




Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich



Pulsed optomechanics using high-overtone bulk acoustic wave resonators

Master Thesis

Jacob Werner
jawerner@ethz.ch

Hybrid Quantum Systems Group
Departement of Physics, D-PHYS
ETH Zürich

Supervisors:
Tom Schatteburg
Prof. Dr. Yiwen Chu

July 2, 2023

Abstract

In this work, we report on progress towards realizing microwave-to-optical transduction between a superconducting qubit and an optical mode of a fiber, by using a phonon mode of a high-overtone bulk acoustic wave resonator (HBAR) as an intermediary between the qubit and the optics.

We adapt the experimental setup from [1] for pulsed laser operation and work towards optimizing a Pound-Drever-Hall laser locking scheme for locking a laser in pulsed operation with $15 \mu\text{s}$ pulses at a repetition rate of around 1 kHz to an optical mode of a Fabry-Pérot cavity placed in a dilution refrigerator. While locking in pulsed operation, we observe average reflected powers on the order of 20 % of the reflected power observed when the laser is completely off resonance. The feedback parameters can be studied and optimized more in the future, but the implementation of lock performance characterization of every individual pulse already allows for post-selection of acceptable lock performance.

As an application of this locking procedure, we perform an optomechanical storage and retrieval experiment. We excite a mechanical mode of an HBAR placed in a Fabry-Pérot cavity by means of Brillouin scattering and retrieve this excitation by the inverse process. This retrieval process is identical to the optomechanical readout required for microwave-to-optical transduction using an HBAR. We probe the response of two separate mechanical modes by varying the time between storage and retrieval as well as the frequency at which we drive the mechanical modes. We measure a retrieved phonon energy corresponding to 30 % of the value predicted by a coupled classical oscillator model. We further find that the T_1 lifetime depends on the frequency at which the mode was driven. Most notably, we find a maximum T_1 time of $4.9 \pm 0.1 \mu\text{s}$ at a drive frequency detuning of 60 kHz, which is ~ 70 % larger than the lifetime expected from the intrinsic linewidth of the mechanical mode. We hypothesize that this is due to diffraction losses in the HBAR.

Contents

Abstract	iii
1 Introduction	1
2 Pulsed laser locking in a dilution refrigerator	3
2.1 Introduction	3
2.2 Pound-Drever-Hall lock	4
2.3 Experimental setup	6
2.3.1 Overview	6
2.3.2 Selected Equipment and Software	7
2.3.3 Assessing lock quality	7
2.4 Initial tests with slow PID 1	9
2.4.1 Tuning PID parameters for the room temperature cavity	9
2.4.2 Testing lock with cryogenic cavity	10
2.5 Red Pitaya as PID 1	12
2.5.1 Characterizing components	12
2.5.2 Tuning PID parameters using both cavities	13
2.5.3 Long-time lock behavior	16
2.5.4 Discussion and potential improvements	19
3 Optomechanical storage and readout	21
3.1 The optomechanical system	21
3.2 Theory	23
3.2.1 Write pulse	24
3.2.2 Readout pulse	25
3.3 Experimental methods and setup	27
3.3.1 Overview and experimental setup	27
3.3.2 Calculating retrieved phonon energy from the balanced detector signal	29
3.4 Results	32
3.4.1 OMIT spectrum	32
3.4.2 Probing the main resonance	32
3.4.3 Comparison with numerical estimates	36
3.4.4 Probing two mechanical resonances	38
3.5 Discussion	40
4 Conclusion	43
Acknowledgments	45
Bibliography	47

Introduction

Since the onset of what is now often referred to as the "second quantum revolution" during the second half of the 20th century, there has been great interest in developing devices that exploit quantum mechanical entanglement for information processing applications [2, 3]. The ability to transfer quantum states from one such device to another over macroscopic distances is an essential ingredient of many relevant quantum information processing protocols, such as quantum key distribution [4] or loophole-free Bell inequality tests [5]. A natural choice is to use infrared photons as information carriers for this, since they are robust to thermal decoherence even at room temperature due to their high frequencies and can be sent over long distances through optical fibers with moderate losses. Another reason for using infrared photons is the fact that a lot of classical communication is already built around infrared light at 1550 nm wavelength, meaning that there are many commercially available components tailored to these wavelengths, that can also be used for quantum applications. In this work, we are specifically concerned with transferring quantum states between superconducting circuits using infrared photons.

Superconducting circuits engineered to behave as anharmonic microwave frequency oscillators are one of many promising platforms for future quantum information processing applications [6]. Once cooled to sub-Kelvin temperatures in a dilution refrigerator, these circuits are operated as quantum two-level systems that can be used as qubits, usually referred to as superconducting qubits.

To achieve the desired long-distance communication between such superconducting devices, we need to be able to transfer the quantum state of the qubit to infrared photons and vice versa. This process is known as microwave-to-optical transduction [7].

One of the main challenges in microwave-to-optical transduction is bridging the large frequency gap between the microwave and optical domains. While there are many other approaches to this problem [8], the one pursued by this group is to use a phonon mode of a high-overtone bulk acoustic wave resonator (HBAR) as an intermediary system. In this case, the qubit state is first converted into a phonon state of the HBAR by piezoelectric coupling [9] and then this phonon state is read out as an infrared photon by means of Brillouin scattering [10].

In this work we focus on aspects of the optomechanical readout, i.e. the conversion of high frequency phonons in an HBAR to infrared photons. The report is divided into two major parts.

In the first part, consisting of chapter 2, we report on work done towards frequency locking a laser in pulsed operation to an optical resonance of a Fabry-Pérot cavity located in the noisy environment of a dilution refrigerator. We do this by means of a Pound-Drever-Hall lock.

The second part consists of chapter 3, where we perform an optomechanical storage and readout experiment: First, we expose our HBAR to laser light, containing a sideband detuned from the carrier frequency by the phonon frequency. This creates a mechanical excitation in the HBAR via Brillouin scattering. We then readout this excitation by the inverse process: Laser light inelastically scatters with the existing phonons, creating a higher frequency optical sideband, that we subsequently measure. The optomechanical readout performed in this experiment is in principle identical to the readout that needs

1 Introduction

to be performed in a microwave-to-optical-transduction experiment using this system, although we only readout coherent states and do so in the low optomechanical cooperativity regime.

Pulsed laser locking in a dilution refrigerator

2.1 Introduction

The goal of the first part of this project was to frequency lock a laser in pulsed operation to a resonance of an optical Fabry-Pérot cavity located in a dilution refrigerator. For this we adapted the fiber optical setup from [1] to allow for pulsing of the main laser and made some adaptations to the locking scheme for pulsed laser operation.

This pulsed laser locking setup will then be used in a future microwave-to-optical transduction experiment, which we introduced in chapter 1 and is being pursued by the hybrid quantum systems research group [11]. In order to achieve the high electric field strengths required for efficient readout of the intermediary phonon state, the high-overtone bulk acoustic wave resonator (HBAR) is placed in an optical Fabry-Pérot cavity (see [10, 12] and chapter 3). This optomechanical readout requires locking a laser to the aforementioned Fabry-Pérot cavity.

Shining an infrared laser near a superconducting qubit destroys the superconducting state, as the infrared photons carry enough energy to break up the Cooper pairs of the superconductor. This means that the laser needed for the optomechanical readout cannot be on while the qubit is active and further the qubit needs some time to recover the superconducting state after each run of the experiment. Effectively this means that the readout laser needs to be pulsed. Continuously leaving the laser on would also lead to excess heating in the refrigerator.

The exact needed parameters for the microwave-to-optical transduction experiment are still to be determined. Our current estimates are that the laser pulses need to be 15 μ s long for the desired readout efficiency and the recovery time of the superconductor limits the repetition rate of the experiment to around 1 kHz for our experiment.

Pulsing of the laser was achieved by introducing an acousto-optic modulator (AOM) into the setup, which is driven by an arbitrary waveform generator (AWG) that sends radio frequency (RF) pulses to the AOM. When the AOM receives an RF signal, it transmits the input optical signal, otherwise it blocks the incoming light. Thus, the AOM acts as an optical shutter controlled by the AWG. The technical specifications of both the used AWG and AOM easily allow us to achieve the desired optical pulse lengths and repetition rates. The laser needs to be resonant with a mode of the Fabry-Pérot cavity located in a Bluefors dilution refrigerator (which we simply call the fridge). In order to keep the laser resonant with the cavity mode, we utilize a Pound-Drever-Hall locking scheme (PDH lock). The PDH lock, which is described in the next section, produces an error signal (the PDH error signal) quantifying how far off resonance the laser is at a given point in time. The PDH error signal can then be sent to a PID controller, which produces an appropriate control signal, that is sent to the laser controller, which controls the laser frequency and thus holds it on resonance.

While the PDH lock is a very standard technique, there are some challenges due to our setup and the pulsing of our laser. Mechanical vibrations in the fridge constantly change the distance between the mirrors of our cavity, causing its resonance frequencies to fluctuate

accordingly. The goal of the PDH lock is to react to these shifts in the optical resonances in real time in order to keep the laser resonant with the cavity throughout these fluctuations. The mechanical vibrations caused by the fridge far exceed the vibrations experienced by a cavity mounted on a conventional optical table. Indeed, with the current vibration isolation appreciable laser locking is only possible if the pulse tube cooler of the fridge is turned off.

A further challenge is posed by the pulsed operation of the laser. Between pulses, when the laser is off, the PDH error signal contains no information on how the resonance of interest is shifting. Further, due to noise in the photodiodes and other electronics involved in the PDH lock, the error signal is non zero and fluctuating between pulses. Hence there should be no feedback applied in between pulses. To achieve this we have to be able to turn the active feedback on and off on a microsecond timescale, which is achieved by sending a 'marker' signal from the AWG to the PID controllers, indicating when the feedback should be active or not. Additionally, our controller needs to have a high enough bandwidth once it is on, to make appreciable corrections to the laser frequency within a couple of microseconds.

In section 2.2 we give a brief introduction to Pound-Drever-Hall laser locking. In section 2.3 we describe our experimental setup and discuss how to quantify the performance of the locking scheme. Our initial tests using two PID controllers directly integrated in our laser controller are presented in section 2.4. After these tests we decided to replace one of the PID controllers with a faster PID controller on a Red Pitaya. The implementation of the faster controller as well as subsequent pulsed laser locking results are shown and discussed in section 2.5.

2.2 Pound-Drever-Hall lock

The Pound-Drever-Hall lock (PDH lock) is a technique for laser frequency stabilization and holding a laser on resonance with an optical cavity mode. This process is commonly referred to as laser locking. In this section we cover the basic operating principles of a PDH lock. A more comprehensive review of the technique can be found in reference [13]. When light enters a Fabry-Pérot cavity with a frequency that is close to one of the cavity's resonances, the amount of reflected optical power is a measure for how close the laser frequency is to the one of the resonances of the cavity. In the case of a critically coupled cavity, where the input coupling rate is equal to the internal loss rate (see [14]), if the laser is perfectly resonant with the cavity, no reflected power is detected. If the laser frequency is far from a resonance of the cavity, all of the optical power is reflected. The optical power reflected by a Fabry-Pérot cavity, when the optical frequency is swept over a resonance of the cavity, can be seen in figure 2.1a.

One could imagine a simple scheme, where the reflected power is entered into a control loop, which tries to minimize the reflected power by adjusting the optical frequency. However, this scheme has one very large draw back: The reflected power is symmetric about the resonance, meaning if some non zero reflected power is measured, one cannot infer if one is above or below the resonance. Hence the control loop cannot reliably adjust the optical frequency in the correct direction.

The PDH lock circumvents this issue. While the reflected power is symmetric about the resonance, its derivative is not. The derivative of reflected power with respect to optical frequency is positive when the frequency is above the resonance and negative when below the resonance.

Measuring the derivative of the reflected power can be achieved by constantly varying the frequency of the input laser beam a little bit in a controlled, sinusoidal fashion and monitoring if the reflected power oscillates in or out of phase with the laser frequency. We

do this by running the optical input beam through a phase modulator, driven by an RF signal at frequency Ω , where $\Omega \ll \omega$ with ω the frequency of the input beam. To first order, this phase modulation creates sidebands at frequencies $\omega \pm \Omega$. Figure 2.2 shows our experimental setup including all of the components needed for a PDH lock.

The PDH lock can be operated in two different regimes. In the first regime, the LO frequency Ω is much smaller than the linewidth κ of the resonance, in the second regime the opposite holds. We operate in the second regime where $\Omega \gg \kappa$. In this regime, one can consider the sidebands created by the phase modulation to be off resonance and hence they are completely reflected by the cavity.

The reflected signal then consists of the completely reflected sidebands oscillating at $\omega \pm \Omega$, as well as the reflected part of the carrier signal, oscillating at frequency ω . The reflected power, which we measure with a photodiode, then contains a beat note at frequency Ω from the sidebands interfering with the reflected carrier, a beat note at frequency 2Ω from the two sidebands interfering with each other as well as constant and high frequency terms. Going through the theory (see [13]), one finds that for laser frequencies ω close to the resonance, the amplitude of the term oscillating at frequency Ω is proportional to the frequency detuning of the carrier from the cavity resonance for small detunings. Therefore if we take the output of a photodiode that captures the reflected light, mix it with an electrical signal at frequency Ω (effectively multiplying the photodiode output with $\sin \Omega$) and then apply a low-pass filter to this signal, we are only left with the signal proportional to the frequency detuning of the carrier to the resonance of the cavity.

In practice, one mixes the photodiode output with the same signal used to modulate the input beam. In order to ensure the correct polarity of the error signal, the LO output and the photodiode output need to be in phase at the mixer. This can either be achieved by adding a phase shifter in the LO arm before the mixer or more simply by adjusting the frequency of the LO until the two mixer inputs are in phase.

This mixed and filtered output of the photodiode is what we call the PDH error signal. The PDH error signal produced for the room temperature cavity located on an optical table in our lab can be seen in figure 2.1b. The PDH error signal can then be input into a conventional linear control loop, e.g. a set of PID controllers, that minimize the error signal by adjusting the laser frequency.

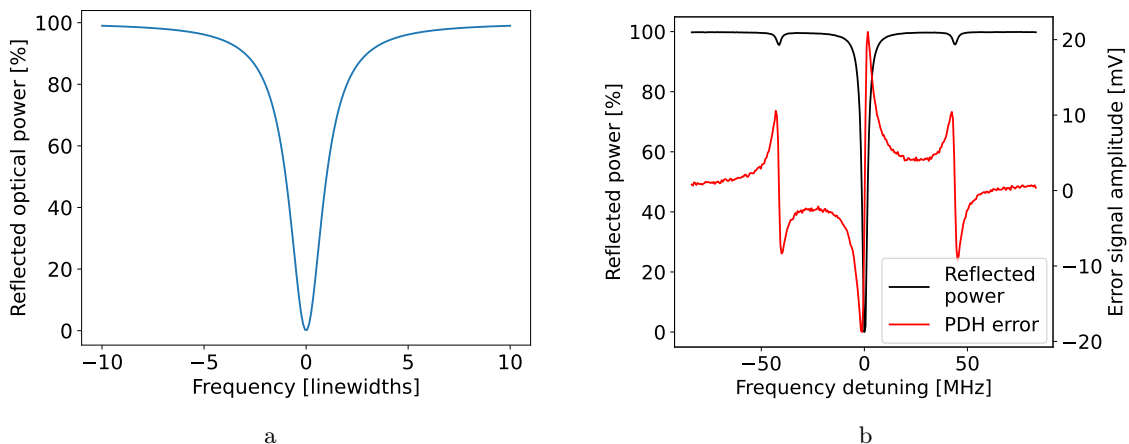


Figure 2.1: **a)** Optical power reflected by a Fabry-Pérot cavity for optical frequencies near a resonance. **b)** Reflected power and PDH error signal as a function of the optical frequency detuning from a resonance of the room temperature cavity. The additional small dips in reflected power appear when the sidebands generated by the phase modulator become resonant with the cavity. The LO frequency is 42.7 MHz.

Note that the PDH error signal is only proportional to the frequency detuning of the carrier when the carrier is very close to the cavity resonance (within a linewidth κ). How-

ever, the PDH error signal maintains the correct sign up to a detuning of Ω , which is useful for the start-up phase of the lock.

2.3 Experimental setup

In this section we give an overview of the experimental setup used for locking a laser in pulsed operation to an optical cavity in a dilution refrigerator by means of a PDH lock and mention the used components and software. We further discuss how to assess the quality of the laser lock.

