
A frequency-stable diode laser system for spectroscopy and trapping of Sr atoms

Nejc Janša



Munich 2016

Ein frequenzstabiles Diodenlasersystem zur Spektroskopie an und Speicherung von Sr-Atomen

Masterarbeit an der Fakultät für Physik
Ludwig-Maximilians-Universität München
Max-Planck-Institut für Quantenoptik

vorgelegt von
NEJC JANŠA
aus Jesenice, Slowenien

München, den 26.08.2016

Abstract

I report on my work done on the new strontium machine in Immanuel Bloch's group at MPQ under the supervision of Dr. Sebastian Blatt, which was a part of my Master program at LMU Munich. My main project was the construction of the 689 nm laser system for laser cooling of neutral strontium atoms. I report on how the diode laser is stabilized using a combination of spectroscopy and a high-finesse cavity. For the detection of the laser light I report on a low-noise photodetector design. To house the high-finesse cavity, I designed a vacuum housing with good temperature and vibration isolation for better stability. In addition, I took up additional projects, which were necessary for the progress in the newly started laboratory. I report on a novel mechanical design for stable and compact mounting of optics, which I used to assemble a module for injection-locking of a laser diode. Last of all, I describe the measurements done on the strontium atomic beam.

Contents

Abstract	v
1 Introduction	1
1.1 Overview	2
2 Cooling Strontium	5
2.1 Strontium Atom	5
2.2 Atomic oven	7
2.3 Zeeman slower	9
2.4 Blue MOT	12
2.5 Red MOT	14
3 Red laser system	15
3.1 Laser setup	15
3.2 Laser locking and stabilization	20
4 Designing low-noise photodetectors	27
4.1 Transimpedance amplifier	27
4.2 Designing the photodetector	32
4.3 Testing the device	35
5 Design of the reference cavity housing	41
5.1 Changes of cavity length	41
5.2 Choice of material and geometry	44
5.3 Thermal properties	54
5.4 Assembling the vacuum chamber	57

6	Setup for stable injection-locking	61
6.1	Injection-locking	61
6.2	New compact mechanical design	62
6.3	Assembling a module	71
6.4	Protection and control circuit	74
7	Measurements of the strontium atomic beam	83
7.1	Absorption in an atomic beam	83
7.2	Experimental setup	85
7.3	Data analysis	87
8	Conclusion	95
	Acknowledgments	97
	Bibliography	99

Chapter 1

Introduction

Ultracold atoms are a great tool for creating quantum-mechanical systems with a high level of control. When a gas of atoms is cooled to the range of 100 nK, the quantum nature of the particles starts to play a role. As the particle wave-functions become macroscopic, particles with the bosonic and fermionic statistics start to behave in a different way.

Bosons will experience a transition to a Bose-Einstein condensate with a macroscopic occupation of the lowest quantum state and fermions will form a Fermi liquid, where all states up to a certain energy are filled. By observing these phenomena we can learn about how our world works on the quantum level. Quantum many body physics is far richer than just a combination of single-body particles, as the famous quotation “more is different” from P. W. Anderson suggests.

The great thing about the cold atom systems is that they can be very well controlled and also observed thanks to the light-atom interaction. The level of the interaction can be precisely controlled by using laser light and changing its intensity and frequency. Furthermore, the internal states of the atoms can be affected using external fields. Using both these effects, we can first of all cool the atoms and then create different potentials to manipulate them.

By creating a lattice potential using a standing wave of light, we can replicate the periodic structure found in crystals. Such a setup allows us to study different models proposed to explain the behavior of real materials. What we have in mind is a type of quantum simulation, as proposed by R. P. Feynman. We can simulate how other quantum systems behave by engineering the Hamiltonian of our atoms [1].

The first experiments with cold atoms were done mostly on alkali metals, due to their simple hydrogen-like electronic structure. Lately, more exotic species of atoms are also being successfully cooled and trapped.

Strontium is a great candidate for studies of quantum many-body physics. The two electrons in the outer orbital offer a number of optical transitions from the ground state, all of which are in the visible regime, where commercial diode lasers are readily available. Furthermore, the techniques for cooling and controlling strontium are al-

ready well explored due to its use in atomic clocks and high precision measurements [2]. The presence of both bosonic and fermionic isotopes allows for observation of different phenomena with little change to the experimental setup.

The goal of our new experiment is to construct a versatile machine for studies of quantum many-body physics with strontium. We want to be able to cool any of the naturally available isotopes to quantum degeneracy and to have the option to load them into an optical lattice. Furthermore, we want to image the atoms in the lattice using high resolution imaging, where there have been a lot of recent advancements in the field [3, 4]. The current systems are typically limited to $\sim 30 \times 30$ sites, but we hope to achieve even larger system sizes by using in-vacuum buildup cavities for the lattice light. By making the lattice larger and more homogeneous, we want to reduce the dependence on the local-density approximation and enable quantum simulation with larger particle numbers and better fidelity.

1.1 Overview

In this master thesis, I report on my work done on the initial setup of the experiment. The main topic of my project was the preparation of the frequency-stable diode laser, which is used in the cooling of strontium. As I joined the team just as the experiment started, I took on additional projects to help with the progress in the laboratory.

In Chapter 2, I review the properties of the strontium atom and the knowledge required to cool the atoms. The atomic beam from an effusive oven is collimated by a microchannel nozzle and a differential pumping tube. Using a Zeeman slower, the atoms will be slowed down until they can be trapped in a magneto-optic trap operating on the broad $^1S_0 \rightarrow ^1P_1$ transition. Finally, they will be loaded into another magneto-optical trap on the narrow $^1S_0 \rightarrow ^3P_1$ transition.

In Chapter 3, I describe the setup of the red laser system. The diode laser linewidth will be reduced using the Pound-Drever-Hall technique which allows the laser to be locked to an ultrastable reference cavity. An absolute frequency reference is provided by a strontium spectroscopy cell. The laser system has to provide several outputs with frequencies that match the isotope shifts of the different isotopes.

In Chapter 4, I continue with the design of low-noise photodetectors used for the PDH lock and for spectroscopy. The photodetector boards can be adapted for different applications and their output was tested to be mostly shot-noise limited.

In Chapter 5, I report on my design of the reference cavity housing. I revise some of the effects that limit the stability of an optical cavity and show how these considerations were taken into account during the design. With proper choice of material,

good thermal and acoustic isolation, we aim for a fractional stability of the cavity length on the order of 10^{-15} . I also describe how the housing was assembled and how the different parts work.

In Chapter 6, I describe a novel optical breadboard design. By using dowel pins for precise mechanical referencing, we built a setup for compact and stable mounting of optics. Using this setup, I designed a compact module for injection-locking a laser diode and I describe how it is assembled. Furthermore, I describe the protection and control circuit I adapted for the monitoring of the used laser diode.

In Chapter 7, I describe my measurements of the strontium atomic beam. The measurement method allows the characterization of the collimating properties of the nozzle. We can also get estimates of the density and flux of our atomic beam.

Chapter 2

Cooling Strontium

In this chapter, I will collect some general information about the strontium atom. Furthermore, I will describe the steps required to laser-cool the strontium atoms to temperatures on the order of $< 1 \mu\text{K}$. The methods required to cool strontium to degeneracy are already well explored [5–7].

2.1 Strontium Atom

Strontium is an alkaline earth metal and is rarely found in its elemental form due to its high reactivity with water and water vapor. Its atomic number is 38 and it sits in the second group and fifth row of the periodic table. It has four naturally occurring stable isotopes with different abundance, shown in Table 2.1. All of the isotopes with an even mass number are bosonic and ^{87}Sr with the odd mass number is a fermion. The statistics of the atom depends on the nature of its constituent particles. An odd number of fermions will make a fermionic particle and an even number of fermions form a bosonic particle. As electrons, protons and neutrons are all fermions and there are 38 electrons in the atom, the statistics for the strontium isotopes depends directly on the mass number, which is the sum of the number of protons and neutrons. Furthermore, only ^{87}Sr has a nuclear spin.

Isotope	Abundance	Nuclear spin	Statistics
^{84}Sr	0.56%	0	Bosonic
^{86}Sr	9.86%	0	Bosonic
^{88}Sr	82.58%	0	Bosonic
^{87}Sr	7.00%	9/2	Fermionic

Table 2.1: Strontium isotope data taken from NIST [8].

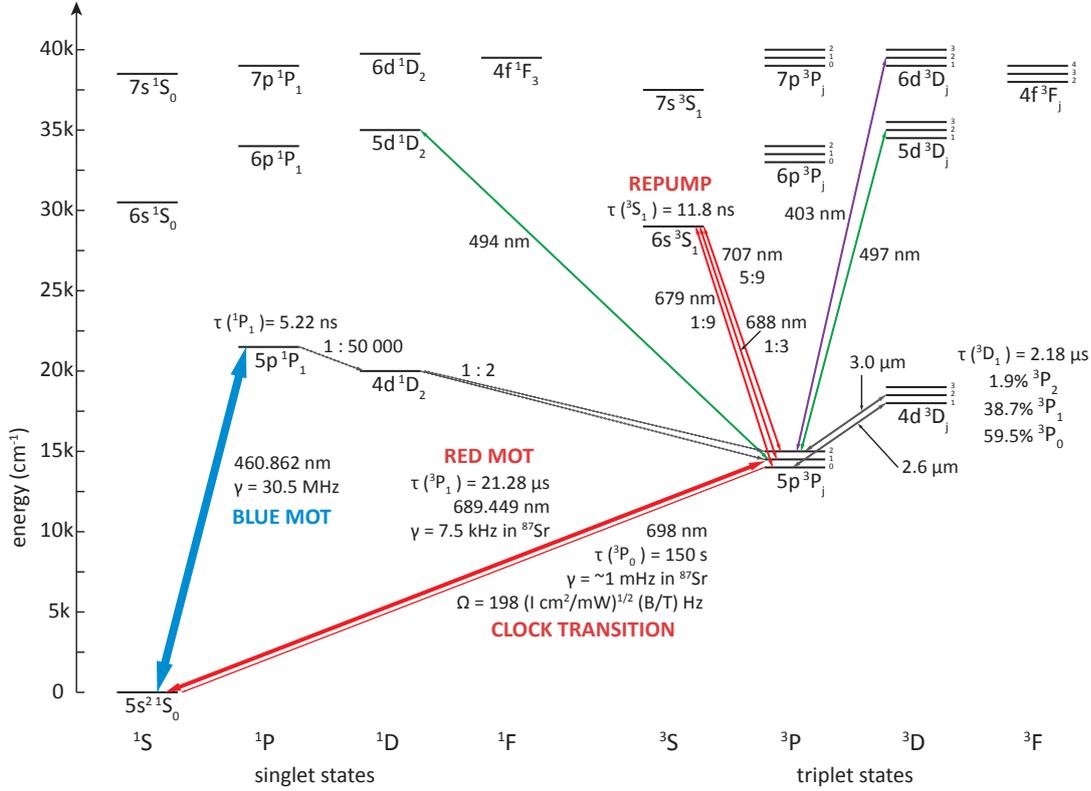


Figure 2.1: The low-lying one-electron excited energy states for strontium. Energy levels are taken from Ref. [10]. Added are the wavelengths of some of the well-explored transitions and the lifetimes of the states [6, 11–13].

2.1.1 Electronic properties

As strontium is found in the second group of the periodic table, it has two electrons in its outer shell. These electrons lead to a rich electronic structure, which is partially depicted in Figure 2.1. The electron pair can form either a singlet or a triplet, depending on the orientation of their spins. The electronic ground state is the singlet $5s^2\ ^1S_0$. Transitions between the singlet and the triplet are dipole forbidden due to the $\Delta S = 0$ selection rule, which makes the transition rates a few orders of magnitude lower. Furthermore, we find a doubly forbidden transition to the $5s5p\ ^3P_0$ state, which in addition breaks the $\Delta J = 1$ selection rule. The transition has a finite transition rate in ^{87}Sr due to small mixing between different hyper-fine states, giving rise to the so-called clock transition [9].

Due to the absence of nuclear spin in the bosonic isotopes, these isotopes do not show a hyperfine structure. The shift in energy levels is present only in the fermionic ^{87}Sr . Here, we should also note that the different masses of the nuclei result in isotope shifts. In Table 2.2 we show the frequency shifts with respect to the electronic transitions of the most abundant isotope ^{88}Sr on which we normally base our spectroscopy.

Isotope	F'	$^1\text{S}_0 \rightarrow ^1\text{P}_1$ (MHz)	$^1\text{S}_0 \rightarrow ^3\text{P}_1$ (MHz)
^{84}Sr		-270.8	-351.49
^{86}Sr		-124.8	-163.81
^{88}Sr		0	0
^{87}Sr	7/2	-9.7	1352.0
	9/2	-68.9	-221.7
	11/2	-51.6	-1241.4

Table 2.2: The detuning from the transitions in ^{88}Sr required to address the different isotopes and hyperfine states [5].

2.2 Atomic oven

To obtain atomic strontium, we first have to transform the metal into a gas. Strontium has a melting point at around $770\text{ }^\circ\text{C}$, but even at lower temperatures, there will be a certain vapor pressure of strontium gas above the surface of the solid. The vapor pressure dependence on temperature can be described by the Clausius-Clapeyron equation, which contains experimentally determined coefficients. For strontium in the temperature range between $361\text{ }^\circ\text{C}$ and $750\text{ }^\circ\text{C}$ the following relation holds [14].

$$p(T) = 10^{10.255 - 8324/T}, \quad (2.1)$$

where p is the pressure in units of Pa and T is the temperature in units of K. We should note that different coefficients for the Clausius-Clapeyron equation can be found in the literature. A comparison of these is plotted in Figure 2.2.

By building an oven that can be heated to high enough temperatures, we can get a gas of strontium with the desired pressure. A small aperture in the oven allows some of the atoms to fly out and we get an atomic beam. If we want the atomic beam to be collimated, we can make the aperture out of long and thin tubes, which form a nozzle. In the ballistic approximation, where the mean free path of the atoms is longer than the dimensions of the tube, only atoms with low transverse velocities will exit the oven. The atoms which collide with the wall, will only slowly diffuse towards the

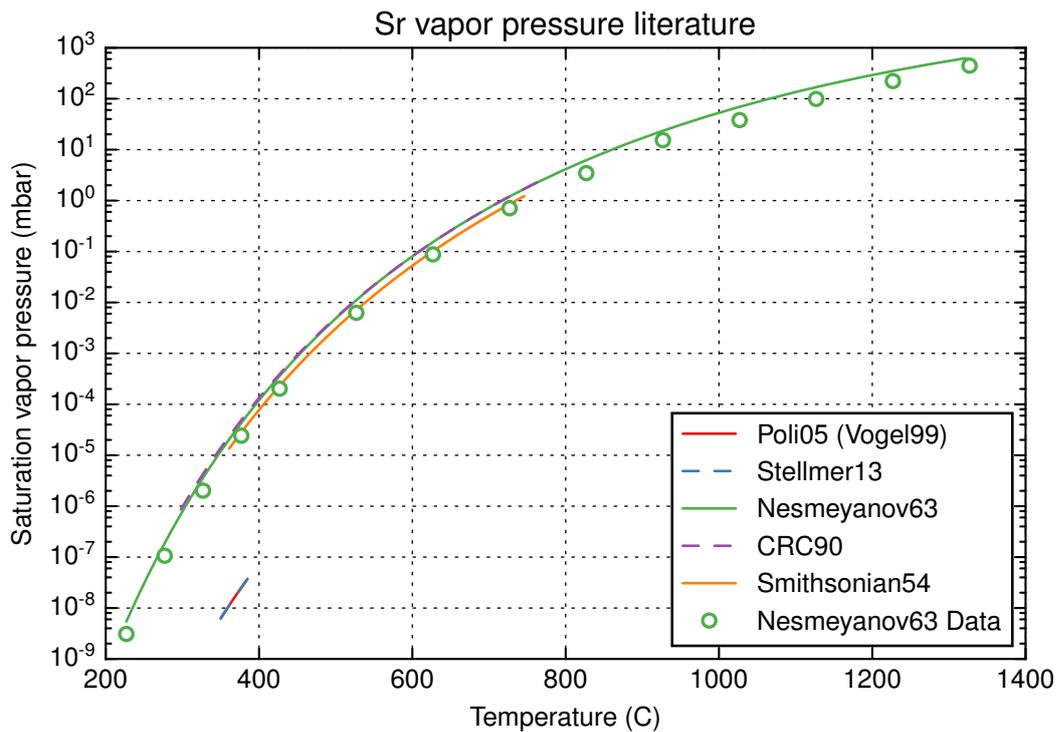


Figure 2.2: A comparison of different fits to the Clausius-Clapeyron equation for the vapor pressure of strontium taken from different literature. Most of the fits seem to agree with the data from Nesmeyanov63 [15]. However, in the temperature range more interesting to us, the fit from Smithsonian54 [14] seems to agree better than the fit from CRC90 [16]. There is also a set of different fit parameters, which seem to originate from an experimental fit on a system with a different geometry [5, 17, 18].

exit of the tube, due to the random direction of flight after the collision. The principles of how such an effusive oven works are discussed in Ref. [19].

The velocity distribution of the atomic beam from an effusive oven is given by the modified Maxwell-Boltzmann distribution [19]

$$f(v) = \frac{1}{2} \left(\frac{m}{k_B T} \right)^2 v^3 e^{-mv^2/(2k_B T)}, \quad (2.2)$$

where v is the velocity of the atom, m its mass, k_B the Boltzmann constant, and T the temperature of the gas in the oven. The mean velocity $\langle v \rangle$ is given by

$$\langle v \rangle = \sqrt{\frac{9\pi k_B T}{8m}}. \quad (2.3)$$

For a well collimated atomic beam, the longitudinal velocity distribution $f(v_z)$ can be approximated by the same distribution.

The modified Maxwell-Boltzmann distribution is shown in Figure 2.3. The atoms have rather large velocities and only a tiny fraction could be loaded directly into the blue MOT. Therefore, we have to first slow the atoms down using a Zeeman slower.

The oven we are using in our experiment was bought from Createc and can be seen in Figure 2.4. Solid strontium is loaded into a cylindrical tantalum crucible. The micro-channel nozzle is inserted into the opening of the crucible and clamped in place. The assembled crucible is then inserted into the oven, where it is surrounded by heating wire to achieve the desired temperature. The outer side of the oven can be water cooled to prevent the vacuum system from getting heated, which would otherwise decrease the quality of the vacuum through increased outgassing. A recent report on a strontium oven using similar principles can be found in Ref. [21].

2.3 Zeeman slower

To increase the number of atoms that can be captured by the blue MOT, we need to decrease the velocity of the faster atoms. Atoms can be cooled using the radiation pressure of a laser beam resonant with an atomic transition.

The radiation force of light with wavelength λ on a stationary two-level atom has the following form [22]

$$F(\delta) = \hbar \frac{2\pi}{\lambda} \frac{\gamma}{2} \frac{s_0}{1 + s_0 + (2\delta/\gamma)^2}, \quad (2.4)$$

where $\hbar k = \hbar \frac{2\pi}{\lambda}$ is the momentum of the photon given by the reduced Planck constant \hbar . The wave-vector of the photon is defined as $k = \frac{2\pi}{\lambda}$. The decay rate of the atomic

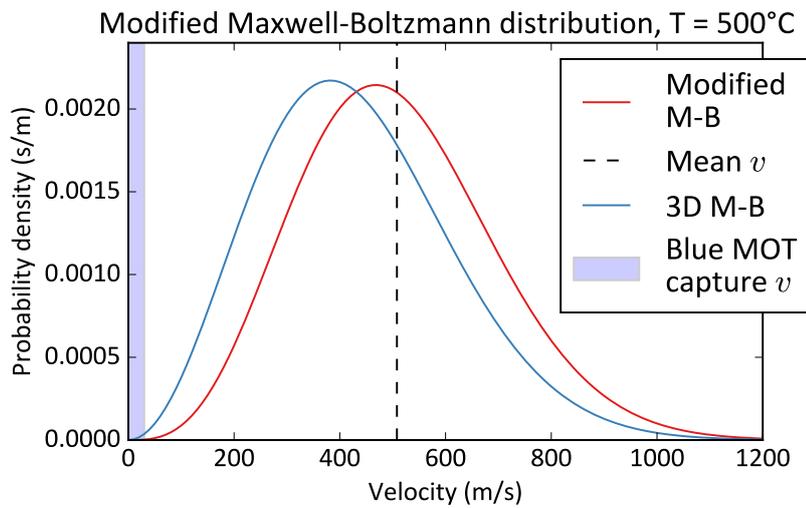


Figure 2.3: Comparison of the modified Maxwell-Boltzmann (M-B) distribution with the standard Maxwell-Boltzmann distribution for a 3D gas. The modified velocity distribution in an atomic beam tends to higher velocities. The plots are made for ^{88}Sr atoms at 500°C . The mean velocity of the modified distribution is marked by the dashed black line. The blue region in the plot shows which atoms can be captured by the blue broadband MOT, which has a capture velocity of approximately 30 m/s [20].

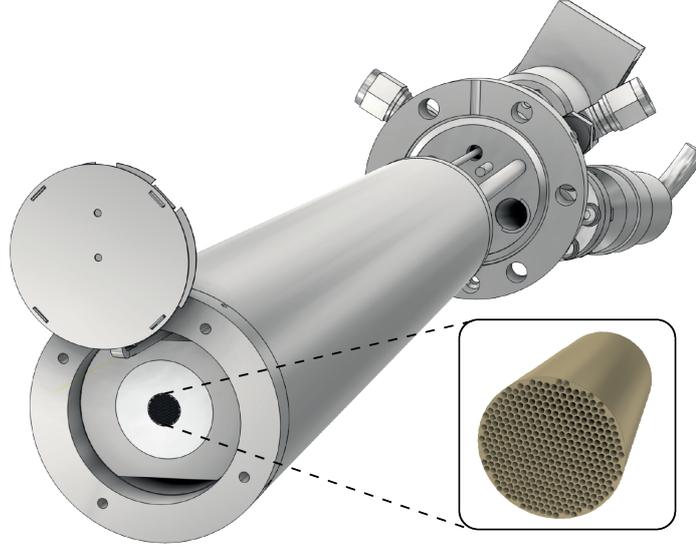


Figure 2.4: The atomic oven from Createc used in our experiment. The custom nozzle is magnified in the bottom right. The nozzle made out of titanium consists of 700 micro-channels 10 mm long and with a 200 μm diameter.

transition is given by γ and δ is the detuning of the laser light from the two-level transition.

In Eqn. 2.4, s_0 is the saturation parameter, defined by $s_0 = I/I_{\text{sat}}$, where I is the intensity of light and the saturation intensity is given by

$$I_{\text{sat}} = \frac{\pi\hbar c}{3\tau\lambda^3}, \quad (2.5)$$

where τ is the lifetime of the excited state.

The saturation intensity of the 461 nm blue transition in strontium is $I_{\text{sat},461} = 40.7 \text{ mW/cm}^2$ and for the 689 nm red transition it is $I_{\text{sat},689} = 0.58 \text{ }\mu\text{W/cm}^2$. Furthermore, the decay rates of the transitions to the ground states are shown in Figure 2.1. For slowing of atoms we want to use the 461 nm transition, as the slowing force will be much larger.

If the atom is moving with a velocity v towards the slowing beam, the force will change due to an effective detuning caused by the Doppler effect:

$$\delta_d = \frac{-\vec{k} \cdot \vec{v}}{2\pi} = -v/\lambda. \quad (2.6)$$

The Doppler effect prevents us from completely slowing down an atom with the

light of a single detuning, as the atom comes out of resonance once its velocity has been decreased. A great solution to this problem is to introduce another effective detuning that is position dependent. The detuning is introduced using the Zeeman effect and a magnetic field that changes along the axis of the atomic beam z . The form of the Zeeman detuning for the 1P_1 transition is

$$\delta_z(z) = \mu_B \frac{2\pi}{\hbar} B(z), \quad (2.7)$$

where μ_B is the Bohr magneton and $B(z)$ is the magnetic field, which points along the z axis pointing along the z axis.

We can optimize the function $B(z)$ for atoms with a certain initial velocity v_i , to decelerate them at a constant force F_0 . The solution is a magnetic profile of the form

$$B(z) = \sqrt{\left(\frac{\hbar v_i}{2\pi\lambda\mu_B}\right)^2 - 2\frac{F_0}{m}z} + \frac{\hbar\delta}{2\pi\mu_B}, \quad (2.8)$$

where m is the mass of the strontium atom and δ the detuning of the laser from the atomic transition. Using this example, we can build a so-called Zeeman slower [22] that slows down all atoms that enter the magnetic field with a velocity below v_i , because at some position z they will get in resonance with the slowing beam. The final velocity v_f after the Zeeman slower can be selected by changing its length.

The selection of the rest of the parameters depends on other considerations, such as how much laser intensity I there is available for the slowing, how large the magnetic fields can be, mechanical constraints, what the divergence of the atomic beam is, and what final velocity we want.

In our experiment, we are planning for a Zeeman slower which can slow down atoms with $v_i = 400$ m/s down to $v_f = 30$ m/s. This result will significantly increase the fraction of the atoms in the atomic beam, which can be captured by the blue MOT.

2.4 Blue MOT

The next stage of cooling is done with a magneto-optical trap. The first stage of the MOT is done on the broadband 461 nm blue transition. The details on how a MOT works, are explained in Ref. [22].

The force used for slowing and trapping of the atoms is again the radiation pressure from Eqn. 2.4, with the additional effective detunings of the Doppler and Zeeman shifts. The chosen magnetic field has a linear gradient, $B(z) = Ax$. Furthermore, we use two counter propagating beams with circular polarization of the opposite helicity, which will result in the opposite sign of the Doppler and Zeeman contributions.

If we want to cool and trap in all three dimensions, we need three pairs of counter propagating beams and a magnetic gradient in all three directions.

The resulting force on an atom along one dimension is

$$F = \hbar \frac{2\pi}{\lambda} \frac{\gamma}{2} \left(+ \frac{s_0}{1 + s_0 + 4(\delta - v/\lambda - 2\pi\mu_B Ax/\hbar)^2/\gamma^2} - \frac{s_0}{1 + s_0 + 4(\delta + v/\lambda + 2\pi\mu_B Ax/\hbar)^2/\gamma^2} \right). \quad (2.9)$$

For small Zeeman and Doppler shifts, compared to the detuning, we can Taylor expand the equation around v and x , obtaining

$$F = -\beta v - \kappa x, \quad (2.10)$$

with the spring constant $\kappa = \frac{2\pi\mu_B\lambda A}{\hbar} \beta$ and the damping coefficient given by

$$\beta = -8\hbar\delta \frac{(2\pi)^2}{\lambda^2} \frac{s_0}{(1 + s_0 + (2\delta/\gamma)^2)^2}. \quad (2.11)$$

If we want the value of β to be positive, we need to select a negative detuning of the laser light, also called red detuning. Now, the force on the atom has a damping coefficient β , which will decrease the velocity of the atoms. Furthermore, we recognize κ as a spring constant, which leads to a restoring force that traps the atoms in the center with the help of the magnetic field gradient.

There are two limits to how low atoms can be cooled using a laser. The Doppler temperature is the expected final temperature of a MOT with the optimal detuning $\delta \approx \gamma$ and is limited by the natural linewidth of the transition γ [22]:

$$T_{\text{Doppler}} = \frac{h\gamma}{2k_B}, \quad (2.12)$$

where h is the Planck constant.

