likely explanation for this modification of the mode structure is the result of a damage of the membrane between the chip and the decoupling shield of the micropillar sample that was induced by a mechanical shock or thermal stress during heat-up or cool-down. The new resonance is very lightly damped, and its low frequency makes it very susceptible to transduce vibrations into cavity length fluctuations. The resonance is therefore even more problematic for the stability of our optomechanical system than the 13-kHz mode that was shown to be a limiting factor in Fig. 6.9.

While unstable zeros of the piezo and electrode transfer functions prevent the use of these actuators for feedback-suppression of the noise of the 6-kHz mode, amplitude-modulation of the light injected into the cavity is expected to provide a viable actuation mechanism also for this mode. The limitation of the simple feedback controller from the previous section was a trade-off between gains at low and higher frequencies. By implementing a more complicated controller loopshape as shown in Fig. 6.14, we can define both the gain and phase of the controller at a finite number of frequencies. The presented controller was designed by combining highly damped pairs of complex-conjugate poles with lightly damped complex-conjugate pairs of zeros near each mechanical resonance. The ratio of damping and the precise frequencies of the poles and zeros allows to fine-tune the gain and phase at the closest mechanical frequency without significantly modifying the controller response at other mechanical resonances. The fine-tuning was performed while monitoring the displacement noise spectrum in real-time. The controller resulting from this iterative tuning procedure was capable to significantly reduce the displacement noise of both the 6 kHz-mode and other mechanical modes. However, the accomplished damping was still insufficient to prevent the optomechanical system from entering a regime of self-sustaining oscillations when a word was spoken out at normal volume near the experiment with injected powers above 50 μW. In this case, the acoustic excitation from the environment caused the system to spend a significant amount of time at detunings where optical and the effective controller (e.g. due to the optical spring effect) feedback gains and phases are significantly changed. At higher laser power, the changes are sufficient to drive the system far enough from the linear regime where our feedback controller is effective. This problem might be reduced by using a laser with higher power as feedback actuator to enable stronger damping of the mechanical modes. However, the power necessary to significantly cool the 6-kHz mode is of the order of 100 μW, which starts to significantly modify the characteristics of the optomechanical system to study (see next chapter). Rather, it is highly preferable to either reduce environmental noise by a suspension system of the measurement cavity inside the cryostat, or to improve the performance of the fast piezo actuator such that the employed feedback actuator does

not compromise our measurements.

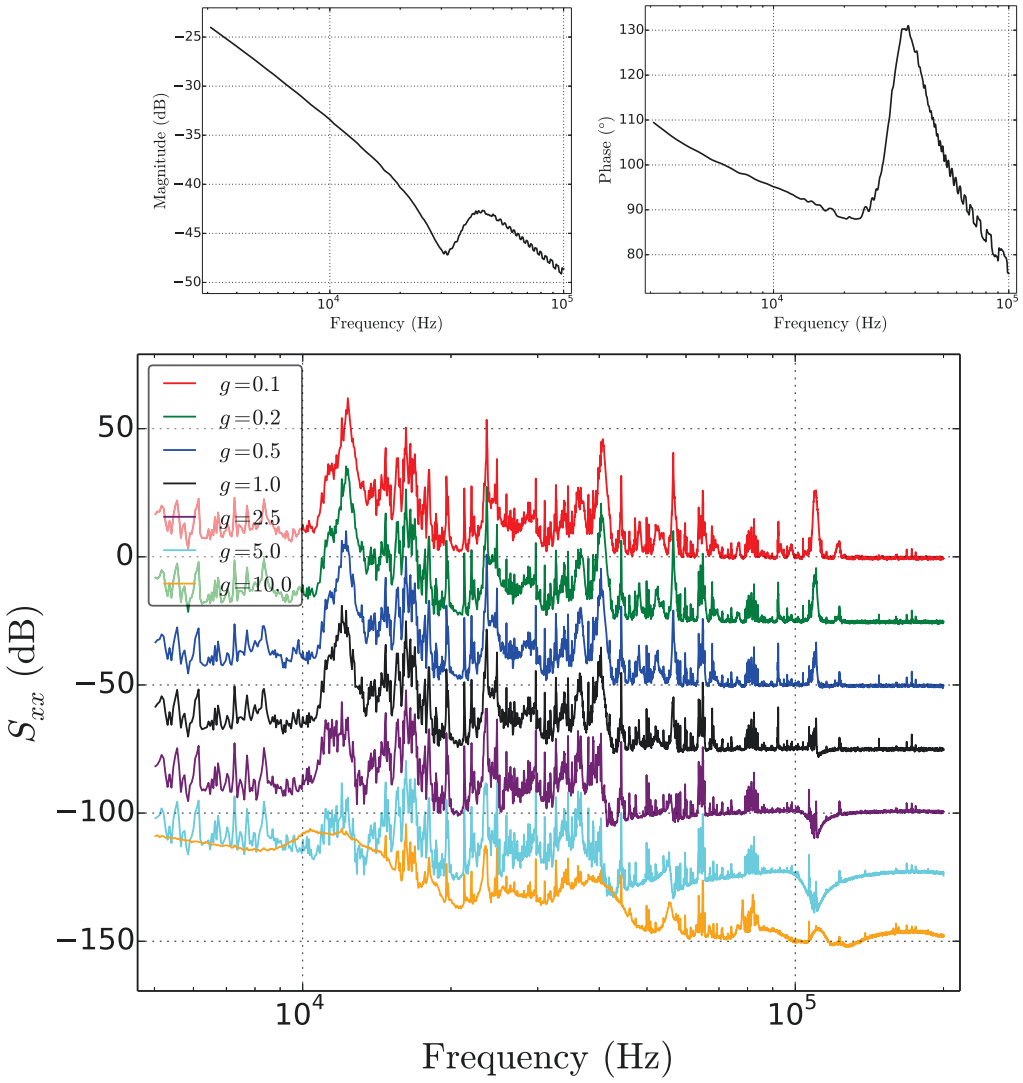


Figure 6.13: Top plots: Magnitude and phase response of the amplitude-modulation feedback controller. Bottom plot: Displacement noise spectra with various feedback gains. Spectra are offset by -25 dB at each increment of the feedback gain for a better visibility.

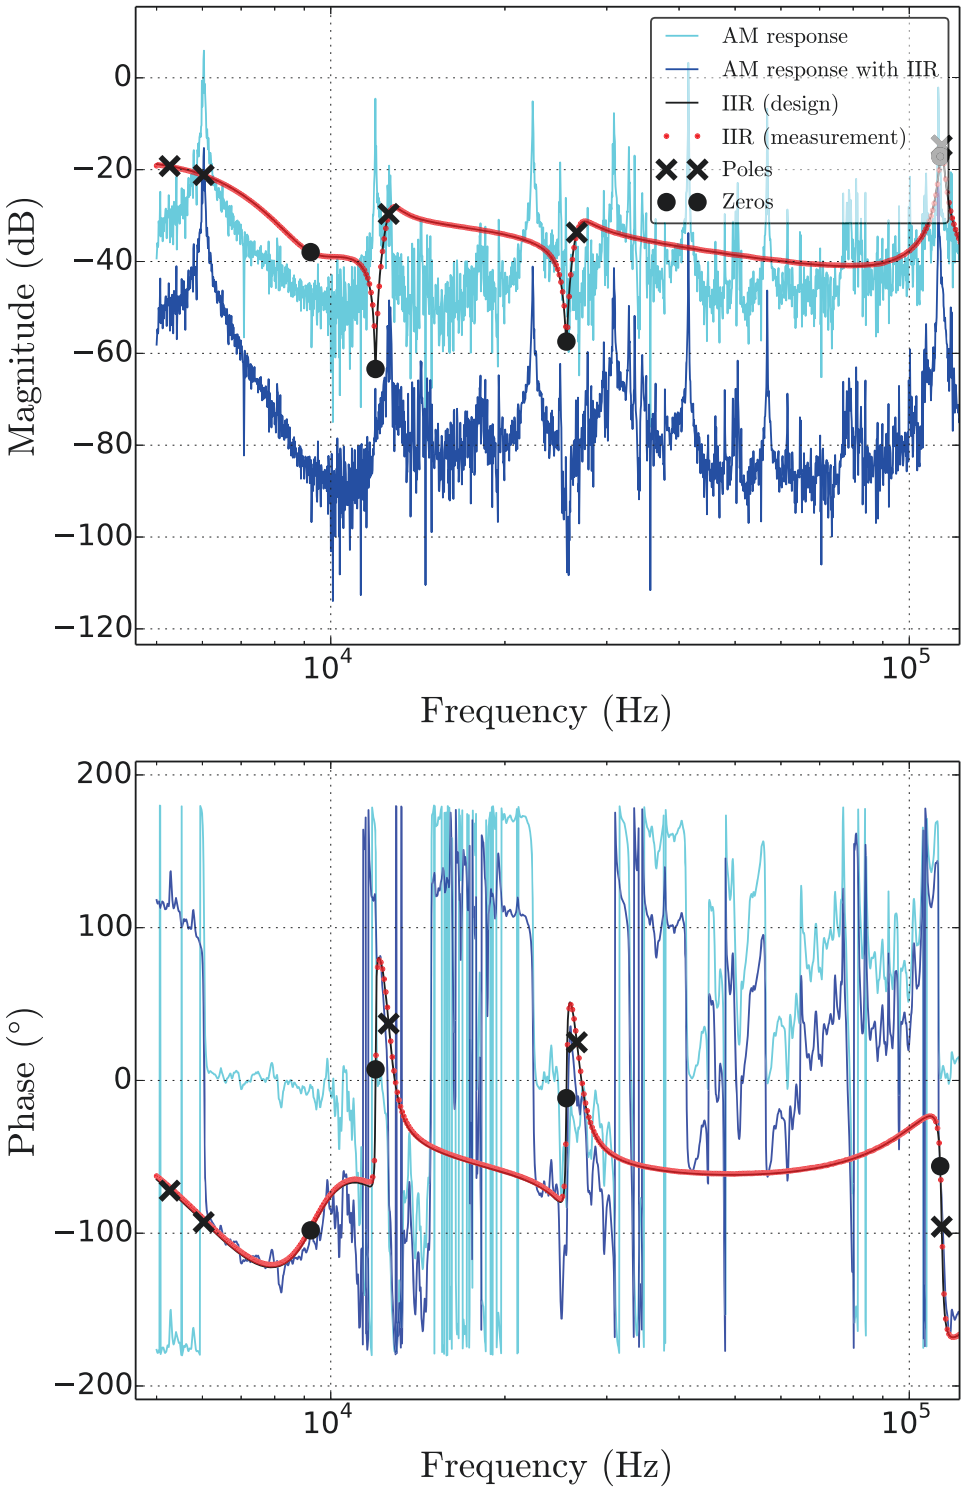


Figure 6.14: Magnitude and phase of the open-loop transfer function of an amplitude modulation of the homodyne laser to the PDH error signal. The designed IIR filter is designed to provide optimal gain and phase at all significant mechanical resonances. The implemented filter that was measured with the internal network analyzer of PyRPL agrees nearly perfectly with the design specifications.

Chapter 7

Cooling close to the quantum ground state and other optomechanics experiments

Mehr Licht!

— Johann Wolfgang von Goethe

This chapter presents a number of optomechanics experiments, both at room temperature and at cryogenic temperature. The room temperature measurements are mainly performed as a pre-characterization of our system and detection calibration. The low-temperature part demonstrates a range of optomechanical effects and constitutes the main scientific results of this thesis. Both OMIT and cooling close to the ground state are performed using a side-of-fringe measurement. The observed spectra are discussed and fully analyzed, taking into account a number of experimental issues. Homodyne spectra are also presented to discuss the results and limitations of a quantum back-action demonstration experiment. Ideas to circumvent the current experimental limitations such as the input mirror noise floor are discussed at the end of the chapter.

7.1 Room-temperature measurements

At room temperature, our optomechanical system has been operated with samples featuring micropillar compression modes with quality factors from 5×10^5 to 3×10^6. More importantly, instabilities were easy to avoid at this temperature due to low parametric gains and low vibrations of the setup. This allowed us to validate our calibration procedure and the convergence of temperature measurements through the Brownian motion of the micropillar compression mode (section 7.1.1). Heating effects from the laser proved to be limited to a fraction of the environment temperature in this temperature region (section 7.1.2). The good stability and autonomy of our setup allow us to conceive a future experiment aiming at the demonstration of quantum backaction effects at room temperature (section 7.3.3).

7.1.1 Brownian motion at 300 K

For the measurement of Brownian motion at room temperature presented in this section, a cavity was assembled with the micropillar sample 20121009C5 and the coupling mirror S6-7. The length of the cavity was estimated to be 98 μm at room temperature with the procedure described in section 3.2.3. With the dilution refrigerator closed and ready for a cryogenic run, one laser beam ("heterodyne laser") was injected into the cavity and all reflected light directed to the avalanche photodiode. A PDH error signal was created by imprinting a phase modulation on the beam at a frequency of 50 MHz and demodulating the measured photocurrent with the same frequency, and an additional 3.6-MHz tone with an electrical power of -21.5 dBm was sent to the phase modulator as a calibration marker. With the cavity locked to the center of the PDH error signal, the RF spectrum of the error signal was acquired with a spectrum analyzer. The acquisition was performed with linear averaging of the electrical power spectrum over 10 traces and by selecting a flat-top filter of 1.0 Hz resolution bandwidth for maximum accuracy in the power estimation. Simultaneously with the frequency band around the mechanical resonance, the spectrum around the calibration peak at 3.6 MHz was also measured[1]. With the calibrated V_π of the phase modulator from Eq. 5.3.2, the voltage amplitude V_0 of the 3.6-MHz tone, and Eq. 5.3.1, the phase modulation depth $\phi_0 = \pi V_0/V_\pi$ of the calibration peak can be computed. From the known cavity length L and laser frequency ν_L and by assuming a flat detection response between the calibration peak and the measurement frequency band, Eq. 2.2.51 allows us to find the equivalent displacement corresponding to the calibration peak, and thereby convert the acquired power spectrum in the measurement frequency band into displacement units. Fig. 7.1 shows a zoom of the calibrated spectrum around the resonance frequency of the micropillar compression mode.

To further process the data, a model consisting of a Lorenzian with an additional incoherent background term is used to fit[2] the spectrum:

$$S(x) = \frac{\text{scale}}{1 + \left(\frac{x - x_0}{\text{bandwidth}}\right)^2} + y_0\,. \tag{7.1.1}$$

A natural question is whether a spectrum should be fitted in linear or logarithmic scale, i.e. whether the least-sum-of-squares optimization should be performed over the array of differences between measured datapoints and the model, or over the array of differences of their logarithms. Denoting the model function as $S(x)$, the data array as $D(x)$, the resulting linear error as $e_\mathrm{lin}(x) = D(x) - S(x)$, we find for the difference of logarithms $e_\mathrm{log}(x)$ in the case of small relative errors $e_\mathrm{lin}(x)/S(x) \ll 1$

$$e_\mathrm{log}(x) = \log(D(x)) - \log(S(x)) = \log\frac{S(x) + e_\mathrm{lin}(x)}{S(x)} \approx \frac{e_\mathrm{lin}(x)}{S(x)}\,. \tag{7.1.2}$$

[1] We use a MXA N9020A spectrum analyzer from Agilent/Keysight Technologies together with the software Keysight 89600 VSA, which allows to acquire spectra in multiple frequency bands simultaneously as long as they fit into the baseband of the analyzer (typically $10 - 25$ MHz).

[2] All fits are performed within our home-programmed "pyinstruments" environment and employ the function "leastsq" of the Python package scipy.optimize. Through continuous usage, we have accumulated a large (> 50) library of standard fitting functions with associated algorithms to guess the initial fit parameters, which have proven their robustness on different data from multiple experiments. Variable names in this text are identical with the ones in the library.

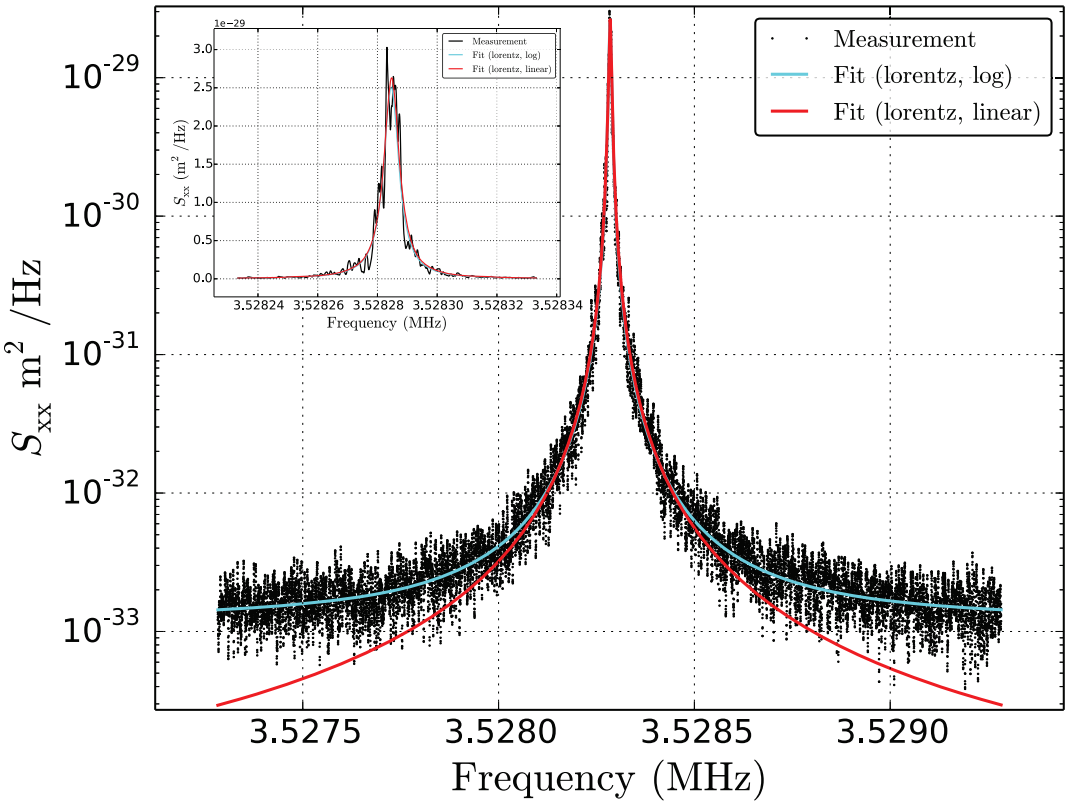


Figure 7.1: Calibrated displacement noise spectrum around the compression mode frequency. Only 2 kHz out of the acquired 5 kHz wide spectrum are shown for better visibility of the central peak. The measurement is fitted with the sum of a Lorentzian and a constant background in linear and logarithmic scales, which results in two different fits. The inset shows a zoom on the central 100 Hz of the main plot in linear scale. Deviations between the measurement and the fits are better visible in the linear plot, in particular the spike on the left of the resonance frequency. These deviations are expected for the fluctuations of the thermal force noise, as the variance estimates in the main text confirms.

We see that for a model function that describes the data with a small relative error, fitting in logarithmic scale with appropriate initial conditions is equivalent to an optimization of a weighted error vector with the weights given by the inverse of the model function. Since our goal is to ultimately obtain an estimate of the magnitude of the flat thermal force noise spectrum, optimizing the weighted errors can be understood as the application of the inverse of the Lorentzian filter (mechanical response) through which we collect information about the force spectrum prior to the optimization. We therefore expect better results with a logarithmic fit as long as the premise of a model that correctly describes the data holds. Fig. 7.1 illustrates this, as the constant background term y_0 does not converge to the measurement background in the linearly weighted fit because the error associated with this background is negligible in comparison with the fluctuations of the data near the more than 10^4 times higher peak. The linear fit of other spectra has resulted in unphysical negative background terms or background terms above the actual background level. While this fault is not discernible in a linear plot of the data and fit (see the inset of Fig. 7.1), the linear plot more drastically shows the power fluctuations over the different spectral bins. We will come back to this observation later in this section.

The above-described measurement was repeated over the course of 8.5 hours and a total of 241 spectra were obtained, out of which 235 spectra were post-selected based on a "waslocked" property that indicates that the reflected intensity never reached 90% of the off-resonant value during the entire acquisition, even though no easily visible anomaly could be found in most of the affected spectra. A thermometer on the mixing chamber plate of the cryostat indicated a decrease of the plate temperature from 298.5 to 290.8 K at a nearly constant rate during the acquisition, which is a result of other cryostat parts already being pre-cooled for subsequent cryogenic experiments. By plotting the observed resonance frequencies versus the cryostat temperature, we find a clean linear dependence with a slope of -263 Hz/K, which is in good agreement with an earlier measurement of -247 Hz/K with another sample in section 3.1.4 and supports the assumption that the sample temperature closely follows the temperature of the mixing chamber plate.

In order to estimate the temperature of the thermal bath that is coupled to the compression mode, we first integrate Eq. 2.1.22 to obtain a simple analogue to Eq. 2.1.14 for a scalar measurement:

$$T_{\mathrm{eff}} = \frac{m_0 \Omega_0^2}{k_B} \left\langle x_{\mathrm{m}}^2 \right\rangle . \tag{7.1.3}$$

While the effective mass m_0 was precisely measured in section 3.1.3, the frequency Ω_0 can be deduced from the measurement with an accuracy of a few parts per million. The displacement variance $\langle x_{\mathrm{m}}^2 \rangle$ can be extracted from the spectra either by computing the combination $\pi \times$ scale $\times$ bandwidth for the linear and logarithmic fits, or directly by numerical integration of the calibrated spectra, possibly accounting for the background level with a numerical estimator or the background level estimated with the Lorentzian fit. While fitting is expected to yield the best results due to the constraints imposed by the underlying model, the numerical methods are expected to be more robust to experimental imperfections such as frequency drifts of the mechanical oscillator or otherwise deformed lineshapes.

T_{eff}	Lorentzian (log)	Lorentzian (linear)
num, no bg	307 ± 22.8	307 ± 22.8
num, num bg	300 ± 22.7	300 ± 22.7
num, fit bg	300 ± 22.7	302 ± 28.0
num, fit bg, weighted	302 ± 23.0	1883 ± 6618
fit	287 ± 21.9	302 ± 28.0

T_{env}	Lorentzian (log)	Lorentzian (linear)
num, no bg	306 ± 7.1	318 ± 51.3
num, num bg	298 ± 6.9	310 ± 50.2
num, fit bg	299 ± 6.9	315 ± 65.7
num, fit bg, weighted	300 ± 5.9	2087 ± 7258
fit	285 ± 5.6	315 ± 65.9

Table 7.1: Mean values and variances of the temperatures deduced by different data-processing methods of 235 spectra. T_{eff} denotes the effective temperature. T_{env} is the environment temperature inferred with Eq. 7.1.6 by estimating Q_0 from the mean damping rate of all curves and inferring Q_{eff} from the fit. The indicated algorithms are: numerical integration of the full spectrum (num, no bg), numerical integration and subtraction of the numerically estimated background level (num, num bg), numerical integration and subtraction of the background estimated by the fit (num, fit bg), numerical integration with background subtraction and weights given by the inverse of the fitted function (num, fit bg, weighted), and estimation of the area of the fitted function from the analytical expression (fit).

The results are shown in Table 7.1. We first discuss the mean values of T_{eff} for numerical estimation, where we note that the photodiode background consistently amounts for an excess temperature of about 7 K if not properly accounted for. All methods that make use of fitted parameters seem to work better when the fit is performed logarithmically, which is expected from the discussion above. The mean temperature deduced from the converged fit in logarithmic scale is about 4% lower than the numerical values, which is most likely a result of temperature drifts of the mechanical resonance frequency (on average 38% of the mechanical linewidth during the acquisition time), rendering the imposed Lorentzian lineshape inaccurate.

While the variance differs little between the different methods, the lowest observed value of 22 K is obtained with the logarithmic fit. To estimate the stability of our setup, we find from [Gav13] that the variance of the temperature estimator converges with the measurement time τ and the mechanical damping rate γ_{eff} as[3]:

$$\sqrt{\langle (T - \langle T \rangle)^2 \rangle_\tau} = \sqrt{2} \frac{\left(e^{-\gamma_m \tau} - 1 + \gamma_m \tau\right)^{1/2}}{\gamma_m \tau} \langle T \rangle \ . \tag{7.1.4}$$

In our case, this simplifies with $\tau = 10\ \text{s} \gg 1/\gamma_m$ to an application of the law of large numbers

$$\sqrt{\langle (T - \langle T \rangle)^2 \rangle_\tau} = \frac{T}{\sqrt{N}} \, , \tag{7.1.5}$$

[3]We have dropped a factor of $\sqrt{2}$ with respect to [Gav13] since the spectrum analyzer we use employs both quadratures of motion for the energy average.

where $N = \gamma_m \tau$ plays the role of the number of samples. Thus, due to the mechanical coherence time of $1/\gamma_m$, an independent sample of the thermal force noise spectrum can only be obtained at the limited rate γ_m. With the parameters of our experiment, Eq. 7.1.4 yields a temperature dispersion of 21.4 K, which is in good agreement with the obtained measurements[4].