2.3.1 Overview

A schematic of the experimental setup is shown in figure 2.2. The output of a continuous wave (CW) tunable laser is pulsed using an AOM controlled by an AWG. The AOM output passes through a phase modulator (PM) needed for the PDH lock (see previous section). Ten percent of the phase modulated light is split off and sent to a photodiode (subsequently called the reference port) in order to monitor the optical input power. The rest of the light is sent through a circulator and finally a polarization controller, before it is sent to the dilution refrigerator containing the optical cavity.

Note that the optical cavity we are using contains a flat-flat quartz crystal, which is needed for the optomechanical storage and readout experiment described in chapter 3. The birefringence of the crystal causes the frequency of the optical resonance to depend on the polarization of the incoming light. More precisely, a single optical resonance is split into two resonances, each corresponding to one of the principal crystal axes. In order to address just one resonance, we tune the polarization of the input beam to be parallel with one of the principal crystal axes.

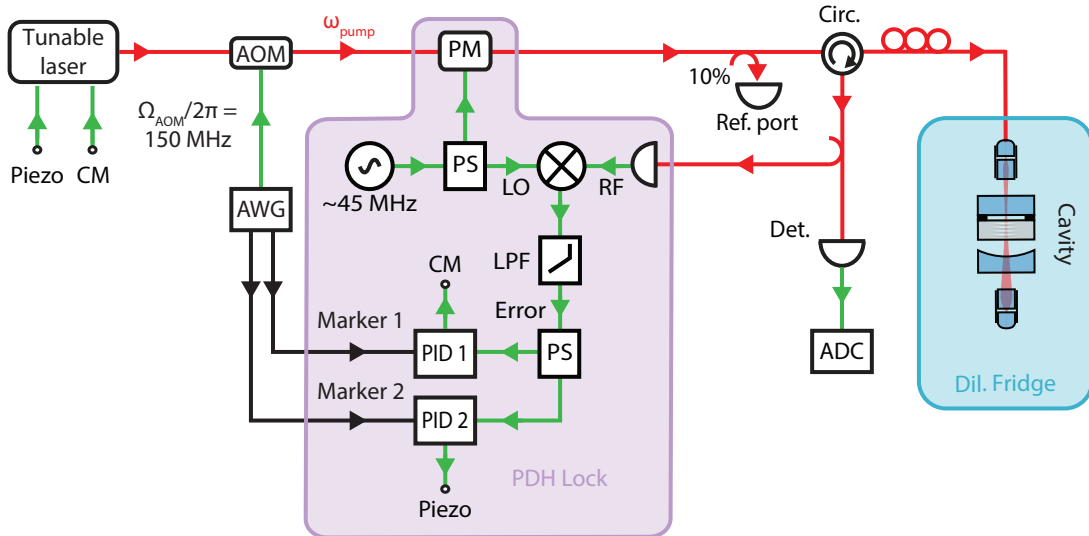


Figure 2.2: The setup used for pulsed laser locking. Red paths are optical fibers, green paths are electrical analog lines and black paths are electrical digital lines. See main text for a detailed description of the setup.

The light reflected from the cavity again passes through the circulator and is then split between two photodiodes. The output of the first photodiode (labeled "Det." in figure 2.2) is sent to an analog to digital converter (ADC) used for data acquisition. This data is used to determine the quality of the lock.

The other photodiode is used for the PDH lock. Its output is processed as described in the previous section. The PDH error signal passes through a power splitter (PS) and is sent to two different PID controllers. PID 1 is a fast PID controller used for high bandwidth current modulation (CM) of the laser. PID 2 is slower and controls the voltage on the piezo element of the laser. This way of adjusting the laser frequency is much slower than CM but grants access to a much wider frequency range. This is needed to correct for larger drifts of the resonance as well as hysteresis of the piezo element, which are too large to be corrected via CM. The separately programmable marker channels of the AWG are used to put the PID controllers on hold between pulses. Putting a PID controller on hold means that the controller holds its last output without making any further corrections. This means that the laser frequency is not further adjusted once a pulse is done and the controllers only resume locking the laser once the next pulse starts. We use two separate marker channels to account for the different reaction times of the two controllers, see section 2.5.1.

In addition to the cavity inside the fridge, we use another cavity placed on an optical table (not shown in figure 2.2) for testing purposes. We refer to it as the room temperature cavity and we call the cavity located in the fridge the cryogenic cavity. The room temperature cavity is similar to the cryogenic cavity except that it does not contain a crystal. We can choose to either send the light to the cryogenic or to the room temperature cavity.

Note that when actually performing experiments involving locking laser pulses, we engage the lock while sending CW light to the selected cavity, before we start pulsing the laser.

2.3.2 Selected Equipment and Software

Here we document the most relevant equipment and software used in this experimental setup.

The laser is a Toptica 1550 CTL tunable laser, creating light around 1550 nm. It is controlled by a DLCpro digital laser controller by Toptica. The AOM is a TEM-150-18-30-1580-2FP by Brimrose, driven by an HDAWG from Zurich Instruments. The phase modulation for the PDH lock is done with an iXblue MPX-LN-0.1. The optical cavity is located within a Bluefors LD400 dilution refrigerator, for details on the cavity and how it was aligned see [1]. Data acquisition was performed with an AlazarTech ATS9870 digitizer, controlled by the QCoDeS software package [15]. For the initial tests of pulsed locking as well as for the optomechanical storage and readout experiment, we used the integrated PID controllers of the laser controller itself for both PID 1 and 2. For the later tests of the laser lock (section 2.5), PID 1 is implemented on a Red Pitaya controlled by the PyRPL software package [16]. For the Red Pitaya, PID 1 is operated simply as a PI controller due to software restrictions. The output of PID 1 is plugged into the Mod DC port directly on the laser head. The photodiode used for the PDH lock was a high bandwidth fixed gain Thorlabs PDA05CF2 photodiode. The photodiode used to record the reflected signal was a Thorlabs PDA10CS2 set to 20 dB gain.

All the optical paths except for the cavity itself are contained within single mode fibers.

2.3.3 Assessing lock quality

We generally use two different methods to assess the quality of the locking scheme. The first method is fairly qualitative: On the screen of the laser controller we can monitor the error signal in real time by eye. If the observed fluctuations of the error signal are much smaller than the amplitude of the error signal during a frequency sweep over the resonance, we can say that we are well locked to the resonance. This is the method we use

to tune PID parameters in section 2.4.1.

A more quantitative method of assessing the quality of the lock is to use what we call the lock level. The lock level is equal to the average reflected power throughout a pulse expressed as a percentage, normalized such that 0% corresponds to the laser being completely on resonance and 100% corresponding to the laser being completely off resonance. Similarly we speak of the instantaneous lock level, when considering the reflected power during a pulse without averaging. Note that our cavity is not critically coupled, therefore there is some residual power reflected, even when the input beam is perfectly on resonance. Thus when the laser is off, as is the case before a pulse, we record a negative instantaneous lock level. This effect is especially pronounced in the cryogenic cavity, where reflections and losses at the vacuum crystal interfaces lead to the modes not all being critically coupled [12].

When averaging the instantaneous lock level (i.e. the measured reflected power during a pulse) to calculate the lock level of a pulse, we exclude the beginning and end of the pulse. This is due to the fact that the reflected power measured during a well locked square pulse exhibits a peak at the beginning and end (see figure 2.3). These peaks can be explained by the cavity dynamics when the rise time of the pulse is much shorter than the average lifetime of a photon inside the cavity. As previously mentioned, when perfectly resonant light gets sent to a critically coupled Fabry-Pérot cavity, no reflected power is measured. This is not due to there being no light reflected by the cavity but rather due to the promptly reflected beam destructively interfering with the light exiting from the cavity. When a square optical pulse is sent to a Fabry-Pérot cavity, some reflected power is measured at the start of the pulse even when the pulse is perfectly resonant with the cavity. While the resonance is building up, there is not yet enough light exiting the cavity to completely cancel out the promptly reflected input beam and hence some reflected power is measured. When the input beam suddenly ends at the end of the pulse, the opposite happens and one again measures a peak in reflected power. The time it takes to build up a resonance is given by the lifetime of the cavity. For our cavity this is around 70 ns, corresponding to a linewidth of $\kappa/2\pi = 2.4$ MHz. Note that the width of the peaks in figure 2.3 agree with the response time of the photodiode, which has a bandwidth of around 1.1 MHz for the used gain setting, meaning that the peaks are likely more narrow, than what the photodiode measures.




Figure 2.3: Instantaneous lock level during a 15 μs long pulse on resonance with the room temperature cavity.

2.4 Initial tests with slow PID 1

Before implementing the Red Pitaya as a fast PID controller (see section 2.5), the first tests were performed with the laser controller as PID 1 and PID 2. We first tune PID parameters (section 2.4.1) and then test the performance of this setup (2.4.2) in pulsed operation. If not specified otherwise all pulses are 15 μs long square pulses and the repetition rate is 1 kHz.

2.4.1 Tuning PID parameters for the room temperature cavity

The first step is to find suitable PID parameters. We do this by minimizing the error signal fluctuations seen on the screen of the laser controller, while sending pulses to the room temperature cavity, as we are adjusting the PID parameters. We can then use the found PID parameters to test the quality of the lock when trying to lock laser pulses to the cryogenic cavity. Note that doing this directly in the cryogenic cavity would have been more tedious and time consuming, since we can only turn the pulse tube of the fridge off (which is required for the lock to work) for a couple minutes at a time, after which the fridge needs to recover.

Recall that for given PID parameters P, I, D and overall gain G , the output V_{out} of a PID controller for given input error signal $\varepsilon(t)$ is given by

$$V_{out}(t) = G \left(P\varepsilon(t) + I \int \varepsilon(t')dt' + D \frac{d\varepsilon(t)}{dt} \right). \quad (2.1)$$

Our starting point is the set of parameters used in [1] for laser locking in CW operation to the same cavity. In [1], laser locking was performed using only the PID 2 controller in the laser controller connected to the piezo element of the laser. The parameters used for CW operation can be found in table 2.1.

PID 2 parameter	
P_2	0.003
I_2	15
D_2	23.3

Table 2.1: PID 2 parameters for laser locking in CW operation. The optimal overall gain depends on the laser power. For example for an optical power of -18 dBm measured at the reference port (see figure 2.2) an overall gain of $G_2 = 1$ is appropriate.

Using just PID 2 for locking of 15 μs long square pulses with a repetition rate of 1 kHz in the room temperature cavity shows that for pulsed operation the performance of the lock is better if D_2 is reduced by roughly a factor of 10. Further, when also using PID 1, we reduce the gains of PID 2 by a factor of 10, since PID 2 only needs to correct for slow drifts and we do not want the slow controller overreacting to errors that PID 1 can correct much faster.

To find suitable PID parameters for PID 1 we first disable PID 2 and find parameters of PID 1 suitable for locking to CW light sent to the room temperature cavity. For this we follow a standard parameter tuning procedure taken from the user manual of the laser controller:

1. Set the integral (I) gain to a non-zero value (all others to zero) and increase the overall gain until the system locks.

2. *Alternate between increasing the proportional (P) gain and the integral (I) gain each until the feedback loop starts to oscillate, then reduce the gains until the oscillation definitely stops. Standard optimization procedures set the proportional (P) gain to about 0.6 of the value where oscillations start.*

3. *Increase the integral (I) gain until the feedback loop starts to oscillate, then increase the differential (D) gain until the oscillation stops. Proceed iteratively Increasing both gains until the oscillation cannot be stopped by further increase of the differential (D) gain. At that point, reduce both gains until the oscillation definitely stops.*

After finding PID 1 parameters suitable for locking a CW laser, we turn on PID 2 with the adapted parameters and begin pulsing the laser. To see if the PID 1 parameters are also suited for pulsed operation we then try increasing and decreasing each PID 1 parameter while monitoring the error signal to see if we can observe an improvement in the lock. The chosen PID parameters for both controllers can be found in table 2.2.

PID 1 parameter		PID 2 parameter	
P_1	10	P_2	0.001
I_1	5	I_2	2.1
D_1	6	D_2	0.2

Table 2.2: PID parameters for laser locking in pulsed operation using the two PID controllers of the laser controller. Parameters were optimized using the room temperature cavity. The optimal overall gains depend on the laser power. For example for an optical power of -20 dBm measured at the reference port (see figure 2.2) overall gains of $G_1 = -3$ and $G_2 = 1$ are appropriate. Note the minus sign for G_1 .

We expect that these parameters have further room for improvement, for the following two reasons: Firstly the parameters were optimized using the room temperature cavity which has a different noise environment than the cryogenic cavity. Secondly the method of optimization (minimizing error signal fluctuations as seen by eye) was only qualitative in nature.

2.4.2 Testing lock with cryogenic cavity

Using the PID parameters from table 2.2 we send pulsed laser light to the cryogenic cavity and measure how well the lock works. Results of sending pulses to the cryogenic cavity for two minutes and measuring the lock level of each pulse can be seen in figure 2.4. The average lock level for this data set is 37 % and the distribution of lock levels (figure 2.4a) can be seen to be very broad. Similar experiments with slightly different PID 1 gains taken on the same day show similar results with the average lock level fluctuating $\sim 2\%$ between subsequent experiments. The moving average of lock levels shows fluctuations in time, but without long-term trend (see figure 2.4b).

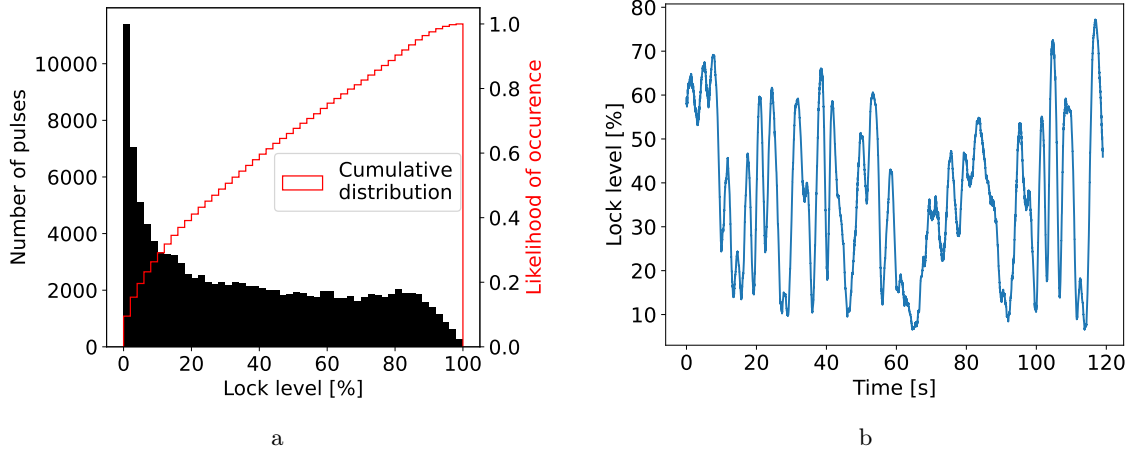


Figure 2.4: Results of two minutes of pulsed laser locking in the cryogenic cavity using the laser controller PID controllers. The used PID parameters can be found in table 2.2, the overall gains were $G_1 = -3$ and $G_2 = 1$, the optical power measured at the reference port was -20 dBm. **a)** Lock level histogram. **b)** Moving average over 1000 pulses.

Looking at individual time traces of the reflected power during a pulse shows the limitation of this setup. Two such traces are shown in figure 2.5. The first trace is an example of the lock making meaningful corrections and the second trace shows the controller over-compensating, causing the lock level to get worse. We can make a couple of observations: First of all, the PID controllers only start making corrections to the laser frequency roughly $3 \mu\text{s}$ after the start of a pulse. This delay comes from the time it takes the PID controllers to be taken out of "hold". For these experiments we did not yet properly understand this delay and simply sent the signal to take the lock out of "hold" at the same time as the RF pulse, which creates the optical pulse. Later (section 2.5) we correct for this delay.

The lock seems to make only two corrections during one pulse and keeps the frequency fairly constant in between these two corrections. This is most likely due to the limited bandwidth of the used PID controllers, which result in this staggered response. While the laser frequency can be adjusted via current modulation with a very high bandwidth (on the order of 100 MHz), the speed at which PID 1 can change the laser frequency is limited by the controller itself, which has a bandwidth of only 30 kHz (see section 2.5.1 for an illustration of this). Note that a controller with a bandwidth of 30 kHz is theoretically too slow to make meaningful adjustments on the timescale of a couple microseconds, as $1/30\text{kHz} \approx 33 \mu\text{s}$. PID 2 also has a bandwidth of 30 kHz however the actual expected bandwidth of the laser frequency response is on the order of 1 kHz, due to the piezo element used to adjust the laser frequency, which is the limiting factor here. This should not be a problem though as PID 2 only needs to correct larger drifts in the laser frequency, which happen on the order of seconds.