The second cooling limit is the recoil temperature. It assumes that the atom is cooled to a complete stop by absorbing the last cooling photon. However, the atom is now in an excited state and has to emit the photon. The recoil of the emitted photon gives the atom a certain energy which can be expressed as a temperature:

$$T_{\text{Recoil}} = \frac{h^2}{mk_B\lambda^2}. \quad (2.13)$$

For the 461 nm transition, the value of the two temperatures is $T_{\text{Doppler},461} = 730 \mu\text{K}$ and $T_{\text{Recoil},461} = 1.02 \mu\text{K}$. Indeed, experimentally the final temperature of the blue MOT is shown to be close to the Doppler temperature [23]. To cool the atoms further, we load them into the red MOT.

2.5 Red MOT

The reason to use the magneto-optical trap on the 3P_1 transition is the much narrower linewidth. The mentioned Doppler and recoil temperatures are $T_{\text{Doppler},689} = 180$ nK and $T_{\text{Recoil},689} = 460$ nK. Indeed, there have been reports of atomic sample temperatures as low as 250 nK [6].

Loading the atoms from the blue MOT to the red MOT is a bit more tricky. The capture velocity of the red MOT is too small to catch the relatively hot atoms of the blue MOT. To overcome this problem, we have to increase the detuning from the transition and artificially broaden the linewidth by adding modulation on the laser light. These two things can be introduced using an AOM, and are one of the requirements for the laser setup discussed in Chapter 3.1. Furthermore, the magnetic field gradient required for the red MOT is quite different than for the blue MOT, so the magnetic field has to be changed rapidly during the loading.

After the atoms are loaded into the red MOT, the MOT has to go through a compression sequence, which allows us to reach the lower final temperatures. The optimal parameters of this compression vary somewhat between different experiments. A recent publication showing good results can be found in Ref. [23]. In this example, the starting intensity of each red MOT beam is $\sim 2000I_{\text{sat}}$, and its frequency is modulated to achieve sidebands at detunings between -150 kHz and -8 MHz with a 20 kHz spacing. In the end, the sidebands are removed, the detuning arrives at -150 kHz, and the beam intensity is reduced to $\sim 0.5I_{\text{sat}}$.

The operation of the red MOT for the fermionic isotope ^{87}Sr is a bit more complicated. Due to the nuclear spin, the 3P_1 state experiences hyperfine splitting into three states with $F' = 11/2, 9/2$ and $7/2$. The large values of F in the ground state and the F' of the excited state give us a series of sub-states with different m_F . The operation of the fermionic MOT becomes quite complicated and a good analysis is found in Ref. [5]. In practice, the solution is to use two MOT beams. The $F' = 11/2$ transition, which is the trapping laser exerts a restoring force on most of the m_F states. Additionally, the $F' = 9/2$ transition is required to stir the population of the m_F states, which allows all of the atoms to be trapped on average.

Chapter 3

Red laser system

As we saw in the previous chapter, the red MOT requires light at 689 nm at several different detunings, corresponding to the different isotopes and the different hyperfine states for fermionic ^{87}Sr . We want our experiment to be versatile, so I planned for a laser system that can provide all of the required frequencies.

Furthermore, the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition has a narrow linewidth of 7.4 kHz. Therefore, we have to take care of the linewidth of the laser and its absolute frequency stability to be able to resolve the transition. A typical linewidth of a grating stabilized external-cavity laser is on the order of 100 kHz. We reduce it using the Pound-Drever-Hall method by locking to a high-finesse cavity. For the absolute frequency reference we want to lock to a spectroscopy signal of a cloud of atomic strontium.

3.1 Laser setup

For the master diode laser we bought a Toptica DL pro laser, which can output up to 30 mW of power at 689.4 nm. This power is not enough for all of the beams necessary for the red MOT. Therefore, we have to amplify the output of the master laser using a tapered amplifier (TA). The schematic of the optic setup of the master laser is shown in Figure 3.1.

First, a small amount of power is split off and sent to the wavemeter. The wavemeter measures the wavelength of the light to approximately 10^{-4} nm precision. It helps us set the master laser to the correct mode and to get it close enough to the atomic transition to see the signal in the spectroscopy cell. The second beam is first frequency shifted by +80 MHz using an AOM (acousto-optic modulator). This shift will effectively offset the master laser -80 MHz below the ^{88}Sr transition. This offset helps us with the preparation of all the necessary frequencies which will be discussed later on. Because our reference cavity has a free spectral range of 1.5 GHz and is not frequency tunable, we might have to offset the next beam by up to ± 750 MHz to match the frequency of the atomic transition. The beam for the reference cavity is offset by double-passing through an AOM. The remaining power is coupled into a fiber and

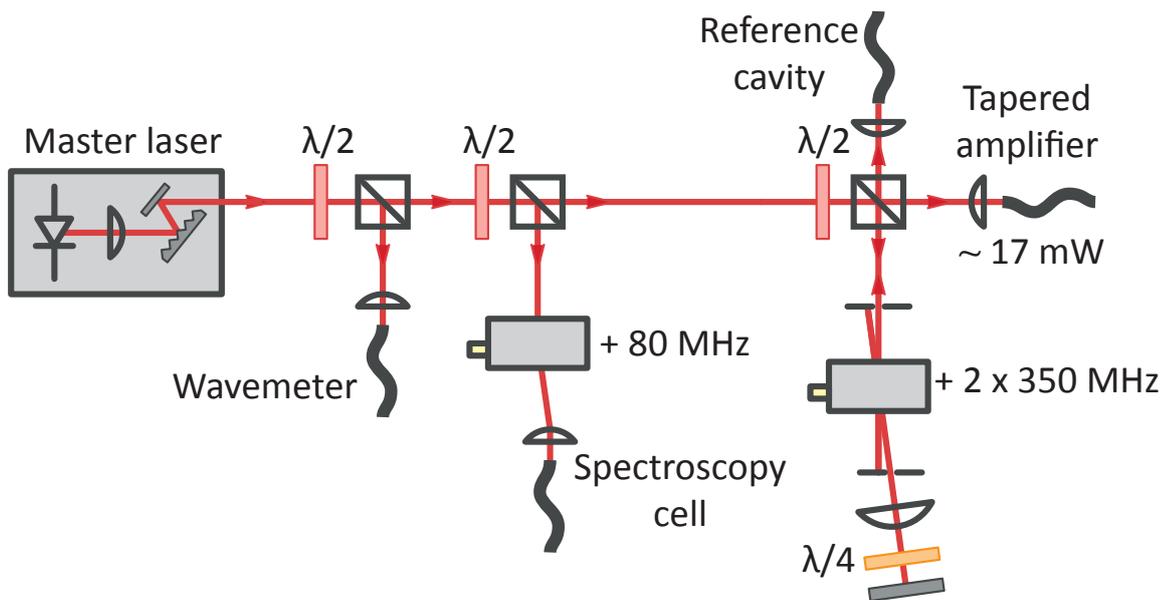


Figure 3.1: The light distribution after the red master laser. The combinations of half-wave plates and polarized beam splitters (PBS) are used to split off a fraction of the laser power. Additional mirrors required for beam alignment and fiber coupling are omitted in the sketch. The grey boxes represent AOMs used for frequency-shifting of the light.

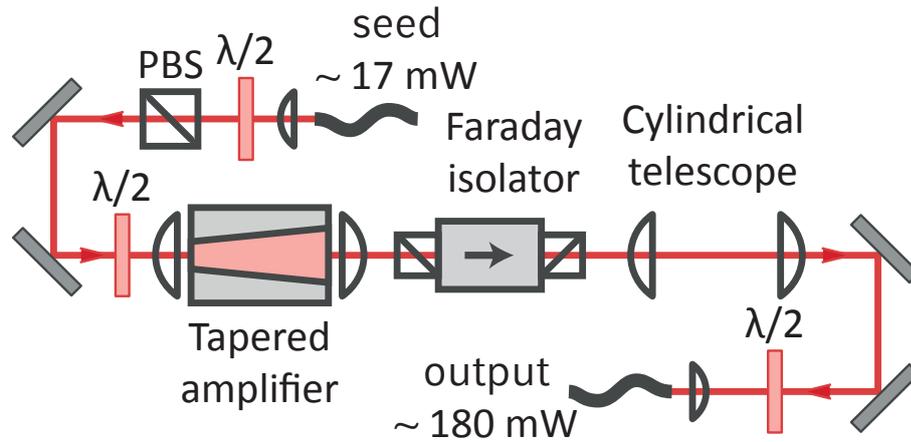


Figure 3.2: The optics setup for the tapered amplifier.

sent to the tapered amplifier, leaving us with around 17 mW of optical power after fiber coupling.

3.1.1 Tapered amplifier

The tapered amplifiers for the red laser system were tested and assembled by our interns. We decided on a modular design, where the TA is a separate module and receives the power through an optical fiber and the output is also coupled to a fiber. This configuration allows easier replacement or repair of the module, without having to realign the rest of the setup. The schematic of the TA module can be seen in Figure 3.2. The polarization of the seed light coming from the optical fiber is cleaned using the half-wave plate and the PBS. It is next coupled into the input aperture of the tapered amplifier chip using two mirrors and a lens. The TA chip acts as a gain medium and amplifies the passing light. The tapered structure prevents the optical power density to become too large and to damage the gain medium. The output of the TA chip diverges differently along the vertical and horizontal axes. The light is partially collimated, along the axis of faster divergence, using the lens just in front of the output. Next, the light passes through a Faraday isolator, which offers approximately 60 dB isolation against light propagating in the opposite direction. The isolation serves to protect the TA chip from back-reflected light. Next, a cylindrical telescope is used to match the beam widths of the horizontal and vertical axes and to collimate the remaining diverging axis. The two mirrors and the lens in front of the fiber allow us to optimize the fiber-coupling efficiency. The last half-wave plate is used to align the polarization of the light to the alignment key of the polarization maintaining fiber.

After fiber coupling we can reliably receive approximately 180 mW of power. Unfortunately, this number tends to decrease over longer periods of operation, as the TA chip ages over time. Therefore, we have two modules prepared so that we can amplify the light in sequence, or to have one TA running while the other can be replaced.

3.1.2 Frequency shifting

The frequencies we require from our 689 nm laser system can be found in Table 2.2. Apart from just shifting of frequencies, we also want to put additional modulation on the MOT beams to artificially broaden the capture velocity of the red MOT. Furthermore, we want to fine-tune the detuning and remotely change the power of each of the beams. I have devised a setup of AOMs, which can be seen in Figure 3.3. I tried to use AOMs, which have a good diffraction efficiency and are commercially available. We decided on using AOMs from Crystal Technologies with nominal frequencies of 80, 110, 200, and 350 MHz. A short guide on how an AOM works and how it can be set up can be found in Ref. [24].

The power in each of the beam paths can be controlled using the half-wave plate and the PBS. By rotating the half-wave plate, we can rotate the polarization of the light. The PBS reflects the component of the light that is vertically polarized and transmits the horizontal component. The turning of the wave-plate allows us to tune the reflected power and allocate the amount of power in the path.

For the ramping of the power, the intensity stabilization and the fast switching-off of a beam, there is a single-pass low frequency AOM in front of each of the fibers. After the diffraction efficiency of the deflected beam into the +1 or -1 order is optimized, we couple the beam into an optical fiber. Afterwards, we can decrease or completely remove the power from the beam by changing the amplitude of the RF signal sent to the AOM. By measuring the intensity in a fraction of the light split off after the output of the fiber, we can feed this signal back to the RF driver and stabilize the output intensity.

The single-pass AOM can also be used for frequency modulation of the beam and for slightly changing the detuning. Both can be achieved by changing or modulating the carrier frequency of the RF signal sent to the AOM. Unfortunately, this type of frequency modulation will result in unwanted intensity modulation, as the angle of the deflected beam changes and it affects the fiber coupling.

The intensity modulation can be suppressed by using the AOM in a double-pass configuration where the first-order beam is retro-reflected [25]. By setting up a lens that has its focus on the crystal in the AOM, we can retro-reflect beams with any angle of deflection without further alignment. As the beam passes through the AOM again, it gets deflected by the same angle and overlaps with the initial beam. The pointing

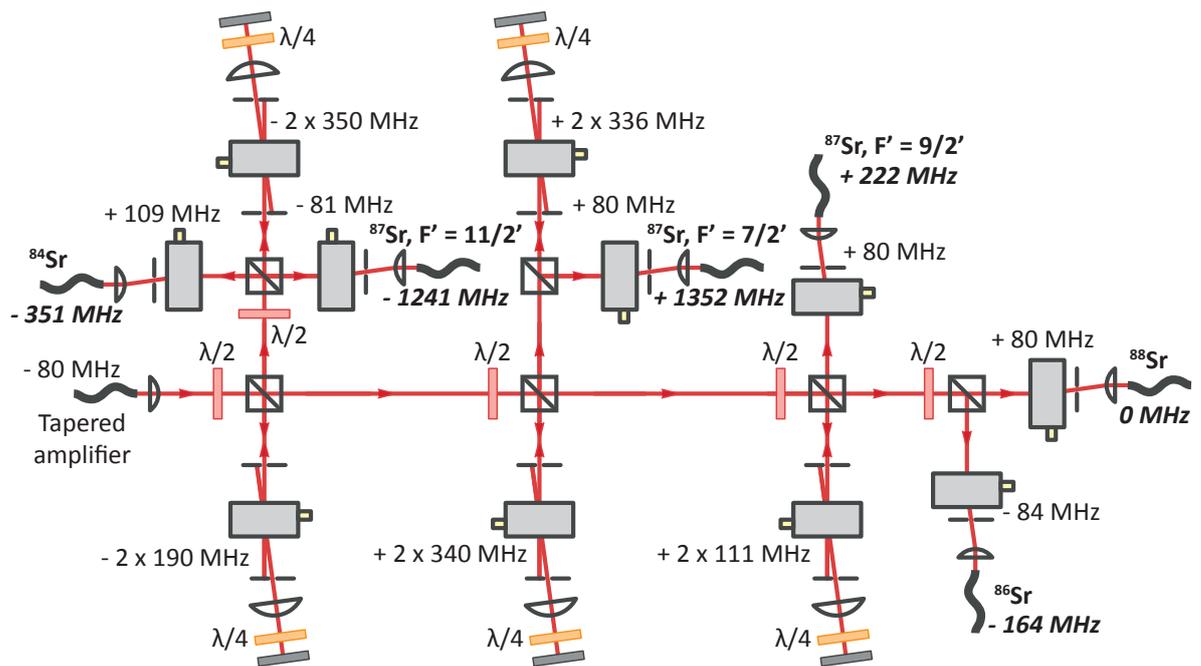


Figure 3.3: A sketch of the optics setup used to provide all the necessary frequencies required for the red MOT of the different isotopes. The additional mirrors and lenses required to align the beam and achieve the optimal beam widths for the AOMs are omitted.

in the beam is no longer sensitive to changes in the shifting frequency. The light from the unwanted deflection orders can be removed using an iris making the total shift after the double-pass twice the carrier frequency of the AOM modulation.

To extract the frequency shifted beam and to prevent it from propagating back into the laser or tapered amplifier, we use the combination of a quarter-wave plate, a mirror, and a PBS. The quarter-wave plate can be rotated until it transforms the linearly polarized light into circularly polarized light. The circularly polarized light will change handedness when reflecting from the 90° mirror. After passing through the quarter-wave plate, it will again turn to linear polarization, but rotated by 90° . This different polarization will now exit a different channel on the PBS compared to where the incoming light came from, allowing it to be fiber coupled.

3.2 Laser locking and stabilization

To stabilize the frequency of a laser, we need to generate a so-called error signal. The error signal should be: zero when the frequency of the laser is correct, positive when the frequency is too low, and negative when the frequency is too high. If we feed this signal into a PID controller, the output of the PID can modify the frequency of the laser by changing its current or cavity length until the frequency is correct again. This control loop is what we consider the locking of a laser. Furthermore, if the error signal has a steep enough crossing through zero, it does not contain too much noise, and if the control loop is fast enough, we can even reduce the linewidth of the laser.

3.2.1 Pound-Drever-Hall lock

A great way to reduce the linewidth of a laser and to give it a stable frequency reference is by using the Pound-Drever-Hall (PDH) lock on a ultrastable reference cavity. The considerations on how to build such a cavity are discussed in Chapter 5. Here, I will try to give an overview on how the PDH lock works, following the derivation found in Ref. [26].

Let us consider what happens to a laser beam when it arrives at an ideal cavity made out of two parallel mirrors. The electric field of the incident beam can be described in its dependence on time t as

$$E_{\text{inc}} = E_0 e^{i\omega t}, \quad (3.1)$$

where E_0 is the amplitude of the field and ω the angular frequency of the light.

An ideal cavity with no losses, made out of two parallel mirrors at a distance L

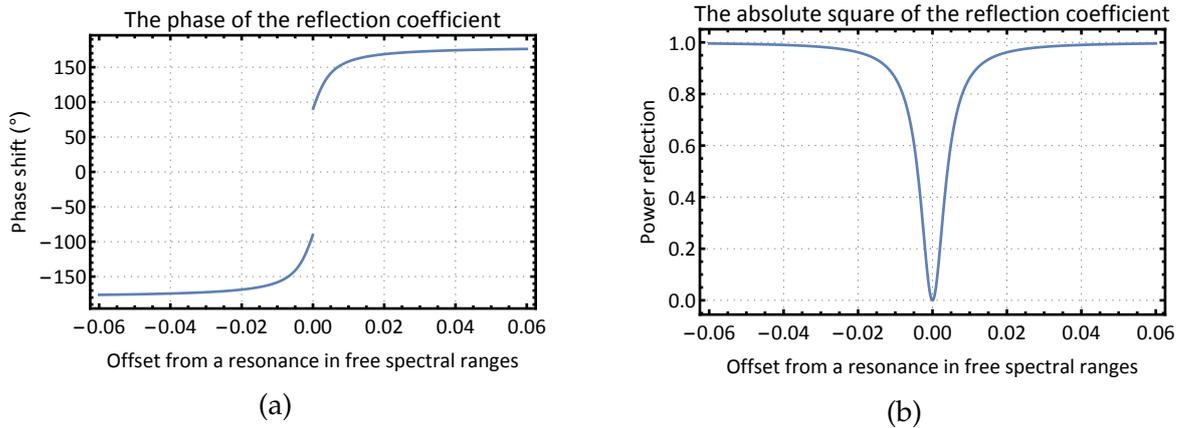


Figure 3.4: The plotted phase of the reflection coefficient $F(\omega)$ (a) and its absolute square (b), which tells us the reflected power. The cavity length for this plot was chosen $L = 10$ cm and $r = 0.998$.

will have the following reflection coefficient [26]

$$F(\omega) = \frac{E_{\text{ref}}}{E_{\text{inc}}} = \frac{r \left(e^{i\omega/\nu_{\text{fsr}}} - 1 \right)}{1 - r^2 e^{i\omega/\nu_{\text{fsr}}}}, \quad (3.2)$$

with the amplitude reflection coefficient r and the free spectral range (FSR) defined as $\nu_{\text{fsr}} = c/2L$. The above equation can be derived by taking the initial electric field and writing a set of equations for the electric fields reflected and transmitted through the two mirrors. Care has to be taken not to forget the π phase shift of reflected fields and the $\frac{\omega L}{c}$ phase shift, which the field obtains when propagating the distance between the mirrors. The final expression is simplified in the numerator using Taylor expansion and assuming $(1 - r) \ll 1$.

The reflection coefficient is a complex number, so it consists of a phase and an amplitude. These are plotted in Figure 3.2. We can imagine trying to lock the laser using the reflected power from the cavity, but unfortunately there is no information on which direction from the resonance we currently are. However, this information is contained in the phase. The problem is that photodetectors cannot resolve the phase of light at optical frequencies, but only its power. The phase could be extracted using interferometry, but as we will see, there is an even better solution.

At this point, we should also define the linewidth of a cavity which can be measured by the full-width half-maximum of the power reflection function. For high

reflectivity of the mirrors it can be approximated by

$$\nu_{\text{fwhm}} = \frac{\nu_{\text{fsr}}}{\mathcal{F}} \approx \frac{1 - r^2}{\pi r}, \quad (3.3)$$

where $\mathcal{F} \approx \frac{\pi r}{1 - r^2}$ is the so-called finesse of the cavity.

By phase modulating the light before sending it to the cavity, we can add sidebands to the carrier frequency. The modulation can be done using an electro-optic modulator (EOM), which adds sinusoidal modulation to the phase of the incoming electric field

$$E_{\text{inc}} = E_0 e^{i(\omega t + \beta \sin(\Omega t))}, \quad (3.4)$$

where β is the modulation index and Ω the frequency of the modulation. In the case of an EOM, Ω is the frequency of the RF signal we send to the EOM and β is proportional to the voltage amplitude of the signal.

The last equation is also a generatrix for Bessel functions J_k , but to keep our calculations simple, we will cut the sum off after the first term [26]:

$$\begin{aligned} E_{\text{inc}} &= E_0 e^{i\omega t} \left(J_0(\beta) + \sum_{k=1}^{\infty} J_k(\beta) e^{i\Omega t} + \sum_{k=1}^{\infty} (-1)^k J_k(\beta) e^{-i\Omega t} \right) \\ &\approx E_0 \left(J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega + \Omega)t} - J_1(\beta) e^{i(\omega - \Omega)t} \right). \end{aligned} \quad (3.5)$$

Therefore, we see that the electric field of the light now contains several components: the carrier at frequency ω and the two sidebands at frequencies $\omega \pm \Omega$. The different frequency components will also affect the reflected beam from the cavity, giving us the electric field of the reflected beam

$$E_{\text{ref}} \approx E_0 \left(F(\omega) J_0(\beta) e^{i\omega t} + F(\omega + \Omega) J_1(\beta) e^{i(\omega + \Omega)t} - F(\omega - \Omega) J_1(\beta) e^{i(\omega - \Omega)t} \right). \quad (3.6)$$

The phase of the reflection coefficient therefore enters the phase of the carrier and the sidebands. When these arrive at a photodetector, the photodetector will be able to detect the interference between the different components, as it measures the power or the absolute square of the total electric field. The fast optical frequencies will be averaged out giving a DC power P_{DC} , but the beating at the difference frequency Ω between the carrier and the sidebands will remain.

The part of the signal from the photodiode, which is relevant in the case where Ω is much larger than the cavity linewidth, is of the following form

$$P_{\text{DC}} + 2|E_0|^2 J_0(\beta) J_1(\beta) \text{Im} [F(\omega) F^*(\omega + \Omega) - F^*(\omega) F(\omega - \Omega)] \sin(\Omega t) + \mathcal{O}(2\Omega t), \quad (3.7)$$

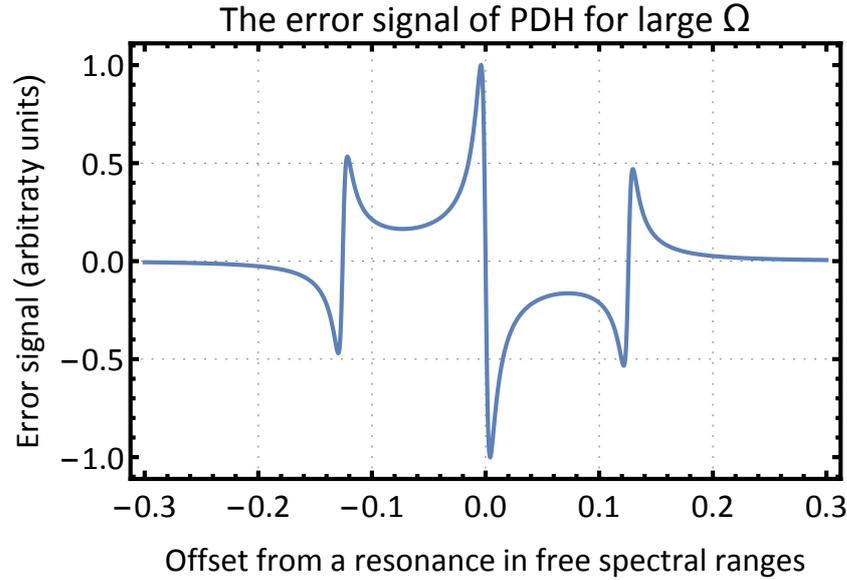


Figure 3.5: The error signal obtained for a high modulation frequency Ω . For this plot I chose $\Omega = 2\pi 30$ MHz while the linewidth was 2 MHz.

By using a mixer, we can multiply the signal with a sinusoidal of the same frequency Ω , taken from the same oscillator that drives the EOM. We can see what this mixing will do by looking at the following trigonometric relation:

$$(c_0 + c_1 \sin(\Omega t)) \sin(\Omega t + \phi) = \frac{c_1}{2} \cos(\phi) + c_0 \sin(\Omega t) - \frac{c_1}{2} \cos(2\Omega t + \phi). \quad (3.8)$$

The resulting mixed signal contains a DC component with an amplitude depending on the complex reflection coefficient and some higher AC components. By using a low pass filter, with a corner frequency below Ω , all the higher components can be removed from the signal. The remaining DC part turns out to have all of the the properties we demanded from our error signal. It is plotted in Figure 3.5.

3.2.2 Optics setup for PDH lock

The optics setup I used to replicate the calculated error signal on a test cavity can be seen in Figure 3.6. The light from the master laser was coupled out of the fiber and aligned with the home built EOM. I used a half-wave plate to set the correct polarization for the optimal operation of the EOM. The Glan-Taylor polarizer is used to make sure that the light only contains the correct component of polarization, as it is specified to have an extinction ratio of $> 100\,000 : 1$. We want the polarization to be clean, to reduce the unwanted intensity modulation caused by the EOM. The uncontrolled

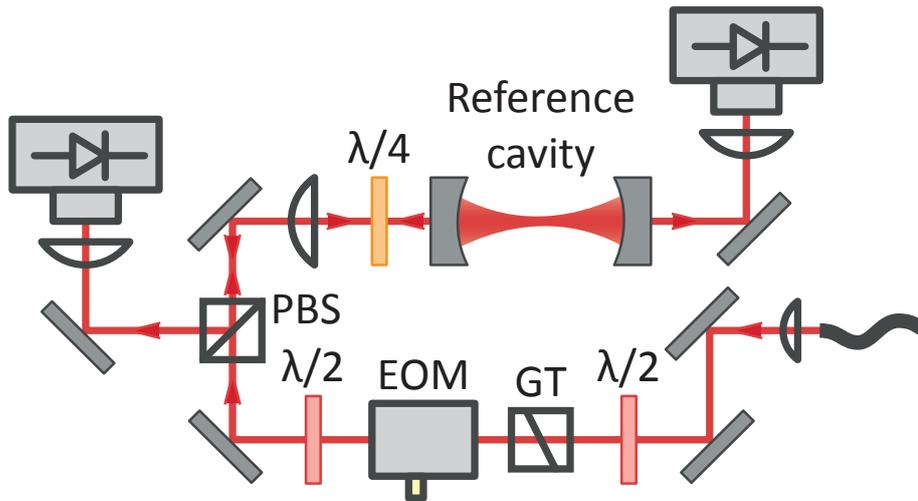


Figure 3.6: The optics setup with the reference cavity used for the Pound-Drever-Hall locking technique. To save space around the reference cavity, the Glan-Taylor (GT) polarizer and the electro-optic modulator (EOM) can be implemented before the optical fiber.

intensity modulation of the light would otherwise introduce a time varying offset in the error signal, which could randomly move around the lock point of the laser [27].