An estimation of the quality factor from the log-scale Lorenzian fits yields $Q_{\text{eff}} = (578 \pm 44) \times 10^3$, i.e. fluctuations of 8 % over the 235 spectra. At first, these quality factor fluctuations were attributed to optical feedback that can provide excess damping or heating to the mechanical oscillator while only introducing an insignificant fluctuating force given by the radiation pressure force spectrum. With this assumption and by estimating the intrinsic (in the absence of feedback) quality factor to be $Q_0 = 574 \times 10^3$, we can deduce the temperature T_{env} of the micropillar environment (temperature of the bulk quartz) as

$$T_{\text{env}} = T_{\text{eff}} \frac{Q_0}{Q_{\text{eff}}} \,. \tag{7.1.6}$$

The obtained values are shown at the bottom of Table 7.1. The linear fits yield imprecise values for Q_{eff} and therefore lead to much larger variances for T_{env} than for T_{eff}. In contrast, the logarithmic fit results in environment temperatures with up to four-fold smaller variance than the variance of the effective temperature, well below the limit imposed by Eq. 7.1.4. To explain this deviation, we recall that when computing the environment temperature, even with numerical methods, we divide each effective temperature by the quality factor Q_{eff} obtained by the fit. Fig. 7.2 shows that these two variables are strongly correlated, which results in the obtained compression of the temperature distribution.

To explain the observed correlations, it was suggested that optical feedback should explain the observed fluctuations of the quality factor. With the injected optical power of 500 nW and mean detunings of at most 2 % of the optical linewidth, we estimate the maximum quality factor variation from optical feedback to be 3 %, which is not sufficient to explain the observations. Together with the Gaussian distribution of the extracted quality factors, this suggests that the quality factor measurement by Brownian motion underlies the same statistical constraints as the effective temperature, where we have also observed a relative variance of 9%. By forcing a fixed intrinsic quality factor Q_0 that can be a large systematic source of error, we are able to remove part of these fluctuations from the extracted temperatures and reduce the variance. To test this hypothesis, the logarithmic fitting procedure was repeated with a constant setting of the quality factor, i.e. Q_{eff} was no longer a free parameter. From these fits, we extract an *effective* temperature $T_{\text{eff}} = 285 \pm 5.5$ K, which is identical with the mean value and variance obtained for the *environment* temperature with an unconstrained fit.

We draw two conclusions from our observation:

1. If the quality factor of a mechanical resonance is independently measured or

[4]We use $\tau = 10$ s for the measurement time since 10 averages with a resolution bandwidth of 1 Hz were acquired. Instead of 1 s, the acquisition of each average took about 3.8 s due to window function constraints of the spectrum analyzer. While Eq. 7.1.5 applies to our measurement, Eq. 7.1.4 does not due to the dead time between successive averages. For the latter to work, only 1 average should be acquired while setting the resolution bandwidth to $1/\tau$.

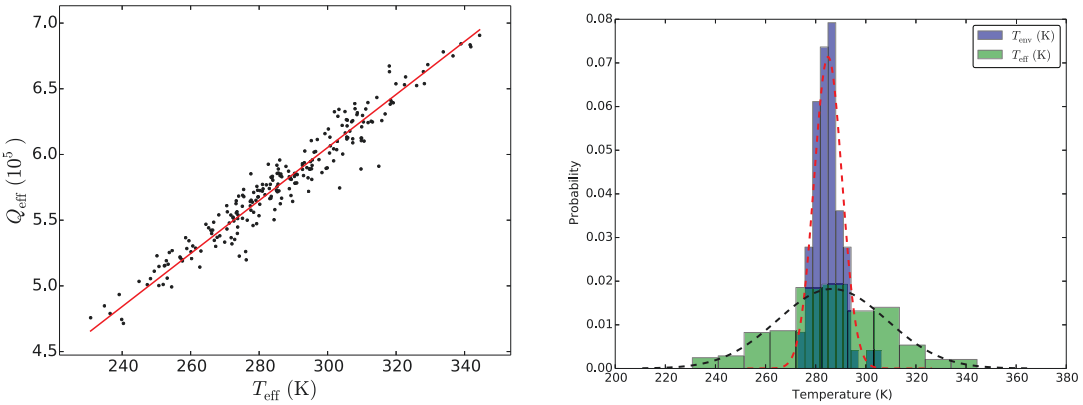


Figure 7.2: Left: Effective quality factor Q_{eff} versus effective temperature T_{eff}. Each datapoint corresponds to the values extracted from a lograrithmic Lorentzian fit of one spectrum. The red line is a linear fit that is constrained to intersect the origin. Effective quality factor and temperature are highly correlated. Right: Histogram of the obtained effective temperatures T_{eff} and environment temperatures T_{env} from 235 spectra. The dashed lines are Gaussians with a mean and standard deviation of the data.

otherwise precisely known, the effective temperature can be estimated faster by using this knowledge than by a naive measurement. In this case, deriving the environment temperature with Eq. 7.1.6 from the effective temperature yields the same result as extracting the effective temperature from a constrained fit. Using Eq. 7.1.6 has however the advantage of allowing numerical methods to be used for the spectral integration when imperfect spectral signatures are observed.

2. If strong dynamical backaction effects are present, we are forced to use Eq. 7.1.6 in order to deduce the environment temperature. The variance that can be obtained is below the thermal limit imposed by Eq. 7.1.4. However, this improvement of the averaging time comes at the cost that the intrinsic quality factor must be assumed constant or accurately known from a model. Any relative error in this knowledge is directly transferred to an identical error of the environment temperature.

The work presented in this section could be extended further: it remains to derive a theoretical expression for the measurement-time dependent estimator of the mechanical quality factor (or equivalently the autocorrelation time), and thereby to analytically confirm that the observed correlations are indeed expected. We note that by writing the estimator of the mechanical quality factor as a term proportional to the estimate of the logarithm of the autocorrelation function, the calculation should be possible in a straightforward manner analogous to the calculations of expectation values in [Gav13, Ver10, GVK12].

7.1.2 Thermal resistance

One of the largest obstacles to the observation of quantum effects at any temperature (aside from the thermal displacement noise of the micropillar) is heating of the sample by laser absorption. From the employed optical power and estimated frequency-shift of 2.5 kHz in the measurement of Fig. 3.1.3, together with the frequency-dependence on temperature inferred from Fig. 3.12, we can obtain a first estimate of the thermal resistance of a micropillar in air at room temperature. Assuming an absorption of the order of 2% for the gold-coated micropillar sample, which is consistent with finesse measurements of cavities constructed with gold-coated samples, we estimate the heat resistance to be about 12×10^3 K/W. This value is already orders of magnitude above the value of 43 K/W reported in [Kuh13] that was obtained with a different sample. The simplicity of the measurements does not leave much room for experimental errors. We thus search an explanation in the difference of the samples:

- Sample 20121002C9 (Fig. 3.1.3) was gold-coated and had a vacuum quality factor of 1.5×10^6. The absorption happens predominantly on the gold-coated surface, and the high quality factor suggests a thin membrane, yielding the high thermal resistance of 12×10^3 K/W.

- The sample in [Kuh13] had a quality factor of the order of 10,000 and was then fully (i. e. also on the membrane) coated with a dielectric mirror, which reduced the quality factor to 3,000. The thermal resistance is expected to be much lower since both the thicker membrane (estimated from the lower Q) and the mirror coating on the membrane contribute to heat conduction. The absorption was estimated from the absorption dip of a Fabry-Perot cavity, though it is unclear whether the micropillar surface or other parts of the cavity absorbed the power. Both differences thus contribute to the lower estimate of 43 K/W.

Since both measurements were performed at ambient pressure in air, we can conclude that the heat resistance of air is larger than the smallest observed value of 12×10^3 K/W, and that the membrane thickness is the main factor determining the heat resistance.

7.2 Cooling experiment

When our optomechanical system is cooled to liquid helium temperatures or below, we observe three main differences with respect to the room temperature behaviour:

- the thermal force noise spectral density is obviously much lower,

- the quality factor of the compression mode, but also of other low-frequency modes, is much higher, and

- the level of vibration is higher than at room temperature.

The higher vibration level in dilution-mode operation is strongly related to the amount of ^{4}He in the reservoir, and can be kept at a manageable level by performing the experiment with a nearly empty reservoir. While the higher quality factor of the compression mode was expected (see section 3.1.4), the weakly damped low-frequency

modes easily jeopardize the lock stability. The sample 20121009C5, with which all experiments in this section were carried out, has the following modes (see section 6.3 for details):

- the drum mode of the micropillar at 112 kHz,
- two tilt modes of the micropillar at 41 and 56 kHz,
- a group of modes in the region from 6 to 30 kHz that are likely drum and tilt modes of the ensemble of micropillar and decoupling shield.

The micropillar drum mode is the low-frequency mode with the highest parametric gain and is unstable at blue detunings just like the like the compression mode. This mode can therefore be simultaneously cooled with the compression mode for red detunings. If instabilities related to parasitic electrical feedback (see section 4.1.4) from the analog PDH error signal are avoided either by the careful selection of the sign of the error signal slope, or by a low-gain feedback loop acting through the antenna actuator, the optomechanical system is found to be stable up to injected powers of the order of 25 μW on the red side of the resonance.

After a short description of the measurement setup in section 7.2.1, section 7.2.2 presents measurements of optomechanically induced transparency, which show a significant modification of the optomechanical cavity response to an optical phase modulation. The damping rates extracted from these measurements agree well with estimations of the cooling rate from ringdown measurements. Section 7.2.3 presents the measurements of Brownian motion in this regime. At the maximum injected power of 25 μW, we have observed our mechanical oscillator at a mechanical occupation number corresponding to only 20 phonons. Unfortunately, due to heating by laser absorption, the accessible parameter regime with this detuning remains about one order of magnitude short of the one where quantum backaction dominates the thermal force noise. Absorption heating is therefore discussed in detail in section 7.2.4. Even at the lowest incident powers of the order of 100 nW, a parametric instability of the compression mode of the micropillar was observed. Operating the system near resonance with a PDH detection was insufficient for a reliable quality factor estimation due to the impossibility to exclude dynamical backaction effects from the residual detuning uncertainty. Section 7.2.5 describes a method that allowed us to reliably estimate the mechanical quality factor with incident powers more than 100 times above the parametric instability threshold power.

7.2.1 Measurement setup

For the measurements described in this section, only one laser beam was used. The setup was modified with respect to Fig. 5.2 by placing the nonresonant phase modulator in series with the amplitude modulator to provide a way to calibrate the measured displacement spectra. The beam was injected through the "Laser 1 input fiber" in Fig. 5.2 and measured by the "Laser 2 APD". At the time of the experiment, an imperfect version of this APD (Thorlabs APD110C) was used. Subsequent characterization of the detector revealed that the DC output voltage of the diode was not linear but closely displayed a square root dependence on the incident optical power. Fortunately,

it was found that the AC-response up to 10 MHz remained flat for all powers with a gain proportional to the DC-gain. We confirmed that the photodiode's high-voltage source had a too high impedance, causing a drop in bias voltage proportional to the photocurrent which in turn reduced the avalanche gain. The drop in avalanche gain allowed to operate the photodiode with incident powers up to 50 μW, which is more than ten-fold above the expected saturation voltage.

The DC-signal from the photodiode was used as an error signal to lock the laser to the red side of the optical resonance. While a nominal detuning $\Delta = \kappa/2$ was chosen, the effective detuning was deduced to be $\Delta = 0.62 \times \kappa/2$ from the nonlinear photodiode response[5]. For incident powers ranging from 50 nW to 25 μW, the photodiode darknoise, optical sweep traces, the error signal of the locked cavity, ringdown measurements, response measurements to a phase modulation (OMIT) and both broad-band spectra and $5 - 20$ kHz-wide spectra around the compression mode resonance were acquired. The spectra and frequency response measurements were calibrated with a phase modulation peak in an identical way as the one described in section 7.1.1.

7.2.2 Optomechanically induced transparency

We first discuss the measurements of optomechanically induced transparency (OMIT). Data were acquired with red laser-cavity detuning by recording the transfer function from a phase modulation of the injected beam to the photocurrent of the reflected beam with a network analyzer. A few of the obtained traces are shown in Fig. 7.3. The spectra show that the cavity response to a phase modulation is strongly altered near the resonance frequency of the mechanical oscillator. The deep minimum in the cavity response to a phase modulation inspired the name optomechanically induced transparency for this type of measurement [WRD$^+$10].

To understand this effect, we can see from Eq. 2.3.14 how the amplitude quadrature of the reflected field p_{ref} depends on the input field's phase quadrature q_{in} :

$$\frac{p_{\mathrm{ref}}(\Omega)}{q_{\mathrm{in}}(\Omega)} = R_{12}(\Omega) + R_{15}(\Omega)L_{12}(\Omega)\frac{2\hbar\omega_{\mathrm{L}}}{c}|\bar{\alpha}|\chi_{\mathrm{eff}}(\Omega)\,. \tag{7.2.1}$$

The frequency over which the cavity coefficients R_{12}, R_{15}, and L_{12} change significantly is of the order of κ, and for $\gamma_{\mathrm{eff}} \ll \kappa$, these coefficients may be assumed constant near the mechanical frequency. Thus, the response can be modeled as the sum of a complex Lorentzian χ_{eff} with a coherent background term $Re^{i\phi}$ that depends on the cavity linewidth and the detuning. Fits with this model are also shown in Fig. 7.3, and are in perfect agreement with the measurement.

The effective mechanical damping rates and effective frequencies extracted from these fits are shown in Fig. 7.4, along with rates from additional ringdown measurements and from the Brownian motion spectra that will be presented in section 7.2.3. All three methods yield similar results, overlapping within the error bars of each other.

[5]The setpoint for each power was computed by PyRPL to correspond to the voltage expected for a detuning of $\Delta = \kappa/2$ with a linear photodiode. The power law found for the photodiode DC-voltage-response V as a function of incident power P was $V \propto P^x$ with $x = 0.48$. This allowed for an easy conversion of the ratio of two powers $P_{1,2}$ from two voltages $V_{1,2}$ as $P_1/P_2 = (V_1/V_2)^{1/x}$ and the subsequent computation of the effective detuning.

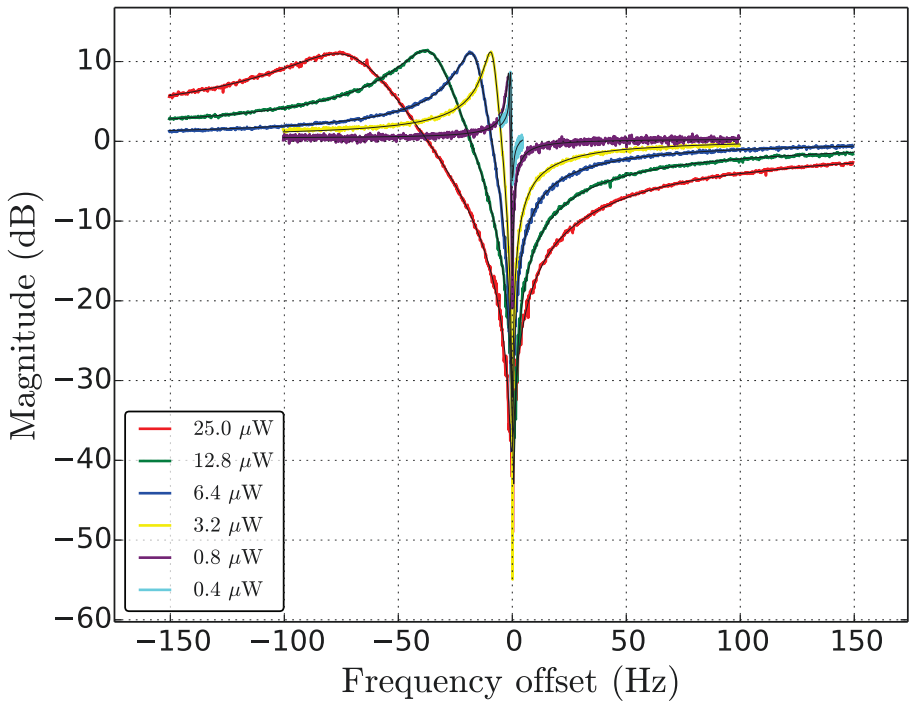


Figure 7.3: Optomechanically induced transparency (OMIT) measurements. During the acquisition with the laser detuned to the red side of the optical resonance, the optical response to a phase-modulation sweep was measured with a network analyzer. The frequency axis indicates the difference from the resonance frequency of the unperturbed mechanical oscillator at 3.578313 MHz.

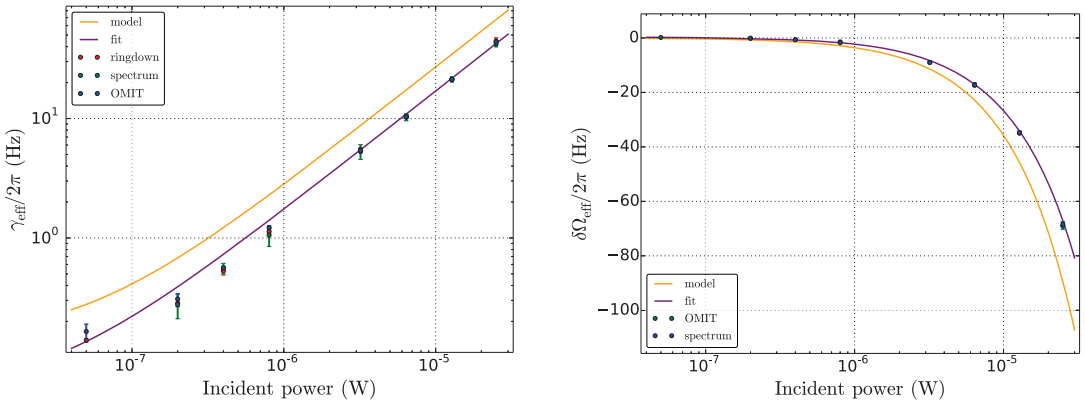


Figure 7.4: Left: Damping rates γ_{eff} obtained by ringdown, spectrum and OMIT measurements for injected powers from 50 nW to 25 μW with a red-detuned laser. Error bars indicate the variance of the available $2-25$ individual measurements with each method at each power. The power dependence from an independent characterization of the system (model) is plotted with a fit of the data. The parameters of the model and fit are shown in Table 7.2. The deviation at low powers can be explained by an incorrect detuning estimation. Right: Effective mechanical resonance frequency shift $\delta\Omega_{\mathrm{eff}} = \Omega_{\mathrm{eff}} - \Omega_0$ versus injected power for the same measurements and models as in the left plot in linear-logarithmic scale. The unperturbed resonance frequency is $\Omega_0/(2\pi) = 3.578313$ MHz. Both plots are indistinguishable from a straight line if plotted in linear scale.

	$L\,(\mu\mathrm{m})$	T (ppm)	P (ppm)	$\mathcal{F}\,(10^3)$	$Q_0\,(10^6)$	$\Delta/(\kappa/2)$	$P_{\mathrm{eff}}/P_{\mathrm{in}}$
Model	58	35	35	90	25	0.62	0.88
Fit	58	35	35	90	70	0.80	0.70

Table 7.2: Parameters of the model (no free parameters) and fit (Q_0, Δ, and P_{eff} adjustable) used to simultaneously deduce the optical spring and damping rates shown in Fig. 7.4. P_{eff} denotes the optical power that is in the spatial mode of the cavity when it is reflected by the input coupler. In contrast, P_{in} is the injected power that is measured outside the cryostat before the injection periscope.

Both the effective damping rate and the frequency-shift display a linear dependence on the optical power, as shown by the good agreement between the linear fits and the measurement. Only the datapoints for the lowest incident powers are slightly off the linear fits, which is easily explained by the difficulty to accurately determine the locking setpoint from the noisy reflection signal in this power region. The fits were performed on the ensemble of the frequency-shifts and damping rates with a model of our optomechanical system. The deduced parameters are shown in Table 7.2. The same model fed only with independently measured parameters is also shown in the figure and the table. While the fit was optimized by adjusting the detuning and optical power, the expected cooling and optical spring depends strongly on the loss and transmission of the optical cavity. We believe that the most likely explanation for the small mismatch between fit and model is a change of the coupling mirror transmission between room- and cryogenic temperature due to an etalon effect between the uncoated backside and the highly reflecting front side (see section 3.2.3, from where transmission variations of the order of 30 % can be expected). A more detailed characterization of the ratio between transmission and loss (only their sum is known from a finesse measurement) is, in principle, possible at cryogenic temperature by measuring the cavity finesse as a function of laser wavelength over at least one half of the etalon's free spectral range, and could be carried out by a future experiment[6].

The only parameter that cannot be determined with good precision from the fit is the intrinsic mechanical quality factor, because even at the lowest injected power, significant dynamical backaction effects were observed. Any value of the quality factor in the range from $10 - 100$ million is in reasonable agreement with the measurement. We will discuss this issue further in section 7.2.5.

7.2.3 Radiation-pressure cooling close to the quantum ground state

We next turn to an analysis of the Brownian motion observed during the measurement run. A few of the measured spectra are shown in Fig. 7.5. A major issue for the measurents at the lowest powers was the small mechanical linewidth, which required resolution bandwidths as low as 50 mHz and consequently long averaging times of the order of 15 minutes per acquired spectrum. With these settings, Eq. 7.1.4 predicts statistical uncertainties near 10% for the estimation of the effective temperature from

[6]The next version of the experiment will furthermore employ coupling mirrors with transmissions of the order of 80 ppm, with which an estimation of the ratio between transmission and loss from a measurement of the resonant reflectivity of the cavity is possible with good accuracy. The closeness of the transmission and loss in this experiment made this method too imprecise to be of practical use.

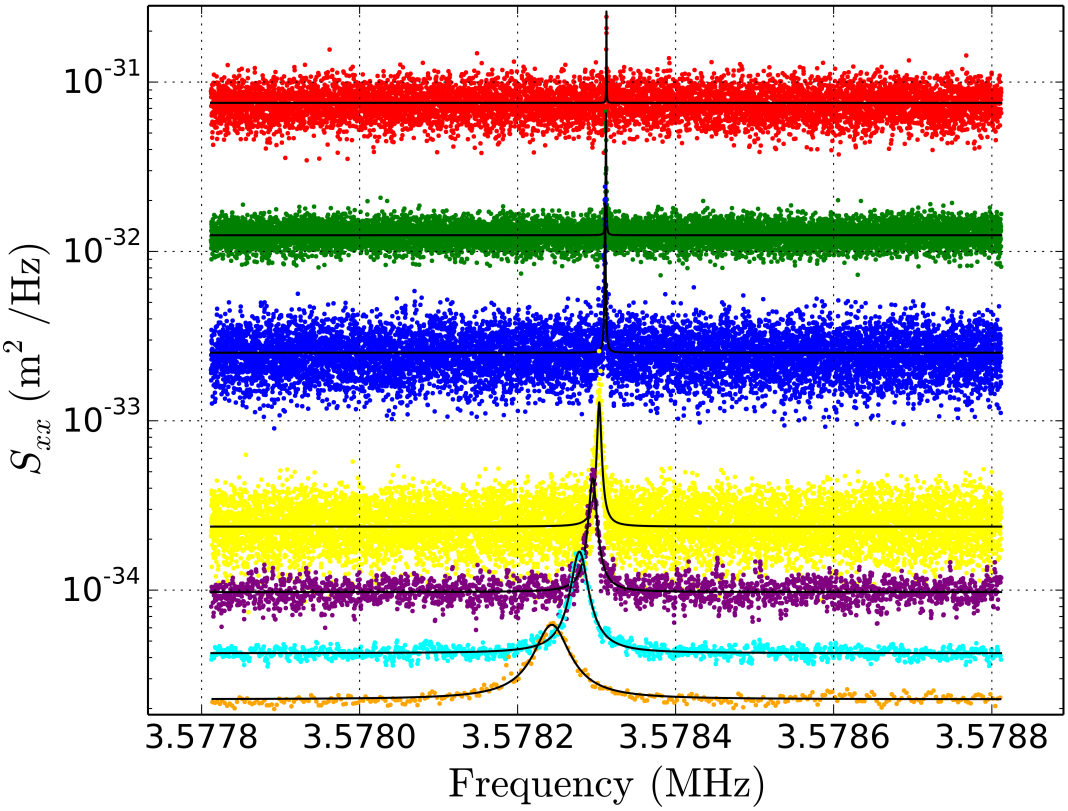


Figure 7.5: Displacement spectra of Brownian motion around the compression mode frequency with a red-detuned laser ($\Delta = 0.62 \times \kappa/2 \approx 2.5 \times \Omega_0$), acquired with direct detection. The incident powers were, from top to bottom: 0.2, 0.4, 0.8, 3.2, 6.4, 12.5 and 25.0 μW. From the bottom (top) spectrum, we deduce an effective temperature of 3.4 mK (110 mK), which is equivalent to an occupation number of 20 (642) thermal phonons. Since our optomechanical system operates in the non-resolved sideband regime, the chosen laser-cavity detuning allows nearly maximum cooling rates at the given incident powers and causes the cooling rates to depend minimally on detuning fluctuations. The thereby enhanced and stable mechanical damping rates permit the best estimation of the effective and environment temperatures for a given incident power.