Finally, we see two very different responses between a pulse that begins near resonance and one that starts further off resonance. The instantaneous lock level during the pulse shown in figure 2.5a starts at around 17 % and the lock manages to bring the laser completely on resonance by the end of the pulse. Ideally this should happen within the first couple of microseconds but this behavior is promising. In contrast, figure 2.5b shows almost the opposite. Initially the instantaneous lock level is at roughly 50 % and the initial response of the lock brings the laser further off resonance than it already was. Since the PDH error signal is still linear in the frequency detuning at an instantaneous lock level of 50 % (see figure 2.1b) this most likely means that the controller over corrected, resulting in a rise in reflected power. The two selected traces are taken from the same data set that was used to create figure 2.4 and show typical behavior for traces starting at their respective instantaneous lock levels.

We conclude that the response for pulses that start close to being resonant is a bit too weak while the response to a pulse further off resonance is too strong. This kind of trade off is fairly typical for PID controllers. While a lower gain might lead to a better average performance, the bandwidth of PID 1 is clearly too small for the lock to properly work with such short pulses.




Figure 2.5: Instantaneous lock level measured during two pulses, while using integrated PID controllers of the laser controller. **a)** The controllers make corrections to the laser frequency, bringing the laser on resonance. **b)** The controllers over correct, causing an initial rise in reflected power.

2.5 Red Pitaya as PID 1

Due to the bandwidth limitations of the PID controllers in the laser controller, we now use the PID module of a Red Pitaya controlled by the PyRPL software package [16] as our PID 1 controller. The PID controller on the Red Pitaya has a much higher bandwidth and should in principle far outperform the previous setup. PID 2 remains the same as before.

2.5.1 Characterizing components

Similarly to the PID module in the laser controller, the PID on the Red Pitaya also comes with a "hold" function that can be toggled with a DC signal sent to a designated input. If a high Voltage is detected the controller is put on hold and the output is frozen at its current value. If a low Voltage is detected the controller runs normally. Since the manufacturers of both of the used PID controllers do not specify the reaction time between the controller receiving a DC signal telling the controller to be put on "hold" and the controller actually being put on "hold", we perform a simple test to find these numbers. This test also illustrates the large discrepancy in bandwidth between the PID module on the Red Pitaya and that of the laser controller.

First we test the reaction time and bandwidth of PID 1. We apply a 50 MHz sinusoidal signal to the error signal input of the Red Pitaya and set $G_1 = 1$, $P_1 = 1$, $I_1 = 0$, $D_1 = 0$ and record both the input error signal as well as the output of PID 1 on an oscilloscope. For our chosen error signal and PID parameters we expect to retrieve the same sinusoid that we input. Further, we periodically put the lock on hold with a marker from the AWG and also record this marker. The result of this test can be seen in figure 2.6a.

The controller reacts roughly $0.1 \mu\text{s}$ after the controller is taken out of "hold". After this, the output of the controller follows the error signal, demonstrating that the PID controller of the Red Pitaya is capable of reacting to changes in the error signal on a sub microsecond timescale, which is comparable to the timescale on which the cavity reacts to external

changes, given by the inverse linewidth of ~ 70 ns. The amplitudes of the error signal and the PID output don't exactly match. However this is not relevant for us and we did not further investigate this discrepancy. When putting the controller on hold we observe the same delay of around $0.1 \mu\text{s}$.




Figure 2.6: **a)** Response of PID 1 (Red Pitaya) to being taken out of hold while receiving a 50 MHz sinusoidal error signal. **b)** Response of PID 2 (laser controller) to being taken out of hold while receiving a 10 kHz sinusoidal error signal.

Performing the same test for PID 2 but with a 10 kHz signal and PID parameters $P = 1$, $I, D = 0$, $G = 10$ (we increased the overall gain to get a stronger response), we observe a different behavior. The PID controller in the laser controller takes roughly $3 \mu\text{s}$ to react to being taken out of "hold". Further, the controller is too slow to properly follow the error signal. The results can be seen in figure 2.6b. Due to this longer reaction time, we anticipate this delay when using the lock by sending the signal to take PID 2 out of "hold" before the AWG sends the RF pulse to the AOM to create the optical pulse.

A further delay that needs to be taken into account is the time that passes between the AWG sending an RF pulse to the AOM and the detection of a reflected optical signal from the cavity, allowing the locking mechanism to make corrections to the laser frequency. This delay mainly comes from the AOM itself. The AOM starts to transmit light once an acoustic standing wave has built up in its internal crystal. For the AOM we used, this was measured to take on the order of 400 nanoseconds. Further, there are delays introduced by the optical paths as well as the photodiodes detecting the reflected signal. When measured in our setup, the total delay between sending an RF pulse with the AWG and measuring a rise in the reflected signal is around 500 ns, the exact number depending on whether the room temperature or the cryogenic cavity is used.

Finally, we time the marker channels such that the PID controllers are taken out of hold around 300 ns after the pulse started. This is to account for cavity dynamics. As discussed in section 2.3.3 and seen in figure 2.3, the reflected signal of a pulse that is on resonance with a cavity mode has peaks at the beginning and end of the pulse. Thus the PDH error signal does not represent the detuning of the cavity during these peaks. To make sure the controllers do not react to these peaks we time the controllers to be taken out of hold 300 ns after the the initial rise in measured reflected power.

2.5.2 Tuning PID parameters using both cavities

After characterizing PID controllers reaction times to being put on "hold", the next step is to find suitable PID parameters. We follow a very similar procedure to section 2.4.1 but try to be a bit more quantitative in our choice of PID 1 parameters.

We decide to set $D_2 = 0$. This can be motivated in two ways. Firstly, the role of PID 2

is to correct for larger drifts in the laser frequency. These do not happen suddenly so the differential gain D_2 should not be needed. Secondly, the role of the differential gain D_2 is to react to sudden changes in the error signal and these are corrected for by PID 1 on a much faster timescale than PID 2 is able to, due to the much higher bandwidth of PID 1. Finally, for similar reasons we reduce the proportional gain P_2 in relation to I_1 a little bit compared to the parameters for CW operation and also reduce the overall gain, since we want the faster PID 1 to perform most of the corrections. The chosen PID 2 parameters for pulsed laser locking can be found in table 2.3.

PID 2 parameter	
P_2	0.0001
I_2	2.1
D_2	0

Table 2.3: PID 2 parameters for laser locking in pulsed operation when using the Red Pitaya as PID 1. The optimal overall gain depends on the laser power. For the typical laser power used for the tests (around -18 dBm measured at the reference port) an overall gain of $G_2 = 1$ is appropriate.

To find good PID parameters for PID 1, we perform tests in the room temperature cavity. We follow standard procedure for calibrating PID parameters (see section 2.4.1) and start by leaving the proportional gain P_1 at zero and increase I_1 until we get a decent lock. The way the controller module is implemented in the PyRPL software package, the overall gain is fixed at $G_1 = 1$ and there is no option for differential gain (i.e. $D_1 = 0$). Note that during all of these tests PID 2 is also on with the parameters from table 2.3. Contrary to section 2.4.1, we now use the lock level to quantify the lock performance while tuning PID parameters. Results from two sweeps of I_1 in the room temperature cavity can be found in figure 2.7. From the results of these two sweeps we choose $I_1 = 15000$ for now.




Figure 2.7: Results of locking a pulsed laser to the room temperature cavity while varying I_1 . For each value of I_1 we plot the average and standard deviation of the lock level of 30'000 pulses. PID 2 was on during these measurements with the parameters from table 2.3 and a gain of $G_2 = 1$. The optical power at the reference port was -17.6 dBm.

Leaving the integral gain of PID 1 fixed at $I_1 = 15000$, we begin to turn up the proportional gain P_1 and observe its effect on the performance of the lock. Results of two such sweeps performed in the room temperature cavity can be seen in figure 2.8. The results merely show that the lock is close to perfect in the room temperature cavity for a fairly wide range of P_1 . However, we then also perform a sweep of P_1 in the cryogenic cavity (see figure 2.9), which leads us to settle on the PID 1 parameters found in table 2.4.




Figure 2.8: Results of locking a pulsed laser to the room temperature cavity while varying P_1 . For each value of P_1 we plot the average and standard deviation of the lock level of 30'000 pulses. PID 2 was on during these measurements with the parameters from table 2.3 and a gain of $G_2 = 1$. The optical power at the reference port was -17.6 dBm. Note that for $I_1 = 15000$ and $P_1 = 0$ we measured an average lock level of 0.8 ± 0.7 % (see figure 2.7b).

PID 1 parameter	
P_1	10
I_1	15000

Table 2.4: PID 1 parameters for laser locking in pulsed operation for typical laser powers (around -18 dBm measured at the reference port). For different laser powers the parameters have to be scales accordingly, as the PID 1 module has no overall gain setting.




Figure 2.9: Results of locking a pulsed laser to the cryogenic cavity while varying P_1 . For each value of P_1 we plot the average and standard deviation of the lock level of 30'000 pulses. PID 2 was on during these measurements with the parameters from table 2.3 and a gain of $G_2 = 1$. The optical power at the reference port was -18 dBm.

We also did some tests of the lock in the cryogenic cavity with the same PID 2 parameters as in table 2.3 but with $P_1 = 15$ and $I_1 = 7500$ that actually produced the best results. If these are simply better PID 1 parameters or if it was some other, potentially external, factor is unclear at this point, see section 2.5.3.

2.5.3 Long-time lock behavior

Having found suitable PID parameters with the new PID 1 controller, we now study the lock performance over longer time scales without changing any parameters during single measurements in order to get better statistics and to see if there are any long term trends. As before, when not specified otherwise we are using $15\ \mu\text{s}$ long square pulses at a repetition rate of 1 kHz.

General performance

First we present results from running the lock in pulsed operation in the cryogenic cavity with the PID parameters from tables 2.3 and 2.4 over a duration of two minutes. The results are shown in figure 2.10. The distribution of the lock levels is less broad compared to a similar measurement done using the slower PID 1 (see figure 2.4a) and we see no clear temporal trend of the lock levels. The average lock level for this measurement was 27 %, showing a clear improvement in the quality of the lock compared to using the laser controller for PID 1.

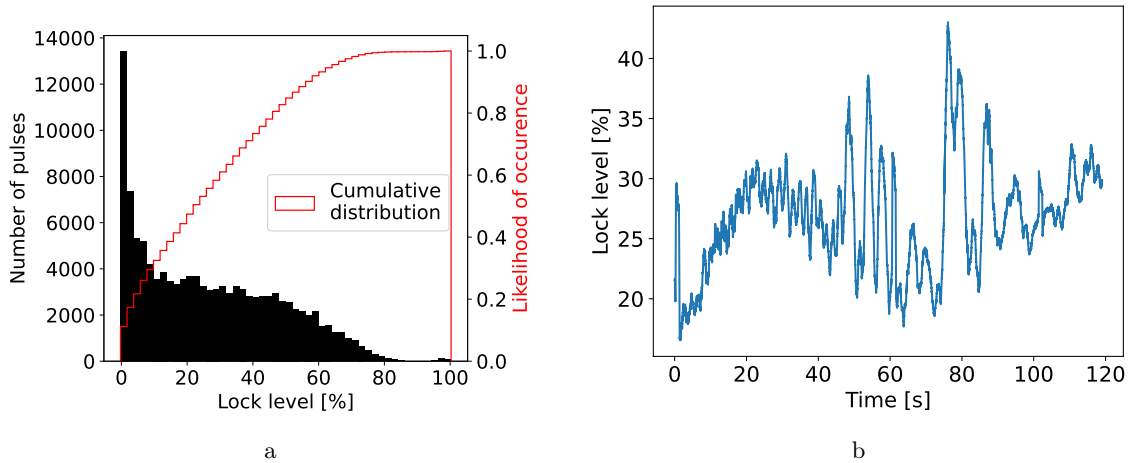


Figure 2.10: Results of locking a pulsed laser to the room temperature cavity over the course of two minutes. The used PID parameters can be found in tables 2.3 and 2.4. The overall gain of PID 2 was $G_2 = 1$, the optical power at the reference port was $-17.9\ \text{dBm}$. **a)** Lock level histogram. **b)** Moving average over 1000 pulses.

Further, we can record the reflected power during a pulse to inspect the effect of the lock. An example of such a time trace can be seen in figure 2.11. At the beginning of the pulse the lock makes a fast correction during the first couple of microseconds and then slows down substantially and does not manage to bring the laser entirely on resonance. Note that the initial peak here differs from the one shown in figure 2.3 and actually comes from the lock making a correction. The peaks seen in figure 2.3 only occur when the laser is almost on resonance to begin with, which is corroborated by the fact that we see no peak in reflected power at the end of the pulse in figure 2.11.

In general, the reflected power during a pulse is a lot smoother when the Red Pitaya is used for PID 1 as opposed to the slower laser controller. However as one clearly sees in figure 2.11, the lock is not quite doing enough to ensure fully locked pulses, see section 2.5.4. Since the chosen I_1 is close to the maximum settable value and in fact the output of the Red Pitaya maxes out quite regularly during pulses (for further discussion of this issue see section 2.5.4), we tried increasing the gain of PID 2 in the hopes of improving the results but this did not have a significant impact.




Figure 2.11: Reflected power during a pulse, while locking with the Red Pitaya.

Turbo pump noise

The main noise source in the fridge when the pulse tube is off is the turbo pump, which creates vibrations at a frequency of around 820 Hz. If this noise is not corrected for by the lock, we should in principle see oscillations in the lock level of subsequent pulses at this frequency. To test this we need to increase the repetition rate, since a sampling rate of 1 kHz is too low to properly detect an oscillation at 820 Hz. More precisely, we expect oscillations at 1640 Hz, as the lock level does not differentiate between being blue or red detuned from the resonance.

To test this, we increase the repetition rate to 10 kHz and measure the lock level of 120'000 pulses. The results from this test are shown in figure 2.12. The lock level over time shows clear oscillations (figure 2.12a) and the discrete Fourier transform (DFT, figure 2.12b) exhibits clear peaks at multiples of 820 Hz with the peak at 1640 Hz being the highest.ⁱ We can conclude that the turbo pump is still a significant error source, even with the lock enabled.




Figure 2.12: Results from locking pulses at a repetition rate of 10 kHz for 12 seconds. **a)** Lock levels during the first 10 ms. **b)** Spectral analysis of the lock levels. The y-axis is normalized such that the maximal value is 1. For clarity the component at 0 Hz is not shown. It is roughly five times the height of the peak at 1640 Hz.

ⁱIf anyone is interested in founding a startup to build terrible (and expensive) microphones using this technique, please write me an email.

Best lock performance

At this point it is worth mentioning that the performance of the lock is greatly influenced by external factors, such as people talking and walking around in the lab, what people are doing in other rooms of the building, how the air conditioning is running, etc. Since we cannot control for all of these factors, it is difficult to compare data sets from different days. Nevertheless we now present the measurements with the best lock quality observed. For these measurements the PID 1 parameters were chosen slightly differently, somewhat at random. Namely we used $P_1 = 15$ and $I_1 = 7500$. Further due to a typo in the measurement code the slower PID 2 turned on roughly $2 \mu\text{s}$ after the faster PID 1. It is unclear if these two aspects led to the better lock quality or if it just was a calm day in the lab. The individual traces of reflected power during the pulses do not look qualitatively different from 2.11. The results of observing the lock level in pulsed operation can be seen in figure 2.13. The average lock level for this data set was 16 % and two further repetitions of the same experiment on the same day had average lock levels of 18 % and 20%.