After the EOM, the polarization of the light is again adjusted by the half-wave plate to be transmitted through the PBS. The light is mode matched to the cavity by a lens and the two adjustable mirrors. The approximate position and focal length of the lens can be calculated using Gaussian optics calculations. In practice, the position and the focal length are then slightly adjusted until the light is coupled well to the cavity. I found it of great help to put a beam camera behind the cavity and to look for transmission. The frequency of the laser or the length of the cavity have to be swept over at least a free spectral range, to get the chance to see a mode. Once at least a little bit of light can be seen on the camera, the beam-walking with the mirrors becomes much easier. The effect of moving the mirrors can be seen on the camera and effectively you want all the blobs to come to the center and form a single Gaussian mode. Further optimization can be done by observing the transmission on a photodiode and increasing the power in the principal Gaussian mode.

To record the reflection from the cavity, we add the quarter-wave plate which will in combination with the reflection on the mirror allow the polarization to turn and reflect on the PBS. The light is focused on a low-noise photodiode with enough bandwidth to record the signal at the frequency of modulation Ω . I used a photode-

tor of my own design as discussed in Chapter 4. The one used here was soldered in configuration C with an AC port.

For the signal processing I used the PDH locking module on the Toptica laser (PDD 110 F). The PDD module already contains all of the necessary electronics and knobs to obtain the error signal. The local oscillator output is amplified to drive the EOM and the AC signal from the photodiode is connected to the photodetector input on the module. After some playing with the phase and amplitude knobs and by adding a low pass filter, the error signal looked qualitatively the same to the calculated one in Figure 3.5.

Laser locking was also achieved by connecting the error signal to the Toptica fast analog linewidth (FALC 110) module. I achieved a laser lock rather resistant to vibrations, but did not analyze and optimize the characteristics of the lock. The test cavity was not very stable and the parameters will most likely change once the lock is attempted with the ultrastable cavity. Unfortunately, I did not manage to finalize the building of the cavity housing discussed in Chapter 5, due to long delays with delivery of the parts.

3.2.3 Modulation transfer spectroscopy

The lock to the reference cavity using PDH is a great way to keep the frequency of the laser constant and to reduce the linewidth. However, a problem we find is that the cavity cannot offer us an absolute reference. The length of the cavity will slowly change over time, no matter how well it is built.

For the absolute referencing to the atomic transition a great way is to do spectroscopy on the desired transition. To generate an error signal centered at the atomic transition, we want to use the technique called modulation transfer spectroscopy. A good explanation of the technique can be found in Ref. [28].

In principle, the technique and signal processing is quite similar to the PDH technique. However, the phase information does not come from the reflection of a cavity but from the four-wave mixing process on the optical transition, which is explained in Ref. [29]. For four-wave mixing, we need to do spectroscopy with two different beams, the pump beam and the probe beam. To obtain the third frequency, we use an EOM to put sidebands on the pump beam. In the four-wave mixing process, the modulation is transferred onto the probe beam, which then also obtains sidebands. The interference of the carrier and the sidebands in the probe beam are then observed on the photodetector. The error signal is obtained the same way as in PDH, by mixing the photodetector output with the signal used to drive the EOM and by isolating the DC signal with the help of a low-pass filter.

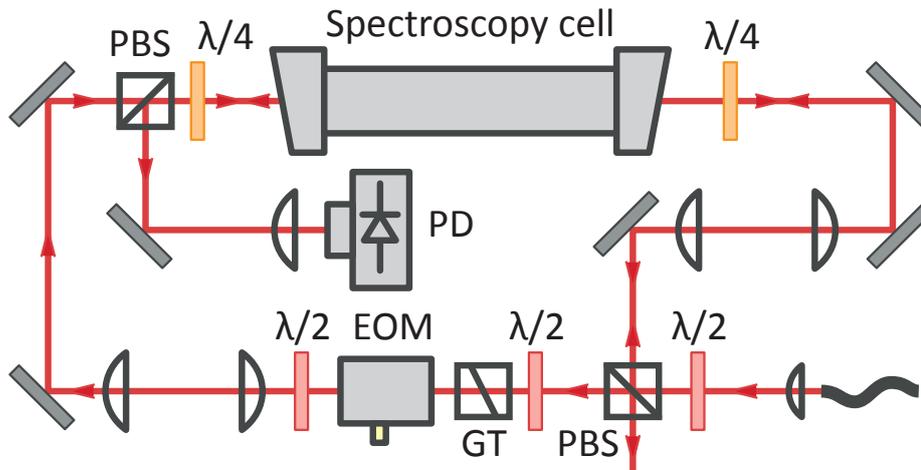


Figure 3.7: The optics setup for the modulation transfer spectroscopy.

The obtained error signal can be used to control the AOM that frequency shifts the light sent to the reference cavity, as seen in Figure 3.1. By using a slow integrator, the frequency of the AOM can slowly follow any drifts of the cavity, keeping the laser referenced to the absolute reference of the atoms. Unfortunately, this control setup has not yet been tested, as I ran out of time because of working on the other projects.

The optics setup required for modulation transfer spectroscopy can be seen in Figure 3.7. Light coupled from the master laser is split into the pump and probe beams using the half-wave plate and the PBS. The pump beam polarization is cleaned using the Glan-Taylor (GT) polarizer and modulated with the EOM. The next half-wave plate sets the polarization to allow the light to pass through the PBS in front of the spectroscopy cell. A telescope is used on both the probe and pump beams to expand the beam to use the whole aperture of the spectroscopy cell. The quarter-wave plates are used to turn the linear polarization of the pump and probe beams to a circular polarization along the spectroscopy cell. The second quarter-wave plate is used to get back to the initial linear polarization. The amplitude of the error signal can be optimized by turning the quarter-wave plates. The now modulated probe beam is reflected on the PBS and collected on our photodetector.

One of the advantages of using the counter propagating probe and pump beams is that they will only be simultaneously resonant with atoms in the zero velocity class. Therefore, our spectroscopy signal does not experience Doppler broadening, caused by the thermal Maxwell-Boltzmann velocity distribution in our atomic cloud.

Chapter 4

Designing low-noise photodetectors

For laser frequency stabilization and linewidth reduction, we have to detect the laser light as reliably as possible. Any noise introduced by the detection method will increase the final linewidth of the lock.

Therefore, we wanted to build photodetectors that we can trust and for which we understand and can quantify the sources of noise. Before we look in more detail into the sources of noise, we need to understand how the light is detected on a photodiode and how a transimpedance amplifier generates a voltage from the photodiode current.

4.1 Transimpedance amplifier

The heart of the photodetector is the photodiode. A more in-depth description on how a photodiode works can be found in the first chapter of Ref. [30]. For this discussion, it is enough to know that the photodiode in the photoconductive mode, which means it is unbiased or reverse biased, acts as a current source producing a current

$$I_{pd} = \eta P, \quad (4.1)$$

where η is the responsivity, which depends on the wavelength of the light, and P is the power of the light collected on the photodiode. Furthermore, we should not forget about the capacitance of the photodiode, C_D which we will see has a large influence on the bandwidth of the photodetector. To reduce C_D and thus increase the bandwidth we usually apply a reverse bias across the photodiode.

Since most of the electronics that we use processes signals as voltages, we want to convert the current from the photodiode into a voltage. The conversion is possible by using only a resistor, but in this case the photodiode's response becomes nonlinear, because of the impedance it has to drive. A better solution is the use of an operational amplifier (op-amp) in the transimpedance amplifier configuration. We can see how a transimpedance amplifier is constructed in 4.1(a). The photodiode is connected to the inverting input of the amplifier, a feedback resistor R_f is put on the negative feedback arm and the non-inverting input of the op-amp is connected to ground.

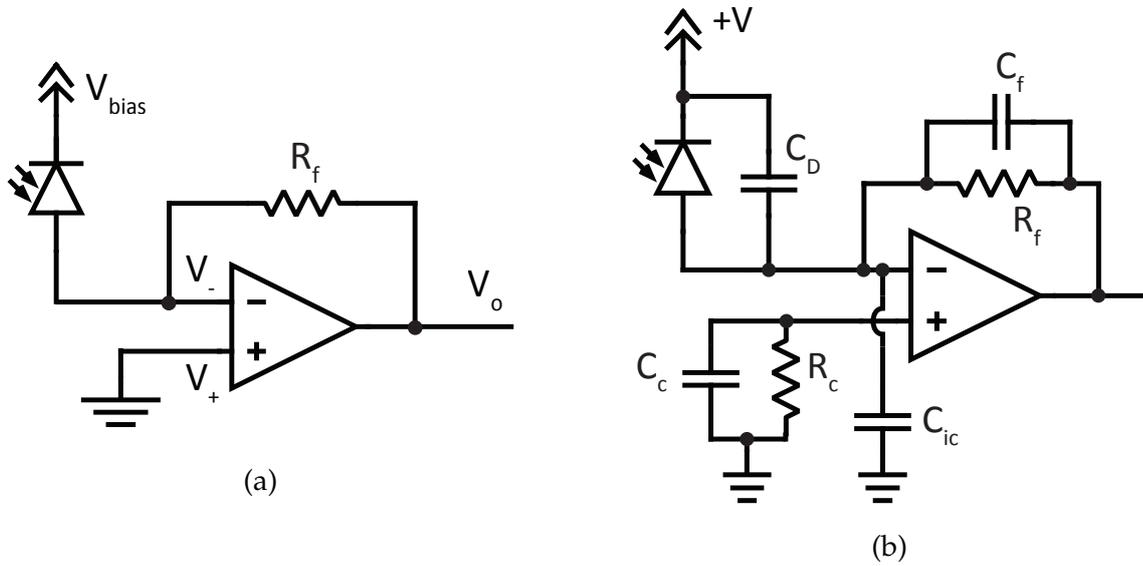


Figure 4.1: The schematics of an ideal transimpedance amplifier (a). In Figure (b) the capacitance of the photodiode, capacitance of the op-amp, and the compensation elements are added.

The idea behind an ideal op-amp is that it greatly amplifies the difference between the voltage of the inverting input pin V_- and the non-inverting input pin V_+ . The gain G is ideally considered to be infinite, so the output of the op-amp

$$V_o = G (V_+ - V_-) \quad (4.2)$$

will in practice shoot up either to the positive or negative supply voltage for any difference between the input signals. As we are using a negative feedback loop, this output is able to modify the signal at the inverting pin. We can show what happens by writing Ohm's law for the feedback resistor R_f

$$V_- - V_o = R_f * I \quad (4.3)$$

and by inserting the expression for V_o . After expressing V_- from the equation and then expanding the final fraction by $1/G$ we get the following result:

$$V_- = \frac{R_f I}{\frac{1}{G} + 1} + V_+ \approx V_+ \quad \text{for } G \rightarrow \infty. \quad (4.4)$$

Taking G to be large, the voltage at the inverting input is forced to the same voltage as at the non-inverting input.

In our case, we have grounded the non-inverting pin, resulting in a so-called virtual ground at the inverting pin $V_- = 0$. As this pin is also one of the pins of the photodiode, the photodiode is always kept at a constant bias voltage, eliminating the non-linear effects we would otherwise observe.

Next, we should look at the output of the transimpedance amplifier V_o . All the current generated in the photodiode I_{pd} has to flow through R_f , because the inputs of the op-amp have a much higher impedance. Therefore, we can rewrite Eqn. 4.3 using $V_- = 0$

$$V_o = -I_{pd}R_f. \quad (4.5)$$

We notice that by selecting the resistance R_f we are also determining the conversion ratio between I_{pd} and the output voltage.

Unfortunately, when building the circuit we have to consider that we are using real components. The op-amp can draw finite input currents, the inputs have a parasitic capacitance C_{ic} and the gain bandwidth of the op-amp is finite. Here, we should also include the capacitance of the photodiode C_D . The schematic of a real transimpedance amplifier can be seen in Figure 4.1(b).

The current drawn into the inverting input steals some current that should be flowing through R_f . This effect results in an offset in the output voltage. To counter this effect, we need to make sure that the non-inverting pin sees the same impedance as the inverting pin. Putting a compensation resistor R_c before the non-inverting input with the same value as R_f will compensate for the offset we would otherwise measure. At the same time, we add a capacitor C_c parallel to R_c to ensure that the non-inverting input is properly grounded for high frequencies.

A feedback capacitor C_f is added to improve the stability of the op-amp by phase compensation. A detailed derivation on how the compensation works and on the optimal value for the feedback capacitance can be found in Ref. [30]. However, for experimental purposes it is usually better to determine the value by trial-and-error, while measuring the output spectrum of the photodetector. A good starting value is $C_f = \frac{C_D}{2}$. An example of determining the correct value will be discussed in a later section and can be found in Figure 4.6.

4.1.1 Bandwidth

There are different effects that can limit the bandwidth of a transimpedance amplifier. We will consider only the most straightforward limit caused by the capacitances of our real components. These capacitances are the limiting factor for large transimpedance gain. At low gain another bandwidth limitation arises from the limited bandwidth of the op-amp interplaying with the phase compensation in the feedback loop [30].

The parasitic capacitance around the inverting input of the op-amp acts as a bypass to R_f for high-frequency signals. The bypass can be understood as a RC low-pass filter with the cutoff frequency given by

$$\nu = \frac{1}{2\pi RC} = \frac{1}{2\pi R_f(C_f + C_D + C_{ic})}, \quad (4.6)$$

where C_{ic} is the parasitic input capacitance of the inverting pin of the op-amp.

4.1.2 Noise sources

When analyzing the transimpedance amplifier, we are able to consider several sources of noise. The derivation in this section follows Ref. [31]. Let us look in detail at some of the sources that have the largest contributions.

The first and most fundamental source of noise is the shot noise, which occurs because the photons are quantized. The photons arrive at the photodiode at discrete times, following a Poissonian distribution. The electrons that get excited on the photodiode will follow the same distribution. In this case, the standard deviation in photon and thus electron number is proportional to \sqrt{N} and the noise is white for all frequencies. We can write the current noise power spectral density as

$$S_i = \sqrt{2qI_{pd}}, \quad (4.7)$$

where q is the unit charge and I_{pd} the DC current from the photodiode. This current noise is transformed to voltage noise by the transimpedance resistor R_f . As different sources of noise are additive in power or voltage square, we will be interested in the mean squared voltage power spectral density of shot noise

$$\overline{S_v^2} = 2qI_{pd}R_f^2. \quad (4.8)$$

The second source of noise we consider is the Johnson noise of the feedback resistor R_f . This noise arises from the thermal movement of charge carriers in the resistor. A derivation and explanation of this type of noise can be found in [32]. The final expression for the mean squared voltage power spectral density is

$$\overline{S_v^2} = 4k_bTR_f, \quad (4.9)$$

where k_b is the Boltzmann constant and T the temperature of the resistor.

The next two important sources of noise are simply properties of the selected op-amp and can normally be found by looking at the datasheet. The corresponding mean squared voltage power spectral density is

$$\overline{S_v^2} = e_n^2 + i_n^2R_f^2, \quad (4.10)$$

where e_n is the op-amp's input noise voltage power spectral density and i_n the input noise current power spectral density.

To get an idea on how these sources of noise compete with each other, we have to select the operational amplifier and photodiode we will be using for our amplifier. A good candidate for the op-amp is the wide-band and low noise LMH6609 from Texas Instruments. The photodiodes that were used for testing were S5973 from Hamamatsu. The calculated mean squared voltage power spectral densities for these components are plotted in Figure 4.2.

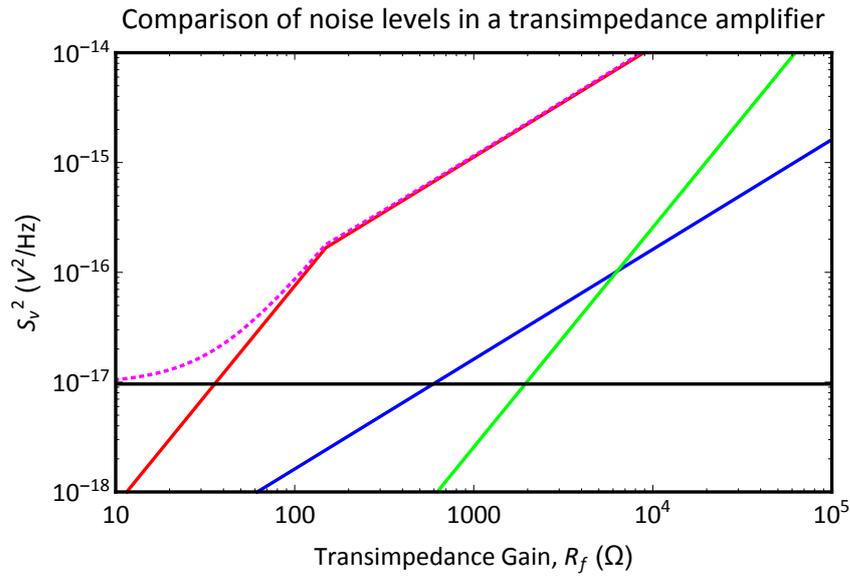


Figure 4.2: A comparison of noise levels for LMH6609 and S5973 photodiode, depending on the selected transimpedance gain R_f . We are assuming the maximal power of light we can put on the photodiode without damaging it or saturating the op-amp output. The total noise level (dotted, magenta) is composed of the contributions from shot noise (red), Johnson noise (blue), op-amp current noise (green) and op-amp voltage noise (black).

The current I_{pD} for the calculation was chosen as the maximal current before the laser power damages the photodiode P_{max} , or at higher gains, the maximal current before hitting the voltage swing limitation of the op-amp V_O .

$$I_{pd} = I_{max} = \min \left[\eta P_{max}, \frac{V_O}{R_f} \right], \quad (4.11)$$

where η is the responsivity of the photodiode for the selected wavelength.

If we have enough laser power available, the shot noise dominates over the other sources of noise for practically all interesting values of the transimpedance gain.

4.2 Designing the photodetector

Now we have some insight on how a transimpedance amplifier works. However, just the amplifier itself is not enough to make a usable photodetector. The device itself needs to deliver its signal to other devices, it has to be able to accept a supply voltage, protect the board against a faulty input and to regulate down the voltage according to the needs of the integrated circuits. Furthermore, we wanted the board to be versatile, to be able to choose between AC or DC outputs, to use different photodiodes and different gains.

The latest version of the electronics circuit can be seen in Figures 4.3 and 4.4. Figure 4.4 shows how the supply voltage is distributed and regulated. First, there are power diodes right after the voltage inputs, to protect the board against wrong polarity of the supply voltage. It can easily happen that a supply cable is faulty or soldered incorrectly, so the diodes will block the current from flowing in the wrong direction and protect the rest of the board. Next, a LC filter is added to cut off high frequency and reduce the amount of noise that can enter the board. Afterwards, low noise voltage regulators are used to regulate down and smooth out the voltage to ± 10 V and ± 5 V which is needed for the operational amplifiers. The capacitors values are taken from the datasheets to assure low-noise behavior.

Next, let us look at the signal path shown in Figure 4.3. Furthest to the left, we see that there is room for three different positions of the photodiode with different choice of bias voltage. The jumpers JP1 and JP2 allow us to change the bias voltage, which makes the board usable with different kinds of photodiodes. After the jumpers, there is another LC filter that suppresses and decouples AC signals. The different photodiode positions D1, D2 and D3 allow us to both choose a positive or a negative reverse bias and also to be able to use two photodiodes simultaneously and amplify the difference of their signals.

After the photodiode, comes the transimpedance amplifier we have already discussed. However, we have three possible configurations of the transimpedance amplifier, which allow us to choose whether the outputs are AC-coupled or DC-coupled. These configurations are controlled by the boxed symbols, R1, R4, R5, R6, R7, R8 and R17. The values needed are written on the schematic, but the general idea is that with option A, the detector is outputting two equivalent DC signals, with option B, the signal path is separated before the transimpedance amplifier, allowing us to choose different gain for the AC and DC signals. With option C, we separate the signals after

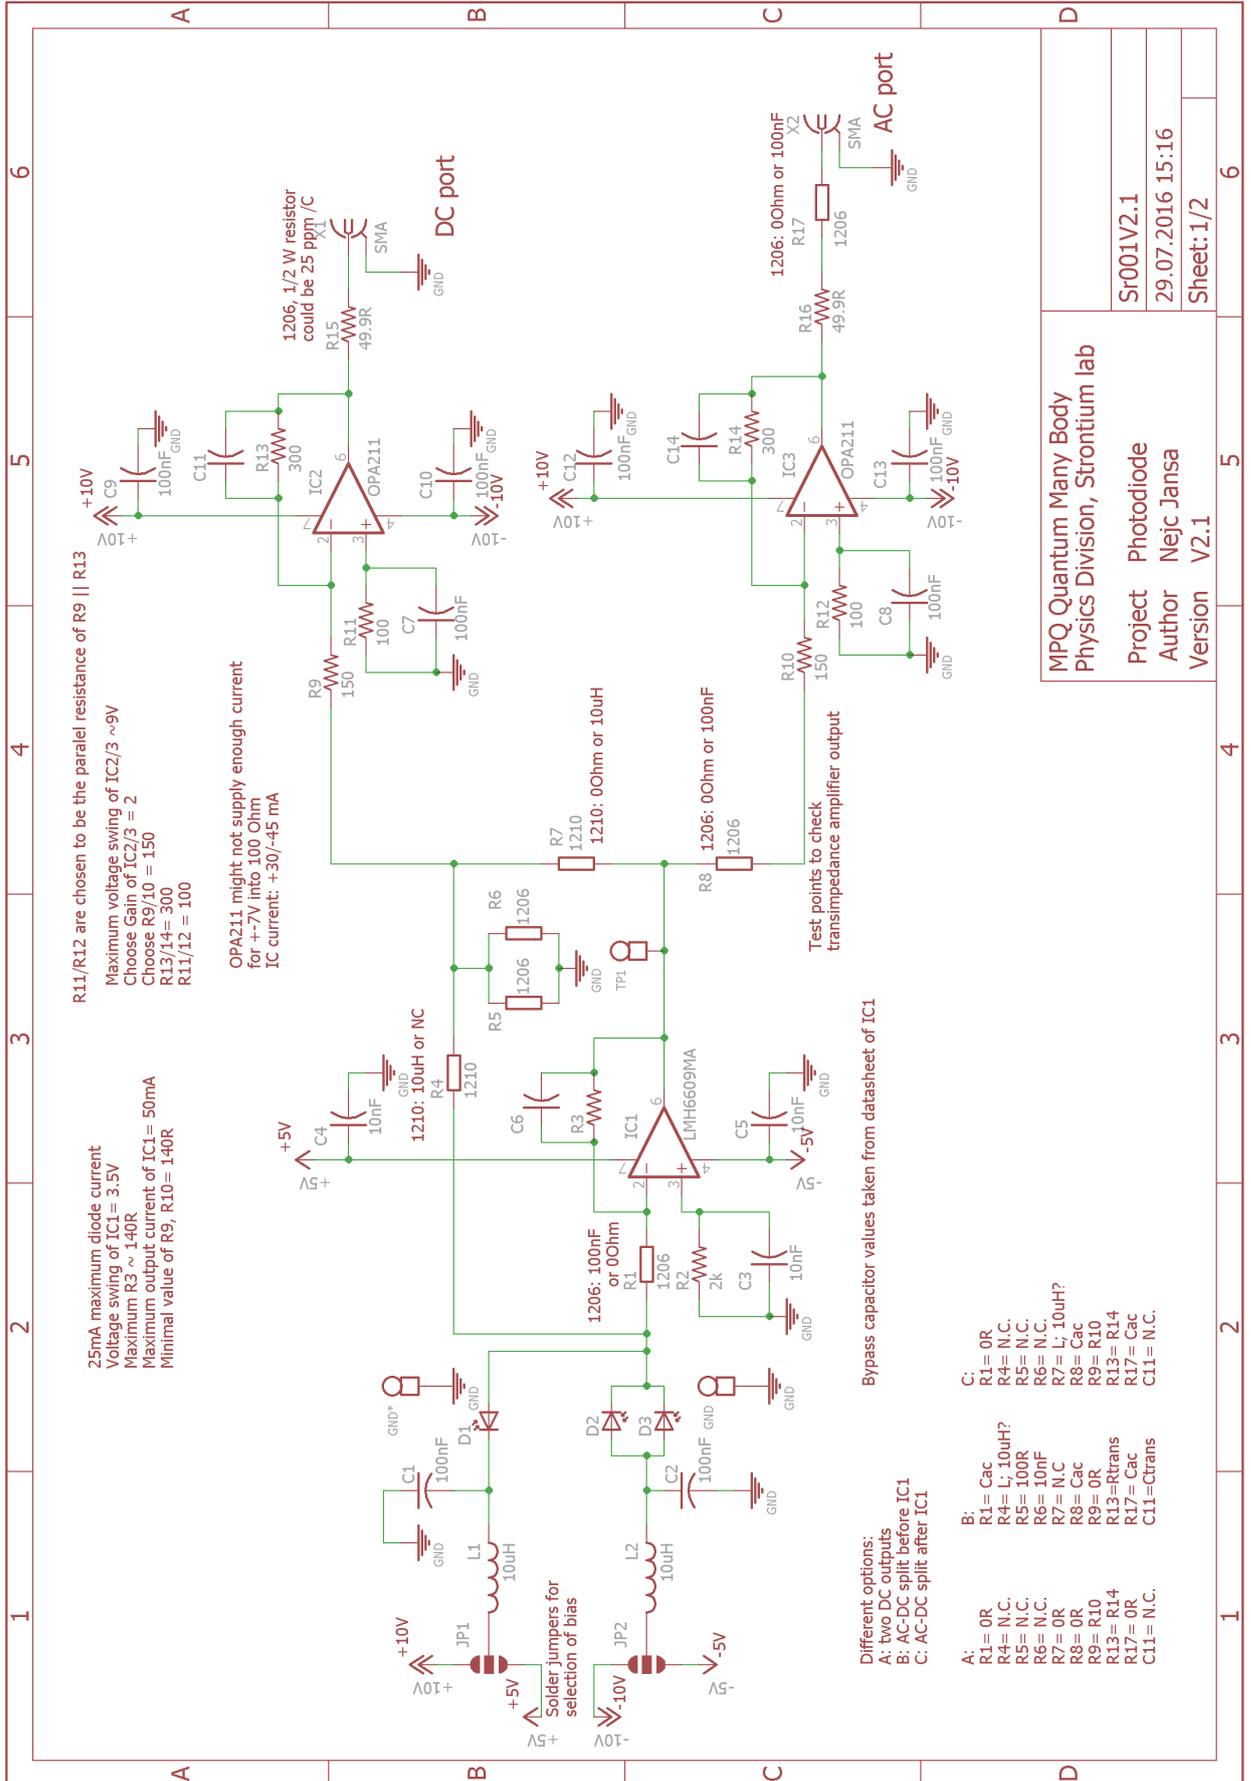


Figure 4.3: Schematics of the photodetector, signal paths.