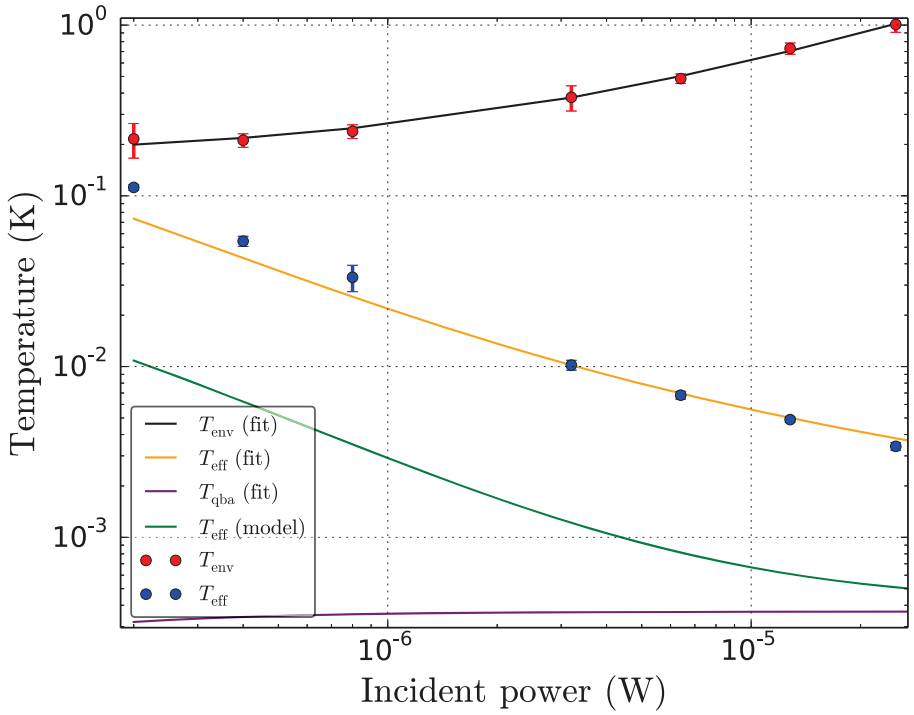


Figure 7.6: Measured effective and environment temperatures versus incident power P for a red-detuned laser. The plotted dots are the mean values extracted from $2-13$ spectra per power, the error bars indicate the corresponding variance. The power-law fit of the environmental temperature (T_{env} (fit)) yields a dependence as $T_{\mathrm{env}} \propto P^{0.68}$. This function is combined with the parameters of the fit from Table 7.2 to yield a good explanation of the effective temperature (T_{eff} (fit)). The expected cooling limitation from quantum backaction (T_{qba} (fit)) is about a factor of 10 below. A model only using the model parameters in Table 7.2 and the assumption of a constant sample temperature of 50 mK (T_{eff} (model)) greatly underestimates the observed effective temperatures.

a spectrum. At higher powers, the averaging times could be shortened more than 10-fold due to the strong increase of the effective mechanical linewidth due to dynamical backaction. The peak signal-to-noise ratio (SNR) of the acquired spectra decreases from about 5 at the lowest power to about 2 at the highest power. This nearly constant SNR is consistent with the assumption that the detection noise floor is limited by a dark noise that scales approximately in proportion to the avalanche gain of the photodiode, such as bulk dark current [McI66]. Lorentzian fits performed in logarithmic scale yield a good description of the data (see figure). For the spectra with the narrowest bandwidth, the resolution bandwidth of the spectrum analyzer was of the order of the mechanical linewidth. To test for possible deviations of the Lorentzian lineshape due to this spectrum analyzer filter, fits were also performed with a function given by the convolution of a Lorenzian plus incoherent background and the flat-top filter of the spectrum analyzer. The fit parameters obtained in this way varied by less than 20 % from standard Lorentzian fits.

The temperatures deduced from all Lorentzian fits can be found in Fig. 7.6. Throughout the measurement, the cryostat was close to its base temperature near 50 mK. The spectrum measured with the lowest power of 0.2 μW corresponds to an effective temperature of 110 mK, while highest employed power of 25 μW lead to the lowest effective temperature of 3.4 mK, which corresponds to a thermal occupation number of 20 phonons. Other optomechanical systems that have been cooled to comparable or lower phonon occupation numbers have employed mechanical oscillators that are about 3 orders of magnitude lighter than our resonator [RDW$^+$11, PYK$^+$15, PPK$^+$16, UML$^+$15]. The ability to achieve large cooling rates despite the large mass of 33 μg of our mechanical oscillator is a result of a number of optimizations: the exceptional mechanical quality factor, the operation inside a dilution refrigerator, the maximum-possible momentum transfer between the mechanical system and the optical field due to the Fabry-Perot implementation of our optomechanical system, and the large optical finesse.

The effective temperatures plotted in Fig. 7.6 only approach the temperature limit imposed by quantum back-action very slowly (as $P^{-1/3}$) due to the deduced increase of the environment temperature. The fit of the deduced values for T_{env} yields

$$T_{\mathrm{env}} = \frac{Q_0}{2.5 \times 10^7} \left(160\,\mathrm{mK} + \left(\frac{P_{\mathrm{in}}}{30\,\mu\mathrm{W}} \right)^{0.68} \right) , \tag{7.2.2}$$

where the dependence on the assumed value of the intrinsic quality factor Q_0 has been stated explicitly. We will attempt to model the heating effect in the next section. To be able to compare the model with the observed dependence, we must deduce the fraction of injected power that contributes to the heating. From the effective detuning, the resonant reflectivity $R_0 = 0.15$ of the cavity and the square root of the round-trip loss of 15% through the cryostat windows, we estimate that a power of $P_{\mathrm{abs}} \approx 0.50 P_{\mathrm{in}}$ is absorbed by the cavity as a function of incident power. While the direct absorption of the micropillar mirror is expected to amount only to about 3% of the absorbed power, we expect that a significant fraction of scattered light (up to 50%) is also dissipated in the vicinity of the micropillar and thus contributes to its heating.

Reference	k (W/m/K)	C_v (W.s/kg/K)	Range of validity
[ZP71]	$100\ T^3$	$5.7 \times 10^{-4}\, T^3$	$0.2 - 10$ K
[GA81]	$15\ T^3$	$5.6 \times 10^{-4}\, T^3$	$0.1 - 3$ K

Table 7.3: Literature values for the thermal conductivity k along the optical axis and the specific heat capacity C_v of quartz at low temperature. For temperature-dependent values, the temperature T must be inserted in units of K. The thermal conductivity in reference [GA81] was limited by boundary scattering at the sample edges.

7.2.4 Heating of the micropillar by absorption

Independently of the assumed intrinsic quality factor Q_0, Fig. 7.6 and Eq. 7.2.2 show a power-dependence of the environment temperature very close to $P^{2/3}$. In this section, we will attempt to model this effect.

Table 7.3 lists a few measurements of the thermal properties of quartz at low temperature from the literature. While the values for the specific heat capacity are in good agreement, the thermal conductivity significantly varies across the measurements. This variation is to be expected because heat transport in quartz at low temperature is ballistic, i.e. the mean free path of the phonons that transport heat across the bulk material are predominantly scattered by the sample boundaries instead of being scattered by lattice imperfections or other phonons. The thermal conductivity must therefore be estimated from the sample geometry. From the Debye model [AM76], we expect the thermal conductivity k to be

$$k = \frac{1}{3}\rho C_v c l\,, \tag{7.2.3}$$

where C_v is the specific heat capacity, c the sound velocity and l the phonon mean free path. While we can find C_v from Table 7.3 and ρ and c from section 3.1.3, we will estimate the mean free path from the dimensions of our sample with the assumption that the thermal conductivity is limited by the membrane. The distance of a line from a point at the center of the membrane of thickness d to the upper or lower membrane surface at an angle ϕ with respect to the plane is $l_{\mathrm{geo}}(\phi) = d/(2\sin\phi)$. As l diverges for angles near zero, we assume that the free path length is furthermore limited by the lateral sample size $l_{\max} = 10$ mm, such that $l^{-1} = l_{\mathrm{geo}}^{-1} + l_{\max}^{-1}$. By integrating $l(\phi)$ over all possible angles, we find for the mean free path numerically

$$l = \frac{1}{\pi}\int_{-\pi/2}^{\pi/2} \frac{1}{l_{\max}^{-1} + 2\sin(\phi)/d}\mathrm{d}\phi \approx 2.4\, d\,. \tag{7.2.4}$$

The prefactor of 2.4 varies by only 40% if the ratio $l_{\max}/d = 10^3$ is varied by an order of magnitude in either direction, and may therefore be assumed to be constant in our approximation. With the numerical values introduced above, we find for the thermal conductivity

$$k \approx 8 \times 10^{-2} \frac{d}{10\,\mu\mathrm{m}} T^3\,(\mathrm{W}/(\mathrm{m.K}^4))\,. \tag{7.2.5}$$

We proceed to estimate the sample temperature assuming that the thermal resistance is dominated by the membrane. A simple analytical model analogous to the one

in [Kuh13] may be derived by approximating the membrane as a ring of thickness d, inner radius[7] $r_{\mathrm{micropillar}} \approx 115\,\mu\mathrm{m}$ and outer diameter $r_{\mathrm{frame}} \approx 2.6\,\mathrm{mm}$. The thermal gradient $\mathrm{d}T/\mathrm{d}r$ at the radial distance r from the center of the disc is related to the absorbed power P_{abs} that is conducted from the pillar edge at $r_{\mathrm{micropillar}}$ to the quartz chip at r_{frame} as

$$P_{\mathrm{abs}} = -2\pi r d k(T) \frac{\mathrm{d}T}{\mathrm{d}r}\,. \tag{7.2.6}$$

After inserting Eq. 7.2.5, the equation can be solved by integrating from the inner to the outer disc radius over a term proportional to $r^{-1}\mathrm{d}r$ on one side and over a term proportional to $T^3\mathrm{d}T$ on the other side. We eventually obtain

$$T_{\mathrm{micropillar}} = \left(\left(\frac{d}{10\,\mu\mathrm{m}}\right)^{-2} \frac{P_{\mathrm{abs}}}{0.4\,\mu\mathrm{W}} K^4 + T_{\mathrm{frame}}^4\right)^{1/4}\,. \tag{7.2.7}$$

By neglecting the cryostat temperature T_{frame}-term due to the T^4-dependence, we find that in our parameter regime, the micropillar temperature should scale with the fourth root of the absorbed power and the inverse square root of membrane thickness.

In addition to the analytical model, we have performed a finite-element simulation of the temperature field at various powers with the same material parameters. Fig. 7.7 shows the obtained average temperature over the volume of the micropillar. A power-law fit of the simulation results yields the expected $P_{\mathrm{abs}}^{1/4}$ depedence with a pre-factor within 10% to the one of the analytical model.

While we have experimentally measured environment temperatures that scale approximately with $P_{\mathrm{abs}}^{2/3}$, both models of temperature conduction yield a $P_{\mathrm{abs}}^{1/4}$-law due to the factor of T^3 in the thermal conductivity. While a T^2-dependence was observed in reference [GA81] for neutron-irradiated α-quartz, even this assumption would only yield a $P_{\mathrm{abs}}^{1/3}$-law. We conclude that the most likely explanation for the experimental observation is that, contrary to the assumption used in figure 7.6, the intrinsic quality factor is not independent on laser power. This dependence may either arise due to an asymmetry of the sample geometry induced by the unilateral irradiation and thermal expansion, or due to the elevated temperature of the micropillar. Asymmetric thermal expansion of the two half-pillars is expected to cause the compression mode of the pillar to have a non-zero displacement at its center. The central displacement is coupled to the surrounding structure by the membrane and can easily become the dominant source of clamping loss [Kuh13, Zer17]. The temperature profile in Fig. 7.7 shows that the micropillar surface deviates from the average temperature of the micropillar volume by about 0.4 K[8]. We expect that the length asymmetry induced by this temperature difference is at least two orders of magnitude below the asymmetries that we expect to have occured at room temperature with constant quality factors of the order of 10^6, and therefore exclude asymmetry as a cryogenic loss candidate. Therefore, a temperature-dependence $Q_0 \propto P_{\mathrm{abs}}^{-0.43} \propto T^{-1.72}$ of the intrinsic quality factor of our samples could explain our observation. An alternative explanation for

[7] The inner and outer radii are chosen such that the circumference of the rings coincides with the corresponding circumference in the triangular sample structure.

[8] This value is likely an overestimate because the thermal conductivity through the micropillar is expected to occur ballistically.

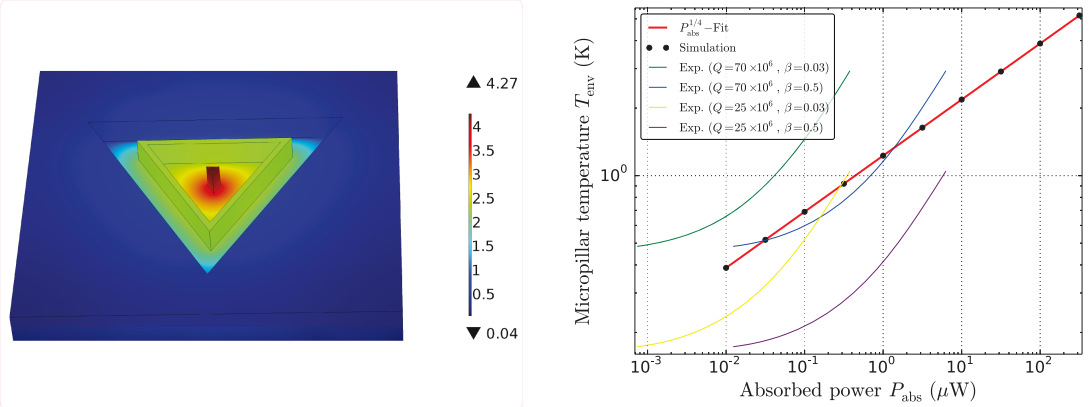


Figure 7.7: Left: Geometry and obtained temperature field from a COMSOL simulation for a membrane thickness of 10 μm and a power of 100 μW absorbed by the round mirror on top of the micropillar. The color bar indicates the temperature in K. Right: Mean temperature of the micropillar versus absorbed power for a membrane thickness of 10 μm. The simulated values are perfectly fitted with the function $T/(1\,\mathrm{K}) = (P_{\mathrm{abs}}/0.44\,\mu\mathrm{W})^{1/4}$, which is remarkably close to the analytical expression in Eq. 7.2.7. Additionally, the fit of the experimental data from Eq. 7.2.2 is plotted along with different assumptions of the intrinsic quality factor Q_0 and the fraction of intracavity power absorbed by the micropillar β. While simulation and experimental data are roughly in the same temperature range, their different power dependence is obvious from the plot.

the quality factor dependence on the optical power is the optomechanical modification of the mechanical susceptibility: the optical spring effect of up to 70 Hz, but also the resonance broadening from optical damping, are expected to modify the overlap of the compression mode susceptibility with mechanical modes of the support structure that have higher mechanical damping. From Fig. 3.12, we estimate that a frequency-shift of the compression mode of the order of 3 kHz with respect to the support structure mode frequencies leads to a change of the compression mode quality factor between about 10^6 and its maximum value. The observed optical spring achieves shifts of the order of 2% of the frequency-shift needed for a full period of quality factor modulations, and may therefore be responsible for a significant modification of the intrinsic quality factor which is about $10-100$ times larger than the presumed minimum value of 10^6. The proposed mechanism could be confirmed experimentally by mapping the environment temperature as a function of the optical spring shift induced by an external feedback loop. If this test should prove successful, it is reasonable to optimize the intrinsic quality factor by using a passive or feedback-induced optical spring in order to shift the mechanical frequency to its optimal value[9].

[9]Frequency-shifts of the order of 3 kHz are expected to be realizable without injecting excessive noise into the system with incident powers of the order of 1 mW. This enables us to cover an entire period of quality factor modulations.

7.2.5 Mechanical quality factor

The discussion in the previous section suffered from the large uncertainty in the estimate of the intrinsic quality factor of our sample. A trustworthy estimate of the quality factor was impossible even at the lowest incident powers due to strong dynamical backaction effects. This problem may be resolved by performing the ringdown measurement with even lower powers, of the order of 10 nW, with the drawback that the power-dependence of the quality factor cannot be studied. In this section we present an alternative solution that enables to use incident powers of the order of $10 - 100\,\mu$W while strongly suppressing the effect of dynamical backaction in the quality factor estimation.

Fig. 7.8 explains the measurement procedure. Dynamical backaction is avoided during the measurement by two effects: a strongly reduced effective intracavity power by only allowing the laser to become resonant with the cavity for a fraction of the piezo sweep period, and an effective cancellation of dynamical backaction effects by forcing the laser to spend equal amounts of time on either side of the optical resonance. In principle, the mean intracavity power can be arbitrarily reduced for a given incident power by increasing the piezo sweep amplitude, or by performing the sweep with a modified waveform such as a square wave. The level of cavity length vibrations induced by the dilution fridge operation in the measurement shown in Fig. 7.8 was of the order of the optical linewidth. Driving the mechanical oscillator to amplitudes well above this level makes our method robust against such vibrations.

Fig. 7.9 shows a few quality factors measured with the described method. The quality factor values around 70×10^6 are scattered within 10% of the mean value. In the measurement, small residual dynamical backaction has occured when the resonance was not well-centered with respect to the piezo sweep range. This is a result of the sinusoidal waveform that was employed for the piezo sweep[10], and could be attenuated further by locking the piezo offset voltage to the time of passage of the mechanical resonance. This type of lock has been implemented before, but is expected to be slightly complicated by a changing optical resonance peak during the ringdown.

Since the measurements in Fig. 7.9 were performed at slightly different temperatures, it is tempting to conclude that the quality factor of our sample improves by a factor of 3 when cooling from 4 K to dilution refrigerator temperatures. However, both measurements were not acquired in the same cryogenics run and the different quality factors may therefore be caused by changes of the mechanical clamping of the sample due to thermal hysteresis-related stress. The low-frequency mode spectrum of the sample significantly changed between both acquisitions (a mechanical mode seems to have moved from 12 to 6 kHz) and therefore suggests that a change of the membrane geometry due to thermal cycling or a vibrational shock is responsible for the observed quality factor difference. A future measurement of the quality factor of this sample over a range of cryogenic temperatures may clarify this ambiguity and is furthermore required to prove the assumption at the end of section 7.2.3 that a temperature-

[10]The time difference that the laser spends on either side of the resonance is proportional to the second derivative of the piezo displacement. While a voltage ramp should result in a zero second derivative, a low-pass filter before the piezo prevents the voltage ramp from being transduced into a displacement ramp.

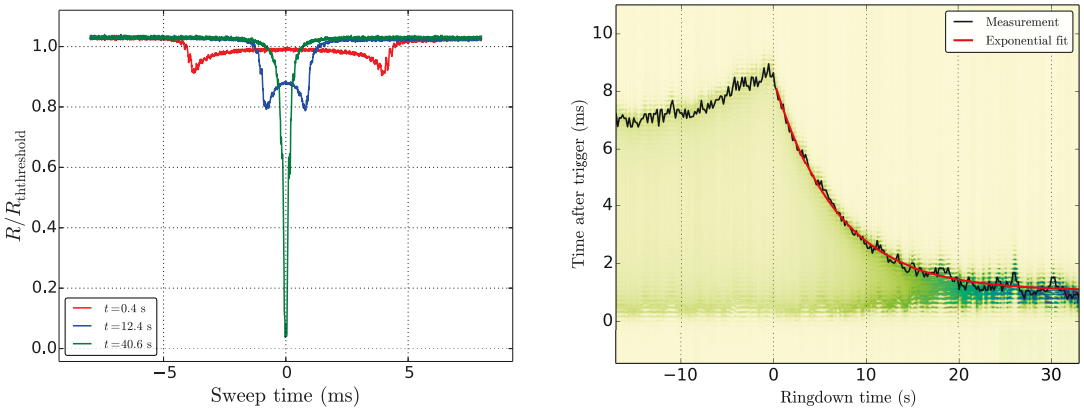


Figure 7.8: Left and right: Cavity length sweep ringdown measurement. The compression mode of the micropillar is driven to displacement amplitudes of the order of 10 cavity linewidths. At $t = 0$, the drive is turned off and the displacement amplitude decays at the intrinsic damping rate. Left: Reflection from the measurement cavity at different times after the RF drive was turned off while slowly sweeping the cavity length with a sinusoidal voltage at 20 Hz. The cavity resonance dip is strongly broadened due to the mechanical oscillation of the compression mode. The maximum distance of the acquired reflection values below a threshold (the ordinate axis is scaled in units of this threshold) is a good estimate of the displacement amplitude. Right: The black line shows the estimated displacement amplitude with the procedure described for the left plot during the mechanical ringdown. The red exponential fit visibly yields a good representation of the tendency in the noisier data trace. The data was recorded by setting the oscilloscope trigger slightly below the off-resonant reflection voltage such that two resonance dips were acquired per sweep period. The color-plot in the background shows the measured cavity reflection. The vertical axis encodes the time relative to the last oscilloscope trigger event. The horizontal axis indicates the absolute time during the ringdown measurement. The fit yields a quality factor $Q = 7.7 \times 10^7$. The vibration-enlarged optical linewidth can be read off from the black trace once the ringdown is complete, and is estimated to be about 1.1 ms. This indicates that the maximum displacement amplitude of the mechanical oscillator is of the order of 8 optical linewidths (about 40 pm). The slow rise of the displacement amplitude before the drive is turned on is likely caused by a slow drift of the relative frequency between the mechanical oscillator and the RF source, since the mechanical linewidth of 44 mHz still corresponds to a fraction of only 4×10^{-10} of the 125 MHz-clock of the RF source.

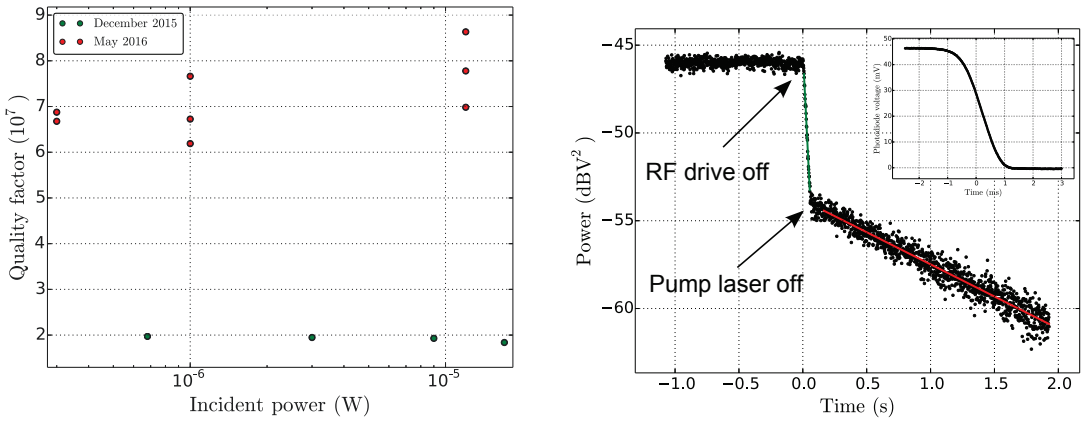


Figure 7.9: Left: Mechanical quality factors measured with variable incident power for sample 20121009C5 in different cryostat runs. The measurements from December 2015 were performed at 4 K, while the ones in May 2016 were carried out at a cryostat temperature of 41 mK. Right: Mechanical ringdown measurement during which a 20 μW pump laser was turned off by a shutter. At time 0, the RF tone that drives the mechanical oscillator is turned off. The presence of the strong red-detuned laser beam increases the effective damping rate. From the slope of the green trace, an effective quality factor of 6.8×10^5 is deduced. Next, the laser beam is turned off by a shutter. The optical power transmitted while the shutter is closed is shown in the inset. Without the optical damping provided by the pump beam, the mechanical displacement decays much slower, corresponding to an effective quality factor of 2.7×10^7. No significant deviation from linear ringdowns is observed when the pump beam is turned off, as would be expected for a temperature-dependent quality factor.

dependent quality factor may explain the observed excess power dependence of the estimated environment temperature.

In another experiment, it was attempted to directly observe the quality factor change due to laser heating. The corresponding measurement is shown in the right plot of Fig. 7.9. For this measurement, a ring-down measurement was performed while a weak probe laser and a stronger pump laser were simultaneously locked to the cavity. During the mechanical ringdown, a shutter was used to turn off the pump laser and thereby to observe the change in quality factor associated with the power change. While a first look at the figure suggests a drastic quality factor change, the dominant observed change can be explained by the modification in the dynamical backaction cooling rate caused by turning off the pump laser which was not locked exactly on the optical resonance. The change in the dynamical backaction cooling rate is expected to occur at a rate comparable to the optical bandwith, which is about ten times larger than the mechanical resonance frequency. A temperature-related change might occur much more slowly. To estimate the associated time-constant, we can write the thermal energy stored in the micropillar at the moment when the pump laser is turned off as $E = \rho V C_v T$, where V denotes the micropillar volume. Eq. 7.2.7 has already established the relation between the power P_{abs} conducted through the membrane and the micropillar temperature T. By realizing that $P_{\mathrm{abs}} = -\mathrm{d}E/\mathrm{d}t$, and writing the temperature-dependence of all variables explicitly in the equation, we obtain a differential equation that equates the time-derivative of the micropillar temperature to a term proportional to the temperature. The solution is an exponentially decaying temperature with a decay rate

$$\gamma_T = \frac{0.4\mu\mathrm{W}}{4\rho V C_v/T^3} = 2\pi \times 430\,\mathrm{Hz}\,. \tag{7.2.8}$$

The rate is comparable to the rate at which the shutter turned off the beam (see inset of Fig. 7.9), and therefore not clearly detectable with the scheme involving a shutter. However, the presented estimation of the temperature decay time is very useful in designing a possible future scheme for this measurement: by injecting an amplitude-modulated pump beam with a modulation frequency comparable to the thermal decay rate, an experimental value of the temperature decay rate and the temperature-related quality factor change may be deduced from the phase-lag and amplitude of the expected ringdown-rate modulation in a future measurement[11].

A previous estimation of a much faster temperature decay rate of the order of 100 kHz [Kuh13] was based on the assumption that the temperature decay rate should be associated to the thermalization within the micropillar volume. Within the micropillar volume, we indeed expect ballistic heat transport to result in thermalization rates of the order of the mechanical frequency.

[11] Such an experiment may encounter two limitations: the measurement bandwidth (which is defined by the span of the spectrum analyzer) must be wide enough to allow the modulation to be detected while providing a reasonable signal-to-noise ratio, and excessive detuning fluctuations due to heating or radiation-pressure effects from the amplitude modulation must be effectively avoided. The second effect could be cancelled by a feed-forward scheme that can be tuned before the ringdown is performed.