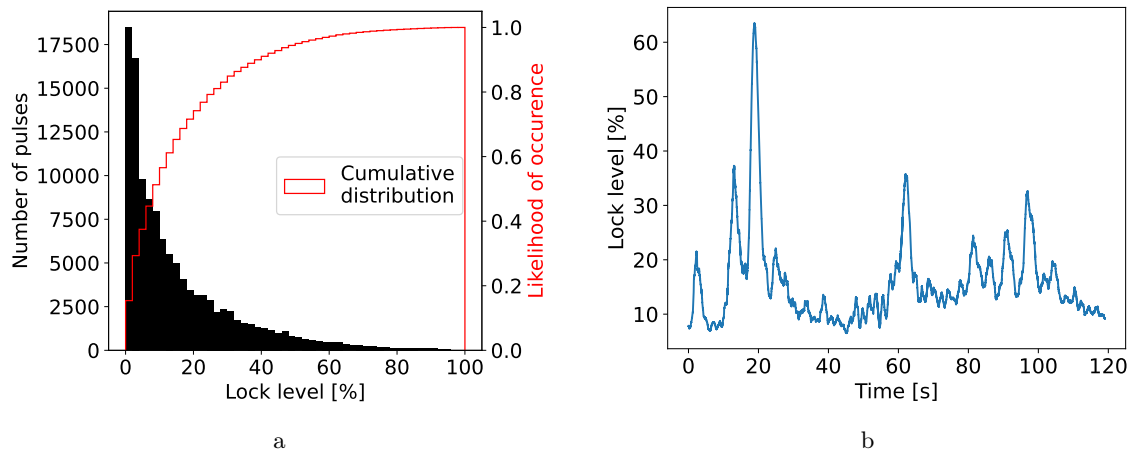


Figure 2.13: The best performance of the lock we observed during this project. We only observed this quality of the lock on one day. **a)** Lock level histogram. **b)** Moving average over 1000 pulses.

On a different day we used the same PID parameters as in the measurement shown in figure 2.13 but using the timing of lock hold described in 2.5.1 and compared it to a measurement taken directly afterwards using the PID 1 parameters from table 2.4. The lock performed very similarly for both sets of PID parameters. Due to the limited time available for this project, we have not yet been able to further study what might be responsible for these measurements having such good results but the adjusted timing of the two PID controllers could very well have made a difference.

Finally we show that the quality of the lock shows some improvement, when the repetition rate is increased. Results of a measurement showing this can be found in figure 2.14. This data was taken on the same day of the good lock performance mentioned above and also with the same slightly different PID 1 parameters.




Figure 2.14: Lock performance as a function of the time τ_{wait} between pulses (i.e. changing the repetition rate). Each data point is the average and standard deviation of 30'000 pulses.

2.5.4 Discussion and potential improvements

In this section we discuss the results of using the Red Pitaya for PID 1 and outline some steps that can be taken in the future to potentially improve the overall performance of the lock.

Due to the short pulses of around $15 \mu s$ required for the proposed microwave-to-optical transduction experiment, the PID controllers and in particular PID 1 needs to have a high bandwidth as it is responsible for making frequency corrections on a microsecond timescale. We have seen (figure 2.6a and 2.11) that the Red Pitaya is capable of doing this in a controlled manner and this leads to a decent performance of the locking mechanism. However there is still a lot of room for improvement and some unresolved technical issues that we would like to address.

Firstly there is the issue of finding optimal PID parameters. While we have performed some sweeps of PID parameters in the cryogenic cavity, these sweeps were fairly sparse. I.e. only a couple different PID parameter values could be tested in each measurement and we only got feedback after the measurement was completed. Ideally we would like to be able to tune the PID parameters while we are pulsing and get real-time, quantitative feedback of the lock performance while doing this. This way we could potentially find better suited parameters for locking a pulsed laser in the noise environment of the fridge. Alternatively if these measurements could be performed while the pulse tube of the fridge is on, we could perform many more and longer measurements, allowing us to more systematically explore the parameter space of this problem.

A further, related avenue worth exploring (especially due to the good lock seen in figure 2.13a) is the timing of taking the PID controllers out of hold. Maybe the two PID controllers cooperate better if the fast PID 1 is allowed to make some initial corrections, before the slower PID 2 is enabled.

Finally there are technical issues with the Red Pitaya itself. The range of the output channel of the Red Pitaya is merely ± 1 V. This voltage range is enough to sweep the laser frequency over the entire resonance but not much further. Due to this limited range the output of the Red Pitaya regularly maxes out, meaning that PID 1 is not making any corrections until the much slower PID 2 makes a correction that leads to the output of the Red Pitaya coming away from ± 1 V. This would happen less often if one would amplify the output of the Red Pitaya. For comparison, the output of the PID controller on the laser controller we used as PID 1 in section 2.4 has a range of ± 4 V.

Another technical aspect with the Red Pitaya which is still unclear, is how well we set the lock point (i.e. which Voltage should the error signal be held at). Before each experiment, we record the PDH error signal on the Red Pitaya with its internal oscilloscope and use this

to set the lock point. However the error signal is very noisy, even for the room temperature cavity, as can be seen in figure 2.15 and taking steps to improve this might prove fruitful.

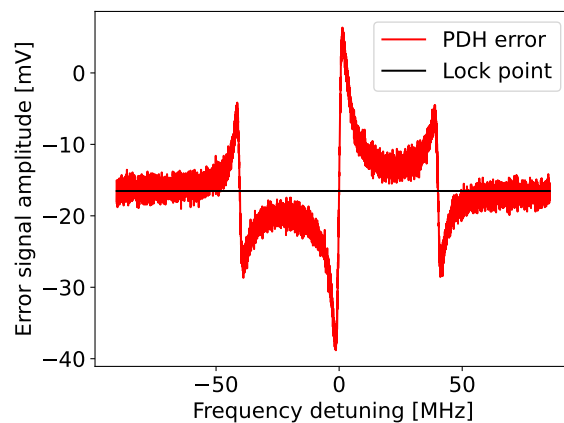


Figure 2.15: Error signal as recorded by the oscilloscope on the Red Pitaya during a frequency sweep over a resonance of the room temperature cavity.

Optomechanical storage and readout

In this chapter we report on the optomechanical storage of an infrared pulse in a high frequency (~ 12.6 GHz) phonon mode of a high-overtone bulk acoustic wave resonator (HBAR) and the subsequent readout thereof. High-overtone refers to the fact that the phonon modes we address have longitudinal mode numbers on the order of $\sim 10^4$. The optomechanical readout of high frequency phonons, i.e. the conversion of phonons to infrared photons, when performed at the single phonon to photon level, is an essential component of the microwave to optical transduction experiment introduced in chapter 1 and further described in [11]. This experiment also constitutes a first application of the pulsed laser locking scheme discussed in chapter 2.

We create coherent phonon states by means of stimulated Brillouin scattering and subsequently read them out after a variable wait time. A very similar experiment has previously been performed using whispering gallery modes of a fused silica microsphere resonator [17, 18].

This chapter is structured as follows: In section 3.1 we describe the optomechanical system consisting of an HBAR placed within an optical Fabry-Pérot cavity. Section 3.2 contains a simple theoretical model of the experiment based on the semi-classical treatment of two optical modes coupled to a single mechanical mode. In section 3.3 we give an overview of the performed experiment and describe the experimental setup, further we relate the measured detector output voltage to the retrieved phonon energy. In section 3.4 we present our experimental results and then discuss them in section 3.5.

3.1 The optomechanical system

In this section we describe our optomechanical system and the Brillouin scattering mechanism, by which the mechanical and optical modes interact. We further briefly discuss the effect of placing a flat-flat quartz crystal in a Fabry-Pérot cavity on the mode spacing.

Our optomechanical system consists of an HBAR, which in our case is a flat-flat, 5 mm thick, z-cut quartz-crystal, placed within a Fabry-Pérot cavity (the cryogenic cavity from chapter 2, see figure 3.1). This system is then cooled to cryogenic temperatures in a conventional dilution refrigerator. This increases the coherence length of phonons in the quartz crystal well beyond its spatial extent, turning the flat-flat crystal into a mechanical Fabry-Pérot cavity, where phonons are reflected at the quartz-vacuum interface [19]. The optomechanical system is the same one used in [1] and very similar to the ones in [12, 19].




Figure 3.1: Drawing of the cryogenic cavity. Light from an optical fiber is focused into a Fabry-Pérot cavity, containing an HBAR made up of a flat-flat quartz crystal. The output is refocused into the output fiber by another GRIN lens. (Source: [1])

The optomechanical interaction in this system takes place between a phonon mode of the HBAR and two optical modes of the cavity, which are separated by the frequency of the phonon mode. The electric fields of the two optical modes form standing waves in the Fabry-Pérot cavity. The beat-note between these two standing waves creates a grating in the index of refraction within the HBAR due to the photoelasticity of the quartz. A photon in the higher energy mode can then scatter off this grating and get downconverted to the lower frequency mode, creating a phonon in the process (Stokes scattering). The opposite process, where a photon in the lower frequency mode absorbs a phonon to create a photon in the higher energy mode (anti-Stokes scattering) is also possible. The optomechanical interaction Hamiltonian of this system in the rotating wave approximation (see [12, 14]) is

$$\hat{H}_{OM} = -\hbar g_0 \left(\hat{a}_1 \hat{a}_2^\dagger \hat{b} + \hat{a}_1^\dagger \hat{a}_2 \hat{b}^\dagger \right), \quad (3.1)$$

where \hat{a}_1 (\hat{a}_2) is the annihilation operator of the lower (higher) frequency optical mode, \hat{b} is the annihilation operator of the phonon mode and g_0 is the single photon coupling rate. In our experiment we classically pump the optical mode \hat{a}_1 . In the rotating frame of the pump laser, this corresponds to making the replacement $\hat{a}_1 \rightarrow \sqrt{N_1}$ in the interaction Hamiltonian (choosing the phase to yield real expectation value \bar{a}_1), leading to

$$\hat{H}_{OM,eff} = -\hbar g_{eff} \left(\hat{a}_2^\dagger \hat{b} + \hat{a}_2 \hat{b}^\dagger \right), \quad (3.2)$$

where $g_{eff} = g_0 N_1$ is the effective, cavity enhanced coupling rate and N_1 is the mean intracavity photon number for the pumped mode \hat{a}_1 [12]. This linearized Hamiltonian shows a beamsplitter interaction between the cavity mode \hat{a}_2 and the phonon mode \hat{b} , allowing us to swap excitations from one mode into the other.

In our setup, the linewidth of the optical modes $\kappa \simeq 2\pi \cdot 2.4$ MHz is much larger than the mechanical mode spacing of ~ 625 kHz. This means that in reality we couple to more than just one mechanical mode at once. However, we can assume these different mechanical modes to be independent, because the mechanical linewidth is much smaller than the mechanical mode spacing. The effective interaction Hamiltonian is the sum over multiple phonon modes of interaction Hamiltonians of the form (3.1) and thus for simplicity we only consider one mechanical mode at once.

The frequency difference between the two optical modes being equal to the phonon frequency ensures that the described Brillouin interaction obeys energy conservation. When taking momentum conservation (i.e. wave vector matching) into account, one finds that this interaction only occurs (i.e. one finds an appreciable g_0) for phonon frequencies near

the Brillouin frequency Ω_B , given by

$$\Omega_B \simeq 2\omega \frac{v_a}{v_o}, \quad (3.3)$$

where ω is the frequency of the pump light and v_a (v_o) is the longitudinal speed of sound (light) in quartz and we assumed $v_a \ll v_o$ (see SI of [19]). For quartz at cryogenic temperatures and 1550 nm wavelength light, the Brillouin frequency is roughly $\Omega_B/2\pi \approx 12.64$ GHz. So to get a non-zero optomechanical coupling rate, we need to tune the free spectral range (FSR) of our optical cavity to be equal to the Brillouin frequency.




Figure 3.2: Simulated FSR modulation of six neighboring modes for similar parameters as our cavity. Each line corresponds to one mode pair. The presence of the HBAR in the optical cavity leads to a periodic modulation of the FSR, when changing the spacing between the mirrors. (Source:[1])

A conventional optical Fabry-Pérot resonator exhibits a regular mode spacing with an FSR of $\Delta f_c = c/(2L)$, where c is the speed of light and L is the spatial separation between the two mirrors. Placing the quartz crystal between the two mirrors however changes this. The crystal introduces reflections within the cavity at the quartz-vacuum interfaces and changes the optical path length between the mirrors. This leads to a modulation in the FSR $\Delta f_c^j = f_j - f_{j-1}$ of the cavity between adjacent modes f_j and f_{j-1} . The FSR varies by up to roughly 20 % of the mean FSR. See the SI of [12] for an in depth treatment of this.

If we change the length of the cavity, we observe a similar periodic modulation of Δf_c^j but now as a function of the length change δL . A plot illustrating this can be found in figure 3.2. Since the cavity is subject to a fair amount of noise in the fridge, causing fluctuations in the distance between the mirrors, which in turn causes fluctuations in the FSR, we tune the distance between our mirrors and the frequency of our laser such that there is a mode pair $(j, j-1)$ with $\Delta f_c^j = \Omega_B/2\pi$, with a mode spacing that is to first order insensitive to fluctuations in δL . For more details on adjusting the mode spacing, see [1].

3.2 Theory

In this section we present a simplified model of the performed optomechanical storage and readout experiment. In section 3.2.1 we describe what we call the write pulse, in which two laser tones separated by the phonon frequency drive the phonon mode into a coherent state. In section 3.2.2 we describe what we call the readout pulse, where the phonon state is partially swapped into the higher frequency optical mode by applying the beam splitter interaction of equation (3.2).

3.2.1 Write pulse

First we describe the write pulse. For this, we closely follow the first steps in the derivation of thermometry signals given in the SI of [1]. We start with the Hamiltonian of two optical modes \hat{a}_1 and \hat{a}_2 at frequencies ω_1 and ω_2 coupled optomechanically to a single phonon mode \hat{b} at frequency Ω_m with single photon coupling rate g_0 ⁱ. As mentioned in section 3.1, we include only one phonon mode in the theoretical treatment since they are well separated in frequency. We assume the optical mode spacing to be detuned from the mechanical frequency by $\Delta_m = \omega_1 + \Omega_m - \omega_2$.

During the write pulse, mode \hat{a}_1 is pumped on resonance with frequency ω_1 and amplitude α_1^{in} , which is ensured by a laser lock like the one described in chapter 2. An intensity modulator driven with frequency Ω_{IM} creates a second tone at frequency

$$\omega_{write} = \omega_1 + \Omega_{IM} = \omega_2 + \Delta_{write} + \Delta_m \quad (3.4)$$

with amplitude α_2^{in} , where $\Delta_{write} = \Omega_{IM} - \Omega_m$ is the detuning between the IM drive frequency and the mechanical frequency.

This results in the following Hamiltonian:

$$\begin{aligned} \hat{H} = & \hbar\omega_1\hat{a}_1^\dagger\hat{a}_1 + \hbar\omega_2\hat{a}_2^\dagger\hat{a}_2 + \hbar\Omega_m\hat{b}^\dagger\hat{b} - \hbar g_0 \left(\hat{a}_1\hat{a}_2^\dagger\hat{b} + \hat{a}_1^\dagger\hat{a}_2\hat{b}^\dagger \right) \\ & + i\hbar\sqrt{\kappa_{ext}}\alpha_1^{in} \left(\hat{a}_1^\dagger e^{-i\omega_1 t} - \hat{a}_1 e^{i\omega_1 t} \right) + i\hbar\sqrt{\kappa_{ext}}\alpha_2^{in} \left(\hat{a}_2^\dagger e^{-i\omega_{write} t} - \hat{a}_2 e^{i\omega_{write} t} \right), \end{aligned} \quad (3.5)$$

where we assume both external fields to have the same coupling rate κ_{ext} to the optical modes.

The Langevin equations of motion [20] for the optical modes are

$$\dot{\hat{a}}_1 = -i\omega_1\hat{a}_1 + \sqrt{\kappa_{ext}}\alpha_1^{in}e^{-i\omega_1 t} - \frac{\kappa}{2}\hat{a}_1, \quad (3.6)$$

$$\dot{\hat{a}}_2 = -i\omega_2\hat{a}_2 + \sqrt{\kappa_{ext}}\alpha_2^{in}e^{-i\omega_{write} t} - \frac{\kappa}{2}\hat{a}_2, \quad (3.7)$$

where we assumed both optical modes to have the same decay rate κ and the single photon coupling rate g_0 to be small compared to κ, κ_{ext} and ω_1 (ω_2), allowing us to drop the term $ig_0\hat{a}_2\hat{b}^\dagger$ ($ig_0\hat{a}_1\hat{b}$).