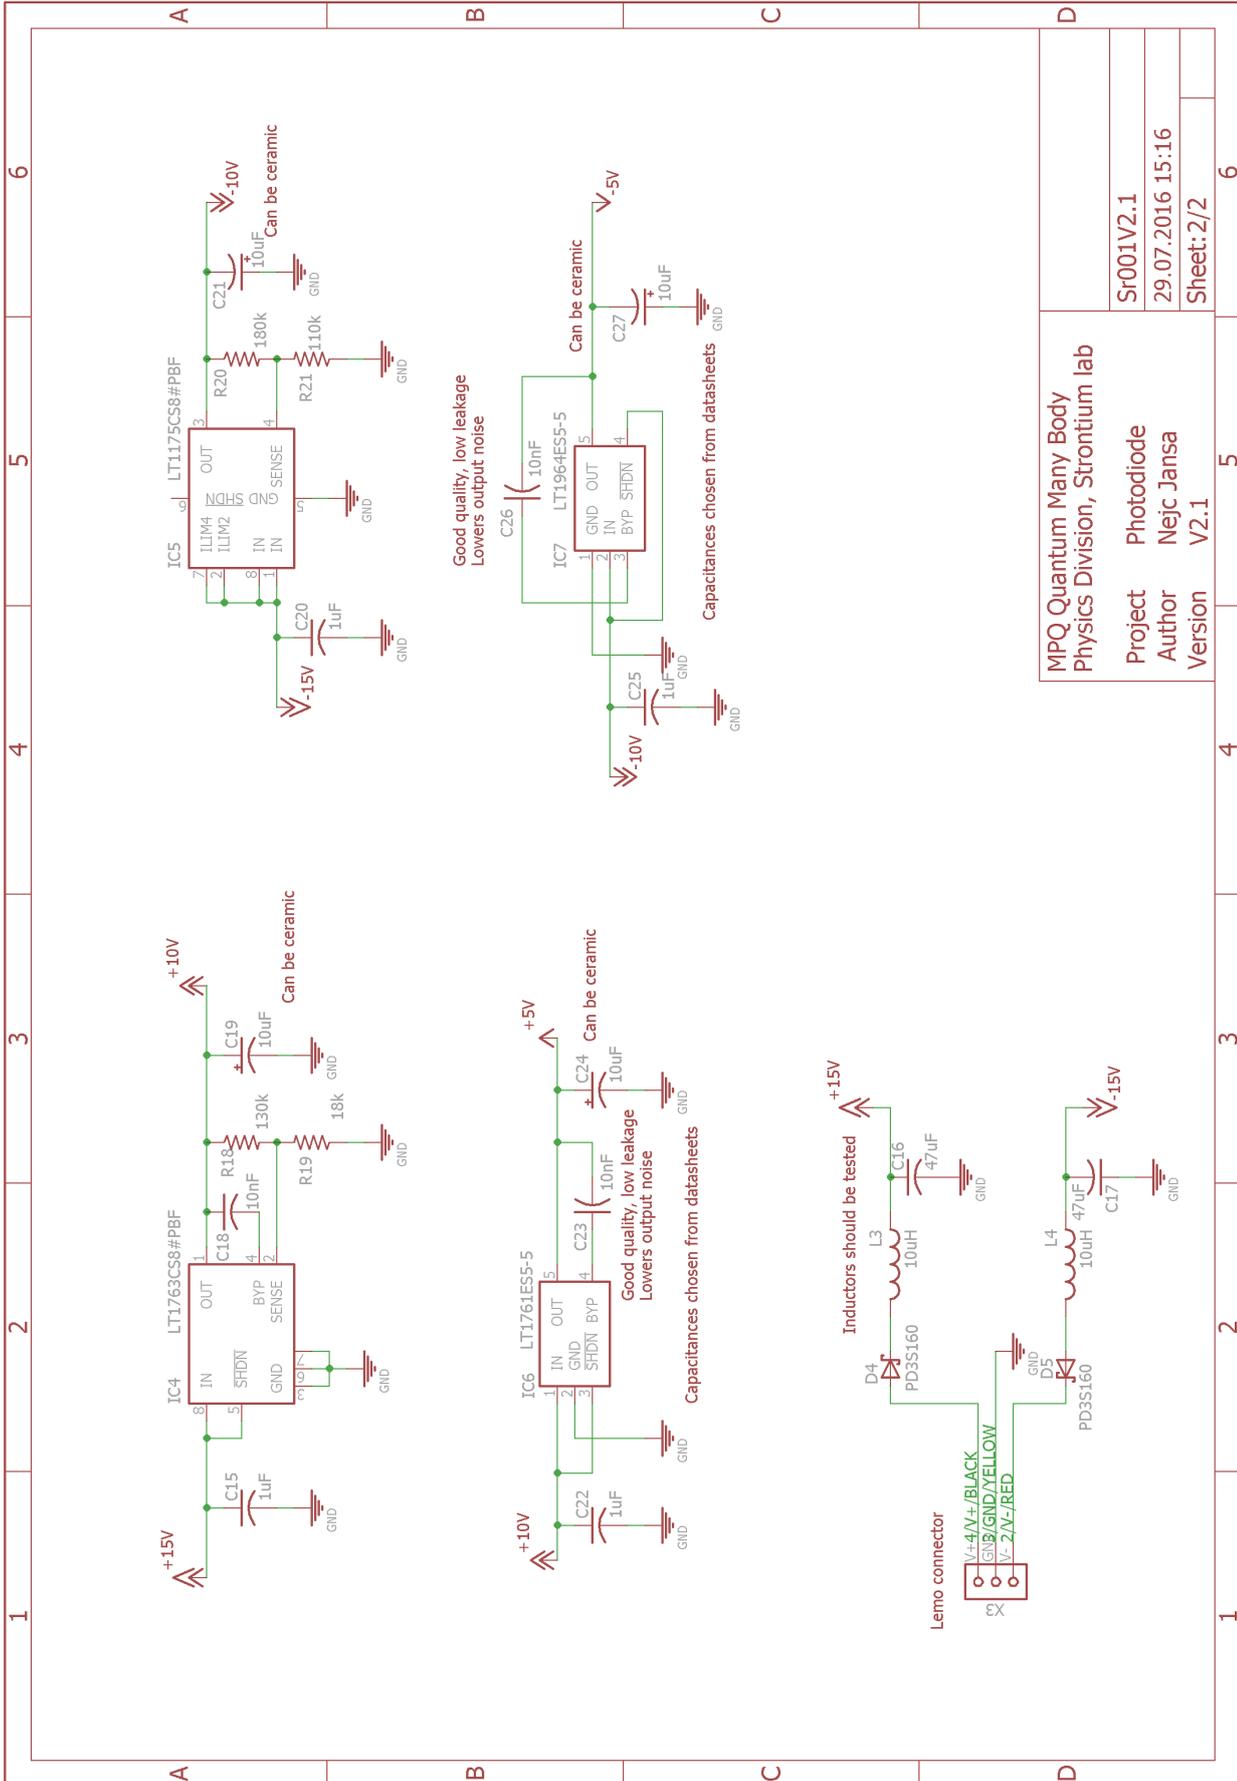


Figure 4.4: Schematics of the photodetector, voltage regulators.

the transimpedance amplifier, which fixes the transimpedance gain to be the same, but seems to perform better when it comes to noise.

To deliver the signal to other devices using coaxial cables, we added a second amplification stage, comprised of inverting amplifiers with gain 2. This amplification brings the signal voltage up to a higher rail voltage and ensures that there is enough current to drive the coaxial cable. The $50\ \Omega$ resistors in front of the connector ensure that the amplifiers are decoupled from any reflections from the coaxial cable. Here, we should note that the choice of $50\ \Omega$ resistors means that the voltage driving a $50\ \Omega$ load will be half of the voltage we would measure on a high impedance oscilloscope. All the op-amps have bypass capacitors close to their supply pins, to protect the IC from ripples.

The design of the printed circuit board (PCB) can be seen in Figure 4.5. The geometry of the board was determined by an existing design, as we want it to be compatible with the existing housings and connectors. The design of the housing allows the board to be easily removed or swapped in an optics setup.

During the routing of the PCB, care was taken to keep the signal paths short and separated from the supply voltages. Large ground planes surround all the other paths to assure an even ground potential and to serve as a heat sink. Furthermore, the components were properly marked and organized to make hand soldering a painless experience.

To reduce unnecessary noise and the change of signal over time, resistors were chosen to be precise and have low thermal coefficients. Care was also taken to choose the correct nominal power of the resistors.

4.3 Testing the device

In this section, we take a look at some of the measurements done on the latest version of the photodetectors. The photodetector was hand soldered using the above discussed components. For the transimpedance gain, we chose $R3 = 3\ \text{k}\Omega$ and the board was soldered in the DC configuration. The capacitor C6 was for now left out. First, we wanted to analyze the transimpedance amplifier. Resistors R7, R8, R16, and R17 were removed to isolate the transimpedance amplifier. The test point TP1 was connected to R17 using a wire and the $49.9\ \Omega$ resistor was soldered in R17 place. This configuration allows us to use the connector of the AC port to measure the direct output from the transimpedance amplifier.

First, I plugged the power connector to the board and tested that all of the voltage regulators are working correctly. Using a voltmeter, I tested all of the supply volt-

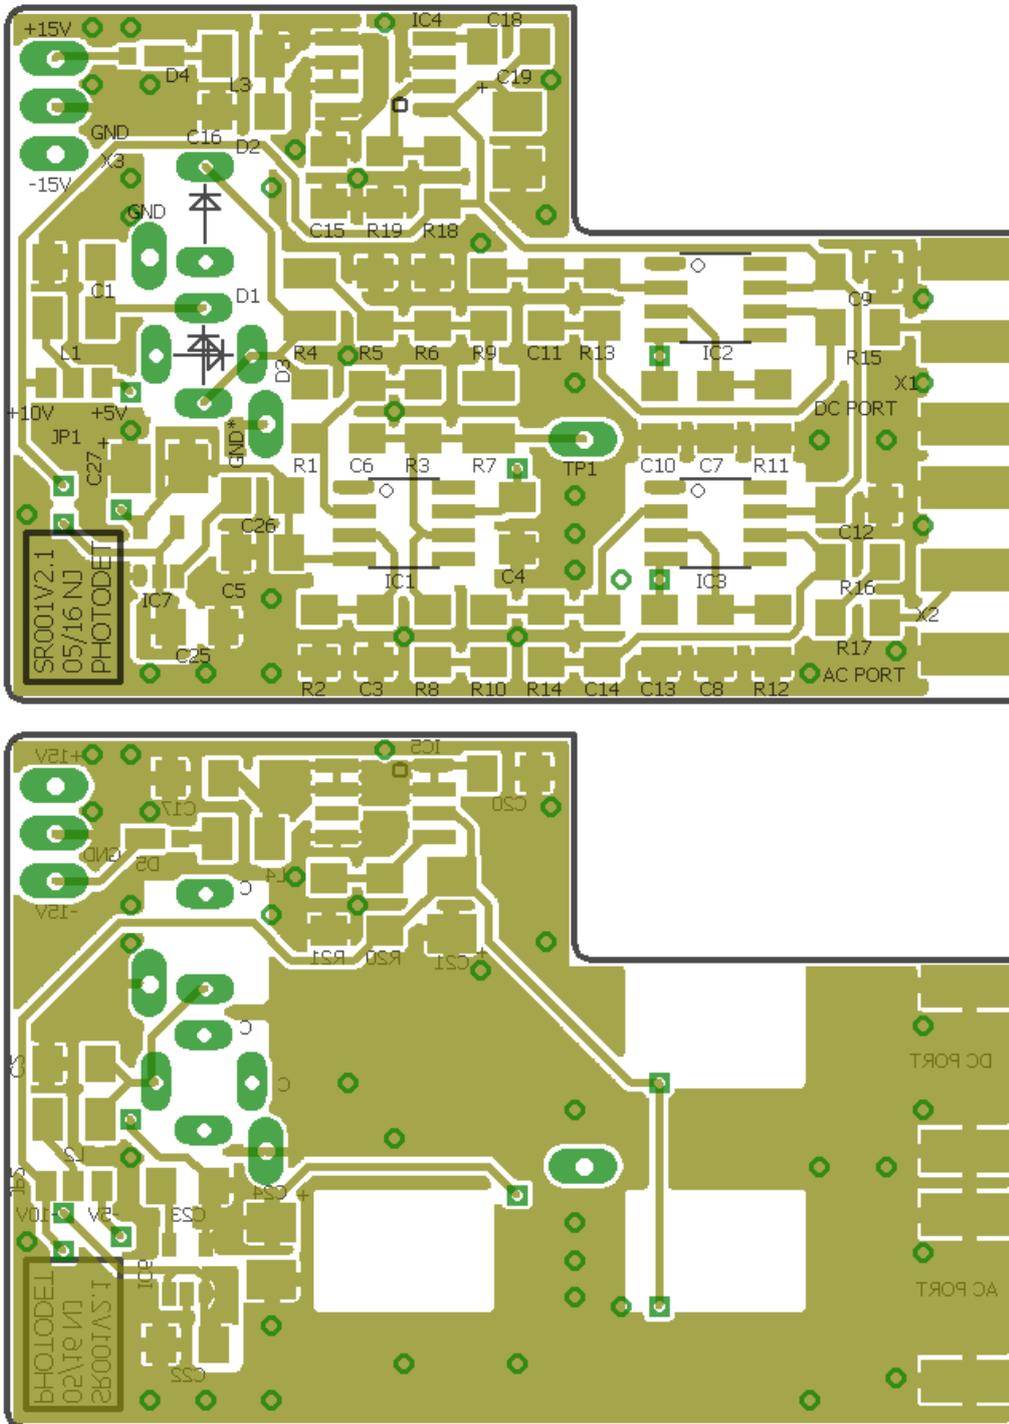


Figure 4.5: Top and bottom sides of the PCB.

ages and they indeed were within 1% of the desired value. Next, the photodetector board was assembled into its case and mounted on an isolating post, to prevent coupling of the ground to the optical table. The aperture of the photodetector case was blocked and I tested the offset voltage of the transimpedance amplifier by connecting the AC port to an oscilloscope. The mean value of the voltage was -0.3 mV, which is reasonable compared to the offset of the selected op-amp.

Next, I tested the spectral response of the transimpedance amplifier using a spectrum analyzer. The spectrum analyzer was set to scan between 0 and 100 MHz with a resolution bandwidth (RBW) of 300 Hz. The reference level was set to -90 dBm and the detection method to rms. Here, I should note that for data comparison we want to normalize the output of the spectrum analyzer to a bandwidth of 1 Hz. Therefore, I divided all the shown data from the spectrum analyzer with the used RBW to give numbers in units of dBm/Hz. First, I recorded the noise floor of the spectrum analyzer with only a connected coaxial cable. The noise floor was at approximately -152 dBm/Hz apart for some small rise in the low and high frequency ends of the scan. Furthermore, we noticed that the spectrum analyzer only works well when running on battery. While it is being charged, a lot of additional noise peaks get added onto the measurement. Unfortunately, later during the measurement the battery ran low and some of the data was taken with additional noise.

Finally, the output of the transimpedance amplifier was connected to the spectrum analyzer and recorded. As I did not yet connect the capacitor C6, the circuit had a resonance at around 42 MHz. The measurement was first taken with the photodiode blocked and afterwards with 3 mW of 461 nm laser light focused on the photodiode. Afterwards, the measurement was repeated for different values of C6. The capacitors with such small capacitance are only available in smaller packages and we bought a few with different capacitance for testing. The measurements can be seen in Figure 4.6.

The effect of changing the feedback capacitor C6 is evident. As we are increasing the capacitance, the resonance peak starts disappearing and moving slightly to lower frequencies. However, it is a bit ambiguous whether we should look at the resonance in the case with laser light or without. In either case, the capacitor with 1.5 pF should be a good choice, as it flattens the response in the case with the light and leaves an approximately 3 dB peak when there is no light. I should also note, that the resonance peak seems to be at a frequency higher than the expected bandwidth of this transimpedance amplifier. The selection of the feedback resistor and the capacitance of the diode and the op-amp would give an estimated bandwidth of 12 MHz. Therefore, the seemingly increased bandwidth could just be an artifact of the selected op-amp (LMH6609).

In Figure 4.6(a), the dashed line corresponds to the Johnson noise of the feedback

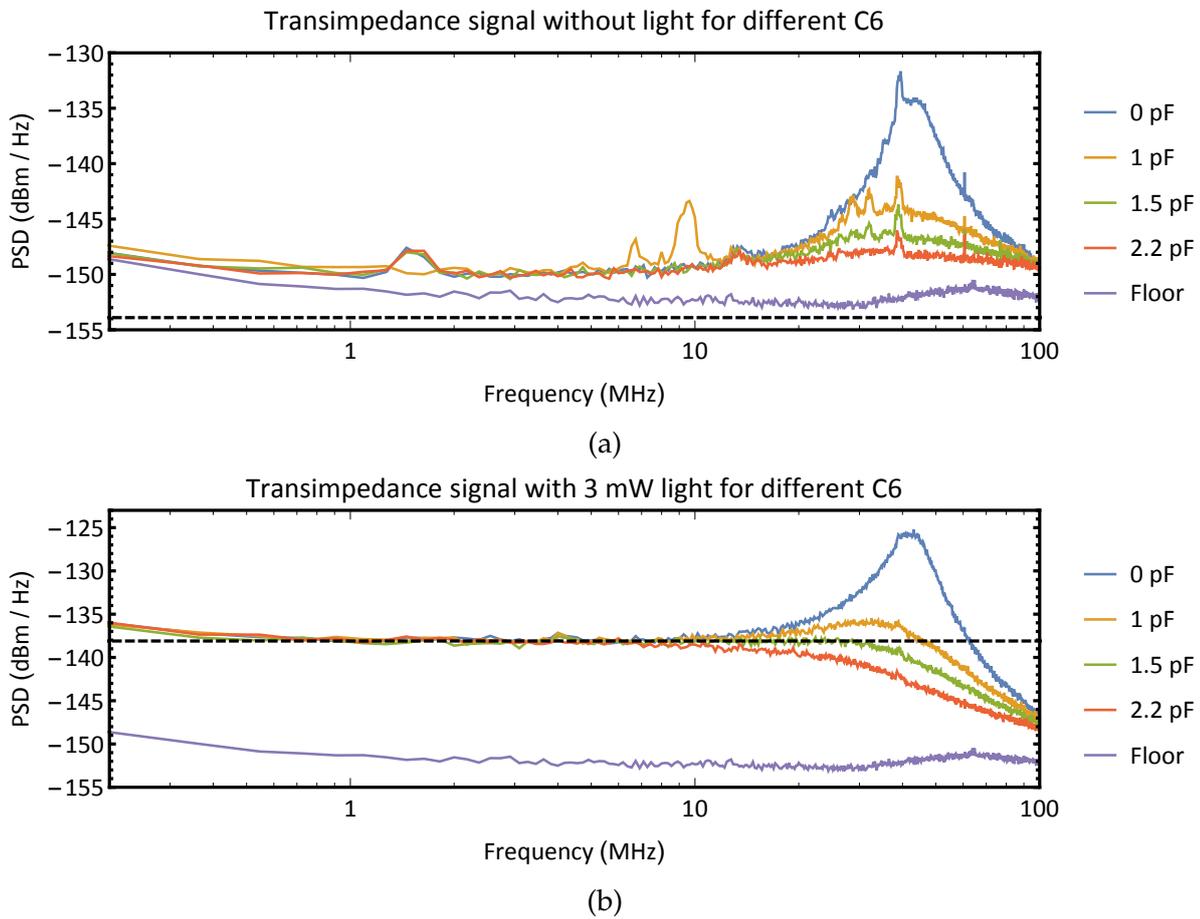


Figure 4.6: Data from the spectrum analyzer measuring the output from the transimpedance amplifier. The noise floor of the spectrum analyzer is shown in purple. The black dashed line shows the theoretically calculated noise level expected for the transimpedance amplifier.

resistor. As the noise level is higher, it can be a consequence of the too high noise floor, or some other unknown source of noise. The observed noise level would correspond to shot noise from 0.2 mW of laser light. Therefore, this detector configuration should be shot noise limited for powers above 0.2 mW. The agreement can be seen in Figure 4.6(b). The calculated shot noise level of -138.1 dBm/Hz agrees well with the data. The mean squared voltage power spectral density from Eqn. 4.8 has to be converted to power spectral density (PSD) using

$$\text{PSD} = \frac{\overline{S_v^2}}{R}, \quad (4.12)$$

where R is the value of the resistor over which the power is measured. Because the transimpedance amplifier is driving both the decoupling resistor (49.9Ω) in front of the coaxial cable and the measurement resistor in the spectrum analyzer (50Ω), we have to be careful about where the power is actually dissipated. As these two resistors act as a voltage divider and we only detect the power on the measurement resistor, where the voltage is halved, we will only measure $\frac{1}{4}$ of the actual power. A quick way to get the correct result is by inserting an effective resistance $R = 4 \times 50 \Omega$. The final result is transformed to units of dBm using the definition

$$\text{PSD(dBm)} = 10 \log_{10} \left(\frac{\text{PSD}}{1 \text{ mW}} \right). \quad (4.13)$$

Chapter 5

Design of the reference cavity housing

If we want to reliably stabilize a laser and reduce its linewidth, we need a frequency reference with the required stability. A great candidate for the reference are high-finesse optical cavities, which transmit light only when the length of the cavity is equal to an integer number of half wavelengths of the light. The great thing about the optical cavities is that their length can be very stable when put in the correct environment.

5.1 Changes of cavity length

There are several effects which can cause changes in the cavity optical length. Therefore, these effects can change its absolute frequency or add unwanted noise and thus increase its linewidth.

We see the changes in the optical length if we look at the frequency of the light that is resonant with the cavity

$$\nu = \frac{Nc}{2Ln'} \quad (5.1)$$

where N is an integer number, c the speed of light, L the geometrical length of the cavity, and n the index of refraction of the medium filling the cavity.

If we differentiate the equation, taking into account that ν , L and n are not constants, and divide it by Eqn. 5.1 we see that the fractional stability of the frequency depends on both the fractional stability of length and on the index of refraction:

$$\frac{d\nu}{\nu} = -\frac{dL}{L} - \frac{dn}{n}. \quad (5.2)$$

5.1.1 Thermal noise

A fundamental limit to the linewidth of the cavity is the thermal noise, which arises from the fluctuation dissipation theorem [33]. The authors of Ref. [34] calculate and show that, with the correct choice of material, the biggest thermal contribution boils down to the coating of the mirrors. There is not much that we can do about this source

of noise, as there are not many alternative material choices for the mirror coating and because we do not want the additional complication of operating at cryogenic temperatures. Therefore, our goal is to try bring any other source of noise down to the level of this limit, which allows the fractional stability of length L or frequency ν on the order of $\frac{\partial L}{L} = \frac{\partial \nu}{\nu} \approx 10^{-15}$. This constraint limits us to the level of 1 Hz stability, but that is enough for the purpose of our experiment. The timescales where we want to address the relevant transitions will be shorter than a second due to the short experimental cycles.

A great collection of knowledge on how state-of-the-art improvements to the thermal noise limit are achieved can be found in Ref. [35].

5.1.2 Thermal expansion

An important contribution to the change of cavity length is thermal expansion. Every real material has a non-zero coefficient of linear thermal expansion (CTE)

$$\alpha_T = \frac{1}{L} \frac{\partial L}{\partial T}, \quad (5.3)$$

where $\frac{\partial L}{L}$ is the fractional change of its length, driven by a change in temperature ∂T .

We can choose materials which are designed for their low CTE, such as ultra low expansion (ULE) glass with a mean coefficient of thermal expansion $\alpha_{T,ULE} = 0 \pm 30 \times 10^{-9} \text{ K}^{-1}$ in the 5 °C to 35 °C range and a zero-crossing in α_T near room temperature or similarly Zerodur with the mean $\alpha_{T,Zer} = 0 \pm 100 \times 10^{-9} \text{ K}^{-1}$ in the 0 °C to 50 °C range.

Unfortunately, the whole cavity is not just a simple rod, but has a three dimensional structure and the mirror substrates are usually from a different material than the spacer. The difference in the thermal expansions of the two materials will result in effects such as the deformation of the mirrors. The authors of Ref. [36] investigate how these effects change the zero-crossing of the CTE using finite element simulations. In the end, the exact zero-crossing also depends on the batch of ULE from which the spacer was made, so it is always useful to measure it for the individual finished cavity.

We should also note, that even as we try to park the temperature of our cavity very near the zero-crossing of the CTE, making it well below 10^{-9} K^{-1} we will have to keep the temperature stable on the order of $< 1 \text{ mK}$ to achieve the desired 10^{-15} fractional stability. Therefore, we have to make a system that is very temperature stable and isolated from the daily temperature fluctuations of the environment.

5.1.3 Vibrations

The next effect that can add on to the linewidth of the cavity, is the vibration or the elastic deformation due to external forces. The materials have a property called Young's modulus, the tensor version of the spring constant, which tells us how much a material deforms under a certain force applied to an area:

$$E = \frac{FL}{A\partial L}, \quad (5.4)$$

where F is the force acting on the object, A is the initial cross section, L the initial length and ∂L the change of length due to the force. The external forces entering our system as vibrations will cause fractional changes of length according to

$$\frac{\partial L}{L} = \frac{F}{EA}. \quad (5.5)$$

The Young's modulus of the already mentioned materials is $E_{\text{ULE}} = 67.9 \times 10^{12}$ Pa for ULE and $E_{\text{Zer}} = 90.3 \times 10^{12}$ Pa for Zerodur, which is quite similar to the ones for metals as aluminum or copper.

By choosing materials with large Young's modulus or taking large cross sections A , we might be able to reduce the fractional length change by a small factor, but what we really want is to damp the vibrations to suppress them by several orders of magnitude. For this purpose rubbery materials are used that can dissipate the energy. Many commercial passive or active isolation stages exist that have great isolation properties. Apart from the vibrations that come through the ground and the table, we also have to ensure some isolation from sound that travels through air.

Another great way to decrease the sensitivity of the cavity is by supporting it in a way, so that even when it is deforming, the net change in cavity length cancels out. This effect is investigated in Ref. [37] for a similar notched cavity design that we are using. Through finite element analysis, it was shown that there are points where we can mount the cavity, for which the sensitivity to vertical accelerations crosses zero.

5.1.4 Pressure changes

Another effect that can alter the optical length of the cavity is the change in the index of refraction of air n . To get a rough estimate on what order of changes to expect from fluctuations in n , we can extrapolate the index of refraction of air for lower pressures. At standard pressure, n for air is given by $n_0 \approx 1 + 3 \times 10^{-4}$. For perfect vacuum $n_{\text{vac}} = 1$ by definition. Neglecting that air is made up from different components and that it is not an ideal gas, we assume a linear relation for the index of refraction

$$n(p) \approx 1 + \frac{3 \times 10^{-7}}{\text{mbar}} p, \quad (5.6)$$

where p is the pressure of air that fills our cavity.

By differentiating, neglecting higher order corrections and using Eqn. 5.2 we arrive at

$$\frac{\partial \nu}{\nu} \approx -\frac{3 \times 10^{-7}}{\text{mbar}} \partial p. \quad (5.7)$$

If we want to keep the fluctuations smaller than those of other contributions, we need to have a pressure stability of roughly 10^{-8} mbar. For such stability, it is necessary to house the cavity in vacuum and pressures in the range $< 10^{-6}$ mbar should be sufficient, as long as the pressure is not fluctuating a lot.

5.2 Choice of material and geometry

Taking into account all the considerations we made about the frequency stability of the cavity, we already have quite some constraints on how the cavity should be housed.