7.3 Towards the observation of quantum backaction

At incident powers above 25 μW, feedback is necessary at any laser-cavity detuning to prevent optomechanical instabilities over the various low-frequency modes. In order to keep the necessary feedback gains low and to maximize intracavity power, detunings closer to resonance were studied. The homodyne detection was employed for these measurements due to its shot-noise dominated noise floor. Section 7.3.1 presents the calibration of our homodyne detection and the measurement procedure, while section 7.3.2 presents the measurement results which demonstrate significant backaction from the laser from quantum and classical noise sources. Section 7.3.3 section discusses the possibility to observe quantum backaction at room temperature with our optomechanical system.

7.3.1 Calibration of the homodyne detection

Despite the phenomenological heating expressed by Eq. 7.2.2, we expect to improve the ratio between quantum backaction and thermal noise by employing higher incident powers approximately as $P^{1/3}$ to $P^{3/4}$, depending on the exact heating mechanism. The desire to achieve optimal sensitivity for arbitrary laser-cavity detunings suggests to employ a homodyne detection for measurements at higher power. A significant obstacle in using the homodyne detection was electrical crosstalk between the feedback loop of the local oscillator phase and the measurement cavity frequency stabilization loop. The solution to this problem has been described in section 5.6. Another difficulty is to reliably calibrate the detection. While the interference fringe amplitude was employed in an interferometric measurement in section 3.1.3, we expect this calibration method to be imprecise for measurements with a Fabry-Perot cavity, because the fringe amplitude obtained from a scan of the local oscillator does not take into account the possibility of different overlaps between the signal beam with the local oscillator beam and with the cavity mode (i.e. the geometrical contribution to quantum efficiency). A better calibration can be obtained by scanning the cavity and local oscillator piezo lengths simultaneously to record the signature of the passage of an optical resonance in the homodyne fringes at different local oscillator phases. While this method is expected to yield a good estimate of the geometrical quantum efficiency, it requires the assumption that the detection sensitivity is linear around the working point once the cavity-laser detuning and local oscillator phase are locked. Due to the large vibrations of our measurement cavity, this assumption is not expected to hold. We therefore resort to yet another calibration method that is a summarized in Fig. 7.10 and a consequence of the already implemented locking scheme of our cavity.

To acquire displacement spectra with the homodyne detection, we start by locking the cavity to a weak probe laser (heterodyne laser) with the PDH method as it was done for the measurements of Brownian motion at room temperature (section 7.1.1). By shifting the frequency of this laser with the AOM setup (section 5.3.2) such that the un-shifted laser frequency nearly coincides with the polarization eigenmode of the cavity that is orthogonal to the mode illuminated by the heterodyne laser, we make the homodyne signal beam resonant with the measurement cavity without any further need of additional feedback loops. This allows for arbitrary detunings and powers of

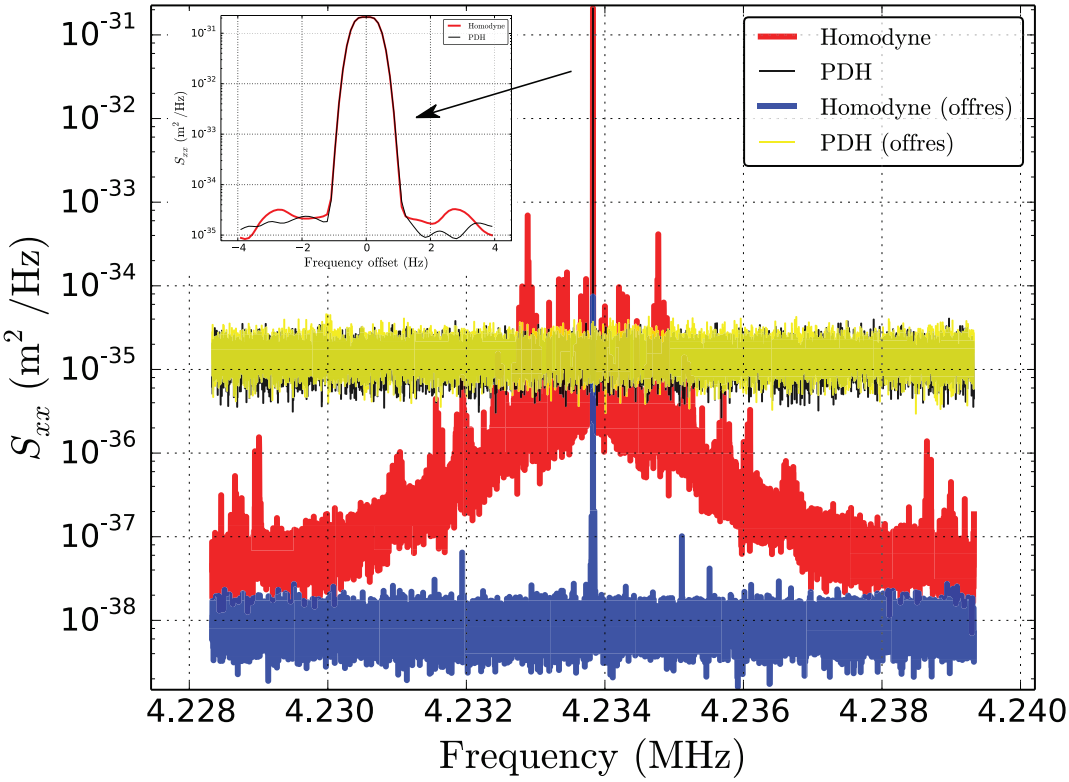


Figure 7.10: Mechanical calibration peaks acquired with an incident power of 300 μW. A mechanical resonance near the measurement frequency band is driven with the antenna actuator and measured simultaneously with the spectrum around the compression mode frequency with both the PDH and the homodyne detections. From the PDH calibration with a phase-modulation peak at 3.6 MHz of known equivalent displacement, the mechanical peak amplitude is known. The homodyne spectra are converted to displacement units such that the mechanical calibration peaks have the same amplitude. Sidebands of the homodyne peak are due to an imperfect local oscillator lock (the presented trace was taken before lock-in technique was used).

the homodyne signal laser. The homodyne signal beam power is then ramped up to the desired value and all feedback loops that are needed to prevent instabilities of the optomechanical system are enabled. A scan of the local oscillator piezo is performed to estimate the fringe amplitude of the error signal for the local oscillator phase, which allows PyRPL to lock the local oscillator phase at the desired value with optimal feedback loop gain. For the measurement of displacement spectra, we use a spectrum analyzer with two baseband channels and the associated software, which allows the simultaneous acquisition of spectra in several frequency bands of both channels. Apart from the PDH and homodyne spectrum around the compression mode frequency, we acquire a peak generated with a phase modulator that is used to calibrate the PDH detection and was already described in section 7.1.1. Furthermore, both the PDH and homodyne signals around a mechanical mode of the micropillar sample near 4.2 MHz are acquired while the mode is driven piezoelectrically by a monochromatic voltage signal applied to the antenna near the micropillar (see section 4.2.4). Since the PDH detection can be calibrated in displacement units from the phase modulation marker, the displacement amplitude of the 4.2 MHz peak during the measurement is known from the PDH spectrum. This allows us to infer a calibration for the homodyne traces from the homodyne measurement of the 4.2 MHz peak. The PDH and homodyne spectra of the mechanical calibration peak are shown in Fig. 7.10.

The presented calibration method has the advantage that the calibration data is acquired simultaneously with the measurement, thereby accounting for sensitivity fluctuations that may arise from modulations of the laser power, laser-cavity detuning or local oscillator phase. Furthermore, the spectra around the calibration peaks contain additional diagnostic information about the measurement: the sidebands around the calibration peak in the homodyne spectrum are absent in the PDH signal, and therefore indicate that the local oscillator phase was unstable during the measurement. This was indeed the case during the acquisition shown in the figure. This problem has motivated the use of an improved error signal for the local oscillator phase, as described in section 5.6. For the calibration method to work effectively, it must be ensured that electrical cross-talk from the generator of the calibration signal to the detection circuit is negligible if the laser is off-resonant with the cavity (see "offres" in Fig. 7.3.1). The frequency of the calibration peak is a compromise between two effects: a calibration peak too close to the measurement frequency band around 3.6 MHz might cause sidebands of the calibration peak to appear in the measurement spectrum, while a calibration peak too far away would risk an imprecise calibration due to frequency-dependent gains of the detection electronics.

7.3.2 Observation of back-action

In this section, we present a number of measurement series with variable homodyne signal beam power. All homodyne measurements were calibrated as described in the previous section. The power of the weak probe laser used for the PDH detection and the cavity-laser detuning lock was of the order of 1 μW for all measurements, and is therefore negligible with respect to the homodyne signal beam. Fig. 7.11 shows the spectra obtained in one of these power series. The measurements with powers up to 50 μW behave as expected for a homodyne detection: the noise floor in displacement units

scales at the inverse signal beam power for the spectra with the lowest powers, and as in Fig. 7.5, the compression mode appears as a Lorentzian peak in the spectrum. All spectra for incident powers above 1 μW had to be acquired with a red-detuned laser with respect to the optical resonance in order to avoid the parametric instability of the compression mode. The detuning had to be a significant fraction (10%) of the optical bandwidth because vibrations can otherwise cause the detuning to fluctuate into unstable regimes. If the optomechanical system nevertheless entered the regime of self-sustained oscillations, for example because the system was perturbed by a clearly audible sound in the laboratory, a stable lock could only be recovered by turning both lasers off and waiting about 15 s until the mechanical oscillation had rung down, or by slowly approaching a weak laser from the red side of the resonance to accelerate the ring-down by dynamical backaction. In order to achieve a reasonably stable lock at the highest injected powers, the shape of the feedback loop that controls the amplitude modulator to suppress low-frequency resonances had to be re-tuned several times during the ramp-up of the signal beam power and optimized for each final power. As a result of these difficulties, the laser-cavity detuning had to be slightly increased for the spectra at the two highest powers in Fig. 7.11 to allow for a lock duration longer than the spectrum acquisition. Since optomechanical cooling can therefore not be avoided in the power series, the Lorentzian peaks in Fig. 7.11 become broader and decrease in peak amplitude for increasing power. As in section 7.2.3, we attempt to quantify the laser-induced heating of the micropillar by computing the environment temperature with Eq. 7.1.6 from the measured effective temperature with an intrinsic quality factor of $Q_0 = 6.8 \times 10^7$, obtained by the measurement in section 7.2.5.

Fig. 7.12 shows a compilation of all power series performed with the homodyne detection, together with some attempts to model the observed behaviour. We first note that measurements acquired in different cryogenic runs mutually agree reasonably well (see the figure caption for a comment on "PDH (5/2016)"), which indicates that the intrinsic quality factor Q_0 was approximately constant despite thermal cycling between the measurements. The model "Heating (expected)" indicates the $P_{\mathrm{abs}}^{1/4}$-dependence expected from Eq. 7.2.7. Howeverm the data agrees much better with the extrapolated phenomenological heating model from Eq. 7.2.2 ("Heating (observed)"). This shows either that the micropillar environment is heated at the theoretically unexpected $P_{\mathrm{abs}}^{2/3}$-dependence, or that the mechanical quality factor degradation is related to the heating-induced temperature-increase. A quality factor degradation due to the optical spring-induced frequency-shift as suggested at the end of section 7.2.4 is very unlikely in light of the presented measurements, which yield mutually consistent estimates of the environment temperature despite a varaiation of the cavity-laser detuning (and thereby the magnitude of the optical spring) from 0.05 to 1 cavity bandwidths $\kappa/2$ among the different measurement series.

At the highest powers, noise in the optical field starts to heat the compression mode. Quantum backaction is expected to significantly contribute (at most about 50%) to the total force noise spectrum (see Fig. 7.12). However, we also expect classical noise in the intracavity field. This can be seen from the excess background in Fig. 7.11 and is also apparent from all other spectra at injected powers above 100 μW: the detection is no longer limited by shot-noise, but instead by a constant displacement background level around $10^{-37}\,\mathrm{m}^2/\mathrm{Hz}$. For any non-zero laser-cavity detuning, these

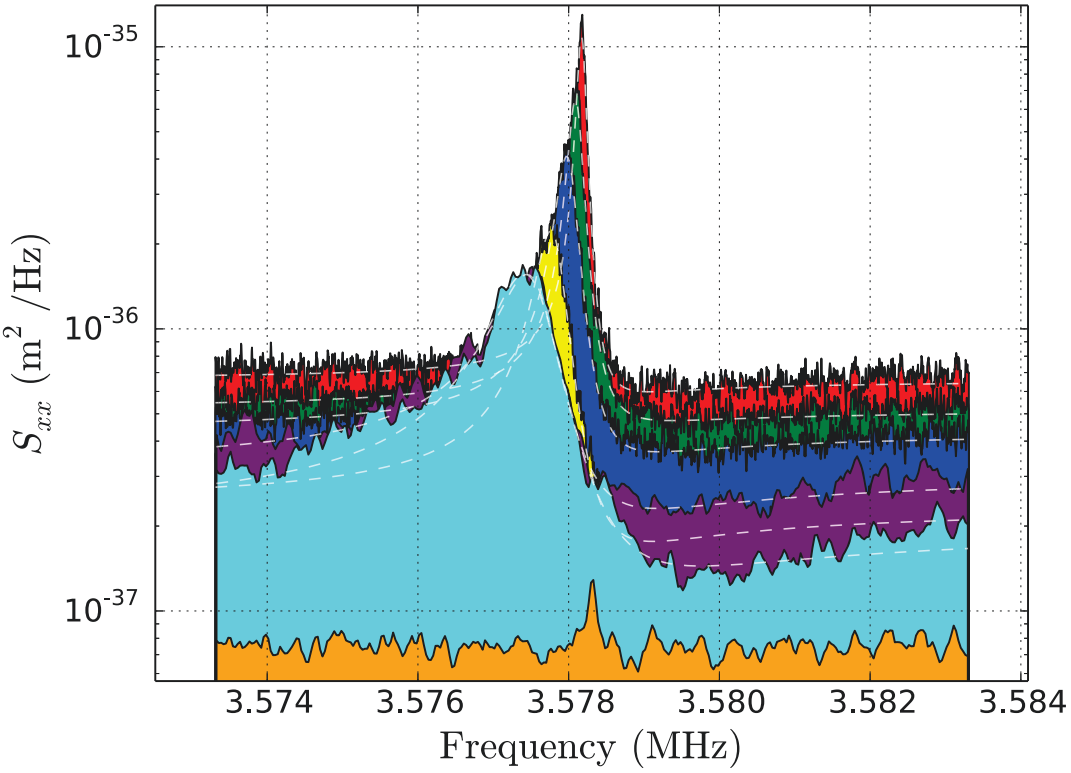


Figure 7.11: Calibrated displacement spectra obtained with the homodyne detection for various powers. The traces correspond to the following powers absorbed by the cavity: 70 μW (red), 110 μW (green), 160 μW (blue), 270 μW (yellow), 360 μW (purple), 450 μW (cyan). The last three traces were acquired with an 8-fold higher resolution bandwidth which explains the difference in frequency resolution. White dashed lines indicate fits of the spectra with a Fano lineshape. The temperatures inferred with Fano-Lorentzians and standard Lorentzians differ at most by 25 %. The orange trace indicates the shot-noise background for the trace in cyan that was acquired with the signal laser off-resonant with respect to the cavity. The noise-floor of the cyan trace is clearly above (160% above) the shot-noise, while the background of the red trace is dominated ($> 95\%$) by shot-noise (the corresponding shot-noise trace is not shown here). The data series corresponds to the label "Homodyne (4/2016)" in Fig. 7.12.

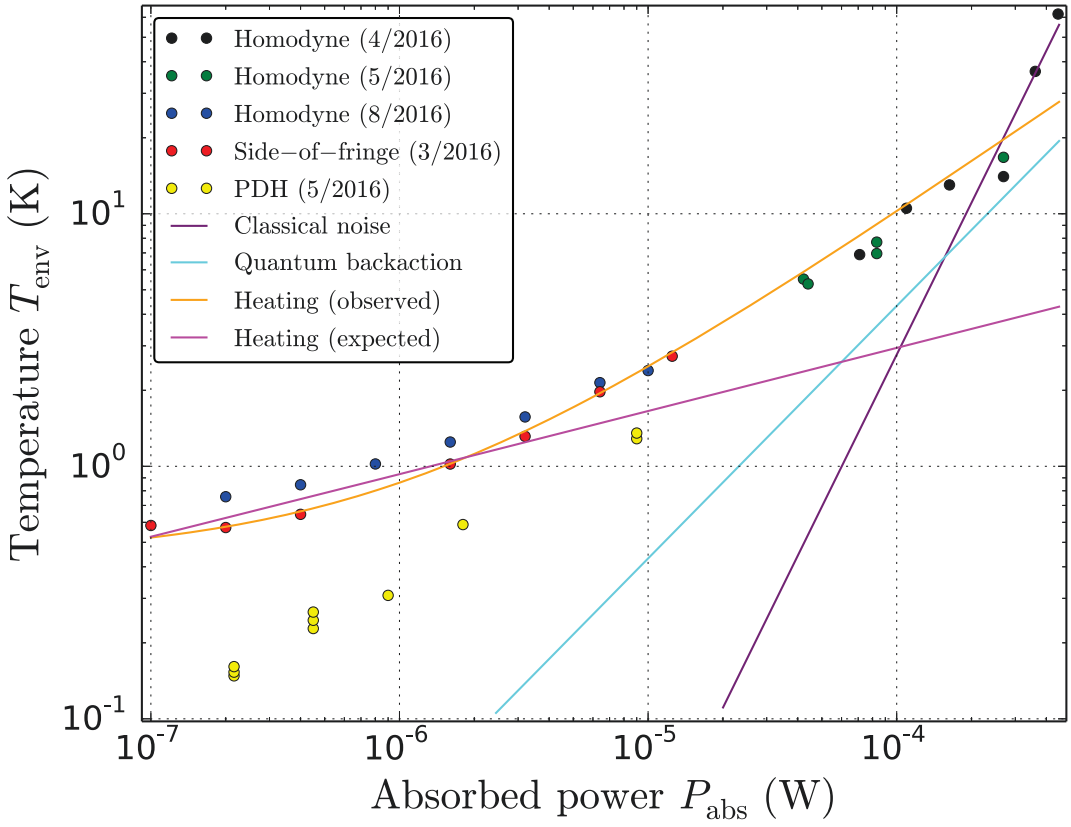


Figure 7.12: Equivalent temperature of the force noise acting on the compression mode for all power series performed on the micropillar sample 20121009C5. The abscissa corresponds to the power that was absorbed by the measurement cavity, which differed with respect to the incident power between different measurement runs due to different alignment, detuning, lock stability and (in the case of "Homodyne (8/2016)") clipping at the cryostat windows. All environment temperatures assume an intrinsic quality factor of $Q_0 = 6.8\times10^7$. The different dots correspond to cryogenic measurement runs between which the sample was heated to liquid nitrogen or room temperature, only "PDH (5/2016)" and "Homodyne (5/2016)" were acquired in direct succession. The significant deviation of "PDH (5/2016)" from the other measurements can be explained by the lack of reliable estimates of the effective damping rate, since the mechanical linewidth was comparable to or smaller than the resolution bandwidth of the spectrum analyzer in this measurement. The full lines indicate various force noise models: "Heating (expected)" is the predicted absorption heating from an analytical model (Eq. 7.2.7) with the assumption that 30% of the power dissipated by the measurement cavity contributes to heating the pillar, "Heating (observed)" is an extrapolation of the fitted environment temperatures of "Side-of-fringe (3/2016)" (Eq. 7.2.2 with Q_0 as above), "Quantum backaction" predicts the force from quantum backaction with Eq. 2.3.25, and "Classical noise" employs Eq. 7.3.1 to model the force from classical noise in the optical field caused by the displacement background of $S_{xx} = 7 \times 10^{-38}\,\mathrm{m}^2/\mathrm{Hz}$ estimated from Fig. 7.11 and a detuning of $\Delta = 0.14 \times \kappa/2$ estimated from the observed dynamical backaction effects in "Homodyne (4/2016)".

displacement fluctuations are transduced into fluctuations of the intracavity power which drive the compression mode through radiation-pressure in the same way that dynamical backaction drives the mode by the radiation-pressure fluctuations induced by its own motion. The heating from this classical noise source can be modeled by comparing the intracavity field fluctuations caused by the mechanical background motion to the fluctuations due to quantum fluctuations of the input fields (Eq. 2.3.25). In the limit of small laser-cavity detunings $\Delta \ll \kappa$ and the unresolved-sideband limit $\Omega_0 \ll \kappa$, we find with Eq. 2.3.11:

$$\frac{S_{II,\,\mathrm{classical}}(\Omega)}{S_{II,\,\mathrm{qba}}(\Omega)} = 2\eta \frac{P_{\mathrm{in}}}{\hbar\omega_{\mathrm{L}}} \frac{\Delta}{\kappa/2} \frac{S_{xx}(\Omega)}{\lambda/(4\mathcal{F})^2} \,. \tag{7.3.1}$$

Importantly, the ratio increases proportionally with the laser power, thereby preventing the quantum-limited operation of an optomechanical system at powers above the threshold power, which corresponds to a ratio of 1. The effect is modeled by the line "Classical noise" in Fig. 7.12 and yields reasonable agreement with the measurements at the highest powers. The corresponding spectra are no longer Lorentzian (see Fig. 7.11), but can be fitted very well with a Fano lineshape (squared magnitude of the sum of a complex Lorentzian and a constant complex background term). This is expected because the pillar is now effectively driven by the displacement noise floor and its response interferes destructively with the noise floor at frequencies above the mechanical resonance for a red-detuned laser. The linear detuning-dependence of Eq. 7.3.1 suggests to avoid the classical heating by tuning the laser closer to the optical resonance. While the heating can indeed be reduced this way, increasing the optical power can obviously not overcome the classical displacement noise floor for the displacement readout. Section 7.4 studies the displacement noise floor in more detail and proposes a way to remove this limitation in a future experiment.

7.3.3 Quantum backaction at room temperature?

Recently, two measurements of quantum correlations in nano-optomechanical systems at room-temperature were reported [SSF$^+$16, PGST16]. While these measurements have undoubtedly pioneered room-temperature signatures of quantum effects in mechanical degrees of freedom, our experiment might extend this endeavour into the regime of more massive systems: for 1 mW of incident power and the parameters of the optomechanical system from section 7.1.1, quantum backaction amounts to about 0.1 % of the thermal noise force spectral density at 300 K, thus produces an excess environment temperature of $\delta T_{\mathrm{env}} = 0.3$ K. To observe such a small temperature shift requires the averaging over about 350 spectra with identical parameters as in section 7.1.1, equivalent to a measurement duration of 4 hours.

The long averaging times require a careful planning of the measurement protocol. The existing two-beam injection setup suggests to use a PDH or side-of-fringe detection with the "heterodyne laser" beam to infer the environment temperature of the compression mode from its displacement spectrum as in section 7.1.1. This prevents the parametric instabilities of both low-frequency modes and the compression mode by appropriate radiation-pressure feedback. The "homodyne laser" should be simultaneously injected with varying power to provide a power-dependent quantum backaction

force. The two-beam difference frequency would be determined with the acousto-optic modulator to coincide up to a fraction of the optical linewidth with the birefringence splitting of the cavity. Simply detecting the reflected power of the pump beam would allow the use of a lock-in technique based on demodulation of a piezo-imprinted cavity length modulation at about 500 kHz. This would lock the two lasers to the cavity with extremely good detuning sensitivity and a small margin for systematic error.

A number of obstacles to such a measurement can be identified and removed by a well-designed measurement protocol:

- **Temperature fluctuations** of the mixing chamber plate have been previously stabilized at room temperature to better than 50 mK without excessive optimization of the temperature controller. This effect of environment temperature fluctuations on the measurement should be negligible, or at least it should be possible to account for this effect in the data analysis with the information obtained from a resistance thermometer near the sample.

- Excess heating by **classical noise** on the optical intracavity field is a more serious limit. While the incident laser beam is shot-noise limited at least up to 10 mW, the thermal displacement background of the measurement cavity is about 30 dB above the shot-noise level at 1 mW. To avoid the transduction of this phase noise into amplitude noise that results in a temperature increase proportional to the square of the optical power, the laser-cavity detuning must be below 1 % of the cavity bandwidth. In previous runs, the laser-cavity detuning has been stabilized to this level over the course of a few hours. The strong optical spring and damping effects of lower-frequency modes (e.g. the 110-kHz mode) can be used as an extremely sensitive error signal to confirm the obtained detunings. The quadratic power-dependence of the measured excess temperature due to this effect should enable us to clearly distinguish this effect from quantum backaction.

- **Absorption of light** by the micropillar sample is expected to result in an excess temperature proportional to the optical power, thus resulting in an identical signature as quantum backaction. The least favourable heat resistance estimate from section 7.1.2 results in a temperature shift of 12 K per mW of absorbed power. Since the absorption of the micropillar surface makes up only 3% of the optical loss, the actual temperature shift might be as low as 36 mK/mW. In the following, we consider the worst case scenario of 12 K/mW. The previously measured temperature dependence of the compression mode frequency (248 Hz/K) allows an estimate of the bulk temperature of the micropillar from the observed resonance frequency to better than 10 mK. For 1 mW and at the above-required maximum detuning of 1 % of the cavity bandwidth, the optical spring effect results in a compression mode frequency-shift of about 30 Hz, equivalent to a temperature shift of 0.1 K when mistaken as a result of the bulk temperature. It is viable to extend the bulk temperature estimation to higher-frequency modes, such as the third harmonic of the micropillar around 9 MHz, which do not suffer from a significant optical spring effect due to their low quality factor. Measuring the effective temperature of the 9-MHz mode can yield additional evidence for

the bulk temperature model, since this mode is expected to experience a much lower ratio of quantum backaction effects to thermal force noise.