Treating the optical modes classically and ignoring any quantum fluctuations in the signal, we replace the operators $\hat{a}_{1/2}$ with their expectation value $\bar{a}_{1/2}$. Solving for these classical amplitudes, we find

$$\bar{a}_1 = \frac{\sqrt{\kappa_{ext}}}{\kappa/2}\alpha_{1,in}e^{-i\omega_1 t} \equiv \alpha_1^{cav}e^{-i\omega_1 t}, \quad (3.8)$$

$$\bar{a}_2 = \frac{\sqrt{\kappa_{ext}}}{\kappa/2 - i(\Delta_{write} + \Delta_m)}\alpha_{2,in}e^{-i\omega_{write} t} \equiv \alpha_2^{cav}e^{-i\omega_{write} t}, \quad (3.9)$$

where $\alpha_{1/2}^{cav}$ is the intra-cavity amplitude of the respective optical mode.

The number of intra-cavity photons $N_{1/2}$ in a given mode is therefore $N_{1/2} = |\alpha_{1/2}^{cav}|^2$. For simplicity, we take $\alpha_{1/2}^{cav}$ to be real valued (the two incident fields are in phase and the phase of $\alpha_{1/2}^{cav}$ can be absorbed into $e^{-i\omega_{1/2} t}$ by a shift in the time t).

ⁱThis is the main simplification of this model. In reality we excite a superposition of eigenmodes of the HBAR with a collective mode profile that matches the Gaussian profile of the optical mode. See section 3.5.

Inserting these results into equation 3.5 and defining the cavity-enhanced optomechanical coupling rate $g_{write} = g_0\sqrt{N_1N_2}$ for the write pulse, we get

$$\hat{H} = \hbar\omega_1N_1 + \hbar\omega_2N_2 + \hbar\Omega_m\hat{b}^\dagger\hat{b} - \hbar g_{write} \left(\hat{b}e^{i(\Omega_m+\Delta_{write})t} + \hat{b}^\dagger e^{-i(\Omega_m+\Delta_{write})t} \right). \quad (3.10)$$

The Langevin equation of motion for the phonon mode is thus

$$\dot{\hat{b}} = -i\Omega_m\hat{b} + ig_{write}e^{-i(\Omega_m+\Delta_{write})t} - \frac{\Gamma_m}{2}\hat{b}, \quad (3.11)$$

where Γ_m is the mechanical decay rate. Treating the phonon mode classically, i.e. making the replacement $\hat{b} \rightarrow \beta e^{-i(\Omega_m+\Delta_{write})t}$ one finds

$$\dot{\beta} = (i\Delta_{write} - \frac{\Gamma_m}{2})\beta + ig_{write}. \quad (3.12)$$

Assuming no initial phonon population ($\beta(0) = 0$), the solution to equation 3.12 is given by

$$\beta(t) = i\frac{g_{write}}{\Gamma_m/2 - i\Delta_{write}} \left(1 - e^{(i\Delta_{write} - \Gamma_m/2)t} \right). \quad (3.13)$$

Thus if the writing pulse is much longer than $2/\Gamma_m$, the initial phonon state right after the write pulse is a coherent state $|\beta_0\rangle$ with amplitude

$$\beta_0 \approx i\frac{g_{write}}{\Gamma_m/2 - i\Delta_{write}}. \quad (3.14)$$

The incident field amplitudes $\alpha_{1/2}^{in}$ needed to calculate the effective coupling rate g_{write} can be related to optical power P_{RP} measured at the reference port (see figure 3.4) via

$$|\alpha_1^{in}|^2 = N_1 = \eta_{in}(1 - \eta_{IM})\frac{P_{RP}}{\hbar\omega_1} \quad (3.15)$$

$$|\alpha_2^{in}|^2 = N_2 = \eta_{in}\eta_{IM}/2\frac{P_{RP}}{\hbar\omega_{write}}, \quad (3.16)$$

where η_{IM} is the fraction of the total optical power contained in the sidebands after passing through the intensity modulator and η_{in} accounts for the conversion between power measured at the reference port to power at the cavity input, including losses. The main losses between the reference port and amount of light coupled into the cavity come from the misalignment of the GRIN lens which focuses the fiber output onto the cavity as well as the performance of the lock.ⁱⁱ

3.2.2 Readout pulse

Next we describe what happens during the readout pulse, where only the lower frequency optical mode \hat{a}_1 is driven but we have a non-zero initial phonon population. For this we follow the derivation from the previous section but assume only the lower frequency optical mode \hat{a}_1 to be pumped classically. This leads to the Hamiltonian

$$\hat{H} = \hbar\omega_1N_1 + \hbar\omega_2\hat{a}_2^\dagger\hat{a}_2 + \hbar\Omega_m\hat{b}^\dagger\hat{b} - \hbar g_{read} \left(\hat{a}_2^\dagger\hat{b}e^{-i\omega_1t} + \hat{a}_2\hat{b}^\dagger e^{i\omega_1t} \right), \quad (3.17)$$

ⁱⁱActually, the performance of the lock should be taken into account by assuming the frequency of the input beam to be detuned from ω_1 . However, for simplicity and since the lock performance is pretty consistent during this experiment, we assume the input to be resonant with the cavity mode and absorb the non perfect lock into the prefactor η_{in} .

3 Optomechanical storage and readout

where $g_{read} = g_0\sqrt{N_1}$ is the cavity-enhanced coupling rate for the readout pulse.

The Langevin equations of motion for the optical mode \hat{a}_2 and the phonon mode \hat{b} are

$$\dot{\hat{a}}_2 = -i\omega_2\hat{a}_2 + ig_{read}\hat{b}e^{-i\omega_1t} - \frac{\kappa}{2}\hat{a}_2 \quad (3.18)$$

$$\dot{\hat{b}} = -i\Omega_m\hat{b} + ig_{read}\hat{a}_2e^{i\omega_1t} - \frac{\Gamma_m}{2}\hat{b}. \quad (3.19)$$

Taking both the remaining optical and phonon mode to be classical as well, by making the replacements $\hat{a}_2 \rightarrow \alpha_2^{cav}e^{-i\omega_2t}$ and $\hat{b} \rightarrow \beta e^{-i\Omega_m t}$, the Langevin equations reduce to

$$\dot{\alpha}_2^{cav} = -\frac{\kappa}{2}\alpha_2^{cav} + ig_{read}e^{-i\Delta_m t}\beta \quad (3.20)$$

$$\dot{\beta} = -\frac{\Gamma_m}{2}\beta + ig_{read}e^{i\Delta_m t}\alpha_2^{cav}, \quad (3.21)$$

taking the form of coupled oscillator equations.

Finally the field amplitude α_2^{out} of optical mode 2 in the output fiber of the cavity is given by

$$\alpha_2^{out} = -\sqrt{\kappa_{ext}}\alpha_2^{in}. \quad (3.22)$$

Note that between the write and readout pulse, the phonon amplitude decays as $|\beta_0| \rightarrow |\beta_0|\exp(-\Gamma_m t_{wait}/2)$. Hence the initial conditions with which one needs to solve the equations of motion are $\alpha_2^{cav}(0) = 0$ and $\beta(0) = \beta_0 \exp(-(i\Omega_m + \Gamma_m/2)t_{wait})$, where t_{wait} is the time between the end of the writing pulse and the start of the readout pulse.

In figure 3.3 we plot the output power in mode \hat{a}_2^{out} given by $\hbar\omega_2|\alpha_2^{out}|^2$ as well as the phonon amplitude $|\beta|$ during the readout pulse for typical parameters of our experiment (see section 3.4.3).




Figure 3.3: Optical power in mode \hat{a}_2^{out} and phonon amplitude $|\beta|$ during the readout pulse for typical parameters of our experiment.

The qualitative behavior of β and α_2^{out} is determined by the value of the optomechanical cooperativity C_{read} given by

$$C_{read} = \frac{4g_{read}^2}{\kappa\Gamma_m}. \quad (3.23)$$

For $g_{read} < \kappa$ as is the case in our experiment, the phonon amplitude $|\beta(t)|$ is well approximated by an exponential decay of the form $\propto \exp(-\Gamma_m(1 + C_{read})t/2)$. The coupling to the optical mode introduces an additional loss channel for the phonons, leading to a faster decay. For $C_{read} < 1$, as is the case in this experiment, we are in the low coopera-

tivity regime. In this case the regular decay channels described by Γ_m are still dominant. For $C_{read} > 1$ (but still $g_{read} < \kappa$) one is in the high cooperativity regime, where a phonon is more likely to swap into a photon rather than decay internally and vice versa. Once the photon is in the cavity however, it is more likely to decay through the optical output port than to swap back into a phonon (as would be the case in the strong coupling regime $g_{read} > \kappa$), allowing for allowing for efficient detectionⁱⁱⁱ.

Integrating the power in the output mode \hat{a}_2^{out} over time yields the energy contained in the readout pulse. Measuring the energy in the readout pulse as a function of the wait time t_{wait} between the write and readout pulse, we can measure the phonon lifetime $T_1 = 1/\Gamma_m$. The measured energy in the readout pulse is proportional to $\int |\alpha_2^{out}|^2 dt \propto \exp(-\Gamma_m t_{wait})$ due to the linearity of the equations of motion. Hence the measured energy follows the same exponential decay as the phonon energy. We will show how to relate the measured detector voltage to the retrieved phonon energy in section 3.3.2, after the description of the experimental setup.

3.3 Experimental methods and setup

In this section we first give an overview of how the experiment was performed. Then we present the experimental setup and mentioned the used equipment. Finally we derive an expression relating the balanced detector output to the retrieved phonon energy.

3.3.1 Overview and experimental setup

Overview

We work with two optical modes $\hat{a}_{1/2}$ with frequencies $\omega_{1/2}$ and $\omega_2 - \omega_1 \simeq \Omega_B$, where the Brillouin frequency Ω_B is close to the frequency of the main mechanical mode of interest. In order to store an optical pulse as an excitation in the phonon mode of the HBAR, we pump mode \hat{a}_1 of the optical cavity and at the same time send an additional weaker optical pulse to mode \hat{a}_2 . This additional write pulse is a sideband of the pump light, created with an intensity modulator (IM) and lasts 30 μ s. During the write pulse, a coherent phonon state is created, as previously described in section 3.2.1. The length of the write pulse was chosen to be much longer than the phonon lifetime of around $\sim 3 \mu$ s in order to create a large phonon population.

After the write pulse we block the laser and wait for some variable time t_{wait} after which we readout our mechanical excitation with a readout pulse. During this pulse, the phonons that remain after t_{wait} scatter with pump photons in mode \hat{a}_1 , creating a population in the higher frequency optical mode \hat{a}_2 as described in section 3.2.2. The readout pulse consists of pumping optical mode \hat{a}_1 for 15 μ s, where we chose the length of the pulse such that most of the phonon energy has either decayed internally or been swapped into the optical mode \hat{a}_2 by the end of the readout pulse. We then detect the optical power that leaves the cavity from mode \hat{a}_2 via balanced heterodyne detection. The measured optical power at frequency ω_2 can then be related to the retrieved phonon energy (see section 3.3.2).

By varying the wait time t_{wait} between write and readout pulse, we can measure the T_1 life time of the mechanical excitation. We further vary the frequency detuning of the sideband in the write pulse to probe the frequency response of the mechanical mode. Since the mechanical FSR is smaller than the optical mode linewidth, we also sweep the frequency

ⁱⁱⁱFor an efficient transduction process, one wants to be in the $C \gg 1$ regime. For our transduction experiment, we aim for $C = 20$.

of the sideband over an entire mechanical FSR in order to see the response of a second mechanical mode.

Experimental setup

The experimental setup can be seen in figure 3.4. The output of a tunable laser is split into two arms, called the pump arm and the local oscillator (LO) arm.

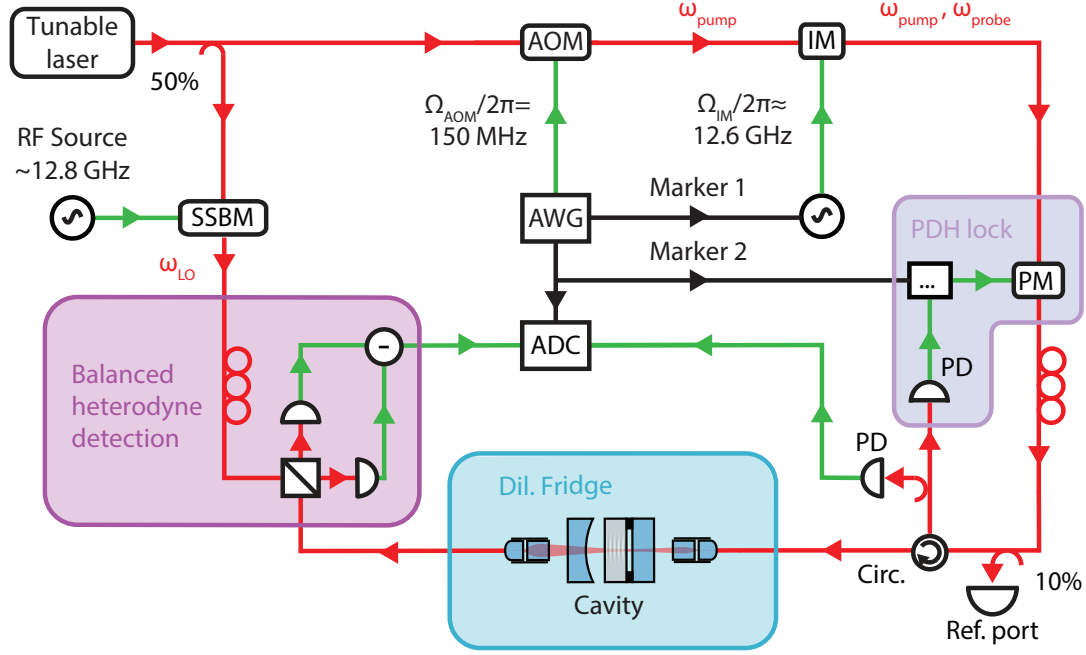


Figure 3.4: Experimental setup for the optomechanical storage and retrieval experiment. Red paths are optical fibers, green paths are electrical analog lines and black paths are electrical digital lines. See main text for a detailed description of the setup.

In the pump arm, the light is pulsed by an acousto optic modulator (AOM), driven by the amplified output of an arbitrary waveform generator (AWG) at frequency $\Omega_{AOM}/2\pi = 150$ MHz. During the write pulse the light is additionally intensity modulated by an intensity modulator (IM) driven by an RF source at frequency Ω_{IM} , close to the Brillouin frequency $\Omega_B/2\pi \approx 12.64$ GHz. This creates sidebands detuned from the pump frequency by the input RF frequency Ω_{IM} . The frequency Ω_{IM} is swept to probe the response of the various mechanical modes. The pump light is then locked to mode \hat{a}_1 of the cryogenic cavity, which is located inside the fridge. Laser locking is performed by a PDH lock (see chapter 2). Similar to the experiment in chapter 2, before entering the cavity, we record the optical power at the reference port and polarize the light to be parallel with one of the crystal axes, due to the birefringence of the crystal. The reflected optical power passes through a circulator and is measured by a photodiode (PD), the output of which we record with an analog to digital converter (ADC) in order to monitor the quality of the lock. In the LO arm, a single sideband modulator (SSBM), driven by an RF source at frequency ~ 12.8 GHz, shifts the frequency of the LO light to ω_{LO} . This light is used as the LO for balanced heterodyne detection of light at frequency $\omega_1 + \Omega_m$ stemming from the readout of phonons: The electric fields of the fridge output and the LO arm are added via a 50:50 beamsplitter, whose outputs are sent to an auto-balanced detector. The auto-balanced detector consists of two high bandwidth PDs of which the outputs are subtracted from each other. The auto-balanced detector uses an internal control loop to balance the power levels

of the outputs of the two PDs to compensate for any beamsplitter or PD gain imbalance. The output of the auto-balanced detector is recorded by the ADC. By postprocessing this data, we can extract the measured optical power coming from mode \hat{a}_2 , see section 3.3.2 for more details.

We choose the frequency of the LO arm such that the difference between ω_{LO} and $\omega_1 + \Omega_m$ is well resolvable with the sampling rate of 100 Megasamples per second we choose for the ADC. We drive the SSBM at 12.78 GHz, this results in $(\omega_{LO} - \omega_1 - \Omega_m)/2\pi \simeq 14$ MHz. Note that the AOM shifts the pump arm frequency by Ω_{AOM} . We further adjust the polarization of the LO arm to match the polarization of the fridge output to allow for maximal interference between the signal and the LO.