Commercial systems are available, such as the Vacuum Housings from Stable Laser Systems. However, due to the high price and the possibility that it can be done equally well or better, we decided to design our own vacuum housing. At the same time, we still bought the commercial cavity, because of the amount of expertise and specialized equipment involved.

Our design included a lot of ideas used in the commercial dual stage housing from Stable Laser Systems and the “JILA cavity” as described in Ref. [27]. The rendered model showing our design can be seen in Figure 5.1.

5.2.1 ULE cavity

As already mentioned, we bought a commercial notched cavity from Stable Laser Systems, more precisely the ATF-4020-4 model machined by Advanced Thin Films. The notched spacer made from ULE has a diameter of 50 mm and length $L = 100$ mm. The length results in a free spectral range $\text{FSR} = \frac{c}{2L} = 1.5$ GHz. The mirrors on a fused silica substrate are planar ($R_1 = \infty$) and concave ($R_2 = 500$ mm), with a reflective coating specified to give a finesse of $> 200\,000$ at 695 nm.

With the cavity we also received a measurement of the zero crossing of the CTE, as discussed in Chapter 5.1.2. The measurement from the received certificate can be seen in Figure 5.2(a). Since the zero crossing is at 49.6 °C, which is fairly above room temperature, we can design the temperature stabilization stages with only heating elements.

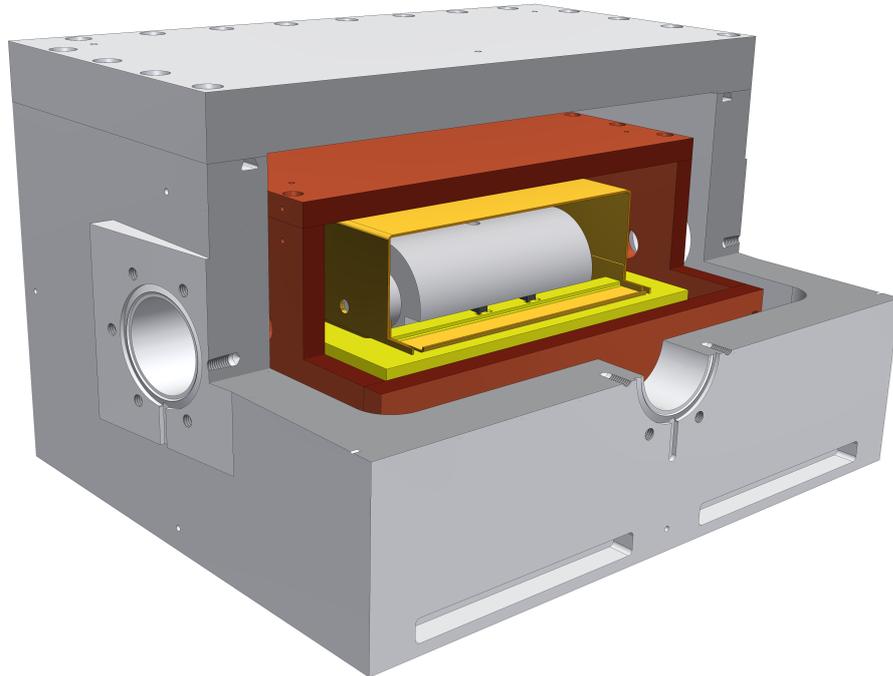
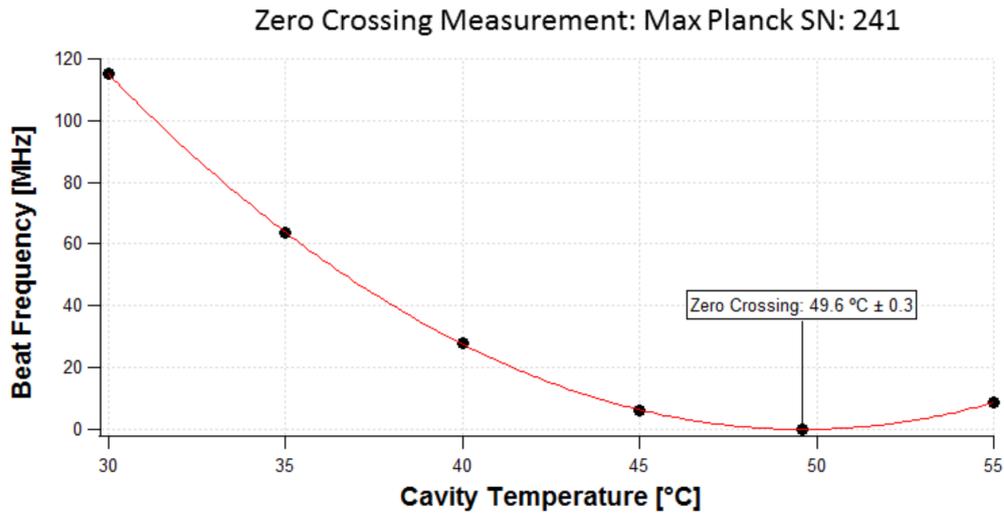


Figure 5.1: A cut showing the components inside the designed vacuum housing. From inside out, we can see the ULE cavity supported by Viton pads on the Zerodur spacer. The cavity is enclosed by a gold coated copper radiation shield. The Zerodur spacer rests on larger Viton pads inside a temperature stabilized copper box. The copper box rests on another set of Viton pads placed in the aluminum vacuum chamber. The chamber has several CF 40 interfaces to connect the viewports on the side, electrical feedthroughs in the back and a pumping arm in the front.

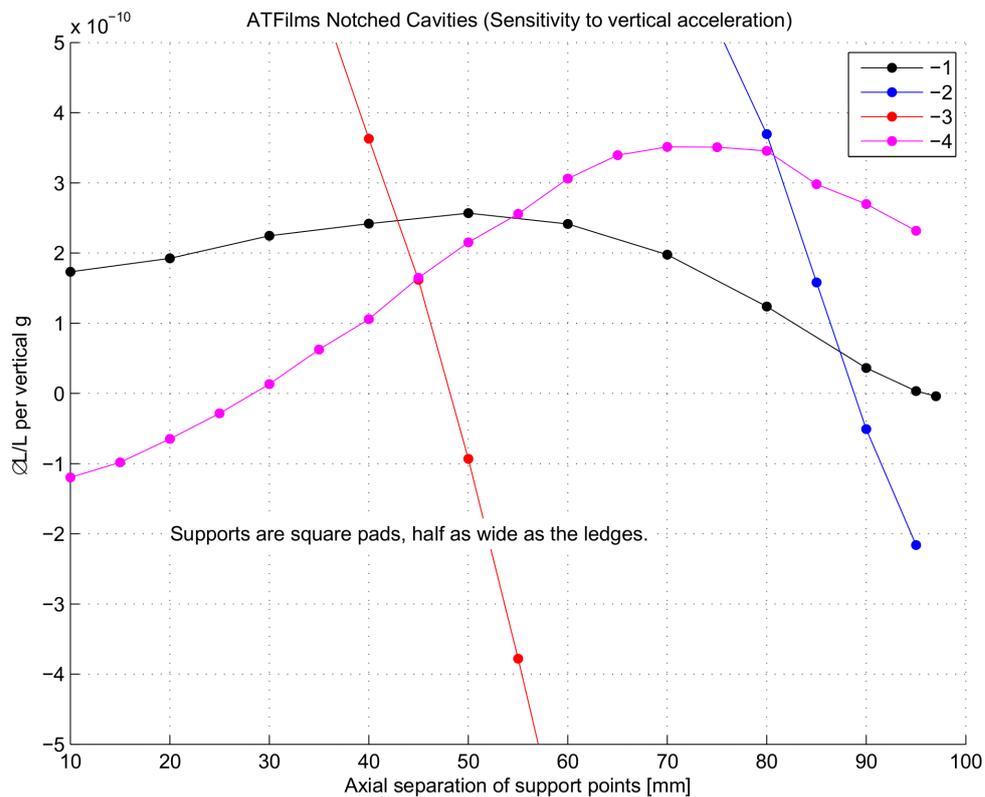
We were also able to receive the information on the mounting points for which the cavity length changes are insensitive to vertical vibrations. The finite element analysis was provided by T. Rosenband from NIST and the results can be seen in Figure 5.2(b). For our model ATF-4020-4, the sensitivity to the position of the pads is small and the zero crossing is at a separation of the support points of approximately 28.5 mm. This value was taken into account in the design of the other parts.

5.2.2 Zerodur spacer

The plate on which the cavity is mounted is made from Zerodur, due to its small coefficient of thermal expansion, $\alpha_{T,Zer} = 0 \pm 100 \times 10^{-9} \text{ K}^{-1}$. The low expansion



(a)



(b)

Figure 5.2: (a) The plot from the certificate of the Zero Crossing Temperature Measurement of ULE Cavity Spacer, provided by Stable Laser Systems. (b) Finite element analysis provided by T. Rosenband, NIST. The pink trace "-4" corresponds to our model.

offers an additional decoupling of the cavity from thermal changes and means that also the position of the pads supporting the cavity will be stable.

The supporting pads are made out of Viton. This material was chosen for its low outgassing properties allowing it to be used in high vacuum. The benefit of a softer material for the support points of the cavity is discussed in Ref. [37], where the authors show that in this case the sensitivity to horizontal vibrations is also decreased. Additionally, the rubber has low thermal conductivity allowing for decoupling of temperature, which is discussed in the later sections.

The 3D model of the Zerodur plate can be seen in Figure 5.3(a). On the top of the plate, there is a small ledge with square extrusions that mechanically reference the position of the supporting pads. The extrusions have an additional spiral groove and “ears” that should prevent air from being trapped under the Viton pad. Removing air is important for high vacuum applications, as a gas bubble forming between the rubber and Zerodur surfaces would act as a virtual leak, degrading the quality of the vacuum.

The raised ledge also serves as a mechanical stop for the radiation shield which is described in the next section.

On the bottom of the plate, there are four more extrusions. These are slots for larger Viton pads, which support the Zerodur plate and at the same time mechanically reference its position. Again, a spiral groove and holes in the edges prevent air from being trapped between the two surfaces.

5.2.3 Radiation shield

The idea behind the radiation shield is to homogenize the thermal radiative links between the ULE cavity and the environment and to act as a passive low-pass filter for changes in temperature. More about how the thermal shielding works is discussed in Section 5.3.

The material was chosen to be copper, due to its excellent heat conductivity $\kappa_{\text{Cu}} = 393 \frac{\text{W}}{\text{m K}}$ and a high specific heat capacity $c_{\text{Cu}} = 385 \frac{\text{J}}{\text{kg K}}$. As the part will have to be in high vacuum, we also have to consider which type of copper we want to use to achieve low outgassing rates. We want the copper to have the oxygen content as low as possible, and have found the copper alloy EN CW008A to be a suitable and available candidate. Furthermore, if we want the shield to have even better radiative properties it can be polished or coated with gold, to get a very low emissivity, down to $\epsilon_{\text{Au}} = 0.02$.

The geometry of the radiation shield can be seen in Figure 5.3(b). It was designed so that the shield covers as much of the solid angle around the cavity as possible,

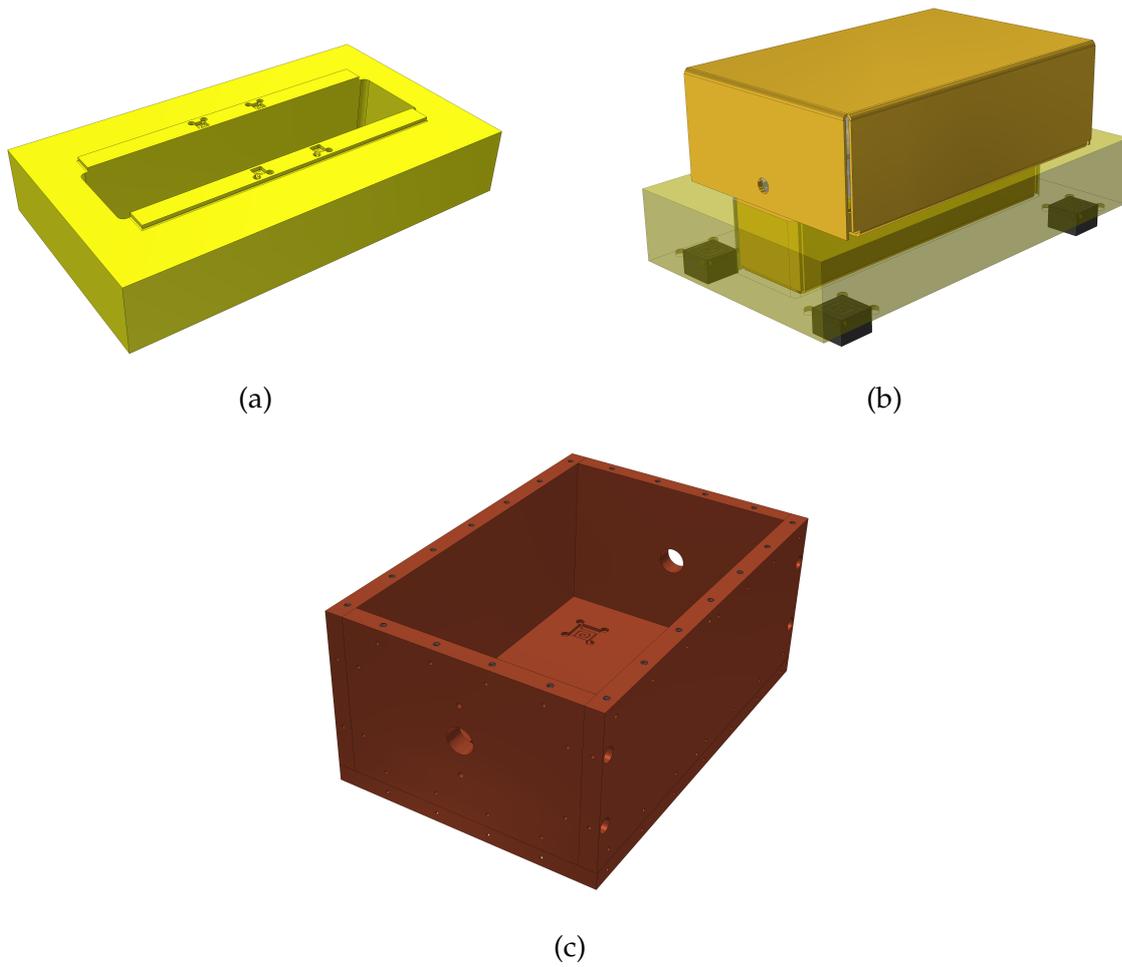


Figure 5.3: A render of the mechanical design for the Zerodur spacer (a), the radiation shield resting on top of the spacer (b) and the copper box without the lid (c).

while still allowing it to be properly mounted on the support pads and making sure the shield cannot crash into the cavity risking damage or scratches. Due to the mounting constraints, the thickness is chosen to be 1mm. Since the system also has to be assembled in a simple way, the shield is made out of two parts, one of which sits under the cavity and the second that covers it from above. The holes in the shield allow for optical access to the cavity inside.

5.2.4 Copper box

The copper box is meant as another radiation shield, but is in addition actively temperature stabilized. The material of choice is again oxygen free copper, and for these parts we even managed to find a higher grade alloy EN CW009A, which has even lower oxygen levels and with it better vacuum and thermal properties.

The geometry was chosen to be a box, due to the simplest mounting and referencing. At first a two-part box was envisioned, from a lid and a hollowed-out block, which would not suffer from bad thermal contacts where sections are screwed together and would be easier to assemble. Unfortunately, the desired high quality material was not available in the correct thickness, so a compromise was made, for all sides to be thin plates. With many screws to hold the pieces together and good quality of the contact surfaces the heat properties should not suffer too much. To prevent air from getting stuck under the screws, additional venting holes were made to release gas from under the head of the screw and from the end of the tapped hole.

For mounting and referencing of the Zerodur spacer and for mounting the box itself inside the aluminum chamber, again Viton pads were used. For this reason, the bottom wall of the box includes the square extrusions and spiral venting grooves.

As we want to temperature-stabilize this box, small holes were added on different points around the outside of the box, that can fit a thermistor head. The thermistors can be thermally contacted to the box using thermal epoxy. Care should be taken not to leave air bubbles in the holes when applying epoxy.

5.2.5 Aluminum vacuum chamber

As we already discussed in one of the earlier chapters, we want the whole assembly to be housed in vacuum. Apart from the pressure fluctuations affecting the stability of the cavity, vacuum also helps with temperature homogeneity and stability, as it removes air currents caused by convection. Furthermore, the absence of air also prevents acoustic vibrations from reaching the cavity from all directions.

Usually, most vacuum chambers are made from stainless steel grades 304 or 316, due to their great vacuum properties. These materials are also used for many of the

commercial vacuum components, employing the conflat (CF) flange sealing technology. Unfortunately, stainless steel does not have good thermal properties considering our application, as the thermal conductivity of these two alloys is rather low $\kappa_{SS} = 16.2 \frac{W}{m K}$. The specific heat capacity is $c_{SS} = 500 \frac{J}{kg K}$.

As we hope to use the vacuum chamber itself as another active thermal heat shield, a better candidate is aluminum. A high grade of aluminum has to be used, that is suitable for use in vacuum and is easily machinable. We used the alloy EN AW 5083 with O3 tempering, which homogenizes the material. The thermal properties for this alloy are $\kappa_{Al} \approx 120 \frac{W}{m K}$ and $c_{Al} = 900 \frac{J}{kg K}$.

The use of aluminum, however, brings some additional problems. Because aluminum is softer than stainless steel and copper, it is not possible to use copper gaskets in the CF seals, as the sealing knife edge would easily get damaged. To avoid the damage, we have to use elastomer gaskets for sealing. The best candidate is Viton, due to its low outgassing and low gas permeability properties. However, even Viton is quite permeable to noble gasses, so using it will limit the vacuum we can achieve to the level of 10^{-8} mbar. By using Viton gaskets, we can machine the CF flange connectors directly onto the chamber and connect any CF parts necessary.

From Figure 5.1 we can get a pretty good idea on what the chamber looks like. Just like on the other pieces, there are extrusions for Viton pads on its inner bottom floor, referencing the copper box to it. The whole chamber has a thickness of 25 mm, giving it a large heat capacity and enough space to machine a sealing surface for the lid. The lid is sealed with a large Viton O-ring, which sits in a trapezoid groove. On the air side of the groove there are several holes for screws, allowing the ring to be compressed evenly between the chamber and the lid.

We can also see some of the CF 40 interface knife edges cut directly into the chamber. The ones for viewports on the sides are cut at a $\sim 7^\circ$ angle, preventing unwanted optical etaloning. Additionally, the chamber has pockets on the outer bottom edges, allowing clamping to the optical breadboard. The bottom side of the chamber is also lifted a bit in the center, so that the contact with the breadboard is defined by the bottom edges. Last of all, some small holes were added on the outside of the chamber for attaching of the temperature sensors.

5.2.6 Other components

In Figure 5.4, we see what the whole cavity housing looks like after assembling. The aluminum chamber sits on a double density optical breadboard made from 20 mm thick aluminum. The custom cut allows it to fit on the vibration isolation stage without blocking the protection nuts that are removed when the stage is floated. It sits on

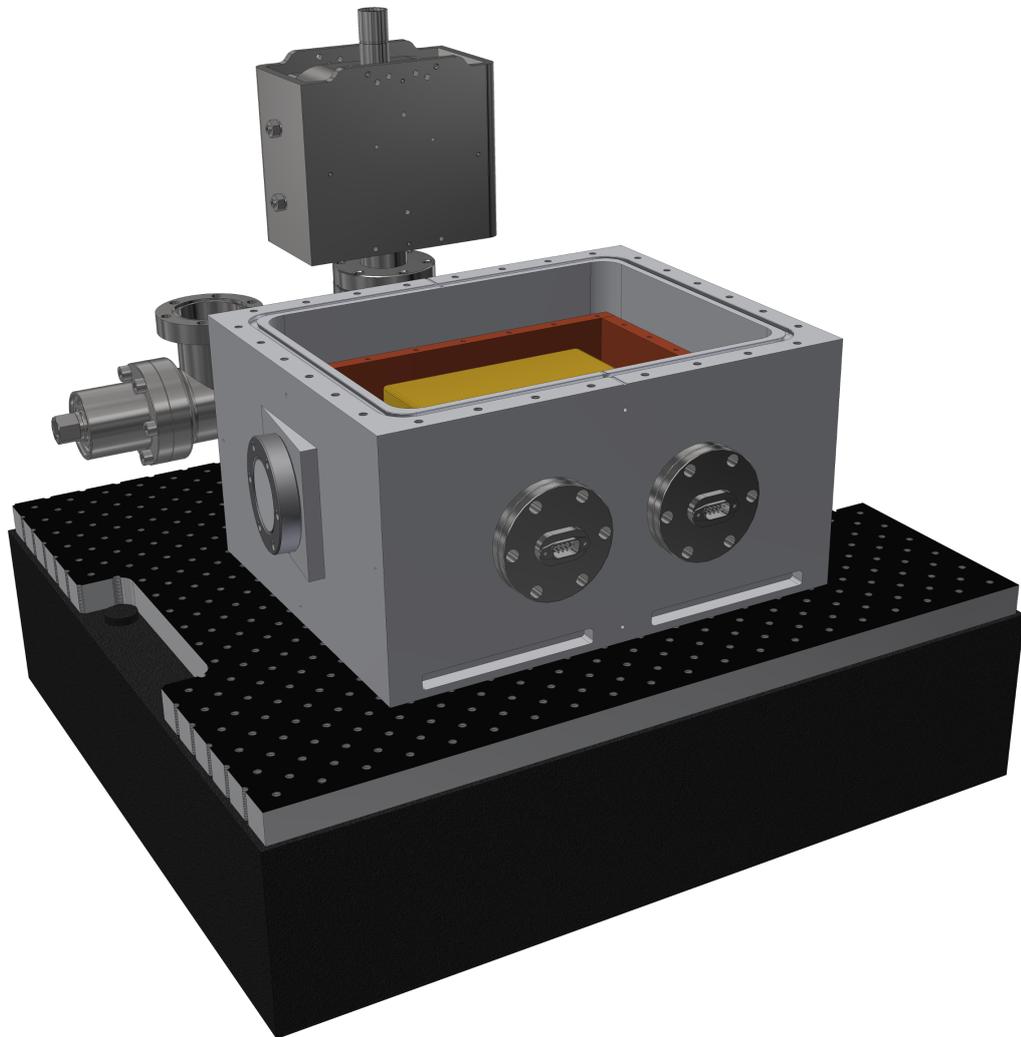


Figure 5.4: The whole assembled vacuum chamber sitting on the optical breadboard and the vibration isolation stage. On the side, the viewports are attached, electrical feedthroughs are in the front and in the back the pumping cross with an 10 L/s ion pump, an angle valve and a vacuum gauge.

four circular pads from Sorbothane, whose hardness and size were chosen to damp the vibrations of the aluminum breadboard. A tool for calculating the correct pads can be found on the Sorbothane website. The breadboard is intended for mounting of the optics necessary for interrogating the cavity.

The vibration isolation stage was bought from MinusK. The chosen model, 125BM-8 had a very suitable size and very low natural vibration frequencies, 0.5 Hz vertically and 1.5 Hz horizontally. An additional nice feature of these stages is that they are completely passive and require no pressurized gas or electricity.

To access the heating components required to temperature stabilize the copper box, two CF 40 electrical feedthroughs are added. These feedthroughs are made to connect D-sub 9 pin cables, allowing to monitor and drive four pairs of heaters and thermistors.

For optical access to the cavity, the two CF 40 viewports are mounted on both sides of the cavity. These viewports have a commercial broadband AR coating and were purchased from Trinos. The reason to use viewports on both sides is that we eventually want to use the cavity to lock two lasers, the 689 nm laser and the 698 nm clock laser. At the same time, the two viewports allow us to observe the transmission through the cavity which is very helpful when trying to mode-match the incoming light.

To pump the chamber and monitor the pressure, a 5-way cross is added on the back. One of the flanges is unnecessary and the commercial 5-way cross could be replaced with a custom one without the bottom arm. To monitor the vacuum level, we used a combination gauge, PKR 361 from Pfeiffer Vacuum, which displays the pressures between 10^{-9} and 1000 mbar.

For rough pumping on the chamber, an angle valve is added to connect to the bellows from a turbo pump. The easy-close all-metal CF 40 angle valve from VAT was used, as its technology allows for an excellent vacuum seal that can be closed numerous times without the need for maintenance or replacements.

The ion pump chosen was the 10SW DI ion pump from Gamma Vacuum. The pump has a differential pumping element, which allows for pumping of noble gasses. One of the limiting factors we expect for the final vacuum pressure is the argon and neon permeating through the Viton sealing. The size of the pump was chosen to be 10 L/s, as bigger pumps would be too heavy and would not fit with the geometry of the chamber. Smaller pumps would not make good use of the pumping conductance from the chamber and would thus result in a lower final pressure.

A lot of information on how to calculate conductance and in general a lot one needs to know about vacuum can be found in Ref. [38]. To find a reasonable pumping speed for our system, a really rough estimation on the conductance from the alu-

minum chamber to the ion pump was made. If we assume the 5-way cross and the wall of the aluminum chamber to be a series of long tubes, we can use the equation for the conductance of a long tube in the molecular flow regime for air at 22 °C [38].

$$C_{\text{tube}} = 121 \frac{\text{m}}{\text{s}} \frac{d^3}{l}, \quad (5.8)$$

where d is the diameter of the tube and l is the length. The inner diameter of our tubes is $d = 37$ mm and the total length of two cross arms plus the thickness of the chamber wall is $l = 151$ mm. These values give us an estimate for the conductance $C \approx 40 \frac{\text{L}}{\text{s}}$. The limited conductance affects the actual pumping speed of the pump, giving us the effective pumping speed

$$\frac{1}{S_{\text{eff}}} = \frac{1}{S} + \frac{1}{C}, \quad (5.9)$$

where S is the nominal pumping speed of the pump. An infinitely large pump would give us a pumping speed of only 40 L/s and the 10 L/s pump gives us almost all of its speed, $S_{\text{eff}} = 8$ L/s.

5.2.7 Heating components

The last parts that have to be included in the vacuum system are the components needed to temperature-stabilize the copper box. All these parts should be vacuum compatible so that we do not unnecessarily degrade the vacuum any further.

The first thing needed is something to measure temperature. A great candidate are the Omega Series 55000 thermistors. These have great stability and come with bare gold plated leads, which should not be a problem for the vacuum. As the thermistor needs to be in thermal contact with the material, it is best to submerge it into small drilled holes filled with thermal conductive epoxy. We are using EPO-TEK's H74 and H77 two component thermal adhesives, due to their low outgassing rates and electrical isolation properties, which prevent the leads from getting shorted.

Next, we need heaters, to add heat to the system and regulate its temperature. To homogeneously heat the surface of the box, we would want to use large heaters to cover most of the box surfaces. A candidate for the desired vacuum levels are Kapton tape heaters, as they are flexible and come in all sorts of shapes. We tested the Kapton Insulated Flexible Heaters from Omega. A problem that might arise here is that the leads may not be vacuum compatible and would in this case need to be replaced.

To connect all these components we also need vacuum compatible wire. The wire should not have an insulation that could outgas too much. Ideally bare wire would be used, but then ceramic beads or spacers have to be used to prevent the wires from getting shorted. We found a Kapton insulated flat cable to be a suitable alternative.