It is realistic to perform the proposed experiment with the sample used in 7.1.1 at a slightly different temperature, where a quality factor of 3×10^6 has been observed before, yielding a 6-fold improvement in the observable temperature shift from quantum backaction. Furthermore, it is possible to use 10 μW of light for the temperature measurement, which should result in at least a 25-fold improvement of the signal-to-noise ratio. Feedback or dynamical backaction cooling with the measurement beam has been demonstrated to reduce the required averaging time by the square root of the signal-to-noise ratio [GVK12], which could yield a significant 100 to 1000-fold improvement in our case at the cost of a small increase in experimental complexity. For the pump beam, up to 1 mW has been successfully injected into current measurement cavities at room temperature. We expect to be able to increase this power at least 5-fold due to a better understanding of the required feedback loops to suppress the low-frequency instabilities from our low-temperature experiments. Lastly, a modulation of the pump beam in the $10 - 100$-Hz range, i.e. well above the thermal relaxation time of the micropillar (of the order of 1 s, with a slowly diverging low-frequency part) and below the feedback-enhanced mechanical damping rate (of the order of 1 kHz), would enable us to drastically reduce the temperature difference due to the absorption of the pump beam at different powers.

While a measurement of quantum backaction effects at room temperature seems at reach even with our unoptimized optomechanical system, a few straightforward improvements are expected to yield at least a 10-fold improvement of the observable signal. The weakest point of the proposed measurement protocol is the linear temperature increase expected from the absorption of light by the micropillar, that must be subtracted from the measured temperature shifts in either a "feed-forward" scheme, a differential measurement scheme involving other mechanical modes, or a lock-in scheme. To avoid this complication, two more alternatives are possible: first, a measurement of the correlations between reflected pump power shot noise and compression mode displacement has been proposed [VTB$^+$09] and realized [PPR13] as a convincing proof of principle of radiation-pressure effects at low temperature. Second, the injection of squeezed light into our optomechanical system that is foreseen for 2017-2018 makes a modulation of the quantum backaction force noise possible without changing any other parameter such as the laser power, for example by rotation of the squeezing angle or by alternate injection of vacuum and squeezed vacuum into the optomechanical cavity. Such an experiment could be used as a first validation of the operability of the squeezing setup at room temperature before progressing to more advanced experiments at cryogenic temperature.

7.4 The noise floor issue

The previous section suggests that the classical displacement noise floor sets a serious limit to the operation of our optomechanical system in the quantum regime. We have therefore studied the effect in more detail. Section 7.4.1 elucidates the origin of the noise floor, while section 7.4.2 presents our approach to overcome the noise floor issue

in future versions of the experiment.

7.4.1 Cavity displacement noise floor

Fig. 7.13 shows the displacement spectrum of the measurement cavity at different temperatures from 0.2 to 300 K, and frequencies from 200 kHz to 10 MHz . The response of the various detectors employed for the measurements is flat to better than 3 dB from 1 to 10 MHz. First, one notes that the noise floor is roughly proportional to the temperature, indicating that the noise originates from a thermal noise mechanism in the measurement cavity and not, for example, from the laser. The large noise level below 2 MHz is expected because the experiment is not optimized for this frequency band. At higher frequencies, large noise peaks occur at 3.5 MHz, 5.9 MHz and 9.0 MHz, which are absent from the cavity that does not contain a micropillar. We therefore attribute these peaks to mechanical modes of the micropillar, specifically to the fundamental and the first and second harmonic compression modes. The plateaus between these peaks have a nearly flat frequency dependence. From a comparison of the room-temperature noise for the cavity constructed with a micropillar and a coupling mirror with the one of a cavity made of two coupling mirrors, we deduce that the dominant contribution of the broad-band noise floor is caused by the coupling mirror (the noise floor with two couplers is about twice as high as with only one coupling mirror). Furthermore, the noise plateau of the cavity with two coupling mirrors starts with a spike near 2.8 MHz, which is also present but less pronounced in the other cavities. This frequency corresponds to the fundamental compression mode of the 1-mm thick silica mirror substrates, which suggests that substrate Brownian motion of the coupling mirror is likely the dominant noise source.

To explain the flat frequency dependence, we have numerically modeled the observed noise. Fig. 7.14 shows a part of the room-temperature spectrum from Fig. 7.13 along with expected noise sources at 300 K. Various thermal noise sources that are known to the gravitational-wave community can be found in [HBD12]. The only relevant known noise sources above 1 MHz are coating thermo-optic noise and substrate Brownian motion. Coating thermo-optic noise is expected to be about two orders of magnitude below the observed noise level. Analytical expressions of substrate Brownian motion [Lev98] are proportional to the frequency-dependent loss angle $\phi(\Omega)$ divided by the frequency Ω. While the standard assumption of structural damping $\phi(\Omega) = \phi_0$ is therefore unable to account for the noise, a loss angle proportional to the frequency yields a flat frequency-dependence of the substrate Brownian motion. We estimate the loss angle of our coupling mirror substrate from our measurement to be

$$\phi(\Omega) = 0.05 \frac{\Omega}{2\pi \times 3.6\,\mathrm{MHz}}\,. \tag{7.4.1}$$

The loss angle of 1/20 at the mechanical resonance frequency seems reasonable for our coupling mirrors[12]. Furthermore, we have attempted to obtain more detailed informa-

[12] In the cavity assembly, about 25% of the back surface of the coupling mirror substrate is glued to a metal support. It is therefore reasonable to expect clamping loss to be the dominant source of mechanical dissipation. Due to mechanical impedance mismatching, only a fraction of the energy of acoustic waves is expected to be lost per reflection by the interface. The observed loss angle 1/20 is therefore within the right order of magnitude.

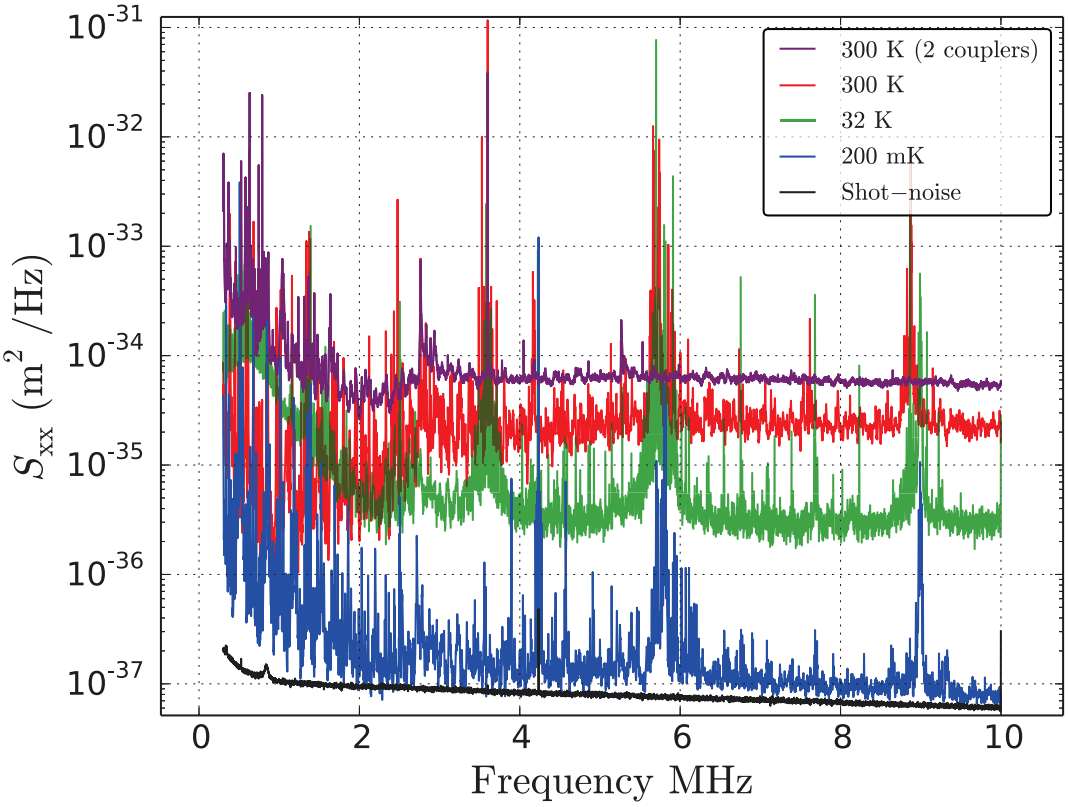


Figure 7.13: Broad-band displacement noise spectra of the measurement cavity with sample 20121009C5 at different temperatures. The measurements were performed by different detection methods (direct detection, PDH, homodyne) and the detection noise floor was subtracted from all spectra. The spectrum of a cavity constructed with two silica mirror substrates (one concave, one flat) of the same type as described in section 3.2.2 is also shown. The measurement at 200 mK was performed with the homodyne detector with 500 μW of incident power and was limited by the shot-noise of the homodyne detector. The shot-noise that was subtracted is plotted along with the spectra to show that the signal-to-noise ratio of the lowest spectrum was larger than one over the entire spectrum. By assuming a linear temperature-dependence of the noise floor, a temperature of 2 K is deduced for the cryogenic spectrum.

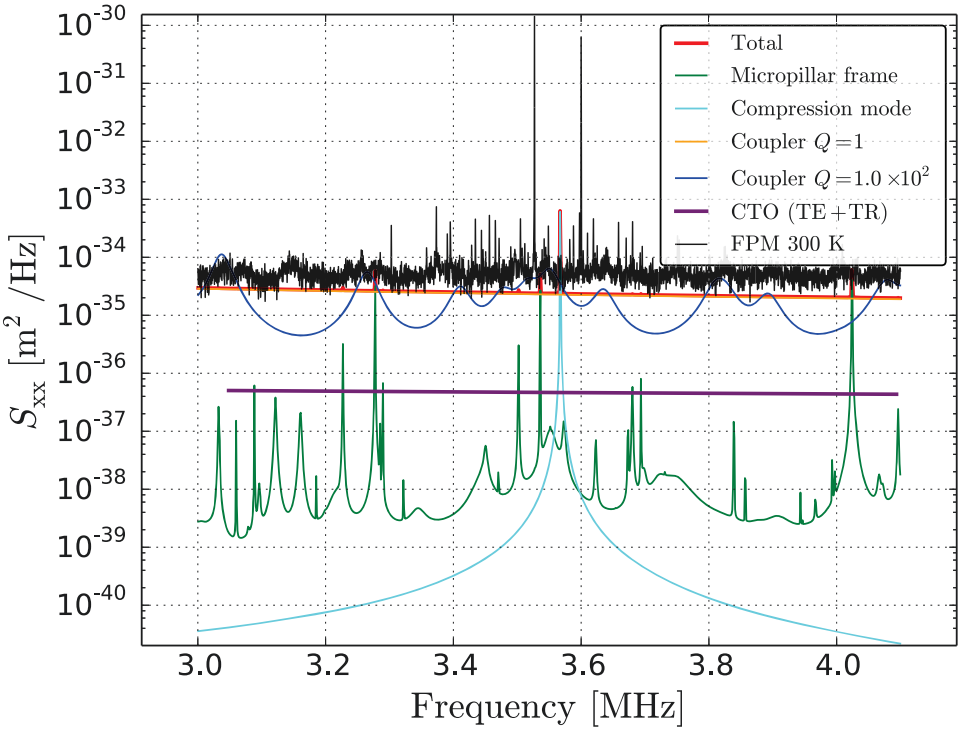


Figure 7.14: Measured ("FPM 300 K") and modeled displacement noise floor at 300 K for a cavity constructed with a micropillar and a coupling mirror. The peak at 3.6 MHz in the experimental trace is a phase-modulation peak, while the peak at 3.58 MHz corresponds to the compression mode of the micropillar. A number of noise mechanisms at 300 K have been modeled. "Micropillar frame": Brownian motion originating from all modeled mechanical modes of the micropillar including its frame, except for "Compression mode", which corresponds to the Brownian motion of the micropillar's compression mode. "Coupler Q=1" and "100": Brownian motion of the coupling mirror substrate assuming a quality factor or 1 and 100 for all modes. "CTO": Coating thermo-optic noise for a single coating, composed of thermo-elastic and thermo-refractive noise. "Total": sum of all previous noise mechanisms, assuming Q=1 for the coupling mirror and two coatings contributing to coating thermo-optic noise. In order to operate the optomechanical system in the quantum regime, a noise floor of the order of $10^{-36}\,\mathrm{m}^2/\mathrm{Hz}$ at room temperature is desirable.

tion about the mechanical mode structure by a COMSOL finite-element simulation. The simulated structure was a free silica plate with dimensions of 10 mm × 10 mm × 1 mm. From the simulated mode shapes, we have calculated the effective mass for a point-sized laser beam at the substrate center with Eq. 2.1.17. The resulting displacement fluctuations could then be computed from the mode frequencies and effective masses with Eq. 2.1.22 by assuming a constant quality factor. The obtained noise spectra are very close to the observed spectrum over a wide range of assumed quality factors and the simulated displacement spectrum is roughly flat over frequency. In the MHz frequency range, the coupling mirror mode density is high enough such that the displacement peaks of different modes overlap significantly for low enough quality factors, thus leading to a near-constant noise floor even for quality factors as high as 100. The analytical model of substrate Brownian motion only takes into account the DC-response of all mechanical modes, and thereby yields a displacement noise spectrum that decreases with frequency. The full simulation of the mechanical mode structure can explain the constant displacement noise spectrum even for a constant loss angle because the resonant contribution of mechanical modes is considered, allowing an increasing density of mechanical states to cancel the decrease of displacement fluctuations with frequency for each individual mode. An identical approach to model substrate Brownian motion was presented in [Wil12] for more macroscopic mirror substrates and found in excellent agreement with the observed noise spectra.

7.4.2 Phononic crystal mirror substrates

While the most likely cause of the cryogenic displacement noise floor was investigated in the previous section, it is not easy to achieve better performance with standard mirrors. Simulations involving bigger (standard 1" size) or smaller (1 mm^3) mirror substrates have yielded very similar noise spectra as our current coupling mirrors: for larger-sized mirrors, the increase in mode mass is compensated for by an increase in mode density, and the inverse occurs for smaller substrates. A more promising approach is the use of a phononic crystal: by periodically patterning the mirror substrate, one can engineer its mode spectrum as to create a band-gap for mechanical modes that encompasses our measurement frequency band. In reference [MEK+07], the three-dimensional acoustic band-structure for thin slabs patterned with holes in a square or honeycomb lattice is computed. The paper shows that the honeycomb lattice works best to implement phononic crystals and demonstrates that for a band-gap to occur, the distance between adjacent holes of the lattice should be of the order of the thickness of the slab. Furthermore, the filling factor, i.e. the diameter of the holes normalized by the maximum possible value, should be as close to 1 as possible. Nevertheless, ratios down to 75% are still expected to allow for a band-gap. We have used these design recommendations and estimated the displacement noise floor for such a structure with the same method as for the simulations in the previous section. Fig. 7.15 presents the results of the simulation with an optimized geometry. From the simulation, we expect a possible noise improvement by up to three orders of magnitude. The design has been tested to be robust (i.e. to yield more than two orders of magnitude better noise performance than our current substrates) against 10% variation in plate thickness, hole spacing and filling factor. For reasonably good performance, at least 96 holes should

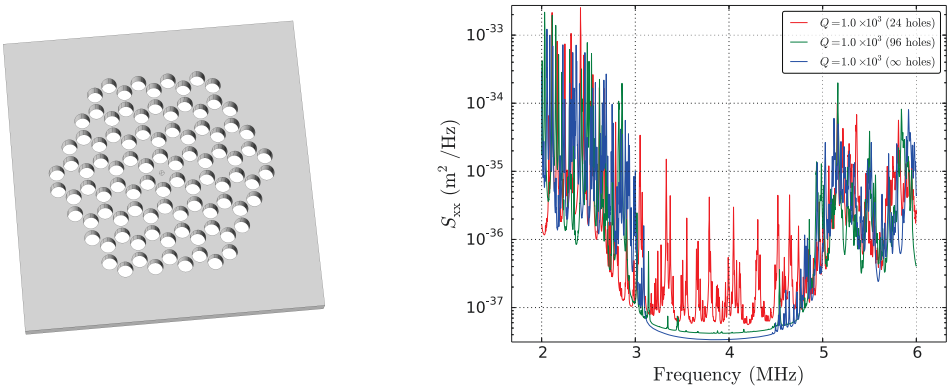


Figure 7.15: Left: Geometry of the designed phononic-crystal mirror substrate. The displayed plate is of dimensions 6 mm × 6 mm × 0.35 mm and contains 96 holes in a honeycomb lattice with a spacing of 0.35 mm and a filling factor of 0.9. The mechanical displacement is simulated at the center of the substrate. Right: Simulated room-temperature substrate Brownian motion spectra for the geometry on the left. Structures with only 24 holes or as many holes as fit into the plate were also simulated.

be patterned. From the figure, one can see that the noise level in band-gap from 3.1 MHz to 4.5 MHz is dominated by the tails of mechanical modes near the band edges. Lower quality factors than a conservatively assumed intrinsic cryogenic quality factor of silica ($Q = 1000$) result in a higher noise floor within the band-gap. The mounting of the phononic crystal mirror is therefore also expected to play a role in the substrate's noise performance.

We expect to fabricate the phononic-crystal substrates by ultrasonic drilling, which is expected to permit a surface roughness of the hole edges of the order of 10 μm and maximum filling factors of about 95%. In order to be able to use the substrates as coupling mirrors in our experiment, the photoablation step needs to be combined with the phononic crystal drilling. While one expects cleaner mirrors by performing the photoablation after the drilling, we believe that the mirror surface can be sufficiently protected during the fabrication of the phononic crystal to allow the opposite fabricatino order. Performing the photoablation before the drilling step has the advantage that any problems related to excessive heating of the phononic crystal structure during photoablation can be avoided, and that the photoablation does not have to be optimized on tediously fabricated phononic crystal samples.

We expect that by using structured mirror substrates, the displacement noise floor of our future measurement cavities will be limited by coating thermo-optic noise at a level around $10^{-39}\,\mathrm{m}^2/\mathrm{Hz}$. If other problems such as instabilities or excessive heating of the micropillar do not worsen in presence of the new coupling mirrors, this should allow to operate our current mechanical oscillator in a regime where quantum backaction dominates other sources of force noise by at least a factor of 10.

Chapter 8
Conclusion

The Book of Nature is written in the language of Python.
— Galileo Galilei, *in a parallel universe*

A large number of topics have been addressed in this thesis. We will briefly summarize the most important points of each chapter in section 8.1, before we discuss future directions of the experiment both in the short term (section 8.2) and in the longer term (section 8.3).

8.1 Summary

After the introduction to this thesis, we have tried to make an original and didactical approach to the presentation of the theory of optomechanics in Chapter 2. It was shown that the mapping of the displacements of the many atoms of a solid onto a single variable, the optomechanical displacement, allows for a great simplification of the dynamics of interest. In particular, it was demonstrated that optomechanical coupling leads to a collocated sensor-actuator pair if photothermal forces can be neglected. Only in this case, a direct application of the fluctuation-dissipation to the effective one-dimensional mechanical system is allowed. The equations derived in the optical section of the chapter lay the foundation for the efficient and complete modeling of linear optomechanical systems that is summarized by Eq. 2.3.14 and Table 2.1 and easily implemented with the Mathematica code in Appendix A.

In Chapter 3, we have discussed the mechanical and optical samples used in this thesis. The fabrication of the mechanical samples is a lengthy process that is not suitable for mass fabrication, but enables to attain exceptional **mechanical quality factors up to 70 millions** at cryogenic temperature. The recipe for coating small-diameter high-reflectivity **dielectric mirrors on a high aspect-ratio structure** that was presented in section 3.1.2 may be of interest outside the scope of this thesis. The expected **effective mass of 33 μg** of our oscillator was experimentally confirmed for the fabricated resonators, which is an important step for inferring effective temperatures from the measured displacements in later chapters. A measurement of the **temperature dependence of the mechanical quality factor** of our samples revealed that despite the design efforts, clamping loss remains the dominant source of mechanical

dissipation in our samples because of a large mode density in the mechanical structure surrounding the micropillar. While the described fabrication of our coupling mirror substrates by photoablation is relatively straightforward, the cleaning procedures of high-reflectivity mirrors with reactive ion etching and "First contact" mentioned in section 3.2.2 might save other experimentalists a tedious trial-and-error procedure. The Python code used to automatically evaluate the finesse measurements in section 3.2.3 constitutes a very practical and robust tool for the various types of optical characterization presented in subsequent sections. A measurement of the **dependence of optical loss on the cavity length** allowed to evaluate the optical quality of our micropillar samples as satisfactory ($P \approx 20$ ppm) and reveals the quantitative contribution of different optical loss mechanisms in our cavities. The construction of a fully-automatized finesse measurement setup allowed to show that datapoints that seemed like outliers in the earlier measurement can be explained by **optical mode degeneracies** at specific cavity lengths shown in Fig. 3.20.

In Chapter 4, we have described in detail the **dilution refrigerator** used for the cryogenic measurements of later chapters in a way that should allow others to successfully operate the cryostat. An inner vacuum chamber that allows to vary the helium pressure independent of the cryostat temperature has been presented end employed to characterize the mechanical damping caused by the presence of either gas or superfluid helium in the micropillar environment. This characterization was done with the possibility in mind to provide additional mechanisms for heat removal from the laser-absorption heated cryogenic micropillar. While gas damping is expected to be inefficient (section 7.1.2), the heat removal rate of superfluid helium could not be studied due to superleaks in the most recent version of the vacuum chamber. However, even the unlikely case that the micropillar is kept at the cryostat base temperature by the heat exchange provided by a **superfluid helium film** would not out-weigh the observed mechanical dissipation caused by the film (section 4.10). The remaining half of the chapter presented the design and implementation of the mechanical structure encompassing the measurement cavity. Valuable information for the experimenter is contained in the cavity alignment protocol in section 4.2.3. While our final measurement cavities do not require the use of a Kapton spacer due to sufficiently low cryostat vibrations and stable mechanical design, the Kapton spacer design presented in section 4.2.1 may be helpful in other cryogenic or room-temperature experiments. The chapter ended with the characterization of the **piezo-electric response** of our mechanical sample to a nearby antenna that modulates the electric field around the sample. The antenna actuator has proved helpful in many other parts of this thesis, for example for ringdown measurements or advanced locking schemes.

Chapter 5 discussed the various optical components of our experiment: the presented design of the filter cavity, the double-pass AOM scheme, the balanced photodiodes and the high actuation bandwidth piezo-mounted local oscillator mirror. The **Raman lasing** observed in an earlier filter cavity prototype (section 5.2.3) demonstrates an important limitation of dielectric mirror coatings that has not been reported anywhere else before. The implemented **automatic feedback loop controllers** for the filter cavity, the homodyne and heterodyne detection, the laser power stabilization, and the AOM difference frequency make the experimental setup completely controllable by a central computer. Beyond simple automatization, the digital controller

implemented with an FPGA allowed to digitally process the error signal of the various feedback loops and enabled further improvements of the setup. In particular, the demodulation scheme for stabilizing the amplitude modulator transmission near zero in section 5.3.1 and for decoupling the homodyne detector lock from cavity length fluctuations in section 5.6 allowed significantly better control of the optomechanical system in later experiments. The enhancement of the homodyne locking bandwidth by an **IIR filter** in section 5.6 shows an improvement which one cannot achieve with currently existing commercial solutions and which can be applied to any feedback loop. Similarly, the generation of a **PDH-like error signal** that is insensitive to modulations of the laser power even for non-zero detunings (section 5.8) constitutes a novel locking technique that has not been achieved with analog electronics before.

Chapter 6 discussed the topic of feedback loops in more detail. Apart from a quick definition of the terminology of control theory, the first part introduced the RedPitaya as a powerful hardware platform for digital servo-loop controllers for analog devices. The presented software package **"Python RedPitaya Lockbox (PyRPL)"** that was originally only developed for the presented experiment contains a large number of tools for the measurement and control in quantum optics experiments, and is therefore interesting in a broader context. Among the presented capabilities are a rapidly interfaced oscilloscope, spectrum analyzer and network analyzer, PID controllers with variable additional input filters, several lock-in modules for digital modulation and demodulation, a frequency-phase discriminator for frequency offset stabilization, and an IIR filter composed of up to 24 biquads. Further software modules allow real-time re-routing of the signal connections between the previously mentioned modules, and can thus perform automatic locking sequences, calibrations of error signals, and gain optimization. The open-source character of all modules allows to adjust the functionality to the specific needs of the experiment. The second part of the chapter described the hardest control problem in this thesis: the suppression of optomechanical instabilities that arise from the low-frequency mechanical mode spectrum of the micropillar samples. We have presented measurements of the low-frequency response of our system to intracavity power modulations, and thereby evaluated the system's stability regime. For the measurement of the transfer function, a **novel error signal** was introduced that allows to avoid the nonlinearities of optical cavity readout in the presence of vibrations of the order of the optical linewidth. With the information obtained by the measurement, a controller was designed that could be implemented with various filters from the PyRPL software package. The final design was shown to prevent the instability of up to **five mechanical modes** at intracavity powers of the order of **25 times above the instability threshold**.