The timing of the various components is coordinated by the AWG using the RF output as well as two marker channels. The pulse sequence executed in every run of the experiment is seen in figure 3.5. Note that for this experiment we use both PID controllers of the laser controller and do not use the Red Pitaya, which is why only one marker channel is needed for the PDH lock. This is mainly due to technical issues that arose with the AWG when using three separate marker channels. Due to the short wait times between pulses (we can run the experiment with a repetition rate of around 10 kHz), the lock performs consistently in this experiment even when using the slower PID controllers of the laser controller.

When starting the experiment, we first send continuous wave (cw) light to the cavity to engage the lock before starting the pulse sequences, since starting the lock during pulsed operation did not work.




Figure 3.5: Pulse sequence executed by the AWG during one run of the optomechanical storage and retrieval experiment. The RF output drives the AOM, creating the optical pulses. Marker channel 1 is responsible for the timing of the IM and marker channel 2 triggers the data acquisition and puts the PDH lock on hold between pulses.

Used equipment

The pump arm was intensity modulated by an Optilabs IM-1550-20-P, driven by a Lab Brick LMS-163 by Vaunix with the pulse modulation option. The SSBM was an iXblue MXIQER-LN-30, driven by a Keysight MXG N5183B signal generator. The heterodyne detection was done by a Thorlabs PDB570C auto-balanced detector. The rest of the setup is identical to the one used in chapter 2 (see section 2.3).

3.3.2 Calculating retrieved phonon energy from the balanced detector signal

After the description of the experimental setup in section 3.3.1, we show how to relate the measured balanced detector voltage to the retrieved phonon energy.

The light from the output fiber of the cavity is mixed with light from the LO arm with frequency ω_{LO} and amplitude α_{LO} in a 2x2 50:50 beamsplitter, whose outputs are fed into a balanced detector. We now calculate the output voltage V_{bal} of this detector from α_2^{out} and relate the frequency spectrum of the output to the optomechanically retrieved energy E_r . We now again assume the amplitudes α_{LO} and α_2^{out} to be complex, as the pump and LO arm are not phase locked and the relative phase does not remain stable between

subsequent runs of the experiment.

The incident field amplitude α_{det1} on the first detector is given by

$$\alpha_{det1} = \frac{1}{\sqrt{2}} \left(\alpha_{LO} e^{-i\omega_{LO}t} + \sqrt{\eta_{out}} \alpha_2^{out} e^{-i(\omega_1 + \Omega_m)t} \right), \quad (3.24)$$

where η_{out} accounts for the efficiency of the GRIN lens that feeds the output mode of the cavity back into the optical fiber as well as further losses in the fibers and optical components between the cavity and the detector. The output voltage of the detector is proportional to the incident optical power, which for the first detector is given by

$$P_{det1} = \frac{1}{2} \left(\hbar\omega_{LO} |\alpha_{LO}|^2 + \eta_{out} \hbar(\omega_1 + \Omega_m) |\alpha_2^{out}|^2 \right) + \frac{\sqrt{\eta_{out}}}{2} \hbar \sqrt{\omega_{LO}(\omega_1 + \Omega_m)} \left(\alpha_{LO} (\alpha_2^{out})^* e^{-i(\omega_{LO} - \omega_1 - \Omega_m)t} + \alpha_{LO}^* \alpha_2^{out} e^{i(\omega_{LO} - \omega_1 - \Omega_m)t} \right), \quad (3.25)$$

where we have neglected the terms oscillating with twice the optical frequencies as these are outside of the detector bandwidth. Defining $\omega_{det} \equiv \omega_1 + \Omega_m - \omega_{LO}$ and writing $\alpha_{LO} = |\alpha_{LO}| e^{i\varphi_{LO}}$ and $\alpha_{2,det} = |\alpha_2^{out}| e^{i\varphi_{2,out}}$ we get

$$P_{det1} = \frac{1}{2} \left(\hbar\omega_{LO} |\alpha_{LO}|^2 + \eta_{out} \hbar(\omega_1 + \Omega_m) |\alpha_2^{out}|^2 \right) + \sqrt{\eta_{out}} \hbar \sqrt{\omega_{LO}(\omega_1 + \Omega_m)} |\alpha_{LO} \alpha_2^{out}| \cos(\omega_{det}t + \varphi_{2,out} - \varphi_{LO}). \quad (3.26)$$

Similarly, the incident field amplitude α_{det2} on the second detector is given by

$$\alpha_{det2} = \frac{1}{\sqrt{2}} \left(\alpha_{LO} e^{-i\omega_{LO}t} - \sqrt{\eta_{out}} \alpha_2^{out} e^{-i(\omega_1 + \Omega_m)t} \right). \quad (3.27)$$

The relative minus sign comes from the phase shift experienced upon reflection within the beamsplitter. The incident optical power on the second detector is therefore

$$P_{det2} = \frac{1}{2} \left(\hbar\omega_{LO} |\alpha_{LO}|^2 + \eta_{out} \hbar(\omega_1 + \Omega_m) |\alpha_2^{out}|^2 \right) - \sqrt{\eta_{out}} \hbar \sqrt{\omega_{LO}(\omega_1 + \Omega_m)} |\alpha_{LO} \alpha_2^{out}| \cos(\omega_{det}t + \varphi_{2,out} - \varphi_{LO}). \quad (3.28)$$

The voltage output by one of the detectors of the balanced detector for optical input power P is given by

$$V = GP, \quad (3.29)$$

where G is the gain of the detector, specified in volts per optical input power. The total output of the balanced detector is given by the difference between the output voltages of the two detectors. In our case this difference is given by

$$V_{det} = V_{det1} - V_{det2} = 2G\hbar \sqrt{\omega_{LO}(\omega_1 + \Omega_m)\eta_{out}} |\alpha_{LO} \alpha_2^{out}| \cos(\omega_{det}t + \varphi_{2,out} - \varphi_{LO}). \quad (3.30)$$

During the measurement time t_{meas} we record the voltage of the balance detector with sampling rate R and store the absolute value $|V_{det}(\omega_n)|$ of the discrete Fourier transform (DFT) of this data, which is given by

$$|V_{det}(\omega_n)| = \left| \sum_{m=-(N-1)/2}^{(N-1)/2} V_{det}(t_m) e^{-i\omega_n t_m} \right|, \quad (3.31)$$

where $\omega_n = n \frac{2\pi}{t_{meas}}$ and $N = t_{meas}R$ is the number of recorded samples, which for simplic-

ity we assume to be odd.

Starting with Parseval's theorem for the DFT we find

$$\begin{aligned}
 \frac{1}{N} \sum_n |V_{det}(\omega_n)|^2 &= \sum_n |V_{det}(t_n)|^2 \\
 &\approx R \int_0^{t_{meas}} |V_{det}(t)|^2 dt \\
 &\approx R \eta_{out} (G\hbar)^2 \omega_{LO} (\omega_1 + \Omega_m) |\alpha_{LO}|^2 \int_0^{t_{meas}} |\alpha_2^{out}(t)|^2 dt \quad (3.32)
 \end{aligned}$$

In the first step we replaced the sum by an integral, since the sampling rate R is much larger than $\omega_{det}/2\pi$. In the second step we inserted the expression 3.30 and made the approximation $\int |\alpha_2^{out}(t)|^2 \cos^2(\omega_{det}t + \varphi) dt \approx 1/2 \int |\alpha_2^{out}(t)|^2 dt$, which is justified as the oscillation in the cosine is much faster than the timescale on which $|\alpha_2^{out}(t)|$ changes appreciably and $t_{meas} \gg 2\pi/\omega_{det}$.

We can now express equation (3.32) in terms of the retrieved phonon energy E_r contained in the output mode \hat{a}_2^{out} of the cavity. The photons in the output mode \hat{a}_2^{out} have energy $\hbar(\omega_1 + \Omega_m)$. Of this energy, $\hbar\Omega_m$ comes from the the mechanical excitation. Thus the retrieved phonon energy is given by integrating $\hbar\Omega_m |\alpha_2^{out}(t)|^2$ from the start of the readout pulse at $t = 0$ for the measurement time t_{meas}

$$E_r = \int_0^{t_{meas}} \hbar\Omega_m |\alpha_2^{out}(t)|^2 dt. \quad (3.33)$$

Using this expression for the retrieved phonon energy, equation (3.32) becomes

$$\frac{1}{N} \sum_n |V_{det}(\omega_n)|^2 = \frac{R \eta_{out} (G\hbar)^2 \omega_{LO} (\omega_1 + \Omega_m) |\alpha_{LO}|^2}{\hbar\Omega_m} E_r. \quad (3.34)$$

As a last step we express $|\alpha_{LO}|^2$ through the optical power in the LO arm via

$$|\alpha_{LO}|^2 = P_{LO}/\hbar\omega_{LO}. \quad (3.35)$$

Thus we find the following expression for E_r in terms of the DFT of the balanced detector voltage:

$$E_r = \frac{\Omega_m}{\omega_1 + \Omega_m} \left(\frac{1}{RN\eta_{out}G^2P_{LO}} \sum_n |V_{det}(\omega_n)|^2 \right). \quad (3.36)$$

The expression in the parentheses is the retrieved energy from the optical mode \hat{a}_2^{out} , the prefactor $\Omega_m/(\omega_1 + \Omega_m)$ scales this energy to be the retrieved phonon energy.

3.4 Results

In this section we present our experimental results. First in subsection 3.4.1 we show an optomechanically induced transparency (OMIT) measurement, used to verify that the optical mode spacing is close to the Brillouin frequency Ω_B and to precisely determine the frequencies of the mechanical modes. Then in subsection 3.4.2, we present experimental data from sweeping Ω_{IM} over the main mechanical mode at frequency Ω_1 for different wait times t_{wait} . In the next subsection 3.4.3, we compare these results with numerical estimates made with the model described in the previous section. Finally in subsection 3.4.4, we extend the sweep of Ω_{IM} to include a second mechanical mode at frequency Ω_2 .

3.4.1 OMIT spectrum

After adjusting the distance between the mirrors in the cavity and the laser wavelength such that we have an optical mode pair separated by the Brillouin frequency and at the displacement insensitive point, we measure the Brillouin frequency for the current conditions. To experimentally determine the mechanical mode frequencies Ω_m of our HBAR modes, we record an optomechanically induced transparency (OMIT) spectrum. First, we lock the pump laser to the lower frequency mode of the optical mode pair. Then we drive the intensity modulator with the output (port 1) of a vector network analyzer (VNA) to create a weak probe tone and feed the output of a 25 GHz-bandwidth photodiode measuring the transmitted optical power back into the VNA (port 2, see [1] for details). This allows us to record the scattering parameter S_{21} as function of frequency, which is proportional to the beat note between the transmitted pump and probe tone.

The OMIT spectrum, seen in figure 3.6, shows a broad optical resonance with two sharp dips located at $\Omega_1/2\pi = 12.643'930$ GHz and $\Omega_2 = \Omega_1 + \Omega_{FSR}$ with $\Omega_{FSR}/2\pi = 625$ kHz. These sharp dips correspond to two mechanical modes. When the VNA frequency matches the frequency of a phonon mode, pump light is upconverted via Brillouin scattering and destructively interferes with the probe light, leading to a dip in S_{21} .




Figure 3.6: OMIT spectrum recorded to verify the location of the mechanical modes and the mechanical FSR.

Seeing that the optical mode spacing is in fact equal to the Brillouin frequency and that we have observed optomechanical coupling between the optical modes and two mechanical modes of the HBAR, we are now ready to perform the storage and readout experiment.

3.4.2 Probing the main resonance

We perform the described optomechanical storage and readout experiment for 40 values of $(\Omega_{IM} - \Omega_1)/2\pi$ ranging from -200 kHz to 190 kHz and 15 values of t_{wait} from $1 \mu\text{s}$ to

8.5 μs . Shorter values of t_{wait} were not possible with the way the data acquisition was set up but the measurement scripts can easily be adapted to allow for this in the future. The experiment was repeated 1500 times for each value of Ω_{IM} and t_{wait} .

During each read pulse we record the discrete Fourier transform (DFT) $V_{\text{det}}(\omega)$ of the output of the balance detector. This spectrum exhibits a clear peak at $\omega = \omega_{\text{det}}$, where $\omega_{\text{det}} \equiv \omega_1 + \Omega_1 - \omega_{LO}$ (see section 3.3.2). This is due to photons from the pump mode at frequency ω_1 getting upconverted to photons at frequency $\omega_1 + \Omega_1$ via anti-Stokes scattering with the phonons in the HBAR (see figure 3.7). The height of this peak depends on both the detuning $\Omega_{IM} - \Omega_1$, which determines how many phonons are created during the write pulse, as well as t_{wait} , which determines how many phonons are still in the HBAR at the start of the read pulse (see section 3.2). The height of the observed peak becomes smaller for larger t_{wait} , due to the smaller remaining phonon population after longer wait times between pulses. All measurements were taken with the same optical power in the read and write pulse (-17.9 dBm CW power measured at the reference port). The error-bars on the DFT amplitude (smaller than the markers in the plot) were taken to be the standard deviation of the DFT amplitude away from the resonance after averaging over multiple runs of the experiment.




Figure 3.7: Amplitude of the DFT of balanced detector signal for different values of t_{wait} with $\Omega_{IM} = \Omega_1$. A constant noise floor has been subtracted from all traces.

We then calculate $\sum_n |V_{\text{det}}(\omega_n)|^2$, where we sum over the frequencies close to the observed peak. Note that there is a second, identical peak at frequency $-(\omega_1 + \Omega_1)$, that we do not show in our plots. We only sum over positive frequencies and then multiply the result by two, to account for the negative frequencies. As shown in section 3.3.2, this quantity is proportional to the retrieved phonon energy, so we refer to it as "retrieved energy [a.u.]" in the following plots.

The behavior of the retrieved energy as a function of the frequency detuning as well as t_{wait} are shown in figure 3.8. The response of the mechanical resonance is not symmetric in the frequency detuning as seen most clearly in figure 3.8b for larger values of t_{wait} . In fact already for the shortest t_{wait} of 1 μs , the retrieved energy as a function of the drive detuning does not follow the Lorentzian profile expected from driving a single mechanical eigenmode coupled to a phonon bath (see figure 3.9). In fact, as we further discuss in section 3.5, we are not driving a single mechanical eigenmode of the HBAR but a superposition of eigenmodes that are close enough in frequency, such that we drive them simultaneously.

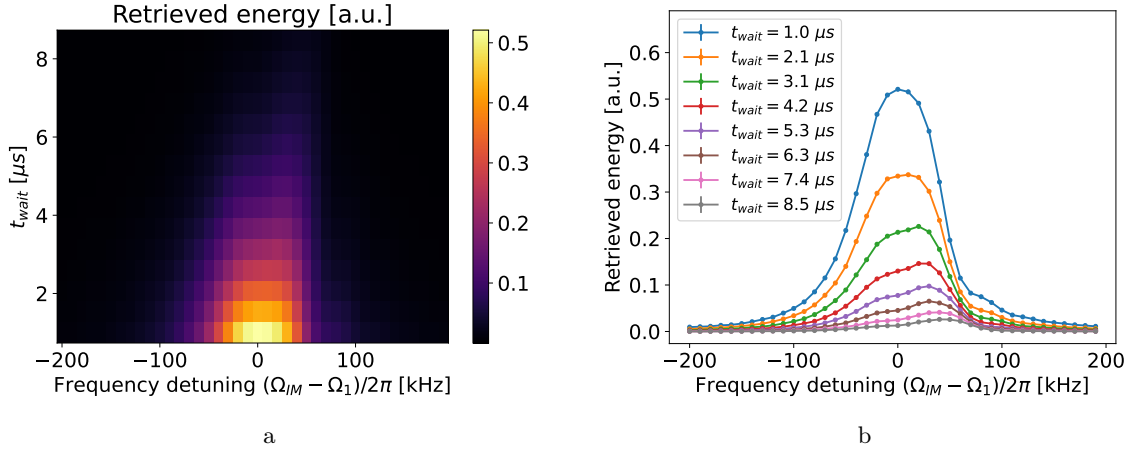


Figure 3.8: Retrieved phonon energy for various wait times t_{wait} between the read and write pulse as well as different IM drive frequencies Ω_{IM} . We can see the response of a phonon mode with frequency Ω_1 .