The materials are UHV compatible and the 9 stranded wire should be convenient to connect to the 9 pinned connectors. As solder should not be used for in-vacuum purposes, the wires can be connected using contact clamps or can be crimped.

5.3 Thermal properties

As we already found out in Section 5.1.2, our cavity is very sensitive to changes in temperature. Our system will be standing on a temperature stabilized optical table, but there is still some coupling to external temperature changes like the day and night cycle and due to other non-constant sources of heat. To reduce the effects of these changes, thermal shields can be built. The authors of Ref. [39] derive a mathematical model for the heat transfer through passive heat shields and the resulting transfer function. It is shown that the thermal shields act analogous to an electronic low-pass filter and that the components can be modeled as linear elements with resistances and capacitances. We will look at a simplified model based on these principles, to get an idea on what materials we should use and to roughly estimate the heat transfer time constants.

5.3.1 RC analogy of heat transfer

In fact, both the transfer of heat and the linear models of electrical circuits can be described by the same linear differential equations. Therefore, it is not surprising that we can find an analogy between the two systems. The difference in electric potential ΔV can be mapped to a temperature difference $\Delta T \hat{=} V$. The electric current or the flow of charge $I = \frac{dq}{dt}$ then corresponds to heat flow $P = \frac{dQ}{dt} \hat{=} I$.

From the correspondence it follows that the capacitance C is defined as both how much charge q a capacitor stores for a certain voltage, or how much heat energy Q a material can store after a certain temperature change:

$$C = \frac{q}{\Delta V} \hat{=} \frac{Q}{\Delta T} = mc, \quad (5.10)$$

where the heat capacitance can also be written as a product of mass m and specific heat capacitance c .

To get an expression for the resistance, we can write the analogy of the Ohm's law

$$\Delta V = RI \quad \Rightarrow \quad R = \frac{\Delta V}{I} \hat{=} \frac{\Delta T}{P}. \quad (5.11)$$

First, we look at the linear heat conduction for a rod with cross section A and

length l

$$P = \kappa \frac{A}{l} \Delta T, \quad (5.12)$$

where κ is the thermal conductivity of the material. By using Eqn. 5.11 we can get the resistance for the linear conduction

$$R_c = \frac{l}{\kappa A}. \quad (5.13)$$

For the radiative link, we look at the Stefan-Boltzmann law, which tells us how much power a black body radiates:

$$P = A \epsilon \sigma T^4, \quad (5.14)$$

where A is the surface area of the body, ϵ is the material's emissivity, T its temperature and $\sigma = 5.67 \times 10^{-8} \frac{\text{W}}{\text{m}^2 \text{K}^4}$ the Stefan-Boltzmann constant.

We are interested in the heat exchange P_{ij} between two cubes where one is set inside of the other. However, deriving the exact equation becomes quite tedious because of the complicated geometry and the series of reflections of radiation due to the emissivity factor. We will skip this step and take the final equation from Ref. [39], where the authors argue that for differences in the size of the two shields, that are not too large, we can consider the cubes to be spheres. Therefore, we can use the expression for the radiative resistance

$$R_{r,ij} = \frac{\beta_{ij}}{4\sigma A_j T_0^3} \quad (5.15)$$

where A_j is the area of the second body, T_0 is the average temperature and β_{ij} is a geometric factor depending on area and emissivity;

$$\beta_{ij} \approx \frac{1}{\epsilon_j} + \frac{1 - \epsilon_i}{\epsilon_i} \left(\frac{A_j}{A_i} \right). \quad (5.16)$$

5.3.2 Heating time constant

Finally, we recognize our components as heat capacitances connected by resistivities of their thermal links. However, as not all properties are well known, such as emissivity of Zerodur and ULE glass, or the actual thermal links between the radiation shield and the Zerodur plate, severe simplifications are made. A schematic of the system as a RC circuit can be seen in Figure 5.5. This model is far from a complete description of the system, but it should be enough to give us a rough estimate on the time scales of the heat transfer through our system. Using this model, we can look at what happens to the temperature of the cavity and the other internal components, when we modulate the temperature of the outside chamber or the copper box.

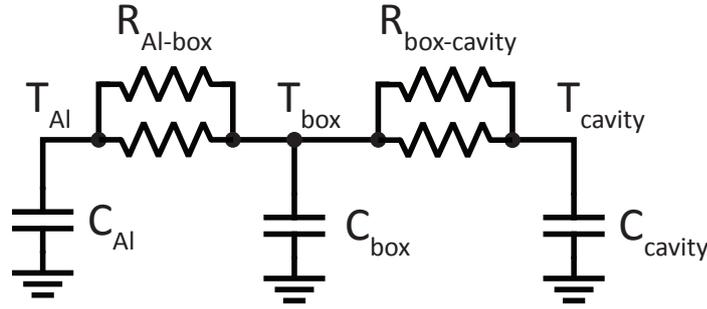


Figure 5.5: A simplified RC diagram for the heat properties of the system. The radiative and conductive thermal links are added in series. The capacitance for the cavity includes the Zerodur spacer and the radiation shield.

In the schematic we recognize two low-pass RC filters between the external temperatures and the temperature of the cavity. If we take the total series resistance of both thermal links

$$\frac{1}{R} = \frac{1}{R_c} + \frac{1}{R_r}, \quad (5.17)$$

we can get an expression for the cut-off frequency of the filter

$$\nu = \frac{1}{2\pi RC'}, \quad (5.18)$$

or a perhaps a more intuitive result, the characteristic time scale $\tau = RC$, where the temperature reaches $1/e$ of the final value.

Because we want to achieve low cut-off frequencies, we need to design the system in a manner such that R and C are large. Looking back at Equation 5.10 we realize why we wanted to select materials that have large thermal capacities and we know that we should make them thick and massive. Here, we should not forget an implicit assumption which we made, namely that the individual pieces are at a homogeneous temperature. This assumption is only true for materials with large thermal conductivity.

We can get values for the masses and areas of the parts with the help of a mechanical CAD software, such as Autodesk inventor. Using these numbers and looking up the thermal properties of the used materials, we can calculate a rough estimate of what the time constant of our system will be. Here, I would like to note that the calculated radiative links can be quite different in reality, as we took numbers for emissivity which can vary by a lot in practice. The values taken are $\epsilon_{Au} = 0.02$, $\epsilon_{Cu} = 0.04$ and $\epsilon_{Al} = 0.10$. Furthermore, the obstruction of the gold plated radiation shield by the

Zerodur spacer is ignored. We assume that it is transparent, although that is not true for the whole spectral range.

The resulting heating time constant is $\tau_{\text{box,cavity}} \approx 70$ h between the copper box and its insides made up of the cavity, spacer and radiation shield. Furthermore, the time constant to the cavity itself is expected to be even longer, as the conduction has to pass through another set of support pads. If that is taken into account, the conductive and radiative links have a comparable value.

The second time constant is $\tau_{\text{Al,box}} \approx 69$ h between the aluminum chamber and the copper box. Here, an interesting fact is that the resistance of the radiative link is much lower than the conductive link. It follows, that almost all heat will be homogeneously transferred through radiation.

Since the time constants are calculated to be almost 3 days long, we can expect that the daily temperature changes of the environment will be greatly suppressed. We can take a look at the transfer function of a single low pass filter to see how much of a 24 h period fluctuation would get transmitted

$$T = \frac{1}{1 + (\omega\tau)^2} = \frac{1}{1 + \left(\frac{2\pi}{24\text{ h}} \times 70\text{ h}\right)^2} = 0.003. \quad (5.19)$$

The result looks quite promising, but the time constants should definitely be double checked by measuring the response of the real system once it is fully assembled.

5.4 Assembling the vacuum chamber

Before the vacuum chamber can be assembled, we have to make sure that the parts are clean enough not to degrade the vacuum level we can achieve. During machining different oils and lubricants are used, that would outgas far too much to be used in high vacuum. Normally, the parts will already be somewhat cleaned but it is hard to rely on them all to be of the same quality, as they can be treated with more or less care, leaving behind fingerprints, oil stains and dust. Since almost every part is made from a different material, I will shortly describe how we decided to treat the components.

5.4.1 Cleaning of components

At all times clean gloves should be worn to prevent leaving fingerprints on the now clean components. Paper-free cloth will prevent any residual dust to be left after wiping things. We use HPLC grade solvents, always freshly opened, which in turn evaporate cleanly without leaving stains. Wrapping things in clean vacuum grade aluminum foil will keep the parts clean and reduce the amount of dust and water that

can stick to them.

The aluminum chamber was first wiped across all surfaces using isopropanol applied to a cloth. Next, the chamber was sonicated in soap for 5 minutes, choosing the soap (Liquinox) that is not too aggressive towards aluminum. Next, all the surfaces were rinsed thoroughly with de-ionized water to remove the soap and then with acetone and isopropanol to dissolve any remaining fat. After that, we rinsed again with de-ionized water to remove the solvents. The parts were left to dry for a few hours and were then wrapped with clean aluminum foil and baked at 120 °C for around two days. The used aluminum is specified not to be operated on temperatures above 180 °C, so we chose a lower temperature to be safe. The baking should reduce the amount of gasses and water adsorbed on the aluminum. After taking the parts from the oven, it was noticed that there are some stains on the surface, resembling water stains. In the future, we should make sure that parts are properly dried, by blowing off the water with high grade dry nitrogen (Grade 6), or just omit the last water rinsing.

The copper parts were similarly wiped with isopropanol, sonicated in soap for 10 minutes and rinsed with water, acetone, isopropanol and again water. Next, they were dried, wrapped in aluminum foil and baked at 100 °C for around two days.

For the plastic parts, the sealing made from Viton and the peek connectors for the electrical feedthrough, the parts were rinsed with only isopropanol and then wrapped and baked at 100 °C. With these materials we have to be careful not to use solvents that are too aggressive, as acetone for example would make the Viton swell and destroy it.

The Zerodur spacer was not sonicated, to prevent accidental damage to the glass due to rubbing against the steel cage of the sonicator. It was instead rinsed over all surfaces with acetone and isopropanol and left to dry on air. It would be a good idea to dry it in the oven at around 100 °C, to make sure everything evaporates from any potential surface cracks in the material. However, due to the manufacturers specifications, saying that cooling the glass faster than it was cooled during the manufacturing process could degrade the thermal expansion properties, we avoided baking it.

The regular stainless steel parts were also sonicated in soap, rinsed with water - acetone - isopropanol - water and then baked at 280 °C for two days. An exception were the viewports, as the glass seal and the coating can get damaged if heated too high or too fast. Thus, the viewports were only wiped with isopropanol and the glass surfaces were cleaned with lens tissues just like sensitive optics. The electrical feedthrough and the angle valve were not sonicated, as it would be hard to get rid of the soap and water from the many edged surfaces. The gauge was also not baked but only wiped, as it cannot withstand high temperatures. The ion pump came vacuum

sealed so cleaning was not necessary.

5.4.2 Assembling the parts

Some care has to be taken when selecting the screws for the assembly. The in-vacuum screws were cleaned by sonicating in isopropanol and then dried on air. The size needed for screwing together the copper box is $M4 \times 10$. The screws can be made from stainless steel, as they will not stick to the copper thread. They also do not have to be specially vented, as that was already taken care of with the venting holes drilled into the copper.

The screws that go directly into the aluminum chamber have to be greased to prevent sticking to aluminum. Silver coating is not a sufficient solution as it might also react with aluminum. Therefore, a small amount of vacuum grease was applied to the thread of the screws using a brush. We used Lithelen, as it has a very low vapor pressure, so it can even be used inside vacuum. The screws used for attaching the CF 40 flanges to the chamber are hexagon head $M6 \times 20$, and the ones for fastening the aluminum lid are cylinder head $M6 \times 30$. For sealing the steel-on-steel flanges on the pumping arm, silver coated $M6 \times 35$ screws were used. Here silver is enough to prevent the steel from sticking.

All the Viton O-rings were also greased with Lithelen. The idea is to apply a layer of grease that will prevent gasses from permeating through the Viton. The grease was first rubbed in all around the ring and then wiped off with cloth, to remove any excess amount. The tightening of the screws around the Viton seals was done in a star pattern and in small increments to make sure that the compression is even. Also, a not too high torque was applied to avoid accidentally damaging the aluminum threads.

The pumping arm was partially pre-assembled as it is hard to reach all of the screws once it is attached to the chamber. A special wrench with a flexible head has to be used to be able to turn all of the screws. Even though the seal with the chamber endured quite some torque from the heavy components, no obvious leak could be detected with a helium leak checker.

After the first test of putting most of the components in the chamber and pumping down with a roughing pump, a rather dissatisfactory vacuum of 5×10^{-5} mbar was reached. To check if the problem lies in the seals and the chamber itself, all of the inside components were removed and the chamber alone was pumped on. The pressure dropped down to 1.7×10^{-6} mbar after 4 days. We realized that the chamber has to be baked to achieve better vacuum, probably because of a porous layer of oxide on the aluminum. After baking the aluminum chamber at 100°C for 3 days under vacuum, and then pumping it down with the ion pump, the final pressure achieved was 4.7×10^{-8} mbar on the ion gauge and 2.7×10^{-8} mbar on the ion pump controller.

This result is in the range that we were aiming for.

After this result, the vacuum chamber was opened several times, and more components were added each time. Now, even after all the components were present in the chamber, apart from the cavity itself, the pressure dropped down to 2×10^{-6} mbar after several days. We expect that after the cavity is put in place and the vacuum chamber is carefully baked again, there should be no problem to achieve a vacuum level on the order of 10^{-7} mbar. To retain a long lifetime of the Ion pump it should not be turned on until the pressure is below 10^{-6} mbar.

Chapter 6

Setup for stable injection-locking

6.1 Injection-locking

In this section I will briefly describe the idea behind injection-locking. A free-running laser diode is normally quite broadband, as the gain medium in the diode allows the emission of photons across a broader spectrum around the nominal wavelength. If we forcibly seed the slave laser diode by injecting external light with a narrow linewidth, the slave diode may start emitting light with the same properties. This effect occurs if the injected light is much stronger than the spontaneously emitted photons of the free-running slave diode. The seeded mode of light will be more likely to get exponentially amplified in the gain medium and it will suppress the other spontaneous modes by drawing all of the available power. Furthermore, the injected light has to be well mode matched to the internal cavity of the slave laser diode, for this effect to happen across the whole gain medium. The alignment gets more sensitive, if we want to use the minimal possible amount of seed light.

The schematic of an injection-locking setup that is used in our design can be found in Figure 6.1. Light from the master laser is coupled through an optical fiber and collimated using a lens. We use a half-wave plate to rotate the polarization of the seed light until it is accepted by the second polarizer of the Faraday isolator. The seed light is coupled into the slave laser diode using the two mirrors. The lens in the slave laser matches to the mode of the laser diode and the half-wave plate in front of it is used to set the correct polarization. The slave laser now emits light that can pass through the Faraday isolator, which at the same time rejects any counter propagating light to protect the diode. A half-wave plate and a PBS are used to split off a fraction of the power for monitoring the operation of the slave laser. The rest of the light is coupled into an optical fiber using two adjustable mirrors.

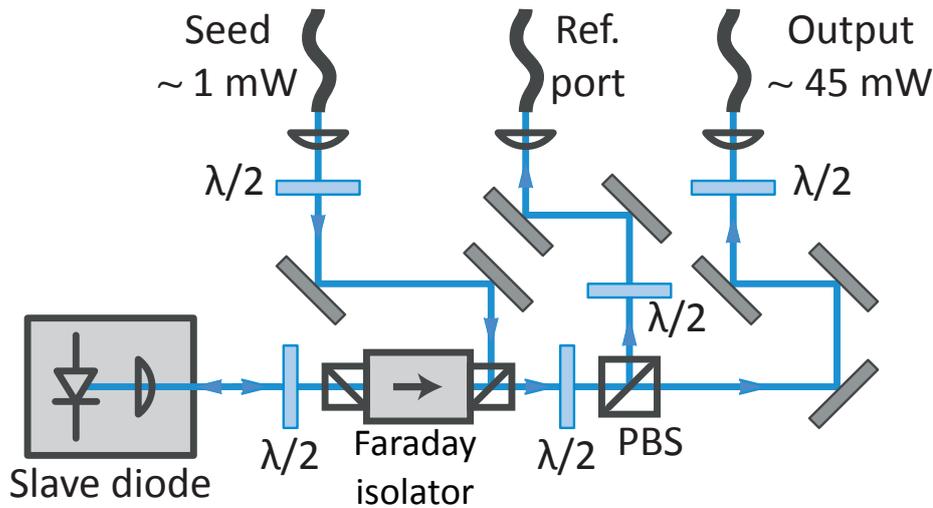


Figure 6.1: The optics setup used for the designed injection lock module.

6.2 New compact mechanical design

Injection-locking is very sensitive to the mechanical stability of the injecting beam. To achieve a stable lock, which does not require constant attention, we wanted to improve on the mechanical stability of our optics setup. By decreasing beam heights and lengths, ensuring temperature stability, good mechanical mounting, and limiting air currents around the beam path, we should be able to decrease the changes in beam pointing over time. Taking these considerations into account, we had an idea for a new compact and modular opto-mechanics setup.

6.2.1 Optical breadboard

A severe limitation in how dense a typical optics setup can be built is usually the size of the clamps and posts used to mount the optics to the optical table. To avoid the use of clamps, we decided to mount the parts differently, by making adapters for the desired optical components and screwing them to a breadboard directly. This method does have a drawback, as parts can only be mounted in discrete positions. However, by careful design, we can reference all the components to a predefined optical path and prevent the beam path from deviating from the hole matrix. To ensure reproducible placement of the mounts on the breadboard, we decided to use dowel pins for precise mechanical referencing. Additional holes for screws ensure that the parts can be fixed in their final position.

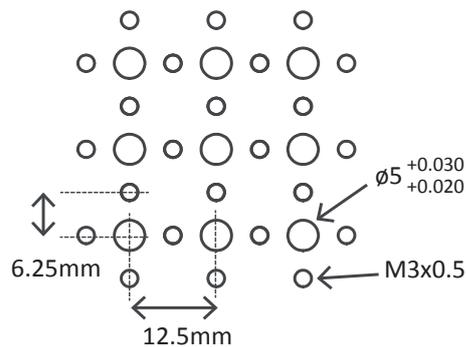


Figure 6.2: The hole pattern for the designed breadboards.

To use both dowel pin holes and tapped holes, we had to design our own hole pattern for the breadboards. After some trial-and-error with the mounts for the desired optical components, I decided on a grid of $\varnothing 5$ mm dowel pin holes separated by 12.5 mm. Between the dowel pins, at a 6.25 mm distance are the M3 tapped holes. A sketch for better visualization can be seen in Figure 6.2. The tolerances were determined after making some test breadboards and parts and seeing what size of dowel pin holes is most practical to work with.

Apart from the precise hole pattern, we also want the breadboard to be temperature stable. For this reason, the material of the breadboard is aluminum, due to its high heat capacity and thermal conductivity. To be able to keep the breadboard at a constant temperature and to sink any heat coming from the laser head, we included a groove for copper tubing, which allows for water cooling. After the tube is bent into the correct shape, it is clamped into the bottom of the breadboard using small clamps. The bottom side of the breadboard can be seen in Figure 6.3(a).

On the bottom side, we also notice the shallow extrusion, that makes sure that the contact with the optical table is defined by the perimeter of the breadboard. The pockets on the short edge of the breadboard allow clamping to the optical table. On one of the long sides, there are two more $\varnothing 10$ mm dowel pin holes for good referencing of the side wall and M6 tapped holes for screwing it in place. It is important for the side wall to be well referenced, as ports for optical fibers will be connected to it.

The side wall of the breadboard can be seen in Figure 6.3(b). The three large holes are for the mounting of Thorlabs's FiberPort collimators. The other holes in the indented extrusion are for electronics connectors that are required for the driving and regulation of the laser diode head. The breadboard and its side wall also have a ledge on the top side that is used to attach plates of acrylic glass, which are used to close the system after it has been assembled. In Figure 6.4, the whole setup is shown, without



Figure 6.3: The bottom view of the breadboard (a) and the top view with the side wall attached (b).

the top and front acrylic plates for clarity.

6.2.2 Optics mounts

As seen in Figure 6.4, there are many different optics components we need for our setup. First, let us look at the commercial components that we have included.

To connect the optical fibers and collimate their outputs, we used Thorlabs's Fiber-Port collimators PAF-X-7-A. These fiber ports are compact, offer good stability, and the 5 adjustable degrees of freedom allow us to collimate the laser beam and also adjust it to the desired optical path. At the same time, the ports can be used for coupling light into a fiber and optimizing the coupling efficiency. One drawback we found with these fiber ports is that the adjustment mechanism is a bit counter intuitive and it takes some practice to make fiber coupling behave the way one wants it to. The focal length of the collimators was chosen from the experience we have with fiber coupling of the light from the Nichia NDB4216E diodes we are using for injection-locking. The fiber ports are mounted directly on the breadboard side wall, using four M2 screws.

The mirror kinematic mounts were chosen to be Thorlabs's Polaris-K05 mounts. The Polaris series offers excellent stability and, furthermore, the K05 model comes with two $\varnothing 2$ mm dowel pin holes on its mounting surface, allowing it to be well referenced in our design. We chose the mounts for half inch mirrors, because they allow us to keep the beam height as low as 20 mm. To place the mirror mount on our breadboard, I have designed an adapter plate, which has a tapped M4 hole and two $\varnothing 2$ mm dowel pin holes on the top for mounting the kinematic mount. Additionally, two $\varnothing 5$ mm dowel pin holes and M3 through holes are added to connect to the breadboard hole pattern. The adapter can be seen in Figure 6.5(a).

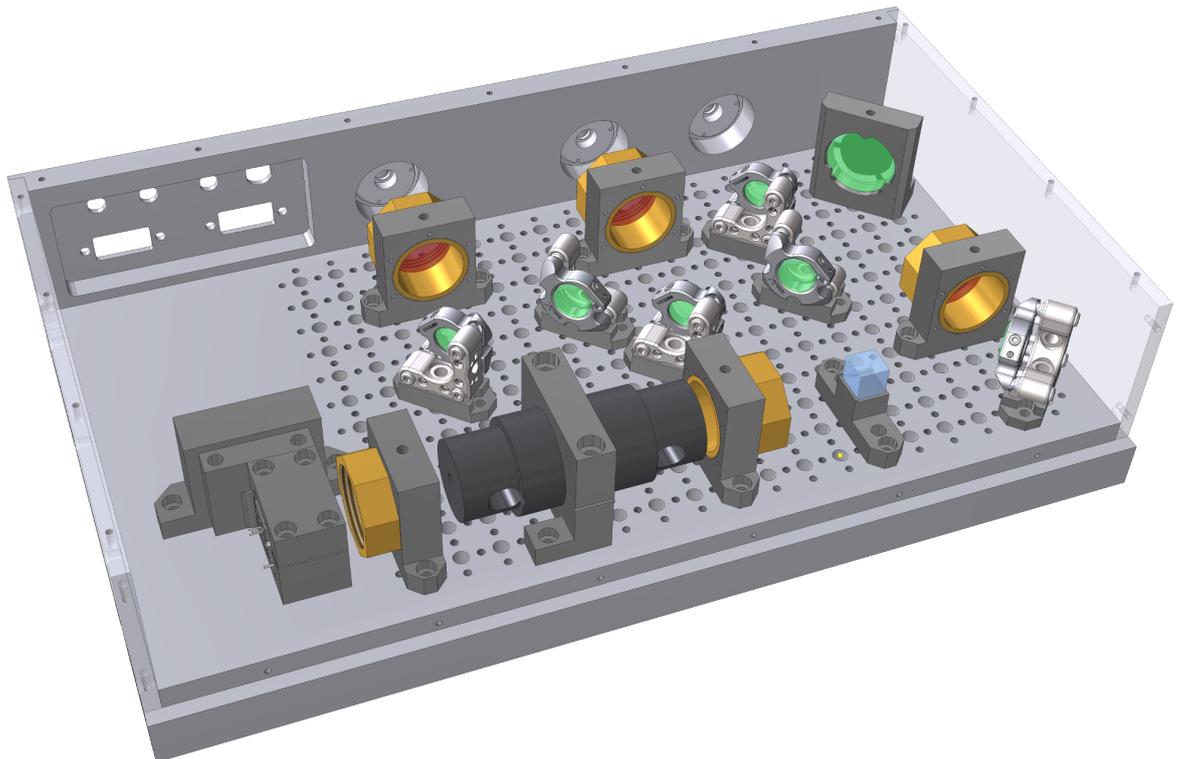


Figure 6.4: The entire injection lock setup. The green glass pieces symbolize mirrors, the red glass half-wave plates, and the blue glass a polarized beam splitter.

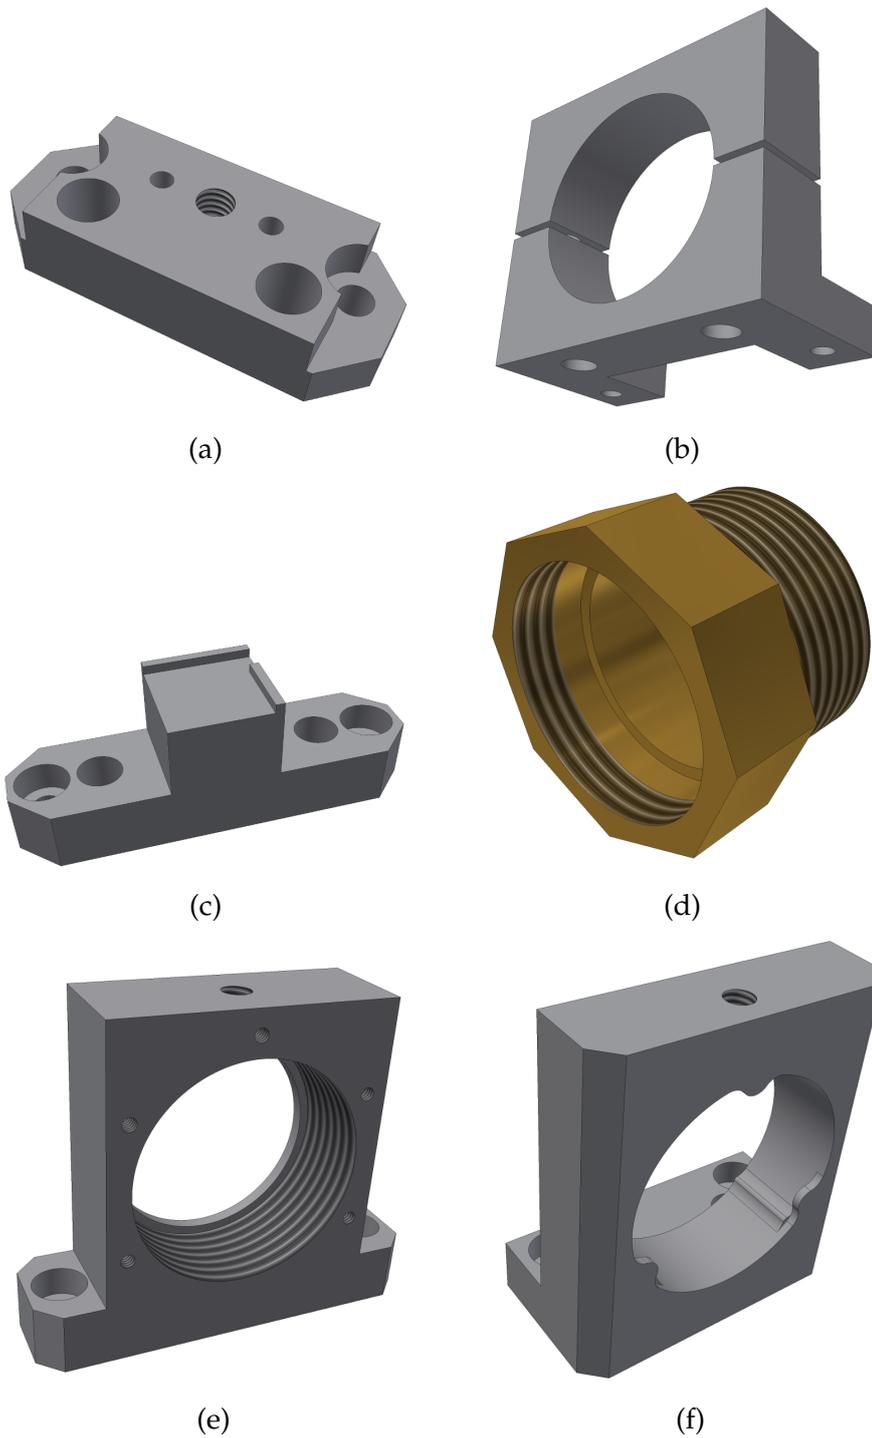


Figure 6.5: The designed mounts for placing the optical components on the custom breadboard. The adapter for Polaris mirror mounts (a), the mount for the optical isolator (b), the 10 mm PBS mount (c), the rotation mount for wave-plates (d), the SM1 threaded mount (e), and the 45° mirror mount (f).