Finally, Chapter 7 presented the actual measurements conducted with our optomechanical system. The **room-temperature measurements of Brownian motion** were used to validate the calibration of our experimental spectra. A study of the variance of the inferred sample temperatures showed a four-fold improvement of the required averaging time over the naive thermal limit that could be achieved by considering additional information about the mechanical properties of our system in the data-analysis. Next, a cooling experiment was performed with our optomechanical system in a **cryogenic environment**. Since our optomechanical system operates in the non-resolved sideband regime, the laser was detuned to the red side of the opti-

cal resonance by one half of the cavity linewidth for this experiment. In a first step, we have used **optomechanically induced transparency** measurements to characterize the effects of dynamical backaction in our optomechanical system at variable incident laser powers, where clear signatures of the modification of the cavity response to a phase modulation by radiation-pressure effects and cooling rates of at least 400 times the mechanical linewidth were observed. The observed effective frequencies and damping rates were in good agreement with independent estimates obtained from mechanical ringdown measurements and from the spectrum of Brownian motion, demonstrating the stability of our setup. Next, we have deduced the effective temperature of the micropillar at various incident powers from measurements of its Brownian motion spectrum. The significant **optical cooling** enabled a reduction of the effective temperature down to only 3.4 mK, which corresponds to a mean thermal occupation number of only **20 phonons** with an optomechanical system of an effective mass that is about three orders of magnitude above the **heaviest optomechanical resonator** [UML+15] that was observed in the quantum ground state up to date. A study of the heating of our sample by laser absorption revealed a stronger power-dependence than expected from an analytical model and finite-element simulations. This dependence can be re-interpreted as a power-dependent degradation of the mechanical quality factor if the assumptions underlying the analytical heating model are assumed to be correct. In a **ringdown measurement** with a method that was shown to be **robust against dynamical backaction effects**, mechanical quality factors up to 7×10^7 were observed with our samples. First steps towards ringdown measurements with temporally modulated mechanical quality factors were presented as well. The next part of the measurements aimed at the observation of quantum backaction effects in our optomechanical system. Our optomechanical system could be operated in a regime where the **quantum fluctuations of the optical intracavity field** are expected to heat the thermal bath of the cryogenic mechanical oscillator to a temperature of **20 K**. However, in this regime another heating mechanism was observed that originates from a **classical noise of the intracavity field** and led in conjunction with absorption heating to a thermal bath at a total temperature of the order of **60 K**. These exciting results motivated the conception of an experiment aiming at the demonstration of **quantum backaction effects at room temperature** with our system. The last part of the chapter presented a study of the temperature and frequency-dependence of the classical optical noise that overwhelmed the quantum backaction in the cryogenic experiment and identified the Brownian motion of the coupling mirror as the dominant noise source. The chapter was concluded with the presentation of the simulation of a new mirror substrate with which we expect a **1,000-fold reduction of optical phase noise from the mirror substrate's Brownian motion** in the near future.

8.2 Short-term prospects

Theory

Intuition is an important ability for an experimental physicist, which usually relies on making the right approximations to understand the dominant mechanisms underlying an experimental phenomenology. We strongly refer the reader unfamiliar with

optomechanics to the references [Coh00, Bri04] that date back to the earlier years of optomechanics for an intuitive explanation of optomechanical coupling, because we have concentrated on a theoretical description of optomechanical coupling with as few approximations as possible in the theoretical introduction in Chapter 2 of this manuscript. As a result of this, Eq. 2.3.14 in conjunction with Eq. 2.1.18 exhaustively describes all effects of linear optomechanics for realistic mechanical structures. These equations quickly become too complex to manually compute the expected experimental signatures without making small mistakes. The Mathematica code in Appendix A is a first step to implement a mistake-free realistic modeling of an optomechanical system. It could however be extended to an **open-source Python software module for modeling optomechanical effects** that is shared and tested by different members of the scientific community. Python modules such as **"SymPy"** allow for an implementation with a minimum number of theoretical assumptions that can easily be verified and altered if necessary. The development of PyRPL has shown that collaborative software development and a broad user-base lead to the rapid eradication of software bugs and to a high rate of functionality improvements, which quickly outpace individual attempts to attain the same results. With **continuous integration** tools such as "Travis CI", a software module for modeling optomechanical effects can be automatically tested on a pre-defined number of example test cases with known results after each change of the software module. A large number of test cases can be taken from the existing articles on optomechanics, making the correctness of the software prediction objective. The resulting ease of availability and absence of doubt of the theoretical modeling of an experiment is expected to foster experimental progress. The long-term generalization of the software module could become the equivalent of "Altium"[1] for quantum optics, which will become necessary for the foreseeable commercialization of quantum optical devices.

Mechanical sample

Undesired mechanical modes of our mechanical sample near the compression mode frequency are a reason for a bad fabrication yield and a source of mechanical dissipation. Additionally, these modes are likely to cause an apparent excessive absorption heating (see the end of section 7.2.4). Furthermore, the fragility and uniqueness of samples with high mechanical quality factors acts as an incentive for the experimenter to postpone straightforward improvements of the measurement cavity to avoid the risk of breaking a good sample. It therefore seems a good idea to search for new designs of mechanical resonators. In the optomechanics community, it has become almost normal to employ **phononic crystals** to drastically reduce the density of modes in the measurement frequency band. Existing work in this direction [Zer17] should be pursued further.

The piezoelectricity of quartz is very practical to drive the mechanical oscillator. However, actuation by an antenna does not constitute a collocated sensor-actuator pair with the optical read-out, and could therefore not be employed successfully as a

[1] Altium provides software tools that allow the design, visualisation, printed circuit board layout and realistic simulation of digital and analog electronic circuits, including issues such as RF-crosstalk and impedance of circuit board connections.

feedback actuator. Since high-quality factor mechanical oscillators are easily driven by environmental noise, working with materials that are **not piezo-electric** will make the experiment easier. Many crystalline materials such as silicon, sapphire, or calcium fluoride are known to feature low mechanical dissipation at cryogenic temperatures [Sch08]. Future mechanical oscillators should be made from one of these materials.

Coupling mirror sample

A design flaw of our coupling mirror samples was exposed in the presented transmission measurement to be an uncertainty in the mirror transmission due to the parallel, uncoated back-surface of the mirror substrate. Future samples should be be **anti-reflection-coated** on the back surface. Attention has to be paid to the compatibility between the AR-coating and other fabrication steps (see Fig. 3.14).

Another improvement that is expected to enable the operation of our optomechanical system deeply in the quantum regime is the development of **phononic crystal coupling mirror substrates** as proposed in section 7.4.2. The fabrication seems straightforward and sufficient margin for experimental imperfections is tolerable to motivate this direction. Preliminary tests have shown that existing coupling mirrors can be thinned down to the required thickness from the uncoated back-side with sandpaper and optically polished again. About 20 coupling mirror substrates with burned anti-reflection coatings on the backside are currently not usable for the experiment and therefore available for the development of the fabrication recipe, such that other more recently fabricated samples can be transformed into phononic-crystal substrates with an optimized fabrication process. The fabrication of phononic-crystal mirrors is compatible with the application of an AR coating in the end.

The optical round-trip loss that we achieve with cavities constructed from a micropillar and a coupling mirror sample is consistently in the range from 30 to 35 ppm. Since the currently employed coupling mirrors transmit about 35 ppm, this results in maximum quantum efficiencies for the detection of the intracavity field of 50%. The already fabricated next generation of conventional coupling mirrors has been observed to feature transmissions up to 110 ppm, which will **improve the quantum efficiency** realistically to about 65 to 70 %.

Cryostat

The heating of the micropillar by absorbed light discussed in section 7.2.4 reveals that a dilution refrigerator is not strictly necessary for our experiments. Using a ^{3}He or even a ^{4}He cryostat in preliminary steps seems to bring about only a small increase of the final micropillar temperature while greatly reducing the experimental complexity and cool-down time. A **test setup in a ^{4}He refrigerator** has already been constructed by Rémi Metzdorff.

The current inner vacuum chamber has **superleaks** and must therefore be improved if it should contain superfluid helium below the λ-point. While it is unlikely that an improvement in the performance of our optomechanical system results from this, the mechanism of **mechanical damping induced by the Rollin film**, which is one possible conclusion from section 4.1.7, is a highly interesting subject of study and should outweigh the small amount of work needed to close the leaks in the IVC.

Measurement cavity

The antenna near the micropillar that was introduced in section 4.2.4 is a very practical tool. However, the possibility that noise and low-frequency cross-talk from the locking electronics is injected into the optomechanical system by this actuator should at least be tested by connecting the **electrode to the cryostat ground potential** in a future cryogenic run. For similar reasons, the fast piezo that was introduced to obtain a higher locking bandwidth should either be grounded completely, or made electrically floating with respect to the cryostat potential in a future run.

In the longer term, either a **fast piezo with better performance** than the current one, or a **mechanical suspension** system to decouple the measurement cavity from cryostat vibrations should be introduced. The expected stability improvement of the measurement cavity lock should translate directly into higher injectable optical powers: the range of detunings over which the feedback loops that suppress the parametric instabilities of low-frequency modes are stable for a given implemented loopshape is currently below the amplitude of the cavity vibrations that originate from helium boil-off in the cryostat reservoir.

Quality factor measurement

The discussion in section 7.2.4 suggests to perform the measurement of the **detailed temperature-dependence of the mechanical quality factor** of our current sample in the cryogenic temperature range to shed light on the mechanism that causes the apparent excessive laser absorption heating. Similarly, a mechanical **ring-down measurement during which the power of a pump laser is modulated**, as suggested in section 7.2.5, should be performed to extract the amplitude and phase-lag of the laser power-dependence of the mechanical quality factor. Another benefit of this measurement is that an experimental estimate of the time-dependence of absorption heating can be obtained, which might enable the operation of our system in a **pulsed regime**.

A more involved study, which is also expected to yield more exciting results, is to conduct a **measurement of the spectrum of mechanical dissipation** with the micropillar mode at cryogenic temperature. This can be achieved by scanning the effective mechanical frequency with a feedback-induced optical spring while tuning the feedback-enhanced mechanical damping rate to a value that corresponds to the desired resolution bandwidth of the loss spectrum[2]. The loss spectrum can be inferred from measurements of the environment temperature with the Brownian motion of the sample during the scan, which makes the proposed method insensitive to small fluctuations of the feedback cooling rate. We expect an improvement of the mechanical quality factor to result from this method.

[2]The required feedback controller is available in PyRPL and can be used if the dynamical range of the inputs and outputs of a RedPitaya (80 dB) is sufficient to prevent the injection of excess noise with the feedback controller.

PyRPL

The work on PyRPL that is described in section 6.2 is an on-going process. We currently share the software with five other research groups who contribute by testing the software, reporting bugs and suggesting or implementing new functionality. Interest and positive feedback from other research groups suggests to pusblish the software as an **open-source project** in the near future to enlarge the possible collaborations to the entire quantum optics community. A vast number of smaller and larger **improvements** are planned and can be found on the project website [Neu16b] in the sections "Issues" and "Milestones". The software is in principle not restricted to the Red Pitaya board as a hardware platform and can be ported to other SoCs with straightforward modifications. In particular, a custom-designed FPGA board with **higher-resolution ADCs and DACs** and the foreseeable release of an upgraded version of the Red Pitaya board based on a **larger FPGA** (Zynq7020) are the most likely options that will enable a performance boost of PyRPL-based laboratory devices.

Room temperature measurement

The **room-temperature measurement of quantum backaction effects** proposed in section 7.3.3 is relatively easy to implement. A successful measurement would prove an exceptional stability of our measurement setup, and can be validated at a later date by repeating the same experiment by replacing the vacuum fluctuations of the laser beam injected into the optomechanical cavity by squeezed vacuum.

Feedback cooling measurement

While not explicitly mentioned, feedback cooling [CHP99] has been employed frequently during the work for this thesis to increase the mechanical linewidth or to suppress the instability of low-frequency modes. However, it remains to be demonstrated that the **quantum ground state** can be reached in the non-resolved sideband regime through feedback cooling [WSP+15]. The corresponding experiment is a straightforward continuation of the experiments in section 7.3.2.

Sideband asymmetry measurement

The **lock-in scheme** of the homodyne detection as described in section 5.6 should be implemented for the heterodyne lock as well. This will allow to perform thermometry through the asymmetry of the Stokes and anti-Stokes sidebands imprinted by the mechanical oscillator on the field reflected from the cavity and lead to a further increase in the confidence level of the measurements of the force noise environment of the cryogenic micropillar compression mode.

8.3 Long-term outlook

The experimental apparatus presented in this thesis offers many possible directions to choose from. The existing automated optical setup is able to inject two low-noise laser beams with amplitude and phase modulations into a dilution refrigerator, and to detect

the reflection by direct, homodyne and heterodyne measurements. Therefore, the setup can be used to characterize various optomechanical systems. As an alternative to the micropillar resonator, current collaborations outside and within our research group suggest the use of silicon wheel resonators coated with dielectric mirrors [MCF+10] or highly reflective photonic-crystal mirrors implemented in silicon nitride membranes [MAK+15]. While the former option is expected to suffer from similar limitations as the micropillar due to the coupling mirror displacement noise, the latter option could immediately allow to operate the experiment in the optomechanical quantum regime due to an effective mass about three orders of magnitude below the micropillar mass. The currently implemented filter cavity is expected to be sufficient to keep laser noise negligible at the required mechanical frequencies and optical powers.

Regardless of the employed mechanical resonator, we plan to replace the homodyne signal laser by a source of bright squeezed light in the near future. The source is expected to feature a frequency-dependent squeezing quadrature. This is realized by combining a frequency independent squeezer based on an optical parametric amplifier with a rotation cavity. Injecting this new laser source into an optomechanical system with identical parameters as for the measurement presented in section 7.3.2 is expected to allow for the demonstration of broadband quantum enhancement of the displacement sensitivity. In particular, Fig. 7.12 enables to estimate that at an incident power of 500 μW, we expect amplitude-squeezed light at the mechanical resonance frequency to reduce the total force noise spectral density acting on the compression mode by up to 15 %, which could become even larger if the incident laser can be successfully tuned closer to the optical resonance to reduce the drive caused by the cavity displacement noise floor. In the measurement with coherent light at 500 μW, shot-noise contributed by about 30% to the total detection noise budget because of the presence of the Brownian motion of the coupling mirror substrate. Squeezing the phase quadrature of the incident laser source at a frequency above or below the mechanical resonance would result in a reduction of the detection noise floor at this frequency by about 12 %. Of course, if the Brownian motion of the coupling mirror is considered a signal instead of noise, the expected noise reduction would be much larger. The experimental improvements listed in the previous section will increase the available quantum-enhancement to a level that can be seen by looking at the displacement spectra without the need of plotting the data in linear scale or excessive data analysis methods. Apart from the quantum-enhancement provided by frequency-dependent squeezed light that finds an optimal trade-off between amplitude and phase fluctuations of the measurement beam, correlations between the two quadratures can push the measurement sensitivity below the level of the standard quantum limit [JR90]. An improved experimental setup is also expected to observe these signatures.

By combining a source of frequency-dependent squeezed light with an optomechanical system, the number of tunable parameters increases from two (detuning and power) to four (detuning, power, squeezing angle and rotation frequency). Recent theoretical work [AZV16] shows that even without a rotation cavity, interesting possibilities such as radiation-pressure cooling down to the quantum ground state in the non-resolved sideband regime arise from the use of squeezed light. It is reasonable to expect more surprises in the even larger parameter space accessible by the upgraded version of our setup.

Appendices

Appendix A

Transfer functions of a Fabry-Perot cavity

A.1 Mathematica code

In order to compute the transfer functions from input quadratures $p_{\mathrm{in}}(\Omega)$, $q_{\mathrm{in}}(\Omega)$, $p_{\mathrm{vac}}(\Omega)$, $q_{\mathrm{vac}}(\Omega)$ and $\delta x(\Omega)$ to the output quadratures $p_{\mathrm{ref}}(\Omega)$ and $q_{\mathrm{ref}}(\Omega)$, the following Mathematica code was used. The corresponding Mathematica notebook susceptibility.nb can be found among the supplementary material of this document.

In[1]:=

```
αin[Ω_]:=If[Ω==0,ᾱin,If[Ω≥0,(pin[Ω]+i qin[Ω]) ᾱin/(2 Abs[ᾱin]),
((Conjugate[pin[-Ω]]+i Conjugate[qin[-Ω]]) ᾱin)/(2 Abs[ᾱin])]];
αvac[Ω_]:=If[Ω==0,0,If[Ω≥0,1/2 (pvac[Ω]+i qvac[Ω]),
1/2 (Conjugate[pvac[-Ω]]+i Conjugate[qvac[-Ω]])]];
δx[Ω_]:=If[Ω==0,0,If[Ω≥0,δxth[Ω],Conjugate[δxth[-Ω]]]];
χc[Ω_]:=κ/(2 (κ/2-i (Δ+Ω)));
α[Ω_]:=χc[Ω] Sqrt[2 F/π] (Sqrt[η] αin[Ω]+Sqrt[1-η] αvac[Ω]
+i Sqrt[η] χc[0] αin[0] (4 F) δx[Ω]/((2 π) λ));
αref[Ω_]:=αin[Ω]-Sqrt[2 π η/F] α[Ω];
F:=2 π c/(κ 2 L);
p[x_,Ω_]:=x[Ω] Conjugate[x[0]]/Abs[x[0]]+Conjugate[x[-Ω]] x[0]/Abs[x[0]];
q[x_,Ω_]:=-i (x[Ω] Conjugate[x[0]])/Abs[x[0]]+i (Conjugate[x[-Ω]] x[0])/Abs[x[0]];
asspt={η>0,η<1,κ>0,λ>0,F>0,ᾱin>0,Δ∈Reals,Ω>0,c>0,L>0};
subst={δx[0]→0};
Print["Sanity checks (should be ):"]
Simplify[ComplexExpand[p[αin,1]]]
```

```
Simplify[ComplexExpand[q[α_in,1]]]
Simplify[(p[α_ref,-Ω]*-p[α_ref,Ω])/.subst,asspt]
Print["Intracavity transfer coefficients:"]
L11[Ω_]= Simplify[p[α,Ω]/.subst,asspt]/.
      {q_in[Ω]→0, p_vac[Ω]→0,q_vac[Ω]→0,δx_th[Ω]→0}
L12[Ω_] = Simplify[p[α,Ω]/.subst,asspt]/.
      {p_in[Ω]→0, p_vac[Ω]→0,q_vac[Ω]→0,δx_th[Ω]→0}
L13[Ω_]= Simplify[p[α,Ω]/.subst,asspt]/.
      {p_in[Ω]→0,q_in[Ω]→0,q_vac[Ω]→0 ,δx_th[Ω]→0}
L14[Ω_] = Simplify[p[α,Ω]/.subst,asspt]/.
      {p_in[Ω]→0,q_in[Ω]→0, p_vac[Ω]→0,δx_th[Ω]→0}
L15[Ω_] = Simplify[p[α,Ω]/.subst,asspt]/.
      {p_in[Ω]→0, q_in[Ω]→0,p_vac[Ω]→0,q_vac[Ω]→0}
Print["Quadrature transfer coefficients: (in agreement with Zhang1995)"]
R11[Ω_]= Simplify[p[α_ref,Ω]/.subst,asspt]/.
      {q_in[Ω]→0, p_vac[Ω]→0,q_vac[Ω]→0,δx_th[Ω]→0}
R12[Ω_] = Simplify[p[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0, p_vac[Ω]→0,q_vac[Ω]→0,δx_th[Ω]→0}
R13[Ω_]= Simplify[p[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0,q_in[Ω]→0,q_vac[Ω]→0 ,δx_th[Ω]→0}
R14[Ω_] = Simplify[p[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0,q_in[Ω]→0, p_vac[Ω]→0,δx_th[Ω]→0}
R15[Ω_] = Simplify[p[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0, q_in[Ω]→0,p_vac[Ω]→0,q_vac[Ω]→0}
R21[Ω_]= Simplify[q[α_ref,Ω]/.subst,asspt]/.
      {q_in[Ω]→0, p_vac[Ω]→0,q_vac[Ω]→0,δx_th[Ω]→0}
R22[Ω_] = Simplify[q[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0, p_vac[Ω]→0,q_vac[Ω]→0,δx_th[Ω]→0}
R23[Ω_]= Simplify[q[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0,q_in[Ω]→0,q_vac[Ω]→0 ,δx_th[Ω]→0}
R24[Ω_] = Simplify[q[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0,q_in[Ω]→0, p_vac[Ω]→0,δx_th[Ω]→0}
R25[Ω_] = Simplify[q[α_ref,Ω]/.subst,asspt]/.
      {p_in[Ω]→0, q_in[Ω]→0,p_vac[Ω]→0,q_vac[Ω]→0}
Print["Quadrature spectra transfer coefficients:"]
Simplify[R11[Ω]*R11[-Ω]*,asspt]
Simplify[R12[Ω]*R12[-Ω]*,asspt]
Simplify[R13[Ω]*R13[-Ω]*,asspt]
Simplify[R14[Ω]*R14[-Ω]*,asspt]
Simplify[R15[Ω]*R15[-Ω]*,asspt]
Simplify[R21[Ω]*R21[-Ω]*,asspt]
Simplify[R22[Ω]*R22[-Ω]*,asspt]
Simplify[R23[Ω]*R23[-Ω]*,asspt]
Simplify[R24[Ω]*R24[-Ω]*,asspt]
Simplify[R25[Ω]*R25[-Ω]*,asspt]
Print["Ratios:"]
```

```
Simplify[R15[Ω] / R12[Ω],asspt]
Simplify[R25[Ω] / R22[Ω],asspt]
```

A.2 Transfer matrix coefficients

For the special case $\eta = 1$, we obtain the following coefficients:

$$L_{11} = \frac{2\sqrt{2}\sqrt{c\kappa}\left|i\Delta - \frac{\kappa}{2}\right|\left(4\Delta^2 + \kappa(\kappa - 2i\Omega)\right)}{\sqrt{L}(\kappa - 2i\Delta)(\kappa + 2i\Delta)(-2i\Delta + \kappa - 2i\Omega)(2i\Delta + \kappa - 2i\Omega)} \tag{A.2.1}$$

$$L_{12} = -\frac{8i\sqrt{2}\Delta\Omega\sqrt{c\kappa}\left|i\Delta - \frac{\kappa}{2}\right|}{\sqrt{L}(\kappa - 2i\Delta)(\kappa + 2i\Delta)(-2i\Delta + \kappa - 2i\Omega)(2i\Delta + \kappa - 2i\Omega)} \tag{A.2.2}$$

$$L_{13} = L_{14} = 0 \tag{A.2.3}$$

$$L_{15} = -\frac{16\sqrt{2}c\Delta\sqrt{c\kappa}\bar{\alpha}_{\text{in}}\left|i\Delta - \frac{\kappa}{2}\right|}{\lambda L^{3/2}(\kappa - 2i\Delta)(\kappa + 2i\Delta)(-2i\Delta + \kappa - 2i\Omega)(2i\Delta + \kappa - 2i\Omega)} \tag{A.2.4}$$

$$R_{11} = R_{22} = \frac{16\Delta^4 + 8\Delta^2\left(\kappa^2 - 2\Omega^2\right) + \kappa^4 + 4\kappa^2\Omega^2}{(2\Delta - i\kappa)(2\Delta + i\kappa)(2\Delta - i\kappa - 2\Omega)(2\Delta + i\kappa + 2\Omega)} \tag{A.2.5}$$

$$R_{12} = -R_{21} = \frac{16\Delta\kappa\Omega^2}{(2\Delta - i\kappa)(2\Delta + i\kappa)(2\Delta - i\kappa - 2\Omega)(2\Delta + i\kappa + 2\Omega)} \tag{A.2.6}$$

$$R_{13} = R_{14} = R_{23} = R_{24} = 0 \tag{A.2.7}$$

$$R_{15} = -\frac{32ic\Delta\kappa\Omega\bar{\alpha}_{\text{in}}}{\lambda L(2\Delta - i\kappa)(2\Delta + i\kappa)(2\Delta - i\kappa - 2\Omega)(2\Delta + i\kappa + 2\Omega)} \tag{A.2.8}$$

$$R_{25} = \frac{8c\kappa\bar{\alpha}_{\text{in}}\left(4\Delta^2 + \kappa(\kappa - 2i\Omega)\right)}{\lambda L(2\Delta - i\kappa)(2\Delta + i\kappa)(2\Delta - i\kappa - 2\Omega)(2\Delta + i\kappa + 2\Omega)} \tag{A.2.9}$$

The stated transfer functions for modulations and noise spectra have been verified to be in agreement with the ones in Ref. [ZPG+95, Sch09].

Appendix B

List of functional micropillar samples

Sample	Q (10^6)	Status at Q-measurement	Actual status
20121009C5	65	low T in cavity	mounted in cryostat
20121009C6	3	demetallized	mirror baked out
20121002C4	2.4	etching finished	etching finished
20140113C6	1.9	demetallized	mirror baked out
20120905C6	1.3	etching finished	etching finished
20140113C1	0.8	etching finished	incompl. demetallized
20140113C7	0.7	mirror coated	mirror baked out
20140113C5	0.5	low T in cavity	mirror baked out
20121009C4	0.5	mirror baked out	mirror baked out
20140113C3	0.5	etching finished	incompl. demetallized
20121002C9	0.3	etching finished	incompl. demetallized
20121002C3	0.15	mirror baked out	mirror baked out
20140114C2	0.12	mirror baked out	mirror baked out
20140114C8	0.1	etching finished	incompl. demetallized
20121002C6	0.1	demetallized	demetallized
20121002C5	0.1	mirror coated	mirror coated
20121009C7	0.06	mirror coated	mirror baked out

Table B.1: Successfully fabricated micropillar samples which are ready to be characterized further or used. For the more promising samples, the quality factor Q was not measured after the coating of the mirror in order to minimize the risk of degradation of the optical properties. The order of fabrication steps is: etching finished - incompletely demetallized - demetallized - mirror coated - mirror baked out - low T in cavity.

Appendix C

Differential photodiode schematic

Figure C.1 shows the implemented photodiode circuit.