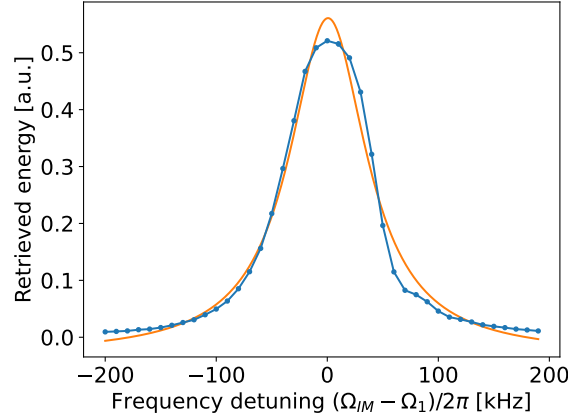


Figure 3.9: Retrieved phonon energy as a function of the IM drive frequency detuning $(\Omega_{IM} - \Omega_1)$ for $t_{wait} = 1 \mu$ s (blue points). We also plot a Lorentzian best fit to the data (orange line), to illustrate the discrepancy between the two.

Looking at vertical slices of the plot from figure 3.8a, i.e. retrieved energy as a function of t_{wait} for fixed frequency detuning, we can analyze the decay of the mechanical mode (see figure 3.10). Ideally, this decay follows an exponential $\propto \exp(-t_{wait}/T_1)$, characterized by the life time $T_1 = 1/\Gamma_m$ of the phonon mode as described in section 3.2.2. However the observed decay only approximately follows an exponential decay best fitted to the data and we find that the observed T_1 is frequency dependent (figures 3.10 and 3.11). Note that the experiments performed in [1] find an intrinsic mechanical linewidth of 54.5 ± 0.4 kHz, corresponding to a T_1 lifetime of 2.9μ s.

Generally, we observe that the retrieved energy follows the exponential for the first $\sim 6 \mu$ s before dropping off below the exponential fit. However, for frequency detunings around 70 to 90 kHz we observe the opposite (see figure 3.10d). For larger t_{wait} , the retrieved energy curve levels out a bit and lies above the best fit exponential decay. For larger frequency detunings we again observe the initial behavior of the retrieved energy curve dropping off (see figure 3.10e). The errorbars on the data points are propagated from the uncertainty on the DFT averages, the uncertainty on the shown T_1 lifetimes are extracted from the fits.

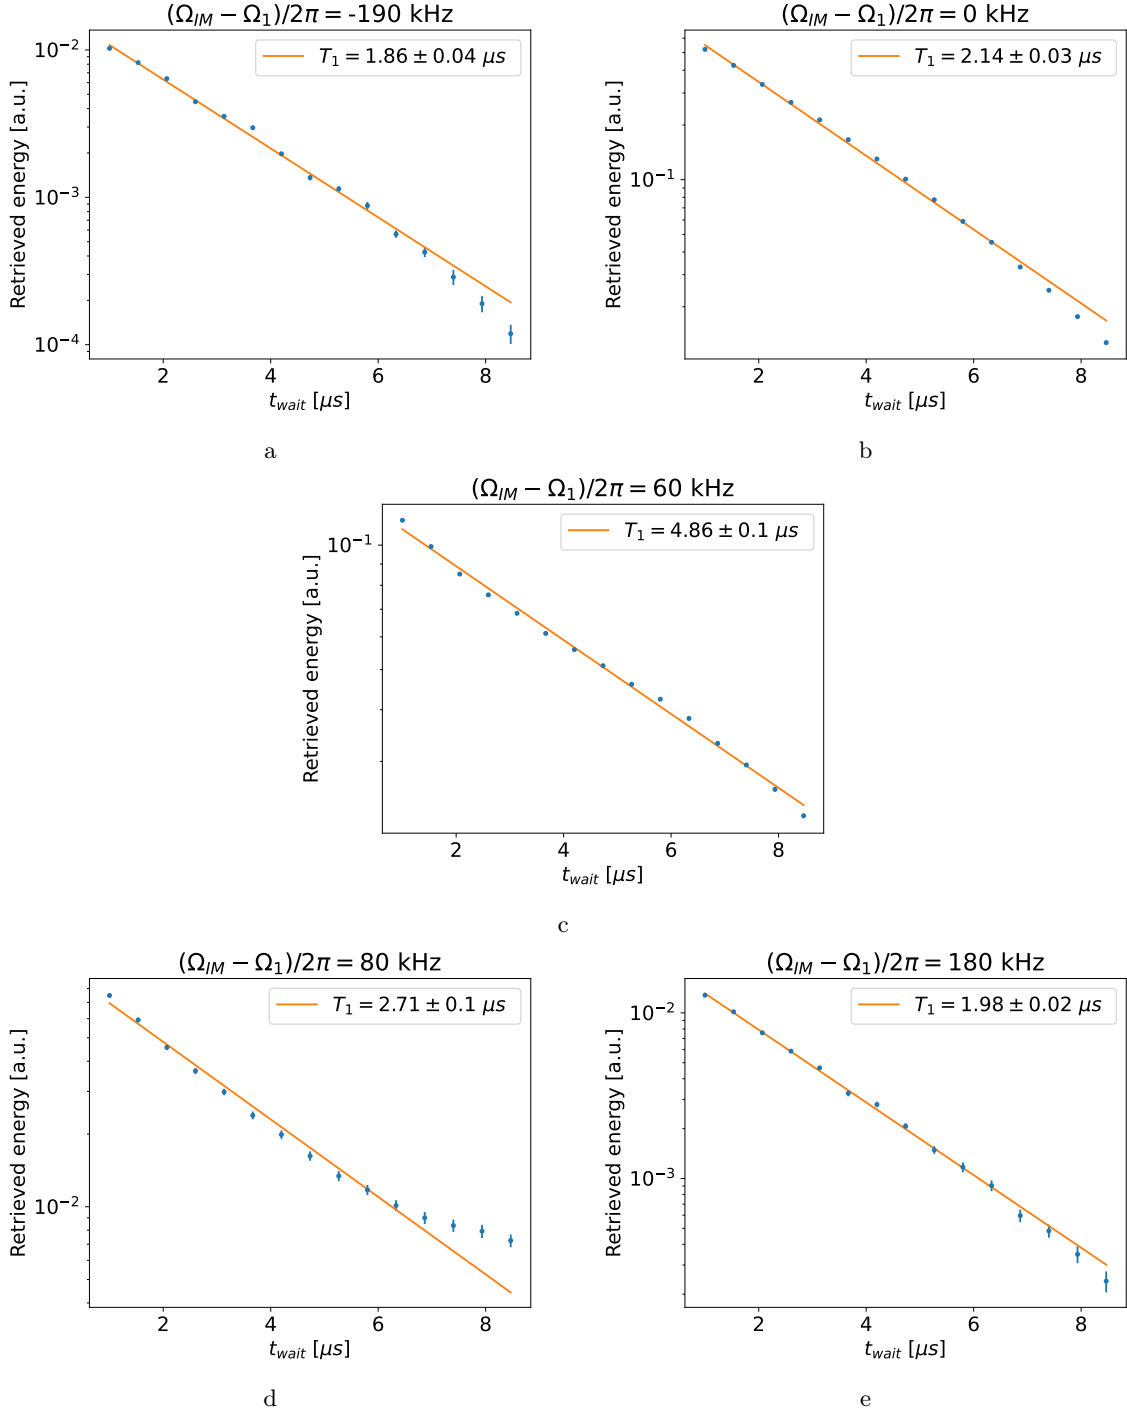


Figure 3.10: Retrieved phonon energy (blue points), as a function of t_{wait} for different values of Ω_{IM} illustrating the different decay behaviors depending on Ω_{IM} . We also plot exponential decays $\propto \exp(-t_{wait}/T_1)$ best fit to the data (orange line).

In figure 3.11 we plot the phonon T_1 lifetime as a function of the frequency detuning. Contrary to what one expects from a single mechanical eigenmode, we find that the T_1 is highly non uniform and peaks at a frequency detuning of 60 kHz (figure 3.10c) with a T_1 more than twice as large as when measured when driving the mechanical mode on resonance. One can also see evidence of this longer lifetime in the evolution of the traces in figure 3.8b. Outside of the peak in T_1 , we measure T_1 lifetimes on the order of 2 μs , which would correspond to a mechanical linewidth of around 80 kHz - quite a bit larger than the 54.5 ± 0.4 kHz that were found in [1]. It is not clear how to relate these varying decay times to a single mode linewidth that we can compare to the linewidth found with

the methods in [1].

Note that optomechanical backaction cannot be the reason for the enhanced or reduced damping, since during the free evolution between write and readout pulses, no light is populating the cavity. Most likely, the varying T_1 times, which are a symptom of the asymmetric response of the mechanical mode, have to do with the fact that we are probing a superposition of eigenmodes of the HBAR. As previously mentioned we further discuss this aspect in section 3.5.




Figure 3.11: Dependence of phonon T_1 on the frequency detuning $(\Omega_{IM} - \Omega_1)$ of the drive.

Here we need to mention a potential calibration issue with the frequency axis in all of these plots. As described in section 3.4.1, we recorded an OMIT spectrum to experimentally determine the mechanical mode frequencies Ω_m . For the OMIT spectrum, the IM was driven by a VNA. For the optomechanical storage and retrieval measurements however we used a separate RF source to drive the intensity modulator. This means that the location of zero frequency detuning in the plots in this section relies on the comparison of two 12.6 GHz signals down to around 10 kHz, implying a relative accuracy of 10^{-6} when comparing these devices, which we cannot guarantee, as we did not use a common clock source for these measurements.

The upshot is that we can conclude that the maximum in retrieved energy (observed at 0 frequency detuning) and the maximum in phonon T_1 (observed at a frequency detuning of 60 kHz) happen 60 kHz detuned from each other, however we cannot make an accurate statement about where these features lie with respect to the dip in the OMIT spectrum.

3.4.3 Comparison with numerical estimates

In this section we list our typical experimental parameters and use them to compare the retrieved phonon energy predicted by the model in sections 3.2.1 and 3.2.2 with the retrieved energy as calculated from the experimental results using equation (3.36) found in section 3.3.2. Note that our theoretical model only accounts for one phonon mode which is why we make the comparison using the data set from section 3.4.2.

A list of typical experimental parameters as well as the various relevant parameters of our setup can be found in table 3.1. These are the parameters dictated by the setup itself and or settings of various devices. Note that we only want to do a rough comparison, which is why we do not give many significant digits the values in the table and rely on order of magnitude similarity for numbers like the single-photon coupling rate.

P_{RP}	-17.9 dBm	ω_1	$2\pi \cdot 193.4$ THz	Γ_m	$2\pi \cdot 55$ kHz
P_{LO}	-8.7 dBm	ω_{write}	$\omega_1 + \Omega_1$	g_0	$2\pi \cdot 8.4$ Hz
η_{in}	4.8	Ω_1	$2\pi \cdot 12.64$ GHz	G	2500 V/W
η_{out}	0.2	Δ_1	200 kHz	R	100 Megasamples/s
η_{IM}	0.02	κ	$2\pi \cdot 2.4$ MHz	N	1401 samples
		κ_{ext}	$2\pi \cdot 1.2$ MHz		

Table 3.1: External parameters for the optomechanical storage and retrieval experiment. Parameters of the cavity as well as the single photon coupling rate g_0 and the mechanical linewidth Γ_m are taken from [1]. For the input coupling efficiency η_{in} we take the value found in [1] and modify it to account for an average lock level of 10 %. η_{IM} was measured using an IM drive power of 13.5 dBm, which is what we use for our experiments.

In table 3.2 we list the effective coupling rates, the initial phonon amplitude β_0 and the optomechanical cooperativity C_{read} during the readout pulse as calculated using parameters from table 3.1. We also list the retrieved phonon energy $E_{r,theory}$ for $t_{wait} = 1 \mu s$ and $t_{read} = 15 \mu s$. To do this we numerically solved the equations of motion for the amplitudes α_2^{cav} and β (see equations (3.20) and (3.21)) and then calculated $E_{r,theory}$ with equation (3.33), again using parameters from table 3.1.

g_{write}	$2\pi \cdot 93$ MHz
g_{read}	$2\pi \cdot 75$ kHz
$ \beta_0 $	2830
C_{read}	0.17
$E_{r,theory}$	$7.65 \cdot 10^{-19}$ J $\approx 5.74 \times 10^5 \cdot \hbar\Omega_m$

Table 3.2: Calculated parameters using parameters from table 3.1, $t_{wait} = 1 \mu s$ and $t_{read} = 15 \mu s$

As can be seen in table 3.2, from our theoretical model we expect to retrieve the energy of roughly 5.74×10^5 phonons. For comparison, at the start of the readout pulse for $t_{wait} = 1 \mu s$ we expect there to be roughly $|\beta_0|^2 \exp(-\Gamma_m t_{wait}) \approx 5.669 \times 10^6$ phonons. The large difference between these numbers is due to the low cooperativity $C_{read} \ll 1$ during the readout pulse, meaning that it is much more likely for a phonon to decay rather than swap its energy into a photon. Note that we could not go to higher cooperativity due to laser power limitations.

Finally, we use our measurement results to calculate a measured value of the retrieved phonon energy using equation (3.36).

For $t_{wait} = 1 \mu s$ and $\Omega_{IM} = \Omega_1$ we measured $\sum_n |V_{det}(\omega_n)|^2 = 0.52$ V² (see figures 3.7 and 3.8b). Plugging this and the required values from table 3.1 into equation (3.36), we find

$$E_{r,measured} = 2.3 \cdot 10^{-19} \text{ J} \approx 1.71 \times 10^5 \cdot \hbar\Omega_m,$$

corresponding to 30 % of the predicted value $E_{r,theory} \approx 5.74 \times 10^5 \cdot \hbar\Omega_m$.

We mainly attribute this discrepancy to errors in estimating the parameters. While we do not believe any single parameter estimate to be off by a factor of 3, it is entirely possible that considerable errors in multiple parameter estimates lead to the theoretical and measured value differing by this much. Notably the single photon coupling rate g_0 varies from mode to mode (see [1]) and the value for the IM efficiency η_{IM} has an errorbar

comparable to its magnitude.

3.4.4 Probing two mechanical resonances

Next we scan the frequency detuning of Ω_{IM} from the Brillouin frequency Ω_B over a wider range than before, in order to see the second mechanical resonance at frequency $\Omega_2 = \Omega_1 + \Omega_{FSR}$ (see figure 3.6). Recall that the mechanical FSR is roughly $\Omega_{FSR}/2\pi \simeq 625$ kHz. Therefore we sweep the frequency detuning from -110 kHz to 748 kHz to see the response of both the mechanical resonance at frequency Ω_1 as well as the mechanical resonance at frequency Ω_2 .

As Ω_{IM} gets close to Ω_2 , we see a second peak appear in the spectrum of the balanced detector output at $\omega \simeq \omega_{det} + \Omega_{FSR}$ (figure 3.12). The frequency resolution of the DFT is merely 71 kHz, resulting in the peaks in the spectrum being separated by ~ 640 kHz. Nevertheless, the second peak is a clear indication that we are creating and reading out phonons in the higher frequency mode at Ω_2 .




Figure 3.12: DFT amplitude for various frequency detunings ($\Omega_{IM} - \Omega_1$) for fixed t_{wait} .

In figure 3.13 we show the retrieved energy as a function of the frequency detuning $\Omega_{IM} - \Omega_1$ as well as the wait time t_{wait} . We can clearly see the presence of the second mechanical mode at $\Omega_{IM} - \Omega_1 = \Omega_{FSR}$, although we get a weaker response from the second mode. It is detuned a bit further from the optical mode spacing and exhibits a slightly lower coupling, as can be seen in the OMIT spectrum (figure 3.6). Note that the optical sideband generated during the write pulse at frequency $\omega_1 + \Omega_{IM}$ is still resonant with the optical cavity even for $\Omega_{IM} \simeq \Omega_1 + \Omega_{FSR}$, as the linewidth $\kappa/2\pi = 2.4$ MHz of the optical resonance is larger than the mechanical FSR $\Omega_{FSR} = 625$ kHz.

The evolution of the retrieved energy for increasing t_{wait} shows the same asymmetric behavior seen before in figure 3.8 in both mechanical modes.