The optical isolator is a custom isolator from Thorlabs, usable with our 461 nm wavelength, with the custom part number IO-5-460-PBS-HI. It is mounted by clamping with two half-moon shapes screwed together with M3 screws. The half-moon shape is of the same nominal diameter $\varnothing 1.20$ inch as the diameter of the isolator, but the asymmetric tolerance ensures that it is always slightly bigger. The 1 mm gap between the two clamps makes sure they can press on the isolator. The bottom clamp includes two $\varnothing 5$ mm dowel pin holes and M3 through holes for mounting to the breadboard. The clamps can be seen in Figure 6.5(b). They mount the optical axis of the isolator on the 20 mm beam height we are using. Unfortunately, the longitudinal position of the isolator has to be adjusted by hand to match the injection port to the hole pattern. To make the alignment easier, a pin hole can be temporarily placed on the breadboard.

To mount the 10 mm polarizing beam splitter (PBS) a simple adapter was designed. The PBS is glued to the central extrusion and the two ledges on the top allow it to be flushed in place and properly centered. The bottom of the part includes the $\varnothing 5$ mm dowel pin holes and M3 through holes for mounting to the breadboard. The mount can be seen in Figure 6.5(c). The thickness of the mount is 12 mm, which is enough to mount the PBS and at the same time uses up only one row of holes. This configuration allows a very compact setup.

For half-wave plates, which need to be rotated to get the desired polarization of light, I have made mounts that are internally threaded with the Thorlabs SM1 thread and brass adapters that can hold and rotate the wave-plate. As the wave-plates only have to be adjusted once and do not distort the optical path by a lot, we expect the referencing by the long and dense thread to be sufficient. After the adapter is turned into the desired position, it can be fixed in place using an M4 nylon-tipped screw. The length of the adapter is chosen so that even after a full turn from the flush position, the part only takes up two rows on the breadboard. The part can be seen in Figure 6.5(d). The front section has an octagonal shape to allow easier turning by hand. The inner SM1 thread can be used to mount the wave-plate using a SM1 retaining ring. Brass is chosen as it is easily machined and the softer thread can turn smoothly against steel without the need for lubrication.

The SM1 threaded mount is compatible with many of the Thorlabs SM1 components such as irises, lens tubes, and retaining rings, and it can also hold 1 inch optics. I have designed two similar parts, one with thread all the way through and one with a ledge on one side, which can be seen in Figure 6.5(e). The ledge serves as a reference plane for 1 inch optics and is more stable than using a retaining ring. If we want to avoid using retaining rings on the other side as well, there are five M2 holes on both sides of the parts, potentially allowing us to lock the part with a back-plate. The mounts have an M4 threaded hole at the top which can be used to fix parts with

a nylon-tipped screw. At the bottom, there are the $\varnothing 5$ mm dowel pin holes and M3 through holes for mounting to the breadboard.

The last of the breadboard parts is the 45° mirror mount, seen in Figure 6.5(f). The front surface of the mirror is referenced to the three tongues at the front of the mount. The tongues can mount the mirror by slight deformation, if we press on the mirror from behind by inserting a wave spring behind the mirror and compress it with a back-plate using M2 screws. This method offers very stable mounting of the mirror and does not stress the mirror surface too much. The mirror height is set by the curved ledges at the bottom of the mirror hole. Alternatively, we can fix the mirror using a M4 nylon-tipped screw from the top. Making the mirror flush with the front surface allows us to use a mirror of different thickness, without changing the design. The rear ledge on the bottom of the mount allows us to mount the mirror to the breadboard using the $\varnothing 5$ mm dowel pin holes and M3 through holes.

All of the above mounts are machined from AISI 303 grade stainless steel. Stainless steel is chosen due to its good mechanical stability and small expansion with temperature. The chosen alloy is non-magnetic which is useful to avoid collisions with the optical isolator and is also easily machinable. After making a test setup of the parts, we chamfered some of the edges allowing the different parts to be placed closer together on the breadboard.

6.2.3 Diode laser mount

To be able to use our new breadboard design for injection-locking, we also had to design a reliable mount for the laser diode. The problem we were facing is that the pointing of the output of the laser diode is specified only to a $\Delta\theta = \pm 3^\circ$ angle. Therefore, we needed to build a mount with enough degrees of freedom, to align the laser beam precisely enough to pass through the aperture of the optical isolator and the PBS up to the first adjustable mirrors. An additional constraint here is that the laser diode has to be temperature-stabilized using Peltier elements. The need to provide good thermal contact for the heat sink seemed to be a big flaw in some of the solutions we were discussing. In the end, we found a good design to overcome these problems while looking at the diode laser system described in Ref. [40], and modified it to our needs.

To estimate how much we need to adjust the beam pointing and position, we look at what happens to the angle error of the diverging beam when it hits the collimating lens. Here, the tolerance to which we can reference the laser diode emitting point with regard to the optical axis of the collimating lens is also important. The way that the laser diode and the lens are mounted is seen in the top-view cross section of the part in Figure 6.6.

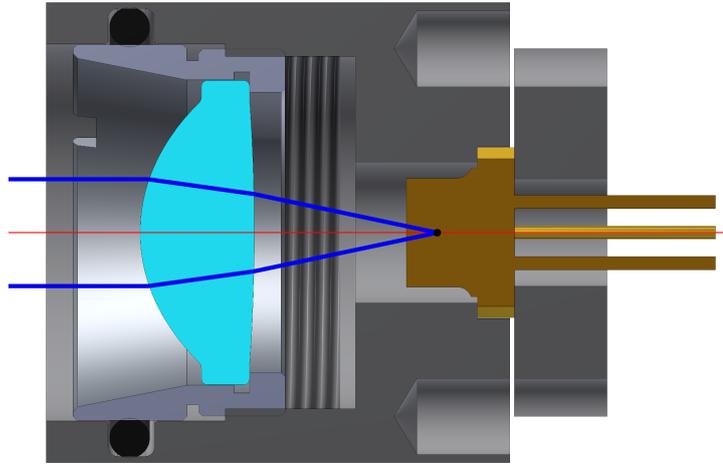


Figure 6.6: The top-view cross section of the designed laser head part. The emitting point of the laser diode is marked with a black dot, the optical axis with the red line and the divergence of the laser light with blue lines.

First, let us take a look at what happens to the error in the pointing angle $\Delta\theta = \pm 3^\circ$, after passing the collimating lens with a focal length $f = 8$ mm. Assuming the emitting point of the laser diode is in the focal point of the lens, the pointing angle will result only in a beam displacement

$$\Delta h = f \arctan(\Delta\theta) = 0.42 \text{ mm.} \quad (6.1)$$

This displacement should not be a limiting factor when trying to align a $\sim \varnothing 3$ mm beam through the $\varnothing 5$ mm aperture of the optical isolator.

Next, we can look at what a displacement between the optical axis of the laser diode and the optical axis of the lens Δs would result in. If we look at the central beam, following the optical axis of the laser diode, parallel to the optical axis of the lens, it will be redirected into the focal point of the lens. The angle of the central beam after the lens is

$$\alpha = \arctan\left(\frac{\Delta s}{f}\right) = 0.87^\circ, \quad (6.2)$$

where we took the maximal displacement allowed by our tolerance $\Delta s = 0.121$ mm and the focal length of the lens is $f = 8$ mm. This angle is not very large, so we will only require a small adjustability to properly align the beam.

As we do not need to adjust our laser head pointing by a lot, it is enough to get the tilt degree of freedom in the up-down direction and the left-right direction. Looking at our design in Figure 6.7, we achieve the adjustability with the freedom from the

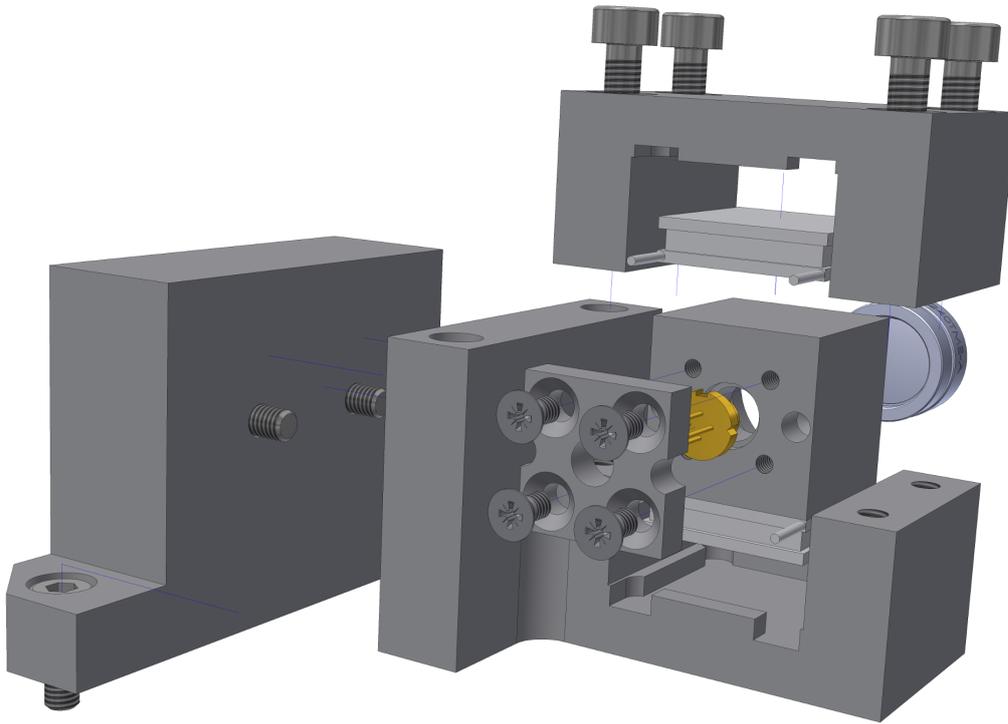


Figure 6.7: The whole assembly for mounting the laser diode. The assembly is made out of four parts, allowing for slight left-right and up-down angular adjustment.

regular tolerance for screw holes. As through holes are always bigger than the actual screws, we are left with some angular freedom before tightening the part in place.

Let us look at the individual parts of the assembly. The design of the laser head was already shown in Figure 6.6. The part is precisely matched to the outline of the used laser diode (Nichia NDB4216E) and the molded aspheric lens from Thorlabs (C240TME-A). The back-plate fixes the laser diode in place using four M2 screws, which cannot be seen in the projection. The two holes that are visible are there for gluing thermistors into the laser head part, allowing measurement of the temperature right next to the laser diode. The laser diode case fits tightly into an extrusion with the same nominal diameter $\varnothing 5.60$ mm and the G7 tolerance. The flat contact surface allows for good angular referencing. The lens is screwed in from the other side into the $M12 \times 0.5$ thread. The length of the thread is chosen so that the focal point of the lens can be moved by ± 1 mm around the emitting point of the laser diode. This degree of freedom allows us to properly collimate the output of the laser diode. Because we want the lens to stay in place after collimation, we added a friction O-ring around the perimeter of the case of the lens. The pressure from the ring and the length of the

lens case should also ensure the lens to be centered and well aligned with the optical axis.

As seen in Figure 6.7, the cubical laser head part is sandwiched between two Peltier elements and fixed using the M3 screws in the clamp from the top. For good thermal contact thermal paste has to be applied on all of the contact surfaces. The Peltier elements sit in small extruded pockets which roughly reference their position. As the Peltier elements have to be heat sunk and the heat has to be conducted through the mount into the breadboard, we chose the material for this assembly to be CW400J, also called “nickel silver”. This alloy has thermal expansion properties well matched to stainless steel, but a somewhat better heat conductivity $\kappa_{\text{ns}} = 30 \frac{\text{W}}{\text{m K}}$.

The bottom part of the clamp is screwed into the mounting part using two M3 screws. One screw is aligned with the central position of the lens, acting as a rotation axis for up-down alignment. The other screw is used to fix the part in place after alignment is done. To make alignment easier, there are two vertical holes with a M4 thread for ball-tipped set screws. These can be used to precisely adjust the up-down tilt using a screwdriver.

The mounting part has a large area of contact with the clamps to improve the thermal contact. The M3 tapped holes allow it to be screwed into the breadboard following the same grid as the other parts. The freedom of the M3 screws is used for left-right alignment of the laser head.

6.3 Assembling a module

This section acts as a guide on how to assemble the next modules for injection-locking and includes some tips and tricks I learned while assembling the prototype.

To assemble the breadboard and its side wall, first insert the $\varnothing 10$ mm dowel pins into the breadboard part. The dowel pins should fit tightly and it is likely necessary to hammer them in using a plastic covered hammer. The pin should not go in all the way to the mechanical stop. Next, slide the wall part over the dowel pins. The wall should fit in a slip fit, so you should be able to put the side of the breadboard on and off by hand. In case of the prototype assembly, the tolerance was too tight, so I had to use a plastic hammer to get the wall in place. Make sure the wall is well flushed with the side surface of the breadboard. Next tighten it in place using $M6 \times 14$ screws. The FiberPorts can be fixed onto the wall using $M2 \times 4$ screws.

Putting dowel pins into the mounts is also a bit of a tricky job. Insert them into parts before attaching any optics to avoid damage or contamination. For the dowel pins, you should ask our technician, as there is a selection of the necessary sizes avail-

able. For all parts, except for the Polaris mount, you should use $\varnothing 5 \times 10$ mm dowel pins. The pins should normally be too tight to put into the part by hand. To make sure the pins are inserted just deep enough, we use a 3 mm thick plate with a drilled hole that loosely fits the dowel pin. First, fix the desired part into a vise and then guide the pin in place using the 3 mm plate. Next, gently tap the pin with a plastic hammer until it starts sinking into the hole. The plate will make sure that the dowel pin will stick out by approximately 3 mm, which works very well with our breadboard. After inserting the pin, make sure that it stays in position and cannot be moved around with the force of your fingers.

If the dowel pin is loose, you will have to glue it in place. Ask in the lab what glue to use. First, you will have to remove the pin from the part. It helps to grab the pin with the vise and then to pull on the part with your hand or by gently tapping it with a plastic hammer. Next, you should insert the pin by just a little and apply a drop of glue and spread it around the perimeter of the pin. Then, gently hammer in the pin again using the plastic hammer and the 3 mm plate. Use a paper towel to soak up any remaining glue on the part and let it dry for around 15 minutes.

The Polaris adapter is assembled similarly, except that you should use the shorter $\varnothing 5 \times 8$ mm pins. Additionally, you should also put in the $\varnothing 2 \times 6$ mm dowel pins on the top side. You can again use the 3 mm plate with a smaller hole to help you insert the pins perpendicularly, but the length will not be set correctly. I managed to set the correct position of the pin by eye and by testing it with the desired Polaris mirror mount. The mirror mount should sit flatly on the adapter.

After all the pins are in position, wipe the parts using isopropanol. Make sure again that the pins do not move when applying small pressure. You should also test if the parts slide smoothly into the breadboard holes. Never use force to get the part in place, as you will get the dowel pins stuck in the part and the breadboard can get damaged. Make sure you insert parts perpendicularly to the holes. When they do not insert smoothly, you can try wiggling them until they properly sit in place. After checking that everything is fine, you are ready to mount the optics.

The PBS has to be glued onto the PBS mount. We are using glue that can be cured with ultraviolet light (Norland's NOA 61). Squeeze a drop of the glue on a mixing tray and then pull it into a clean syringe needle. Avoid contaminating the plastic part of the syringe by sucking in too much glue. Next, squeeze the syringe a little and a small drop of glue should form on the end of it. Apply tiny drops of glue to three edges of the PBS mount. Next, using clean gloves (make sure that there is no glue on them!) unwrap the PBS, grab it only by the edges and then place it on the PBS mount and press it against the ledges. When you are sure that the PBS is flush with the positioning ledges, use the UV curing gun to cure the glue for around 30 seconds for each droplet. After this amount of time, the glue should be pre-cured and the

optics stable enough for use.

The wave-plates can also be glued into the rotation mount. You can glue them by setting the wave-plate in to the front of the mount, applying small drops of glue on three points between the wave-plate and the mount and then curing them. Gluing was necessary in the prototype parts, but the corrected version should have enough thread to fix the encased wave-plates using a retaining ring.

The mirrors are simply inserted into the Polaris mounts and fixed with the compression plate by tightening the correct screw. For the 45° mount you can make the mirror flush with the front wall and then either add the nylon-tipped screw or fix it with the back-plate and the wave-spring.

Next, you need to prepare the laser diode head. Start by getting thermal epoxy (EPO-TEK H74) and with it fill one of the two empty holes on the back side of the laser head part. Prepare it by adding both components on a disposable mixing tray. The specified mass ratio between the two components is 100:3. Stir the components with a disposable mixing stick until the mixture is homogeneous. Next, take a droplet of it and let it flow into the hole for the thermistor. Wipe away any excess using a paper towel. Next, insert an epoxy encapsulated thermistor into the filled hole. Make sure that the leads are well separated to prevent a short. Wipe away any epoxy outside of the hole and be very careful not to get it into the extrusion that fits the laser diode. Check the curing time and temperature from the specification sheet and cure the part in the oven. While heating up, the viscosity of the epoxy changes, so make sure the then liquid epoxy cannot flow all around the part and the oven.

After the part has cooled down, you should sandwich the laser head between the 15 × 15 mm Peltier elements and clamp it with the top and bottom clamp of the assembly. Check the datasheet of the Peltier elements to make sure you are correctly orienting the hot and cold plates. We usually need to heat these laser diodes. Get the syringe with the thermal paste and draw a thin layer into the pocket in the bottom clamp. Insert the Peltier element in the correct orientation and repeat the step when adding the laser head and the other Peltier element and top clamp. Once all parts are in place, try to orient the laser head as straight as possible and then add the M3 × 16 screws. Do not use too much force when compressing the clamp, as the Peltier element can get damaged. After you are done, carefully wipe away any thermal paste that has crept out of the gaps.

Next, you should insert the 1.30 mm thick O-ring. Unfortunately, we could not get the rings in the correct length, so you have to take a bigger one and cut a ~ 36 mm long tube. When testing, inserting a cut O-ring into the groove on the lens's side proved to work well. Next, you can push in the molded asphere C240TME-A and partially screw it in using a spanner wrench. Here I should note that removing the

asphere can sometimes be difficult as it can get stuck due to the friction ring.

When inserting the Nichia NDB4216E laser diode, you should be very careful as the part is expensive and really easy to break. First, make sure you understand exactly what each of the pins is as it is easy to mix them up when the part turns or if you have a mirrored image in mind. Always wear a grounding wristband when handling the laser diodes and do not solder anything when they are attached. You will have to pre-solder wires to a 2×2 socket (Thorlabs S8060-4), keeping in mind the correct pin layout for the laser diode. Once the wires are soldered, carefully insert the laser diode into the slot in the laser head, making sure it fits loosely. If the fit is not loose, the case of the laser diode will deform and you will likely destroy it when trying to remove the diode in the future. There is a tongue in the extrusion to set the orientation of the diode. Fix the laser diode case in place by pressing it with the back-plate using four countersunk $M2 \times 4$ screws. Check that the leads of the diode are not touching each other or the back plate by measuring the resistance between the different pins.

Now, you are ready to put the whole assembly together and mount it on the breadboard. Prepare the electronic board and connectors, making sure the polarities of the connections are correct and that all the leads are long enough. The connectors attached to the breadboard side wall, should be screwed from the side which allows the board to be removed. Connect the laser driver and temperature control modules and first make sure that the temperature stabilizes correctly. Next, slowly turn up the laser driver until the diode starts emitting a beam. Use a spanner wrench to adjust the position of the asphere and collimate the laser beam. Using two pinholes, you can fine-tune the orientation of the laser head by loosening the corresponding screws and adjusting the up-down and left-right pointing. Similarly, the optical isolator should be put in place while aligning its side port output with pinholes.

If the injection lock module should be used in free space, you can drill a hole in the acrylic glass and redirect the beam using the adjustable mirrors. Once the optical setup is complete, add the acrylic panels and tighten them in place using $M2.5 \times 8$ countersunk screws.

6.4 Protection and control circuit

In this section, I report on my design of the printed circuit board (PCB) that is used in the injection-lock modules and also in the tapered amplifier modules. The main reason we wanted to change the previous design used in our group is to add a read-out circuit for the photodiode in the laser diode case. Furthermore, the old circuits by themselves are not compatible with the Thorlabs laser diode drivers we want to use. The project also seemed good for me to start with and learn about PCB design.

Let us jump straight to the schematics and afterwards I will explain what each of the sections of the circuit are for and how they work. The full schematics of the board are seen in Figures 6.8 and 6.9.

6.4.1 Temperature control

The most straightforward part of the circuit is the temperature control part, seen in the bottom left of Figure 6.8. We simply add a male 9-pin sub-D connector, where we can plug in the cable from the temperature controller. The pin layout seems to be rather standardized, at least between Thorlabs and Toptica, whose controller devices we were looking at. The required pins of the sub-D connector are rerouted to the connectors for the Peltier elements and the thermistor. We use Molex KK connectors, which can be easily disconnected, as they are made from pins that are soldered to the PCB and form a head where we crimp the connecting wires.

6.4.2 Current control

The other sub-D connector cable comes from the current controller. In this case we need a female 9-pin sub-D connector to plug the cable in the board. The corresponding part of the schematics is on the top of Figure 6.8. The pin layout of this connector is also compatible between the Thorlabs and Toptica controllers, apart from the different operation of the interlock system. The pins have to be differently connected to the laser diode, depending on its polarity. To make the board compatible with both positive and negative polarities, I added a series of jumpers, JP2, JP3 and JP4. I numbered the pins in a easy to understand way, so that the same placement of the jumper corresponds to one of the polarities. The jumpers effectively select which of the laser diode pins is grounded. Additionally, resistors are added after the pins 6 and 9, protecting the voltage sense inputs of the controller.

The next thing we notice to the right of the jumpers, are a series of bypass capacitors. Their function is to act as a short circuit for high frequency noise and prevent it from reaching the laser diode. Two large capacitance electrolytic capacitors and two smaller ceramic capacitors are chosen. The choice of values is not trivial, as we are deliberately using the non-ideality of the components. I copied the components chosen in the previous design, as they seemed to work in practice. The key feature here is to damp any noise of the current controller or noise that may have been picked up in the long cable.