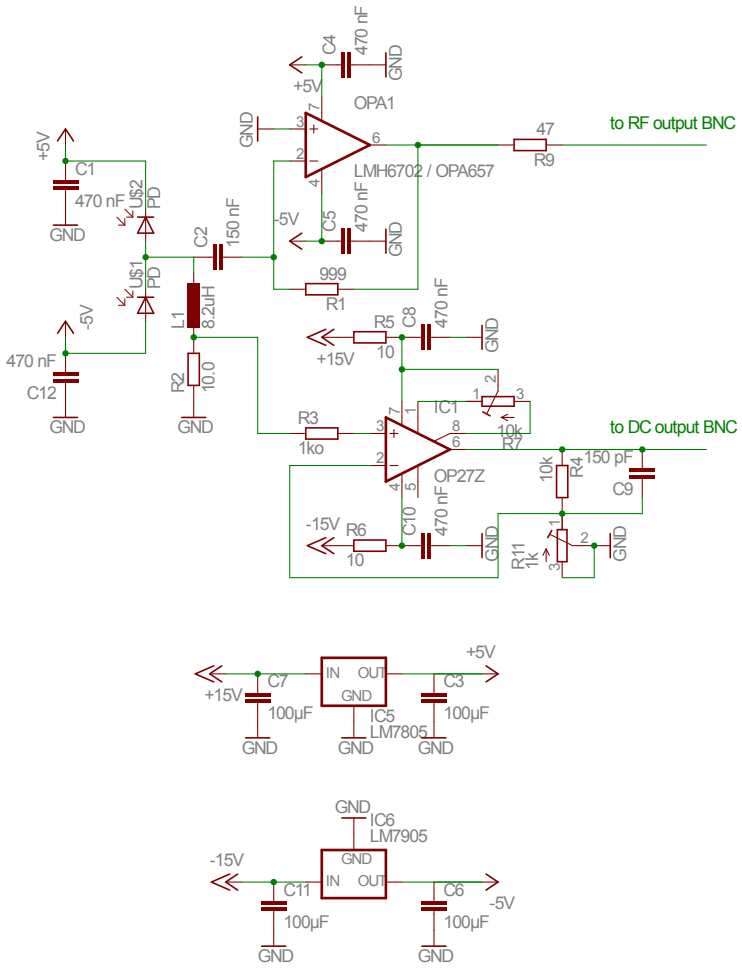


Figure C.1: Electronic circuit of the photodetector.

Appendix D
Cryostat wiring schematics

Figure D.1 shows the implemented wiring inside the cryostat.

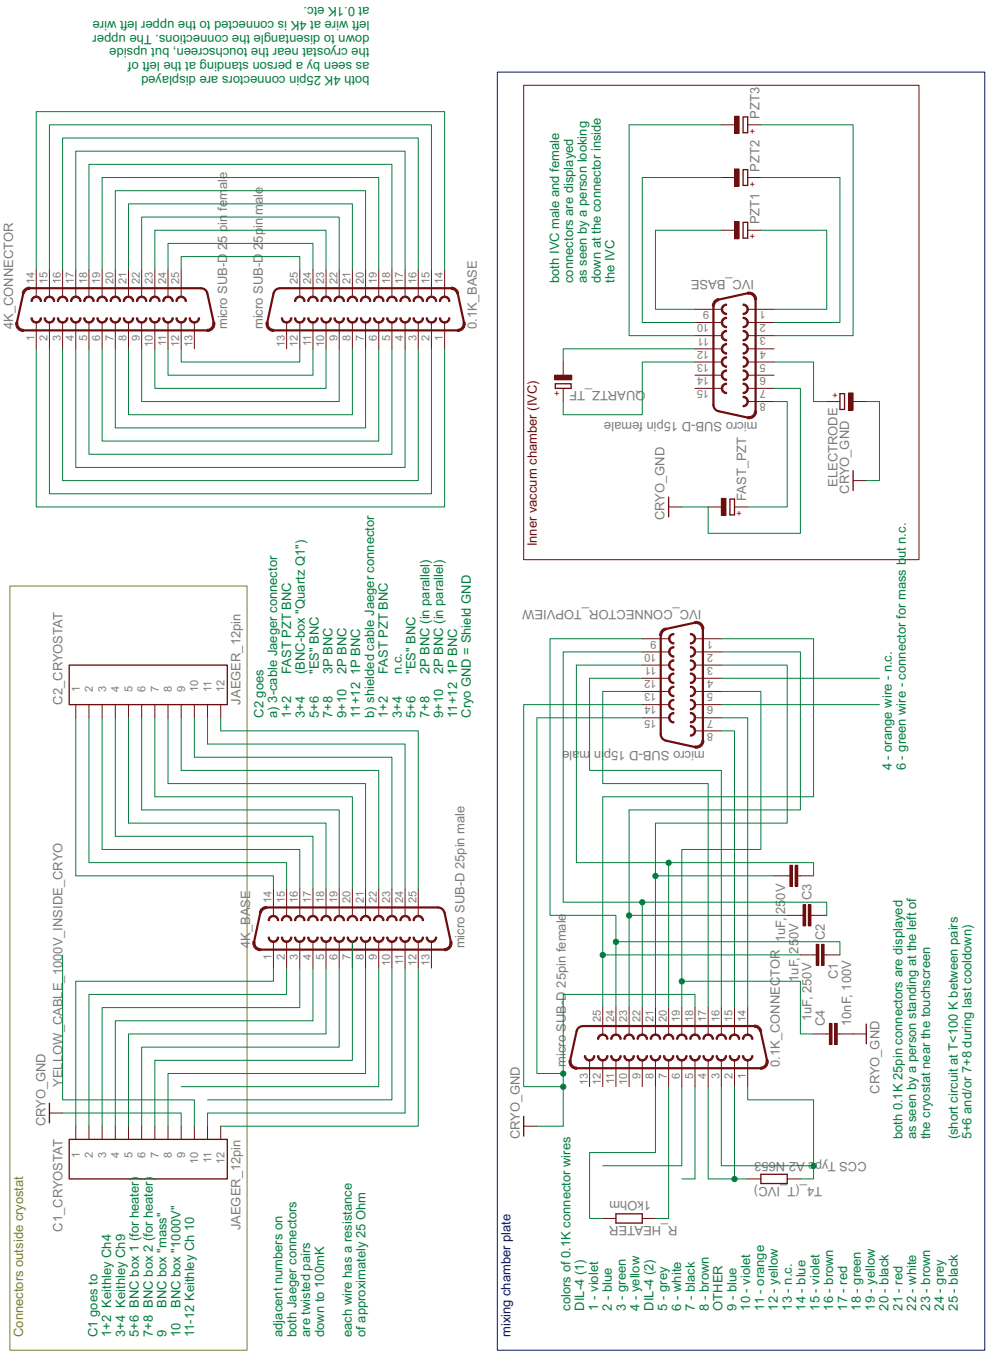


Figure D.1: Schematics of the wiring inside the cryostat.

Appendix E

Operation of the horizontal dilution refrigerator

Here, we summarize the most important steps to successfully operate the dilution refrigerator.

1. Mount and test the sample and associated electrical connections on the mixing chamber plate. Align the free-space laser beam with the open cryostat such that it is minimally clipped by the alignment diaphragm at its foreseen position, which simulates the position of the last thermal shield's window. A cryogenic lens with lateral positioning screws is usually employed for that purpose. Tighten all screws before closing the cryostat. Estimate the reflection coefficient of the beam from the cryostat.

2. Mount all five heat shields. The firsts 3-4 shields can be installed manually, the last ones require the custom-built shield holder to lower the risk of damaging the cryostat. Apply screws to all holes except for the ones marked with an 'x'. After installation of each heat shield, re-optimize the incient beam alignment to prevent losing the alignment completely since the shields' weight tends to cause misalignment. Verify that the reflection coefficient from the cryostat has not been degraded by more than 5 %, otherwise restart the alignment with a corrected alignment diaphragm position.

3. Remove the O-ring from the cryostat case, clean the O-ring and the contacting metal surfaces, apply vacuum grease to the O-ring and close the outer vacuum chamber (OVC) with all screws. Re-confirm that the alignment has not been degraded.

4. Pump down the OVC pressure. We usually pump from 1000 mbar to 10 mbar over the course of 1 hour in order to keep convection and the associated risk of degradation of optical surfaces low. Once the pressure is below 2 mbar, start the turbo pump and monitor the OVC pressure. Confirm that the alignment has not been degraded during the pumpdown.

5. Empty the cold trap if necessary (this step may precede the previous ones): if with all valves closed and the cold trap at room temperature, the cold trap

pressure K5 reads more than 1000 mbar, it is a good idea to remove the trapped air and oil. If the pressure is lower, no emptying of the cold trap is required for one or more further cryostat cycles. To empty the trap, first fill the cold trap dewar with liquid nitrogen. Once the cold trap is cold (K5=0), open the manual valve V11 and launch the cycle "Nettoyage du mélange" in order to circulate the mixture through the cold trap (2-3 cycles are sufficient). After this, bring the mixture back to the reservoir by launching "Rapatriement du mélange". Close the manual valve V11 and confirm that the no mixture was lost (compare values of K6 and K3 to earlier experiments, typically 890 mbar and 70 mbar). Once the turbo has stopped, make sure that the turbo pump pressure P1 reads "<5E-4". With all valves closed, open valves V17 and V6. Confirm that P1 does not rise above "1E-2". Now, remove the cold trap from the nitrogen bath in order to let it heat to room temperature. In the meantime, connect a turbopump stand to the manual outlet V18 (still closed) and start the primary pump. Once the pressure K5 reaches a plateau, the cold trap has reached room temperature and you can start pumping. Therfore, open V18 with the primary pump running and launch the turbo once K5 has reached 1 mbar. Pump at least until P1 reads less than "1E-2" mbar. When finished, close valves 6, 17 and the manual valve 18. Disconnect the pump and apply a blind flange behind V18 for additional security.

6. If the cryostat has not been running for a few weeks, it is a good idea to pump on the dilution tubing while pumping on the OVC in order to remove air that might have leaked into the dilution circuit. To do so, make sure that all valves are closed. Open V10 in order to balance the pressure around the primary pump PP, then close it again. Launch PP, after 10 seconds open V2, V1, V23. Once P1 is below 1 mbar, start the turbo pump and keep pumping for at least 1 day. When finished, close V23, V1, turn off the turbo pump. Once the pump speed is below 10, close V2, stop PP, open V10 and then close it again.

7. Make sure all thermometers are running and that you have a good real-time way to monitor the temperatures.

8. For the cooldown, it is best to start in the early morning. Keep the turbo pumping on the OVC. Make sure a bottle of 100 l of ^{4}He is available. Exchange the bottle head to make it compatible with our setup, connect the compressor plastic tube to it and open the corresponding valve on the bottle head. Make sure that the helium recovery valves of the cryostat are all open. Close the manual valve "Olivier" and set the ^{4}He throughput to 60 l/s. K9 should display about 920 mbar (recovery pressure). Start transferring helium to the cryostat by inserting the transfer cane to the bottle and into the ^{4}He inlet of the cryostat and closing the recovery valve of the bottle. Once K10 falls below 1000 mbar, set the transfer thresholds to 980 and 1050, the transfer rate to 30 l/s and hit "Départ transfert auto". The temperature of the 200 K, 100 K and 4 K plates should start to drop.

9. Fill the cold trap dewar with liquid nitrogen. Once the trap is cold (after 3 minutes), make sure all valves of the dilution circuit (especially V18) are closed,

then open V11. Note the values of K6 and K3 for future reference. Start a new cycle by launching "Nettoyage du mélange" with 2 cleaning cycles and the condensation pressure set to 145 mbar. Within 10 minutes, the cleaning should be finished and mixture should circulate through the cryostat. The 200 K, 100 K and 4 K plate temperatures should now drop more slowly, and the DB and BOUILLEUR temperatures should start decreasing. Keep this mode of operation for about 10 h. When the cryostat BOUILLEUR temperature reaches about 50 K, increase the ^{4}He transfer rate to 60 L/s until the reservoir is full. If this does not occur within 1.5 h, it may help to open the valve "Olivier". During the entire cooldown, make sure not to lose alignment of cryogenic optical parts.

10. If the day is over, bring the mixture back to the reservoir by launching "Rapatriement du mélange". Stop the ^{4}He transfer by selecting "Stop auto transfer" and wait until K10 reaches the helium recovery pressure (about 920 mbar). Then, remove the transfer cane from the cryostat and bottle and make sure the ^{4}He inlet valve on the cryostat is closed properly. On the next day, start transferring ^{4}He as necessary and recirculate the mixture as above by starting from "Nettoyage du mélange" with a full cold trap nitrogen dewar.

11. Once the cryostat's ^{4}He reservoir is full (>14 cm) and all the temperatures of BOUILLEUR, DB and BM are below 6 K, condensation may start. Open the valve "Olivier" in order not to have to bother about the ^{4}He circuit. Click "retour au mode manuel", then "restart auto-run" and select "Inversion". Within 1 minute, condensation starts and should finish after no more than 2.5 hours. The cryostat automatically passes into dilution operation, which you will note by the turbo pump starting and the BM temperature drop below 500 mK. After another 30 minutes, the mixing chamber base temperature (currently about 44 mK) should be reached. The OVC pressure should remain below 1×10^{-5} mbar during the entire condensation process. If the pressure rises suddenly by a factor of two or more, this might indicate a superleak in the dilution circuit.

12. Once in dilution mode, the Jaeger connector "C2" should be replaced by the shielded version if this has not been done before, in order to reduce electrical noise in the cryostat. To reduce vibrations, you may (but do not have to) stop to pump on the OVC for up to 15 days, if you firmly close the OVC valve on the cryostat. The cryostat with a full ^{4}He reservoir and full nitrogen dewar can be left alone for 24 hours before a refill is necessary. The reservoir can be refilled about three to four times with a ^{4}He bottle containing 100 l.

13. When the experiment is to be re-heated, launch "Rapatriement du mélange". Once finished, make sure that no mixture was lost by comparing the values of K6 and K3 with previously noted ones. Close the manual valve V11. Keep valve "Olivier" open to let residual ^{4}He evaporate. The cryostat will automatically heat up over the course of 1-5 days, depending on the OVC pressure. If a fast heatup is desired, mixture may be circulated by going to "Pré-refroidissement" mode, even if no ^{4}He is present in the cryostat. In this case, the cryostat can reach room temperature within one day. Do not break the OVC vacuum before the cryostat has reached room temperature.

Bibliography

[AAA+16] Abbott, B. P., *et al.* "Observation of Gravitational Waves from a Binary Black Hole Merger." *Phys. Rev. Lett.*, **116** (6), 61102 (2016). URL http://dx.doi.org/10.1103/PhysRevLett.116.061102. (Cited on page 1.)

[AAU+09] Anetsberger, G., *et al.* "Near-field cavity optomechanics with nanomechanical oscillators." *Nat Phys*, **5** (12), 909 (2009). URL http://dx.doi.org/10.1038/nphys1425http://www.nature.com/nphys/journal/v5/n12/suppinfo/nphys1425_S1.html. (Cited on page 4.)

[AB09] Adler, S. L. and Bassi, A. "Is Quantum Theory Exact?" *Science*, **325** (5938), 275 (2009). URL http://dx.doi.org/10.1126/science.1176858. (Cited on page 5.)

[ACB+06] Arcizet, O., *et al.* "Radiation-pressure cooling and optomechanical instability of a micromirror." *Nature*, **444** (7115), 71 (2006). URL http://dx.doi.org/10.1038/nature05244. (Cited on page 3.)

[AGR81] Aspect, A., Grangier, P., and Roger, G. "Experimental Tests of Realistic Local Theories via Bell's Theorem." *Phys. Rev. Lett.*, **47** (7), 460 (1981). URL http://dx.doi.org/10.1103/PhysRevLett.47.460. (Not cited.)

[AKM14] Aspelmeyer, M., Kippenberg, T. J., and Marquardt, F. "Cavity optomechanics." *Rev. Mod. Phys.*, **86** (4), 1391 (2014). URL http://dx.doi.org/10.1103/RevModPhys.86.1391. (Cited on page 1.)

[AM76] Ashcroft, N. W. and Mermin, N. D. *Solid state physics*. Saunders College (1976). (Cited on page 195.)

[Arc06] Arcizet, O. *Mesure optique ultrasensible et refroidissement par pression de radiation d'un micro-résonateur mécanique*. Ph.D. thesis, Université Pierre et Marie Curie, Paris VI (2006). (Cited on pages 60, 111 and 116.)

[Asp07] Aspect, A. "Description statistique de la lumière émise par un laser. Largeur de raie, cohérence temporelle." (2007). URL https://www.lcf.institutoptique.fr/content/download/3290/22193/file/OptiqueStatistiquecoursecrit.pdf. (Cited on page 17.)

[Atk59] Atkins, K. R. "Third and fourth sound in liquid helium II." *Phys. Rev.*, **113** (4), 962 (1959). URL http://dx.doi.org/10.1103/PhysRev.113.962. (Cited on page 98.)

[AZV16] Asjad, M., Zippilli, S., and Vitali, D. "Suppression of Stokes scattering in optomechanical cooling with squeezed light." (Mc), 1 (2016). URL http://arxiv.org/abs/1606.09007. (Cited on page 225.)

[Bae08] Baetoniu, C. "Method and apparatus for true random number generation." (2008). URL http://www.google.com/patents/US7389316. (Cited on page 154.)

[BDFF14] Ballmer, S., *et al.* "Comparing Finesse simulations, analytical solutions and OSCAR simulations of Fabry-Perot alignment signals." *arXiv Prepr.* (2014). URL https://arxiv.org/abs/1401.5727. (Cited on page 74.)

[Bec58] Bechmann, R. "Elastic and piezoelectric constants of alpha-quartz." *Phys. Rev.*, **110** (5), 1060 (1958). URL http://dx.doi.org/10.1103/PhysRev.110.1060. (Cited on pages 50 and 109.)

[Bec05] Bechhoefer, J. "Feedback for physicists: A tutorial essay on control." *Rev. Mod. Phys.*, **77** (3), 783 (2005). URL http://dx.doi.org/10.1103/RevModPhys.77.783. (Cited on page 147.)

[BHM+15] Benedikter, J., *et al.* "Transverse-mode coupling and diffraction loss in tunable Fabry–Pérot microcavities." *New J. Phys.*, **17** (5), 53051 (2015). URL http://stacks.iop.org/1367-2630/17/i=5/a=053051. (Cited on page 72.)

[Bla01] Black, E. D. "An introduction to Pound–Drever–Hall laser frequency stabilization." *Am. J. Phys.*, **69** (1) (2001). (Cited on page 27.)

[BMP85] Braginsky, V. B., Mitrofanov, V. P., and Panov, V. I. *Systems with Small Dissipation*. University of Chicago Press (1985). (Cited on page 9.)

[BPBU14] Bahrami, M., *et al.* "Proposal for a Noninterferometric Test of Collapse Models in Optomechanical Systems." *Phys. Rev. Lett.*, **112** (21), 210404 (2014). URL http://dx.doi.org/10.1103/PhysRevLett.112.210404. (Cited on page 5.)

[BPLH13] Batelaan, R. B., *et al.* "Controlled double-slit electron diffraction." *New J. Phys.*, **15** (3), 33018 (2013). URL http://stacks.iop.org/1367-2630/15/i=3/a=033018. (Not cited.)

[BR04] Bachor, H.-A. and Ralph, T. C. *A Guide to Experiments in Quantum Optics*. Wiley (2004). (Cited on pages 18 and 19.)

[Bri04] Briant, T. *Caractérisation du couplage optomécanique entre la lumière et un miroir : bruit thermique et effets quantiques*. Ph.D. thesis, Université Pierre et Marie Curie, Paris VI (2004). (Cited on pages 93 and 221.)

[BYC+10] Briles, T. C., *et al.* "Simple piezoelectric-actuated mirror with 180 kHz servo bandwidth." *Opt. Express*, **18** (10), 9739 (2010). URL `http://dx.doi.org/10.1364/OE.18.009739`. (Cited on page 172.)

[CASN+11] Chan, J., *et al.* "Laser cooling of a nanomechanical oscillator into its quantum ground state." *Nature*, **478** (7367), 89 (2011). URL `http://dx.doi.org/10.1038/nature10461`. (Cited on page 3.)

[Cav81] Caves, C. M. "Quantum-mechanical noise in an interferometer." *Phys. Rev. D*, **23** (8), 1693 (1981). URL `http://dx.doi.org/10.1103/PhysRevD.23.1693`. (Cited on page 2.)

[CCF+04] Crooks, D. R. M., *et al.* "Experimental measurements of coating mechanical loss factors." *Class. Quantum Gravity*, **21** (5), S1059 (2004). URL `http://dx.doi.org/10.1088/0264-9381/21/5/101`. (Cited on page 40.)

[CDG+10] Clerk, A. A., *et al.* "Introduction to quantum noise, measurement, and amplification." *Rev. Mod. Phys.*, **82** (2), 1155 (2010). URL `http://dx.doi.org/10.1103/RevModPhys.82.1155`. (Cited on page 17.)

[CHP99] Cohadon, P. F., Heidmann, A., and Pinard, M. "Cooling of a Mirror by Radiation Pressure." *Phys. Rev. Lett.*, **83** (16), 3174 (1999). URL `http://dx.doi.org/10.1103/PhysRevLett.83.3174`. (Cited on pages 3 and 224.)

[Coh00] Cohadon, P.-F. *Bruit thermique et effets de la pression de radiation dans une cavité optique de grande finesse.* Ph.D. thesis, Université Pierre et Marie Curie, Paris VI (2000). (Cited on page 221.)

[Coh14] Coherent. "Ultra-Narrow linewidth laser : Mephisto." Technical report, Coherent (2014). (Cited on page 115.)

[CW51] Callen, H. B. and Welton, T. A. "Irreversibility and Generalized Noise." *Phys. Rev.*, **83** (1), 34 (1951). (Cited on pages 10 and 37.)

[DBM92] Davidson, M., Bastian, S., and Markley, F. "Measurement of the Elastic Modulus of Kapton Perpendicular to the Plane of the Film at Room and Cryogenic Temperatures." In "Fourth Annu. IISSC Conf.", Fermilab, New Orleans, LA (1992). (Cited on page 101.)

[Deg08] Degallaix, J. "OSCAR, an optical FFT code to simulate Fabry Perot cavities with arbitrary mirror profiles." (2008). URL `http://www.mathworks.com/matlabcentral/fileexchange/20607`. (Cited on page 74.)

[DHK+83] Drever, R. W. P., *et al.* "Laser phase and frequency stabilization using an optical resonator." *Appl. Phys. B*, **31** (2), 97 (1983). URL `http://dx.doi.org/10.1007/BF00702605`. (Cited on page 27.)

[Dió87] Diósi, L. "A universal master equation for the gravitational violation of quantum mechanics." *Phys. Lett. A*, **120** (8), 377 (1987). URL http://dx.doi.org/http://dx.doi.org/10.1016/0375-9601(87)90681-5. (Cited on page 5.)

[DMM+83] Dorsel, A., *et al.* "Optical Bistability and Mirror Confinement Induced by Radiation Pressure." *Phys. Rev. Lett.*, **51** (17), 1550 (1983). URL http://dx.doi.org/10.1103/PhysRevLett.51.1550. (Cited on page 31.)

[EAD64] Everitt, C. W. F., Atkins, K. R., and Denenstein, A. "Third Sound in Liquid Helium Films." *Phys. Rev.*, **136** (1961), A1494 (1964). (Cited on page 98.)

[El 93] El Habti, A. "Physical limitation on the quality factor of quartz resonators." *J. Acoust. Soc. Am.*, **94** (2), 917 (1993). URL http://dx.doi.org/10.1121/1.408193. (Cited on page 40.)

[FC07] Friedt, J.-M. and Carry, E. "Introduction to the quartz tuning fork." *Am. J. Phys.*, **75** (5), 415 (2007). URL http://dx.doi.org/10.1119/1.2711826. (Cited on page 93.)

[GA81] Gardner, J. W. and Anderson, A. C. "Low-temperature specific heat and thermal conductivity of neutron-irradiated crystalline quartz." *Phys. Rev. B*, **23** (2), 474 (1981). URL http://dx.doi.org/10.1103/PhysRevB.23.474. (Cited on pages 195 and 196.)

[Gav13] Gavartin, E. *Optonanomechanical Systems for Measurement Applications.* Ph.D. thesis, Ecole Polytechnique Fédérale de Lausanne (2013). (Cited on pages 183 and 185.)

[GBP+06a] Gigan, S., *et al.* "Self-cooling of a micromirror by radiation pressure." *Nature*, **444** (7115), 67 (2006). URL http://dx.doi.org/10.1038/nature05273. (Cited on page 3.)

[GBP+06b] Gigan, S., *et al.* "Self-cooling of a micromirror by radiation pressure." *Nature*, **444** (7115), 67 (2006). URL http://dx.doi.org/10.1038/nature05273. (Cited on page 61.)

[GCI+12] Goryachev, M., *et al.* "Extremely low-loss acoustic phonons in a quartz bulk acoustic wave resonator at millikelvin temperature." *Appl. Phys. Lett.*, **100** (24) (2012). URL http://dx.doi.org/http://dx.doi.org/10.1063/1.4729292. (Cited on pages 40 and 59.)

[GHSN+13] Gröblacher, S., *et al.* "Highly efficient coupling from an optical fiber to a nanoscale silicon optomechanical cavity." *Appl. Phys. Lett.*, **103** (18) (2013). URL http://dx.doi.org/10.1063/1.4826924. (Cited on page 4.)

[GPS02] Goldstein, H., Poole, C. P., and Safko, J. L. *Classical Mechanics.* Addison Wesley (2002). URL https://books.google.fr/books?id=tJCuQgAACAAJ. (Cited on page 10.)

[GRW86] Ghirardi, G. C., Rimini, A., and Weber, T. "Unified dynamics for microscopic and macroscopic systems." *Phys. Rev. D*, **34** (2), 470 (1986). URL http://dx.doi.org/10.1103/PhysRevD.34.470. (Cited on page 5.)