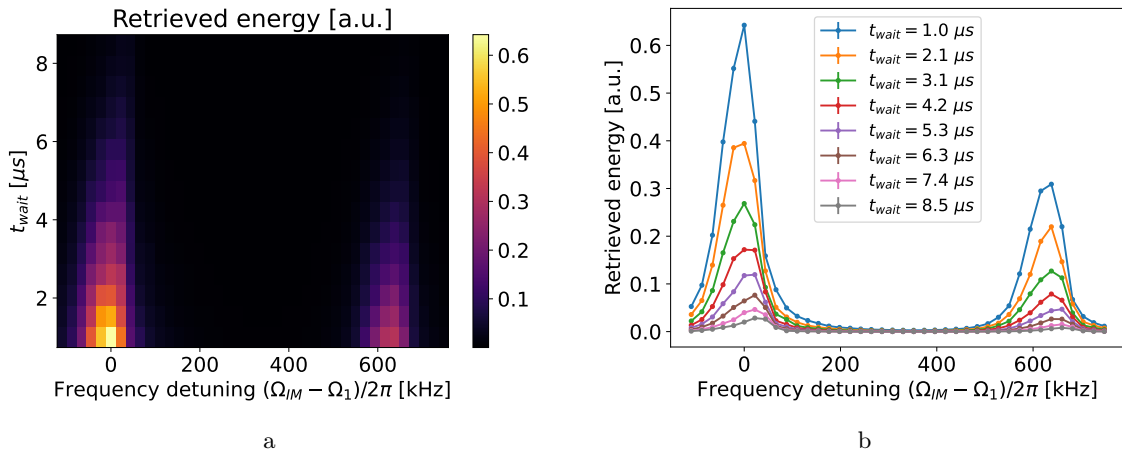


Figure 3.13: Retrieved phonon energy for various wait times t_{wait} between the read and write pulse as well as different IM drive frequencies Ω_{IM} . We can see phonon modes located at frequencies Ω_1 and Ω_2 .

We perform the same analysis of the retrieved energy as a function of t_{wait} for fixed frequency detuning as we did in section 3.4.2. We fit an exponential decay $\propto \exp(-t_{wait}/T_1)$ to the vertical slices of figure 3.13a and plot the retrieved T_1 as a function of the frequency detuning in figure 3.14. We again see that the phonon lifetime T_1 peaks when the mode is driven roughly 60 kHz above the phonon frequency. The second phonon mode exhibits the same behavior.

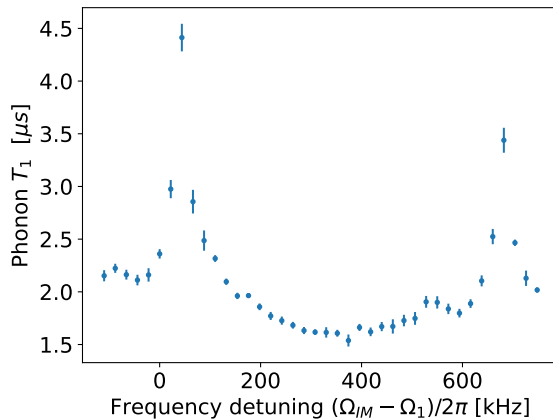


Figure 3.14: Phonon lifetime T_1 as a function of the frequency detuning $(\Omega_{IM} - \Omega_1)$.

Finally, we can also separate the retrieved phonon energy into energy retrieved from phonon mode 1 at frequency Ω_1 and energy retrieved from phonon mode 2 at frequency Ω_2 . Recall that the retrieved energy is proportional to $\sum_n |V_{det}(\omega_n)|^2$, where $V_{det}(\omega)$ is the DFT of the the auto-balanced detector output. We split this sum in half and assign the energy coming from terms in $\sum_n |V_{det}(\omega_n)|^2$ with $\omega < \Omega_1 + \Omega_{FSR}/2$ to mode 1 and the energy coming from terms in $\sum_n |V_{det}(\omega_n)|^2$ with $\omega > \Omega_1 + \Omega_{FSR}/2$ to mode 2. In figure 3.15 we plot the phonon energy retrieved from mode 1 and 2 respectively as a function of the frequency detuning $(\Omega_{IM} - \Omega_B)$ for $t_{wait} = 1 \mu$ s. We can see that this scheme works fairly well for differentiating between the two modes. However we do still see a bump in the retrieved energy from mode 2 at zero frequency detuning. This is due to the "tails" of the peak at ω_{det} in the DFT spectrum reaching out into the upper half of the considered spectrum (see figure 3.12), meaning that we assign some of the retrieved energy from phonon mode 1 to phonon mode 2. The width of the peaks in the DFT spectrum can be explained by the linewidth of the phonon mode and the finite duration of the readout

pulse causing a broadening in the detected frequency.

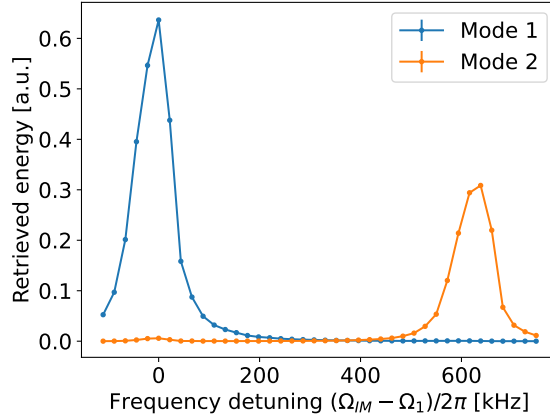


Figure 3.15: Retrieved phonon energy separated by the two phonon modes we excite during the write pulse.

3.5 Discussion

In this section we briefly recap the results of this chapter and discuss some of the unexpected and unexplained features thereof. We further propose an avenue with which one might be able to better theoretically describe the results of our experiments.

With our experimental setup we were able to perform the proposed optomechanical storage and readout experiment and achieve a sufficient signal to noise ratio to get clear signals in the explored parameter regimes. We were able to selectively store an optical pulse in one of two mechanical modes and retrieve it after multiple microseconds. We probed the response of two mechanical modes to being driven with different frequency detunings and characterized their temporal decay. The measured energy retrieved out of the generated mechanical excitations matches within an order of magnitude what we expected to measure based on the simple model in sections 3.2.1 and 3.2.2, based on the semiclassical treatment of two optical modes coupled to a single phonon mode.

However, we did encounter behavior not captured by this model: The observed T_1 lifetimes of the mechanical modes depend strongly on the frequency detuning with which they were driven at (e.g. see figure 3.14) and we see systematic deviations from exponential decay in the T_1 plots. This is also related to the fact that the response (i.e. the retrieved phonon energy) of the mechanical modes is not symmetric in the frequency detuning of the drive and generally does not follow a Lorentzian profile for the used values of t_{wait} . Probably the response would look more like a Lorentzian for shorter values of t_{wait} but for the reason mentioned below, we do not necessarily expect a Lorentzian response even for $t_{wait} = 0$.

The explanation we currently deem the most likely for these effects mainly relies on the fact that the mechanical resonance we see is not a single mechanical eigenmode of the HBAR. Each mechanical mode we see consists of multiple eigenmodes with Bessel function mode profiles. They are not frequency degenerate due to the finite extent of the crystal, however they are close enough in frequency, such that we always excite a superposition of these eigenmodes in our experiment (see SI of [12]). In the following, we sketch an argument that may be able to explain the mentioned properties of our results, although more theoretical work needs to be done to completely verify this.

Our optical modes apply a forcing function to the HBAR with a Gaussian mode profile, meaning that the mechanical excitation we create is a superposition of Bessel modes with a Gaussian mode profile that matches the optical mode profile. This superposition is subject to diffraction loss stemming from the discrepancy in longitudinal and transverse sound

velocities in quartz [12].

In the Langevin equation that we used to describe the evolution of our mechanical mode, losses are modeled by the term $-\Gamma_m/2\hat{b}$, describing the coupling of the mode \hat{b} to a phonon bath with Markovian dynamics. However diffraction loss, which has been studied in HBARs (see e.g. [12] or [21]), is not a Markovian process and is only qualitatively described by the term $-\Gamma_m/2\hat{b}$ in the Langevin equation. Since we expect diffraction to be our main loss channel [1], this might be able to explain why we see deviations from exponential decay in our T_1 plots.

The frequency detuning with which we drive the mechanical mode affects which superposition of mechanical eigenmodes is excited in the HBAR. We suspect that the amount of diffraction loss depends on the specific superposition of eigenmodes. This might provide an explanation for the large variations in phonon lifetimes that we observe. One should be able to verify (or falsify) this theory by calculating what the excited superpositions are and simulating their diffraction losses using the method from [21].

Conclusion

In this work we have made progress towards microwave-to-optical transduction by adapting the experimental setup and PDH laser locking scheme from [1] for the pulsed operation needed for a future transduction experiment. Further, we have demonstrated the creation and subsequent readout of large coherent phonon states by means of Brillouin scattering. The readout of phonon states by means of Brillouin scattering comprises an essential part of microwave-to-optical transduction using a phonon mode of an HBAR as an intermediary between the superconducting qubit state and the optical mode of a fiber.

First tests of the locking scheme in pulsed operation using PID controllers directly built in to the laser controller demonstrated the need for higher bandwidth PID controllers with faster reaction times to turning the active feedback on and off.

We have demonstrated that the PID controller implemented on a Red Pitaya through the PyRPL software package has a bandwidth high enough to make corrections on the timescale of the cavity dynamics and has active feedback that can be turned on and off within $\sim 0.1 \mu\text{s}$. Through tests performed in both the room temperature test cavity as well as the cryogenic cavity, we have found a PID parameter regime that appears to be relatively well suited to deal with the noise environment in the fridge.

Monitoring the quality of the lock during longer measurements on the timescale of minutes, while pulsing at a repetition rate of 1 kHz, has shown average lock levels as low as 17 %, although this is not typical and average lock levels over 30 % were observed in similar measurements. These large fluctuations are probably caused by external environmental conditions (people in the lab, construction in the same building, etc.).

It remains open how the relative timing of turning on the PID controllers affects the lock performance and most likely the PID parameters can be further optimized. A more systematic study of these aspects is still needed. Further, we believe that the performance of the Red Pitaya is currently limited by the limited output range, causing the feedback voltage to get stuck at the minimal or maximal value, hindering further active feedback. Related to this, we have seen that the quality of the lock clearly oscillates at the frequency of the turbo pump, which is the dominant noise source in the fridge, indicating that our locking scheme is not yet doing enough to suppress this noise source and the setup would benefit from further vibration isolation¹.

In the second part of this work we successfully stored and subsequently retrieved optical pulses from a phonon mode of an HBAR through Brillouin scattering. Within an order of magnitude, the retrieved phonon energy matches what we predict using a simplified coupled oscillator model. We probed the response of two phonon modes for different detunings of the drive frequency and characterized the lifetimes of the mechanical excitations we created through this. We found that the lifetimes of the mechanical excitations depend on the frequency detuning at which the mechanical mode was driven and the decay only approximately follows an exponential. We find T_1 lifetimes as high as $4.9 \pm 0.1 \mu\text{s}$ for a drive detuning of 60 kHz, which is 70 % larger than what one expects from the intrinsic linewidth of the phonon modes found in [1]. Related to this, the mechanical modes did

¹Note that we used the same vibration isolation as [1]. The improved vibration isolation stage had not yet been mounted at the time of this project.

4 Conclusion

not exhibit a Lorentzian response for the values of t_{wait} we used. Although based on our discussion in 3.5 this is to be expected, as we were measuring the response for $t_{wait} > 0$. We hypothesize that these effects are caused by diffraction loss, which generally does not fulfill the assumptions which lead to the typical exponential energy decay. The observed mechanical modes are comprised of a superposition of eigenmodes of the HBAR and we suspect that the frequency detuning of the drive affects the exact composition of these superpositions, leading to drive frequency-dependent diffraction losses.

Acknowledgments

I would like to thank Tom Schatteburg for supervising this work and providing me with competent and valuable feedback throughout the entire project. Tom was always willing to take the time to answer "just a couple of short questions" that constantly came up as I was working on this project, thank you!

I would also like to thank Yiwen Chu for providing guidance at crucial points throughout the project and especially for offering me the opportunity to do my master thesis in the HyQu group in a field completely unrelated to anything I had worked on before.

And of course last but not least I would like to thank the entire HyQu group for providing a positive and healthy working environment and always being willing to help out when needed. Such as to explain how to use an outdated digitizer and install even obscurer drivers for it (Marius) or to baby sit me during "pulse tube off" measurements (Tom, Hugo and Max) and of course a special shout out goes to Max and Rodrigo, who graciously let me use their AOM after I fried the AOM intended for this project one fateful afternoon early during my time at HyQu.

Bibliography

1. Doeleman, H. M. *et al.* Brillouin optomechanics in the quantum ground state. <http://arxiv.org/abs/2303.04677> (2023).
2. Ladd, T. D. *et al.* Quantum computers. *Nature* **464**, 45–53. <https://doi.org/10.1038/nature08812> (Mar. 2010).
3. Jozsa, R. *Entanglement and Quantum Computation* 1997. arXiv: quant-ph/9707034 [quant-ph].
4. Bennett, C. H. & Brassard, G. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science* **560**, 7–11. <https://doi.org/10.1016/j.tcs.2014.05.025> (Dec. 2014).
5. Giustina, M. *et al.* Significant-Loophole-Free Test of Bell’s Theorem with Entangled Photons. *Phys. Rev. Lett.* **115**, 250401. <https://link.aps.org/doi/10.1103/PhysRevLett.115.250401> (25 Dec. 2015).
6. Blais, A., Girvin, S. M. & Oliver, W. D. Quantum information processing and quantum optics with circuit quantum electrodynamics. *Nature Physics* **16**, 247–256. ISSN: 1745-2473 (Mar. 2020).
7. Lauk, N. *et al.* Perspectives on quantum transduction. *Quantum Science and Technology* **5**. ISSN: 20589565 (2020).
8. Lambert, N. J., Rueda, A., Sedlmeir, F. & Schwefel, H. G. L. Coherent Conversion Between Microwave and Optical Photons—An Overview of Physical Implementations. *Advanced Quantum Technologies* **3**, 1900077. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qute.201900077> (2020).
9. Chu, Y. *et al.* Creation and control of multi-phonon Fock states in a bulk acoustic-wave resonator. *Nature* **563**, 666–670. ISSN: 14764687 (Nov. 2018).
10. Kharel, P. *et al.* Multimode strong coupling in cavity optomechanics. *arXiv*. <http://arxiv.org/abs/1812.06202> (2018).
11. Schatteburg, T. *Microwave-to-optical transduction using high-overtone bulk acoustic wave resonators* PhD proposal, 2021.
12. Kharel, P. *et al.* High-frequency cavity optomechanics using bulk acoustic phonons. *Science Advances* **5**. ISSN: 23752548 (2019).
13. Black, E. D. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics* **69**, 79–87. ISSN: 0002-9505 (Jan. 2001).
14. Aspelmeyer, M., Kippenberg, T. J. & Marquardt, F. Cavity optomechanics. *Reviews of Modern Physics* **86**, 1391–1452. ISSN: 15390756 (Dec. 2014).
15. Nielsen, J. H. *et al.* *QCoDeS/Qcodes: QCoDeS 0.38.1* version v0.38.1. Apr. 2023. <https://doi.org/10.5281/zenodo.7866908>.
16. Neuhaus, L. *et al.* *PyRPL (Python Red Pitaya Lockbox) — An open-source software package for FPGA-controlled quantum optics experiments* in *2017 Conference on Lasers and Electro-Optics Europe European Quantum Electronics Conference (CLEO/Europe-EQEC)* (2017), 1–1.
17. Fiore, V. *et al.* Storing optical information as a mechanical excitation in a silica optomechanical resonator. *Physical Review Letters* **107**, 1–5. ISSN: 00319007 (2011).

Bibliography

18. Fiore, V., Dong, C., Kuzyk, M. C. & Wang, H. Optomechanical light storage in a silica microresonator. *Physical Review A - Atomic, Molecular, and Optical Physics* **87**, 1–6. ISSN: 10502947 (2013).
19. Renninger, W. H., Kharel, P., Behunin, R. O. & Rakich, P. T. Bulk crystalline optomechanics. *Nature Physics* **14**, 601–607. ISSN: 17452481 (June 2018).
20. Ford, G. W., Lewis, J. T. & O’Connell, R. F. Quantum Langevin equation. *Phys. Rev. A* **37**, 4419–4428. <https://link.aps.org/doi/10.1103/PhysRevA.37.4419> (11 June 1988).
21. Chu, Y. *et al.* Quantum acoustics with superconducting qubits. *Science* **358**, 199–202. ISSN: 10959203 (2017).



Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

PULSED OPTOMECHANICS USING HIGH-OVERTONE BULK ACOUSTIC WAVE RESONATORS

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):

WERNER

First name(s):

JACOB

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the '[Citation etiquette](#)' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date

Zürich, June 18, 2023

Signature(s)

J Werner

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.