[GSA+10] Gorodetksy, M. L., *et al.* "Determination of the vacuum optomechanical coupling rate using frequency noise calibration." *Opt. Express*, **18** (22), 23236 (2010). URL http://dx.doi.org/10.1364/OE.18.023236. (Cited on page 27.)

[GVK12] Gavartin, E., Verlot, P., and Kippenberg, T. J. "A hybrid on-chip optomechanical transducer for ultrasensitive force measurements." *Nat. Nanotechnol.*, **7** (8), 509 (2012). URL http://dx.doi.org/10.1038/nnano.2012.97. (Cited on pages 158, 185 and 210.)

[HBD12] Harry, G., Bodiya, T. P., and DeSalvo, R. *Optical Coatings and Thermal Noise in Precision Measurement*. Cambridge University Press, New York (2012). URL http://dx.doi.org/10.1017/CBO9780511762314. (Cited on page 211.)

[HDB+12] Hunger, D., *et al.* "Laser micro-fabrication of concave, low-roughness features in silica." *AIP Adv.*, **2** (1), 012119 (2012). URL http://dx.doi.org/http://dx.doi.org/10.1063/1.3679721. (Cited on page 62.)

[HHP97] Heidmann, a., Hadjar, Y., and Pinard, M. "Quantum nondemolition measurement by optomechanical coupling." *Appl. Phys. B Lasers Opt.*, **64** (2), 173 (1997). URL http://dx.doi.org/10.1007/s003400050162. (Cited on page 2.)

[HR94] Heidmann, A. and Reynaud, S. "Photon noise reduction by reflection from a movable mirror." *Phys. Rev. A*, **50** (5), 4237 (1994). URL http://dx.doi.org/10.1103/PhysRevA.50.4237. (Cited on page 2.)

[HR06] Haroche, S. and Raimond, J.-M. *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford University Press, 1st edition (2006). (Not cited.)

[HVVT14] Huisman, F. M., *et al.* "Quartz Tuning Forks as Cryogenic Vacuum Gauges." *J. Low Temp. Phys.*, **177** (5-6), 226 (2014). URL http://dx.doi.org/10.1007/s10909-014-1206-5. (Cited on page 93.)

[JR90] Jaekel, M. T. and Reynaud, S. "Quantum Limits in Interferometric Measurements." *EPL (Europhysics Lett.*, **13** (4), 301 (1990). (Cited on pages 2 and 225.)

[JRK+11] Jöckel, A., *et al.* "Spectroscopy of mechanical dissipation in micromechanical membranes." *Appl. Phys. Lett.*, **99** (14), 14 (2011). URL http://dx.doi.org/10.1063/1.3646914. (Cited on page 56.)

[KAB+16] Kaltenbaek, R., *et al.* "Macroscopic Quantum Resonators (MAQRO): 2015 update." *EPJ Quantum Technol.*, **3** (1), 5 (2016). URL http://dx.doi.org/10.1140/epjqt/s40507-016-0043-7. (Cited on page 5.)

[KBD+11] Kuhn, A. G., *et al.* "A micropillar for cavity optomechanics." *Appl. Phys. Lett.*, **99** (12), 121103 (2011). URL http://dx.doi.org/http://dx.doi.org/10.1063/1.3641871. (Cited on page 41.)

[Kep19] Kepler, J. *De Cometis Libelli Tres.* (1619). (Cited on page 1.)

[KL66] Kogelnik, H. and Li, T. "Laser Beams and Resonators." *Proc. IEEE*, **54** (10), 1312 (1966). URL http://dx.doi.org/10.1109/PROC.1966.5119. (Cited on pages 14 and 15.)

[KPJ+11] Kleckner, D., *et al.* "Optomechanical trampoline resonators." *Opt. Express*, **19** (20), 19708 (2011). (Cited on page 3.)

[Kuh13] Kuhn, A. G. *Optomécanique en cavité cryogénique avec un micropilier pour l'observation du régime quantique d'un résonateur mécanique macroscopique.* Ph.D. thesis, Université Pierre et Marie Curie, Paris VI (2013). (Cited on pages 39, 41, 42, 43, 45, 59, 81, 99, 186, 196 and 201.)

[Leg02] Leggett, A. J. "Testing the limits of quantum mechanics: motivation, state of play, prospects." *J. Phys. Condens. Matter*, **14** (15), R415 (2002). URL http://stacks.iop.org/0953-8984/14/i=15/a=201. (Cited on page 5.)

[Lev98] Levin, Y. "Internal thermal noise in the LIGO test masses: A direct approach." *Phys. Rev. D*, **57** (2), 659 (1998). (Cited on pages 11 and 211.)

[LIG09] LIGO Scientific Collaboration. "Observation of a kilogram-scale oscillator near its quantum ground state." *New J. Phys.*, **11**, 073032 (2009). URL http://dx.doi.org/10.1088/1367-2630/11/7/073032. (Cited on page 3.)

[Lin16] Linear Technology. "LTC2145-14/LTC2144-14/LTC2143-14 data sheet." Technical report, Linear Technology (2016). (Cited on page 151.)

[LL58] Landau, L. D. and Lifshitz, E. M. *Course of Theoretical Physics: Theory of Elasticity.* Pergamon (1958). (Cited on pages 12 and 109.)

[LL80] Landau, L. D. and Lifshitz, E. M. *Statistical Physics, Part 1.* Pergamon, 3rd edition (1980). (Cited on pages 10 and 94.)

[Lou74] Lounasmaa, O. V. *Experimental Principles and Methods below 1K.* Academic, London (1974). (Cited on pages 80 and 81.)

[LS14] Lorenzo, L. A. D. and Schwab, K. C. "Superfluid optomechanics: coupling of a superfluid to a superconducting condensate." *New J. Phys.*, **16** (11), 113020 (2014). URL http://stacks.iop.org/1367-2630/16/i=11/a=113020. (Cited on page 90.)

[LWH+99] Lawrence, M. J., *et al.* "Dynamic response of a Fabry-Perot interferometer." *J. Opt. Soc. Am. B*, **16** (4), 523 (1999). URL http://dx.doi.org/10.1364/JOSAB.16.000523. (Cited on page 24.)

[MAK+15] Makles, K., *et al.* "2D photonic-crystal optomechanical nanoresonator." *Opt. Lett.*, **40** (2), 174 (2015). URL http://dx.doi.org/10.1364/OL.40.000174. (Cited on page 225.)

[MCF+10] Marino, F., *et al.* "Classical Signature of Ponderomotive Squeezing in a Suspended Mirror Resonator." *Phys. Rev. Lett.*, **104** (7), 73601 (2010). URL http://dx.doi.org/10.1103/PhysRevLett.104.073601. (Cited on page 225.)

[McI66] McIntyre, R. J. "Multiplication noise in uniform avalanche diodes." *IEEE Trans. Electron Devices*, **ED-13** (1), 164 (1966). URL http://dx.doi.org/10.1109/T-ED.1966.15651. (Cited on page 194.)

[MEK+07] Mohammadi, S., *et al.* "Complete phononic bandgaps and bandgap maps in two-dimensional silicon phononic crystal plates." *Electron. Lett.*, **43** (16), 898 (2007). URL http://digital-library.theiet.org/content/journals/10.1049/el_20071159. (Cited on page 214.)

[MERS+08] Müller-Ebhardt, H., *et al.* "Entanglement of Macroscopic Test Masses and the Standard Quantum Limit in Laser Interferometry." *Phys. Rev. Lett.*, **100** (1), 13601 (2008). URL http://dx.doi.org/10.1103/PhysRevLett.100.013601. (Cited on page 5.)

[MF53] Morse, P. M. and Feshbach, H. *Methods of Theoretical Physics, Part I.* McGraw-Hill (1953). (Cited on page 23.)

[MK04] Metzger, C. H. and Karrai, K. "Cavity cooling of a microlever." *Nature*, **432** (7020), 1002 (2004). URL http://dx.doi.org/10.1038/nature03118. (Cited on page 3.)

[MMGSK08] Murch, K. W., *et al.* "Observation of quantum-measurement backaction with an ultracold atomic gas." *Nat. Phys.*, **4** (7), 561 (2008). URL http://dx.doi.org/10.1038/nphys965. (Cited on page 3.)

[MMM08] Mori, T., Moriwaki, S., and Mio, N. "Mechanical Q-factor measurement of a quartz oscillator at cryogenic temperature." *Appl. Phys. Express*, **1** (7), 0770021 (2008). URL http://dx.doi.org/10.1143/APEX.1.077002. (Cited on page 93.)

[MND+17] Metzdorff, R., *et al.* "Scheme to avoid excess loss due to transverse mode coupling in Fabry Perot cavities." (2017). (Cited on page 74.)

[NBH15] Nowak, K. M., Baker, H. J., and Hall, D. R. "Analytical model for CO2 laser ablation of fused quartz." *Appl. Opt.*, **54** (29), 8653 (2015). URL http://dx.doi.org/10.1364/AO.54.008653. (Cited on page 62.)

[Neu16a] Neuhaus, L. "Adding voltage regulators for the RedPitaya output stage." (2016). URL http://dx.doi.org/10.6084/m9.figshare.4055691.v1. (Cited on page 151.)

[Neu16b] Neuhaus, L. "PyRPL." (2016). URL https://github.com/lneuhaus/pyrpl. (Cited on pages 150, 154 and 224.)

[Neu16c] Neuhaus, L. "Red Pitaya DAC performance." (2016). URL http://dx.doi.org/10.6084/m9.figshare.4055658.v1. (Cited on page 151.)

[OS75] Oppenheim, A. V. and Schafer, R. W. *Digital Signal Processing*. Prentice-Hall (1975). (Cited on page 159.)

[PBVF97] Poirson, J., *et al.* "Analytical and experimental study of ringing effects in a Fabry–Perot cavity. Application to the measurement of high finesses." *J. Opt. Soc. Am. B*, **14** (11), 2811 (1997). URL http://dx.doi.org/10.1364/JOSAB.14.002811. (Cited on page 24.)

[PDV+05] Pinard, M., *et al.* "Entangling movable mirrors in a double-cavity system." *EPL (Europhysics Lett.*, **72** (5), 747 (2005). URL http://stacks.iop.org/0295-5075/72/i=5/a=747. (Cited on page 2.)

[Pen96] Penrose, R. "On Gravity's role in Quantum State Reduction." *Gen. Relativ. Gravit.*, **28** (5), 581 (1996). URL http://dx.doi.org/10.1007/BF02105068. (Cited on page 5.)

[PGST16] Purdy, T. P., *et al.* "Observation of Optomechanical Quantum Correlations at Room Temperature." *arXiv* (2016). URL http://arxiv.org/abs/1605.05664. (Cited on page 208.)

[PHH99] Pinard, M., Hadjar, Y., and Heidmann, A. "Effective mass in quantum effects of radiation pressure." *Eur. Phys. J. D - At. Mol. Opt. Plasma Phys.*, **7** (1), 107 (1999). URL http://dx.doi.org/10.1007/s100530050354. (Cited on pages 8 and 11.)

[Pob07] Pobell, F. *Matters and Methods at Low Temperatures*. Springer, Berlin, Heidelberg, New York (2007). URL http://dx.doi.org/10.1007/978-3-540-46360-3. (Cited on pages 80, 81 and 96.)

[PPK+16] Peterson, R. W., *et al.* "Laser Cooling of a Micromechanical Membrane to the Quantum Backaction Limit." *Phys. Rev. Lett.*, **116** (6), 1 (2016). URL http://dx.doi.org/10.1103/PhysRevLett.116.063601. (Cited on pages 3 and 194.)

[PPR13] Purdy, T. P., Peterson, R. W., and Regal, C. A. "Observation of Radiation Pressure Shot Noise on a Macroscopic Object." *Science*, **339** (6121), 801 (2013). URL http://dx.doi.org/10.1126/science.1231282. (Cited on pages 3 and 210.)

[PRB69] Payne, W. H., Rabung, J. R., and Bogyo, T. P. "Coding the Lehmer pseudo-random number generator." *Commun. ACM*, **12** (2), 85 (1969). URL http://dx.doi.org/10.1145/362848.362860. (Cited on page 154.)

[Pre11] Preumont, A. *Vibration Control of Active Structures , An Introduction.* Springer (2011). URL http://dx.doi.org/10.1023/A:1004398914135. (Cited on pages 8 and 11.)

[PYK+15] Purdy, T. P., *et al.* "Optomechanical Raman-ratio thermometry." *Phys. Rev. A*, **92** (3), 31802 (2015). URL http://dx.doi.org/10.1103/PhysRevA.92.031802. (Cited on pages 3 and 194.)

[PZD10] Phelps, M., Zhang, L., and Dannenburg, R. "Drag Wiping with Methanol vs First Contact." Technical report, LIGO (2010). (Cited on page 65.)

[RDW+11] Rivière, R., *et al.* "Optomechanical sideband cooling of a micromechanical oscillator close to the quantum ground state." *Phys. Rev. A*, **83** (6), 63835 (2011). URL http://dx.doi.org/10.1103/PhysRevA.83.063835. (Cited on pages 3 and 194.)

[red16] "Red Pitaya." (2016). URL http://redpitaya.com/. (Cited on pages 150 and 155.)

[RH89] Reynaud, S. and Heidmann, A. "A semiclassical linear input output transformation for quantum fluctuations." *Opt. Commun.*, **71** (3), 209 (1989). URL http://dx.doi.org/http://dx.doi.org/10.1016/0030-4018(89)90429-X. (Cited on page 19.)

[RHGF92] Reynaud, S., *et al.* "Quantum Fluctuations in Optical Systems." In "Prog. Opt.", volume 30, pp. 1–85. Elsevier (1992). URL http://dx.doi.org/http://dx.doi.org/10.1016/S0079-6638(08)70096-9. (Cited on page 19.)

[RS39] Rollin, B. V. and Simon, F. "On the "film" phenomenon of liquid helium II." *Physica*, **6** (2), 219 (1939). URL http://dx.doi.org/http://dx.doi.org/10.1016/S0031-8914(39)80013-1. (Cited on page 88.)

[RS98] Richardson, R. J. and Smith, E. N. *Experimental Techniques in Condensed Matter Physics at Low Temperatures.* Addison-Wesley (1998). (Cited on pages 88, 90 and 91.)

[RS16] Ryou, A. and Simon, J. "Active Cancellation of Acoustical Resonances with an FPGA FIR Filter." *arXiv Prepr.* (2016). URL http://arxiv.org/abs/1604.04668. (Cited on pages 159 and 160.)

[Sau90] Saulson, P. R. "Thermal noise in mechanical experiments." *Phys. Rev. D*, **42** (8), 2437 (1990). (Cited on pages 9 and 10.)

[Sch26] Schrödinger, E. "Der stetige Übergang von der Mikro- zur Makromechanik." *Naturwissenschaften*, **14** (28), 664 (1926). URL http://dx.doi.org/10.1007/BF01507634. (Cited on page 4.)

[Sch08] Schroeter, A. *Mechanical losses in materials for future cryogenic gravitational wave detectors.* Ph.D. thesis, Friedrich-Schiller-Universität Jena (2008). (Cited on pages 59 and 222.)

[Sch09] Schliesser, A. *Cavity optomechanics and optical frequency comb generation with silica whispering-gallery-mode microresonators*. Ph.D. thesis, Ludwig-Maximilians-Universität München (2009). (Cited on page 231.)

[SDN+06] Schliesser, A., *et al.* "Radiation Pressure Cooling of a Micromechanical Oscillator Using Dynamical Backaction." *Phys. Rev. Lett.*, **97** (24), 243905 (2006). URL http://dx.doi.org/10.1103/PhysRevLett.97.243905. (Cited on page 3.)

[Sei15] Seidler, F. "Digital high bandwidth feedback controller." (2015). (Cited on page 151.)

[Sem12] Semiconductors, N. "DAC1401D125 Product data sheet." Technical report, NXP Semiconductors (2012). (Cited on page 151.)

[SHY+85] Slusher, R., *et al.* "Observation of squeezed states generated by four-wave mixing in an optical cavity." *Phys. Rev. Lett.*, **55** (22), 2409 (1985). (Cited on page 2.)

[Som12] Somiya, K. "Detector configuration of KAGRA–the Japanese cryogenic gravitational-wave detector." *Class. Quantum Gravity*, **29** (12), 124007 (2012). URL http://stacks.iop.org/0264-9381/29/i=12/a=124007. (Cited on page 4.)

[SS98] Sprague, D. T. and Smith, E. N. "Resealable Superfluid Valve." *J. Low Temp. Phys.*, **113** (5), 975 (1998). URL http://dx.doi.org/10.1023/A:1022539914232. (Cited on page 90.)

[SSF+16] Sudhir, V., *et al.* "Quantum correlations of light due to a room temperature mechanical oscillator." *arXiv* (2016). URL http://arxiv.org/abs/1608.00699. (Cited on page 208.)

[Tex92] Texas Instruments. "OA-31 Current Feedback Amplifiers." Technical report, Texas Instruments, Dallas, Texas (1992). (Cited on page 137.)

[Tex02] Texas Instruments. "LMH6702 1.7-GHz Ultra-Low Distortion Wideband Op Amp." (2002). URL http://www.ti.com/product/LMH6702/datasheet. (Cited on page 137.)

[TF88] Thomas, G. B. J. and Finney, R. L. *Calculus and Analytic Geometry*. Addison-Wesley, 7th edition (1988). (Cited on page 161.)

[UBRR15] Uphoff, M., *et al.* "Frequency splitting of polarization eigenmodes in microscopic Fabry–Perot cavities." *New J. Phys.*, **17** (1), 13053 (2015). URL http://stacks.iop.org/1367-2630/17/i=1/a=013053. (Cited on page 76.)

[UML+15] Underwood, M., *et al.* "Measurement of the motional sidebands of a nanogram-scale oscillator in the quantum regime." *Phys. Rev. A - At. Mol. Opt. Phys.*, **92** (6), 1 (2015). URL http://dx.doi.org/10.1103/PhysRevA.92.061801. (Cited on pages 3, 194 and 220.)

[VDW+12] Verhagen, E., *et al.* "Quantum-coherent coupling of a mechanical oscillator to an optical cavity mode." *Nature*, **482** (7383), 63 (2012). URL `http://dx.doi.org/10.1038/nature10787`. (Cited on page 3.)

[Ver10] Verlot, P. *Etude des effets de pression de radiation et des limites quantiques du couplage optomécanique.* Ph.D. thesis, Université Pierre et Marie Curie, Paris VI (2010). (Cited on page 185.)

[Vir] Virgo Collaboration. "Virgo Physics Book." URL `http://wwwcascina.virgo.infn.it/vpb/`. (Cited on page 60.)

[VTB+09] Verlot, P., *et al.* "Scheme to Probe Optomechanical Correlations between Two Optical Beams Down to the Quantum Level." *Phys. Rev. Lett.*, **102** (10), 103601 (2009). URL `http://dx.doi.org/10.1103/PhysRevLett.102.103601`. (Cited on page 210.)

[War60] Warner, A. W. "Design and Performance of Ultraprecise 2,5 mc Quartz Crystal Units." *Bell Syst. Tech. J.*, **39**, 1193 (1960). (Cited on page 56.)

[War92] Ward, R. W. "THE CONSTANTS OF ALPHA QUARTZ." *Piezoelectricity*, **5**, 211 (1992). (Cited on page 109.)

[Wik16a] Wikipedia. "Gaussian Beam." (2016). URL `https://en.wikipedia.org/wiki/Gaussian_beam`. (Cited on page 14.)

[Wik16b] Wikipedia. "Integrating factor." (2016). URL `https://en.wikipedia.org/wiki/Integrating_factor`. (Cited on page 23.)

[Wik16c] Wikipedia. "Partial fraction decomposition." (2016). URL `https://en.wikipedia.org/wiki/Partial_fraction_decomposition`. (Cited on page 161.)

[Wil67] Wilks, J. *The Properties of Liquid and Solid Helium.* Clarendon Press (1967). (Cited on pages 88 and 98.)

[Wil12] Wilson, D. J. *Cavity optomechanics with high-stress silicon nitride films.* Ph.D. thesis, California Institute of Technology (2012). (Cited on pages 100 and 214.)

[Win16] Winchester, M. "Frequency Dependent Squeezed Light in Optomechanical Systems." Technical report, IREU (2016). (Cited on page 163.)

[WRBV+11] Wilson-Rae, I., *et al.* "High-Q nanomechanics via destructive interference of elastic waves." *Phys. Rev. Lett.*, **106** (4), 1 (2011). URL `http://dx.doi.org/10.1103/PhysRevLett.106.047205`. (Cited on page 56.)

[WRD+10] Weis, S., *et al.* "Optomechanically Induced Transparency." *Science*, **330** (6010), 1520 (2010). URL `http://dx.doi.org/10.1126/science.1195596`. (Cited on page 188.)

[WSP+15] Wilson, D. J., *et al.* "Measurement-based control of a mechanical oscillator at its thermal decoherence rate." *Nature*, **524** (7565), 325 (2015). URL http://dx.doi.org/10.1038/nature14672. (Cited on page 224.)

[Xil16] Xilinx. "Zynq-7000 All Programmable SoC Overview." Technical report, Xilinx (2016). URL https://www.xilinx.com/support/documentation/data_sheets/ds190-Zynq-7000-Overview.pdf. (Cited on page 150.)

[Zer17] Zerkani, S. *Fabrication of high Q mechanical resonators with a high reflectivity mirrors for cavity optomechanics.* Ph.D. thesis, Université Pierre et Marie Curie, Paris VI (2017). (Cited on pages 41, 45, 49, 56, 196 and 221.)

[ZNT+17] Zerkani, S., *et al.* "Fabrication of high Q mechanical resonators with a high reflectivity mirrors for cavity optomechanics." (2017). (Cited on page 41.)

[ZP71] Zeller, R. C. and Pohl, R. O. "Thermal conductivity and specific heat of noncrystalline solids." *Phys. Rev. B*, **4** (6), 2029 (1971). URL http://dx.doi.org/10.1103/PhysRevB.4.2029. (Cited on page 195.)

[ZPG+95] Zhang, T. C., *et al.* "Quantum noise of free-running and externally-stabilized laser diodes." *Quantum Semiclassical Opt.*, **7** (4), 601 (1995). URL http://dx.doi.org/10.1088/1355-5111/7/4/015. (Cited on page 231.)

[ZSJ+08] Zwickl, B. M., *et al.* "High quality mechanical and optical properties of commercial silicon nitride membranes." *Appl. Phys. Lett.*, **92** (10), 130 (2008). URL http://dx.doi.org/10.1063/1.2884191. (Cited on pages 3 and 4.)

[ZTZ13] Zienkiewicz, O. C., Taylor, R. L., and Zhu, J. Z. *The Finite Element Method: Its Basis and Fundamentals.* Butterworth-Heinemann, 7th edition (2013). URL http://dx.doi.org/http://dx.doi.org/10.1016/B978-1-85617-633-0.00019-8. (Cited on page 13.)

All references to websites refer to the website state on the 27th of October, 2016.

Leonhard NEUHAUS **December 9, 2016**

Cooling a macroscopic mechanical oscillator close to its quantum ground state

Abstract: In this work, we attempt the experimental demonstration of quantum effects in the motion of a macroscopic mechanical resonator with a mass of 33 μg, about 3 orders of magnitude above the mass of the heaviest system demonstrated so far in the quantum ground state. We have designed, fabricated, and operated an optomechanical resonator at 3.6 MHz, with an optical finesse of 10^5 and a mechanical quality factor near 10^8, embedded in the 100 mK environment of a dilution refrigerator. We present a fully automatized optical measurement setup, including a filter cavity, a homodyne detector, and various feedback controllers implemented in an FPGA with the custom-developed software PyRPL. We have laser-cooled the compression mode of our mechanical resonator to a mean thermal occupation number of 20 phonons. Cooling is limited by the onset of an optomechanical instability of suspension modes with frequencies below 100 kHz. A custom-tailored digital filter to suppress this instability has enabled us to reach a regime where quantum backaction amounts to about 30 % of the total force noise on the mechanical resonator. For even higher ratios in the future, we present the design of a phononic-crystal input mirror with a reduced Brownian motion displacement noise floor.

Keywords: optomechanics, quartz, MEMS, Fabry-Perot, cryogenics, feedback, FPGA, quantum ground state, laser cooling, quantum backaction.

Résumé : Ce travail s'attaque à la mise en évidence expérimentale d'effets quantiques dans le mouvement d'un résonateur mécanique macroscopique avec une masse effective de 33 μg, soit 3 ordres de grandeur au-dessus de celle du système mécanique le plus massif observé à ce jour dans son état quantique fondamental. Nous avons conçu, fabriqué et fait fonctionner un résonateur optomécanique à 3,6 MHz avec une finesse optique de 10^5 et un facteur de qualité mécanique proche de 10^8, inséré dans l'environnement à 100 mK d'un réfrigérateur à dilution. Nous présentons un montage optique complètement automatisé incluant une cavité de filtrage, une détection homodyne et plusieurs asservissements, implémentés dans un FPGA avec le programme PyRPL développé spécifiquement pour cette expérience. Nous avons refroidi par laser le mode de compression de notre résonateur mécanique jusqu'à un nombre moyen d'occupation thermique de 20 phonons. Le refroidissement est limité par l'apparition d'une instabilité optomécanique de plusieurs modes des suspensions, au-dessous de 100 kHz. Un filtre digital particulier pour supprimer cette instabilité nous a permis d'atteindre le régime où l'action en retour quantique contribue à hauteur d'environ 30 % au bruit de force total de l'oscillateur mécanique. Pour atteindre des contributions encore plus importantes à l'avenir, nous présentons la conception d'un miroir d'entrée à cristal phononique, caractérisé par un plancher de bruit de mouvement Brownien réduit.

Mots clés : optomécanique, quartz, MEMS, Fabry-Perot, cryogénie, asservissements, FPGA, état quantique fondamental, refroidissement laser, action en retour quantique.