The series inductors and resistors L1, L2, R5, and R6 are used to decouple the laser diode and the modulation circuits from the capacitors. If we want to add external modulation on our laser diode, we need to make sure it is not also damped by the bypass capacitors. The Schottky diode, right next to the connector for the laser diode,

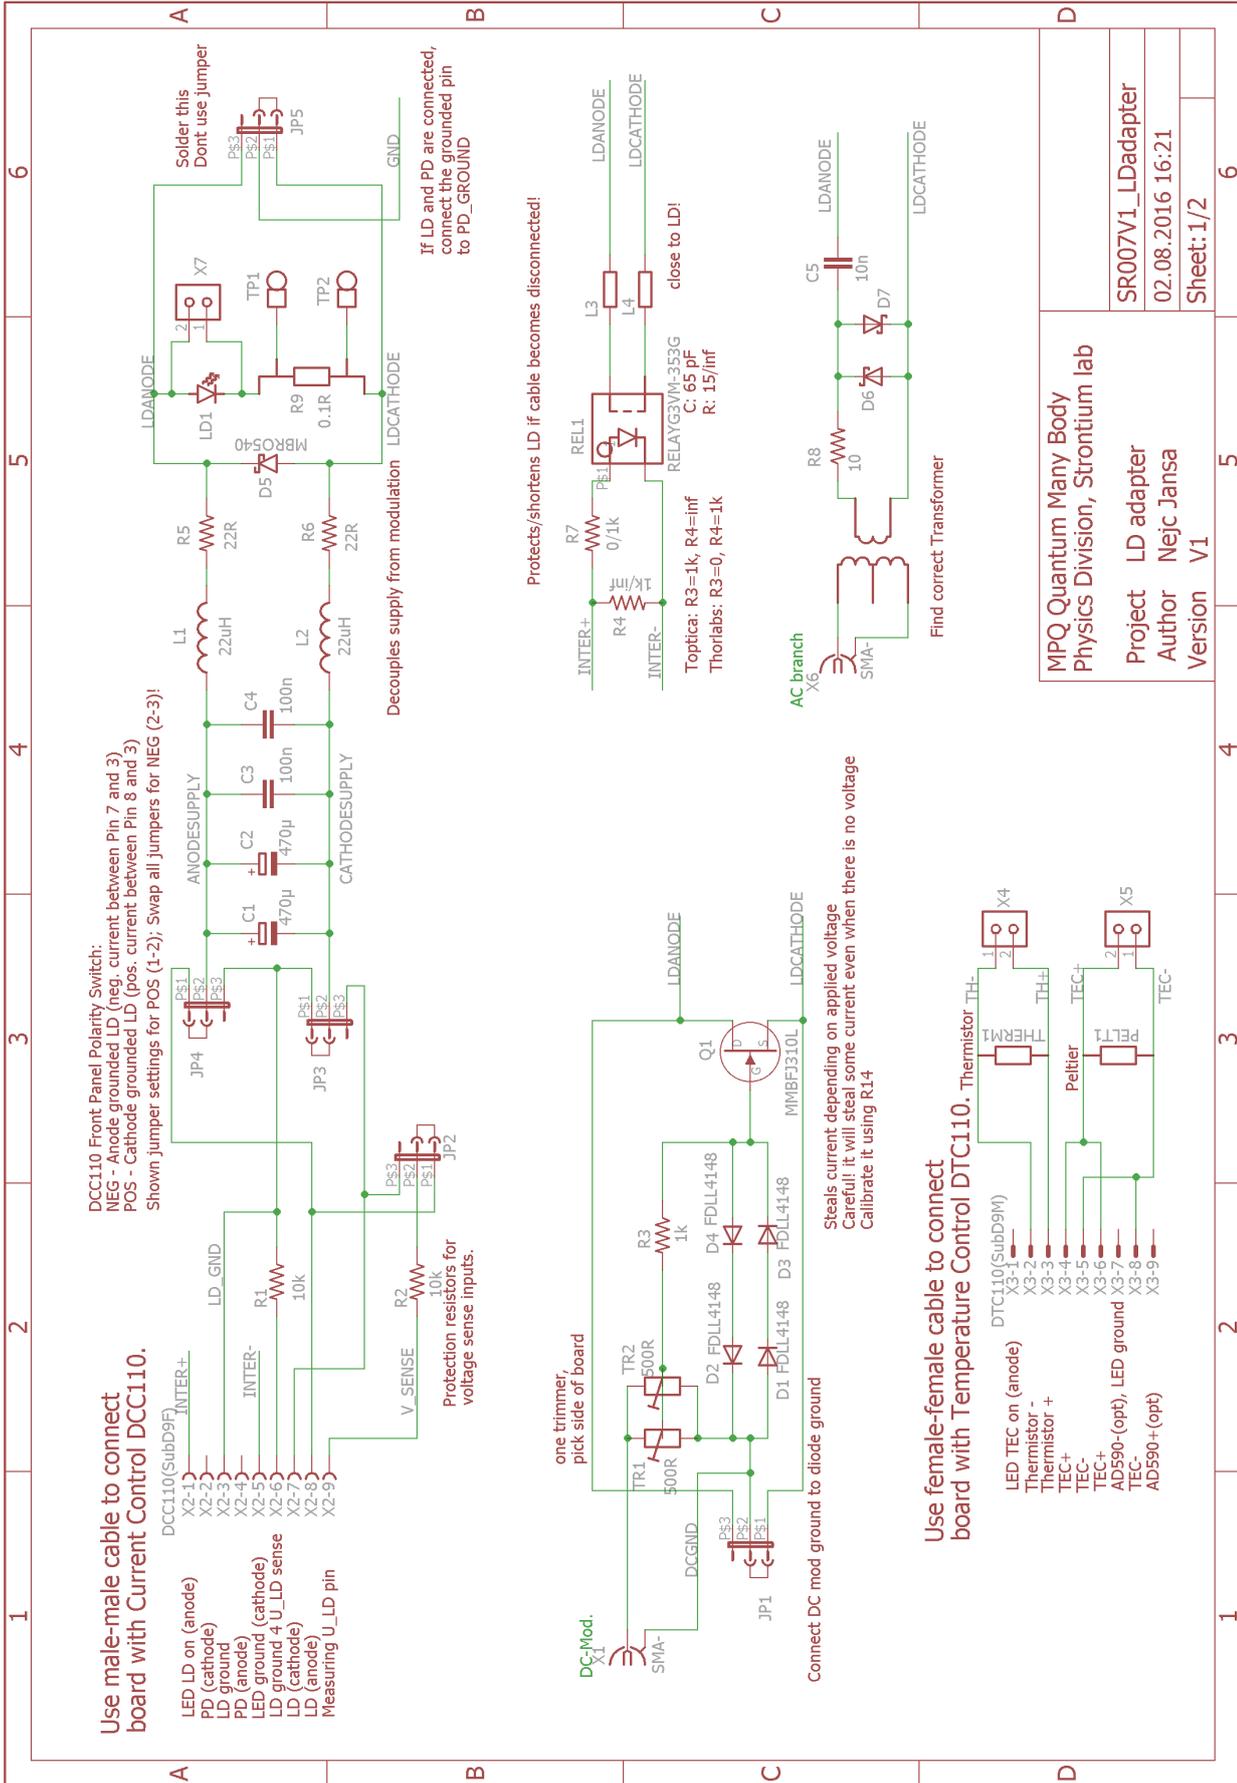


Figure 6.8: The circuit diagram for the laser diode protection circuit.

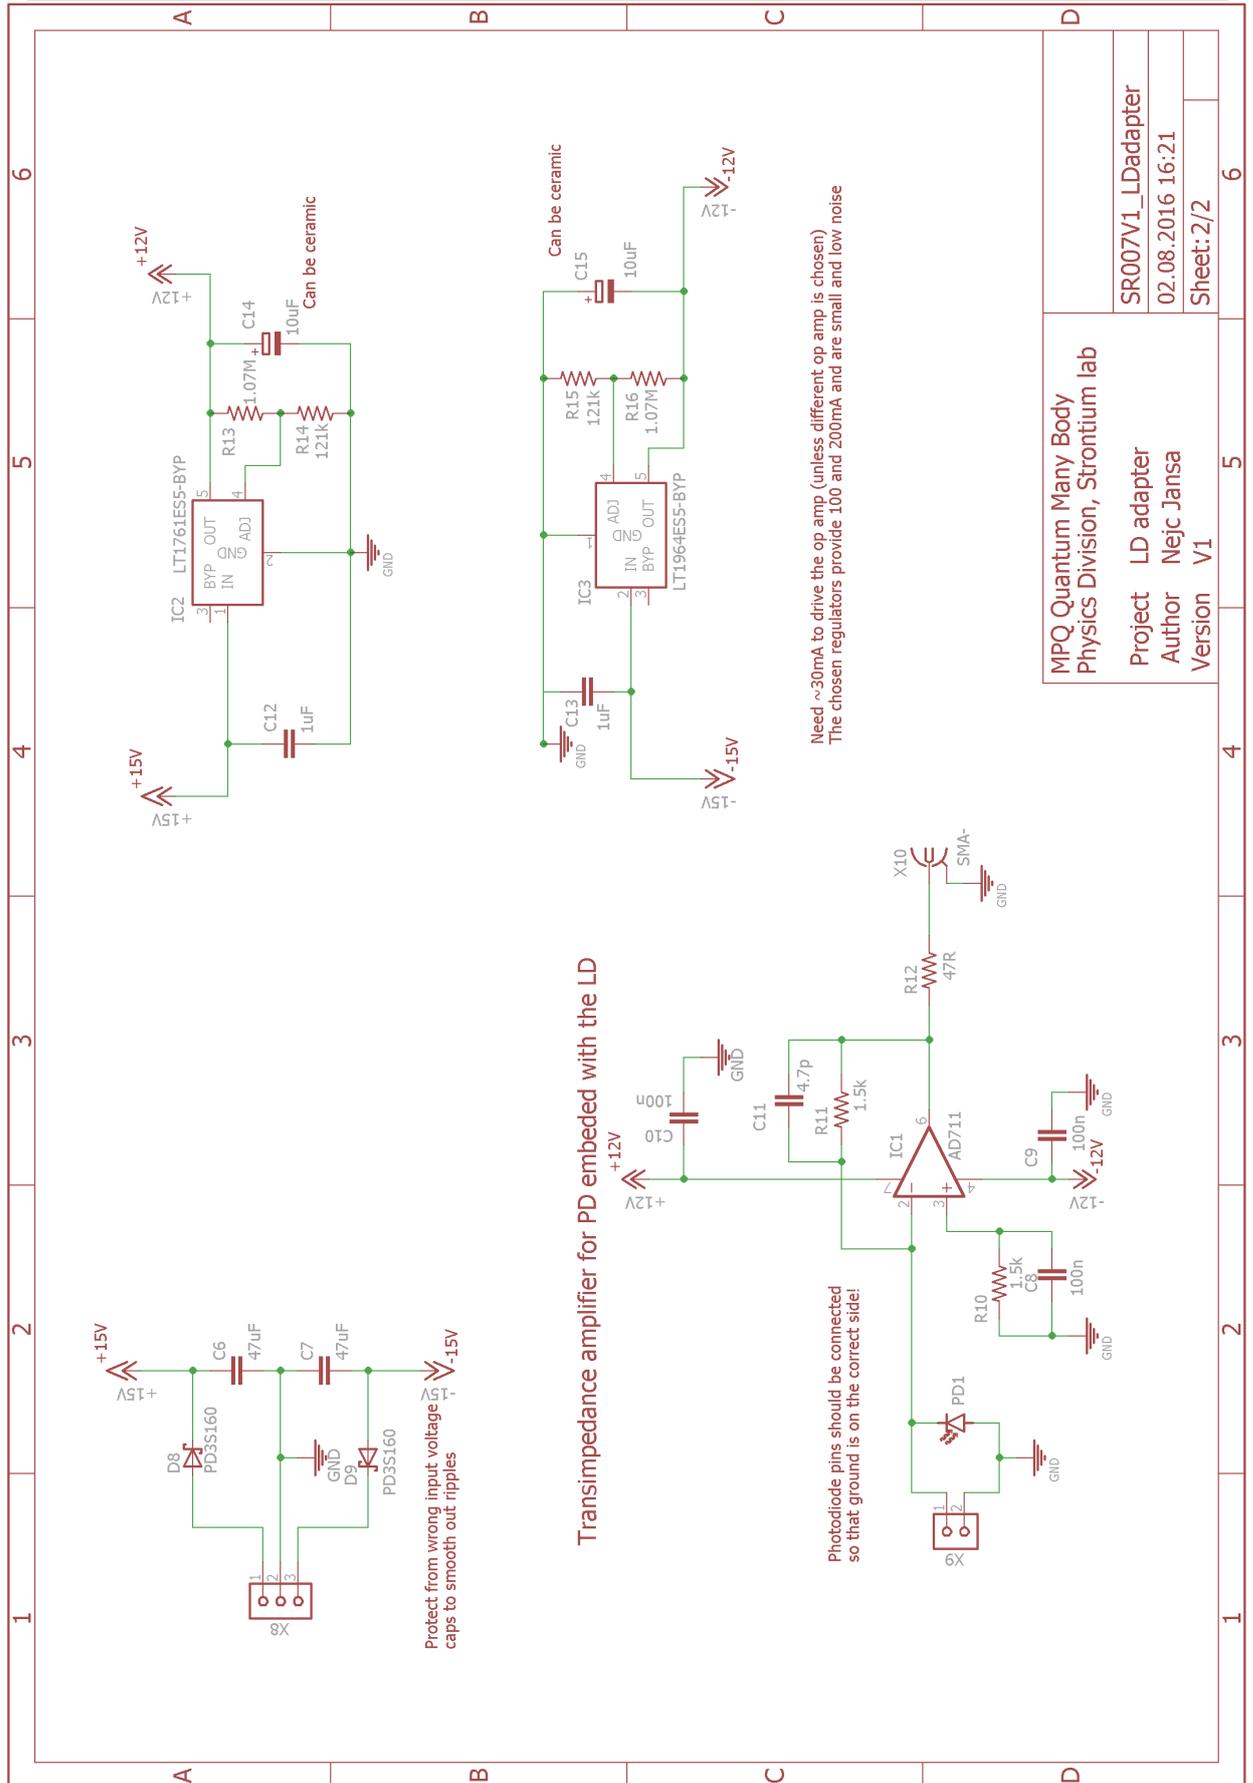


Figure 6.9: The circuit diagrams of the laser diode protection circuit for reading out the value of internal the photo diode in the laser diode case.

protects the laser diode from incorrect driving currents. The sense resistor R9 uses four-point measurement to allow us to precisely monitor the current that is running through our laser diode.

Another form of protection we included in the design is the relay circuit seen in the middle right section of the schematic. The goal is to short-circuit the leads of the laser diode, when it is not in operation. The short-circuit should protect the laser diode from harmful static electricity and any stray voltage which might occur due to photo-voltaic effects in the laser diode. We want to open and close this relay using the interlock logic of the controllers, but a problem arises, because the Thorlabs controller provides much less current in the interlock pins. For this reason, we had to use a solid state relay, which is normally closed and opens when at least 1 mA of current runs through the control pins. Resistances R4 and R7 were chosen and tested to meet the conditions for the relay to operate in the correct way. L3 and L4 are slots for inductors that can decouple the parasitic capacitance of the solid state relay. In the testing done to date, the use of inductors was not necessary and 0 Ω resistors were soldered in their place.

6.4.3 External current modulation

There are two sections of the circuit that can be used for current modulation, the AC and the DC ports. However, due to the high bandwidth of the so called DC port, the AC port is likely to be obsolete.

The AC modulation branch is seen on the right bottom side of the schematics in Figure 6.8. I merely copied the previous design and found space for it on the PCB, in case somebody wanted to use it in the future. However, I did not test it or double-check the components and rather did the tests on the DC branch. The SMA connector allows us to connect the cable carrying the external modulation. The transformer couples the AC signal into the circuit and changes its voltage swing. The diodes D6 and D7 limit the maximal voltage swing across the laser diode. The capacitor C5 transmits the AC and acts as a DC block for the currents driving the laser diode.

The DC modulation circuit diagram is seen in the left middle section of Figure 6.8. Again we need an SMA connector to connect the modulation signal with our circuit. As the signal can be DC, which will offset the current flowing through the laser diode, we need to make sure that the external ground of the connector is properly connected to the ground of our circuit. The connection is done with jumper JP1, which follows the same logic as JP2, JP3 and JP4.

We can achieve the modulation using a JFET N-channel transistor. Here we should note that this type of transistor is symmetric, meaning that the drain and source pins are interchangeable. The transistor conducts a defined amount of current between

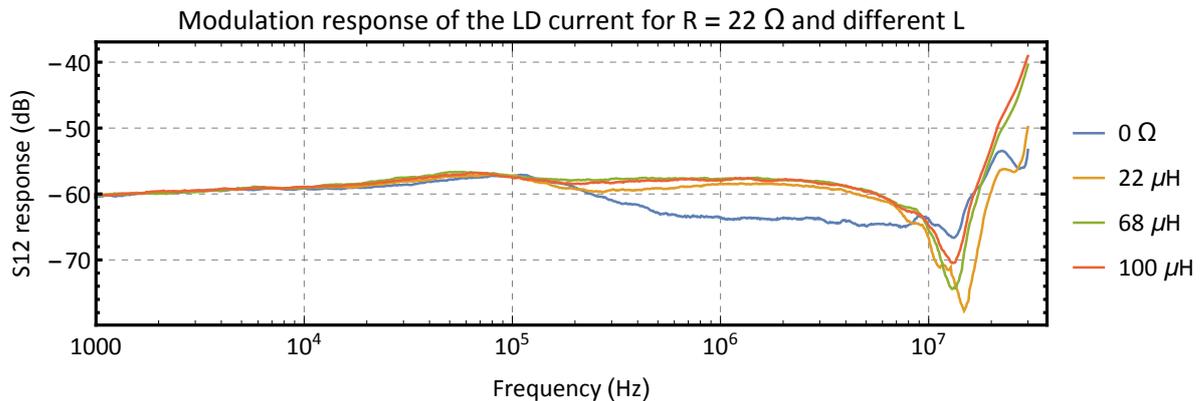


Figure 6.10: Measurement data of the modulation response of the DC port using a network analyzer. The S12 function measured the ratio between the amplitude of the modulation observed across resistor R9 versus the amplitude of the modulation applied to the DC port. A dummy LED was used instead of a laser diode for safety. The measurement beyond the 10 MHz point was not reliable as we did not calibrate the system properly for such high frequencies. To test the higher frequency response, the modulation should be observed on the light emitted from a real laser diode instead.

the drain and source pins, depending on the applied gate voltage. Therefore, we can use this principle to steal a part of the current running through the laser diode and then modulate it by changing the voltage applied to the gate pin. The benefit of this method is that the selected fast JFET MMBFJ310L can operate in the several 100 MHz regime. The bandwidth of our modulation is most likely limited by other parts of the circuit. The measurement that showed that the modulation has a relatively flat response up to at least 10 MHz can be seen in Figure 6.10.

A downside is that even when we are not inputting anything to the DC modulation port, the current through the laser diode will be decreased. Each board should be calibrated by checking the actual current through the laser diode using R9, as the transistors can vary from part to part. A measurement of how much DC current is stolen can be seen in Figure 6.11.

The combination of the trimmer TR1 (TR2) and resistor R3 acts as an adjustable voltage divider, which lets us change the sensitivity of the circuit to the applied modulation. The trimmer total resistance also sets the input impedance of the DC modulation port. The signal diodes D1, D2, D3, and D4 set a limit to the maximal swing of the gate voltage. The choice of diodes probably needs to be revised to optimize functionality with the chosen JFET.

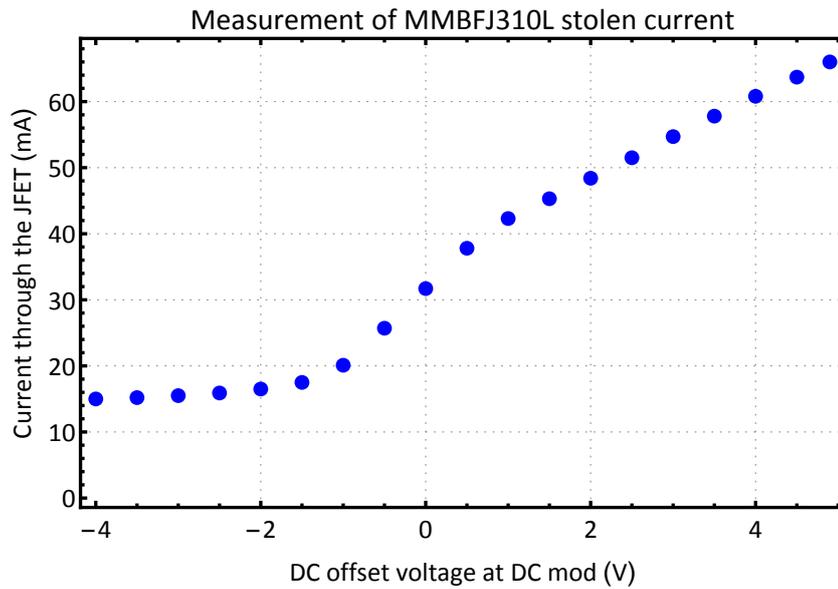


Figure 6.11: A test measurement of the current drawn through the MMBFJ310L JFET. The board was soldered accordingly to Fig. 6.8 and a DC voltage was applied to the DC-mod channel. The laser current driver was set to 80mA and the measured current through R9 was subtracted to get the current through the JFET.

6.4.4 Photodiode readout

To read off the photodiode we use a transimpedance amplifier, which was already explained in detail in Chapter 4.1. The circuits required to make it work on this board can be seen in Figure 6.9. After the Lemo connector X8, we have the two power diodes, preventing wrong currents from flowing through the circuit in case the power supply is connected incorrectly. The two large capacitors smooth out high frequency noise from the supply. Next, the two voltage regulators regulate the ± 15 V supply down to ± 12 V for our op-amp. The layout for the transimpedance amplifier is similar to what we have discussed in the previous chapters.

One thing that one should be careful with when assembling this circuit is the orientation of the photodiode. As we do not use a bias voltage, it is only important to correctly connect the grounded pin of the photodiode. The configuration varies with different laser diode cases, as they can have either 3 or 4 pins and the ground of the laser diode and the photodiode can be joined. In that case, the jumper JP5 in Figure 6.8 should also be soldered accordingly.

The reason why we wanted to include the photodiode readout is that we hoped it might help with the injection of the master laser beam into the slave laser diode. As the photodiode is placed in the case of the laser diode, we can couple in the master beam by optimizing the signal on the photodiode. Unfortunately, it proved that this optimization does not lead to the optimal mode of light for injection-locking. Therefore, it is better to try other methods such as fiber coupling the output of the laser diode that is rejected by the optical isolator.

Chapter 7

Measurements of the strontium atomic beam

In this chapter I will describe how we performed the measurements on the strontium atomic beam. The measurements allowed us to characterize the output of our new atomic oven. We want to determine how well different nozzle designs collimate the atomic beam. The measurements allow us to determine the atomic flux directly out of the oven and the atomic flux in the other sections of the vacuum system.

7.1 Absorption in an atomic beam

To interpret the data from the measurements we are performing, we first have to understand how the absorption in an atomic cloud works. This topic is usually discussed in the context of light-atom interaction, and we will take the basic formulas from Ref. [41].

Absorption of light in an absorptive medium is described by the Lambert-Beer law. It tells us how the intensity of light I changes after passing through an absorptive medium:

$$I = I_0 e^{-\text{OD}}, \quad (7.1)$$

where I_0 is the initial intensity of light and OD is the optical depth of the material given for a cloud of atoms by

$$\text{OD} = \int_0^l \sigma n dx. \quad (7.2)$$

Here, x is the coordinate along which the light passes, l is the length of the material, σ is the scattering cross-section, and n is the number density of atoms.

If we interrogate a cloud of two-level atoms with near-resonant probe light, well below the saturation intensity, the scattering cross section as a function of detuning can be described by

$$\sigma(\delta_0) = \frac{\sigma_0}{1 + \frac{4}{\gamma^2}(\delta_0 + \delta_d)^2}. \quad (7.3)$$

δ_0 is the detuning from the two-level transition given in units of Hz and γ is the natural linewidth of the transition. The resonant scattering cross section is given by

$$\sigma_0 = \frac{3\lambda_0^2}{2\pi} \quad (7.4)$$

and it only depends on the two-level transition wavelength λ_0 .

The effective detuning δ_d , which occurs due to the Doppler effect for atoms with a certain velocity \vec{v} , is given by

$$\delta_d = -\vec{k} \cdot \vec{v} \frac{1}{2\pi}, \quad (7.5)$$

where \vec{k} is the wave-vector of the light and \vec{v} the velocity of the atom.

If we select the probe laser beam to propagate perpendicularly to the atomic beam along axis x , then Eqn. 7.5 simplifies to

$$\delta_d = -kv_x \frac{1}{2\pi} = -\frac{v_x}{\lambda}, \quad (7.6)$$

where v_x is the velocity component along the x axis and the length of the wave-vector is expressed with the wavelength λ of the laser light $k = \frac{2\pi}{\lambda}$.

Additionally, we have to consider that atoms can have different velocities and rewrite the number density as an integral over the velocity distribution:

$$n = \int \frac{dn(v_x)}{dv_x} dv_x = \int \frac{dn(\delta_d)}{d\delta_d} d\delta_d. \quad (7.7)$$

Because of the linear relation between δ_d and v_x , we can substitute them to get the distribution over the detuning, which we will be directly measuring.

Taking these considerations into account, we arrive at a full description for the optical depth of our atomic beam

$$\text{OD}(\delta_0) = \sigma_0 D \int \frac{1}{1 + \frac{4}{\gamma^2}(\delta_0 + \delta_d)^2} \frac{dn(\delta_d)}{d\delta_d} d\delta_d, \quad (7.8)$$

where D is the diameter of the atomic beam. We took into account that our probe beam is much smaller than the atomic beam and crosses through its center. Furthermore, we can assume that the spatial density of the atomic beam is homogeneous as the nozzle is made out of numerous identical channels that homogeneously emit the atoms.

If the velocity distribution does not change much compared to the natural linewidth of our transition, we can approximate the Lorentzian function using a delta function

$\frac{1}{1 + \frac{4}{\gamma^2}(\delta_0 + \delta_d)} \approx \delta(\delta_0 + \delta_d)$ and pull the velocity distribution out of the integral. The remaining integral will give the area underneath the Lorentzian:

$$\text{OD}(\delta_0) \approx \sigma_0 D \frac{dn(\delta_0)}{d\delta_0} \int \frac{1}{1 + \frac{4}{\gamma^2}(\delta_0 + \delta_d)^2} d\delta_d = \sigma_0 D \frac{dn(\delta_0)}{d\delta_0} \frac{\pi\gamma}{2}. \quad (7.9)$$

The normalized intensity of the light, which we can measure in the probe beam, is given by

$$\frac{I}{I_0}(\delta_0) = e^{-\text{OD}(\delta_0)}. \quad (7.10)$$

We can also extract the total density of the atoms:

$$n = \frac{2}{\sigma_0 D \pi \gamma} \int \text{OD}(\delta_0) d\delta_0. \quad (7.11)$$

The expression for the total atomic flux is

$$\phi_{\text{Sr}} = n \langle v \rangle \frac{\pi D^2}{4}, \quad (7.12)$$

where we approximate the mean longitudinal velocity with the mean velocity of the modified Maxwell-Boltzmann distribution which was given in Eqn. 2.3.

Using the derived formulas, we can analyze the absorption data obtained from our atomic beam. The measurements are described in the next sections.

7.2 Experimental setup

The experimental setup is sketched and described in Figure 7.1. The measurement is performed on the oven section of the main strontium machine, which was mentioned in Chapter 2.2. The goal of the measurement is to determine the collimating properties of the nozzle and to estimate the atomic flux after the nozzle and after the differential pumping tube. We first measured the properties of a 3D printed titanium nozzle, to see if the prototype nozzle design performs well. Next, we replaced it with a steel nozzle which is made of a stack of welded hollow needles.

The light for the probe beams comes from the blue master laser and is coupled through an optical fiber that delivers it to the main experimental table. After the out-coupler, the light is split into the two probe beams using a PBS and a half-wave plate. The beams are shaped using a telescope, to achieve a smaller beam-waist of 0.8 mm.

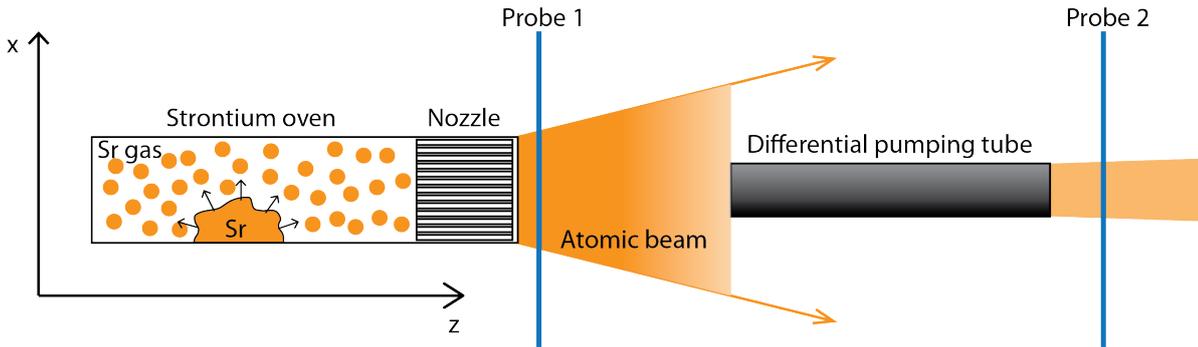


Figure 7.1: The experimental setup for the analysis of the atomic beam. The atomic oven is heated up to increase the vapor pressure of strontium. The nozzle is made out of several thin and long tubes that allow only atoms with velocities well aligned with the z axis to pass. The diameter of the nozzle is $\varnothing 12$ mm and the length of the tubes is 10 mm. The first probe beam is used to interrogate the atoms right after the nozzle. A part of the atomic beam flies through the aperture of the differential pumping tube and is detected using the second probe beam. The differential tube is 130 mm long and has an inner diameter of $\varnothing 5$ mm. The distance from the end of the nozzle to the differential tube is 30 mm.

The first probe beam passes just in front of the nozzle to make the approximations in Eqn. 7.8 valid. After passing through the atomic beam, the probe beam is collected on a photodetector.

The blue master laser is swept in frequency using the Toptica scan generator (SC 110). Data is recorded by connecting the outputs of the two photodiodes to an oscilloscope. The oscilloscope is triggered using the trigger output of the scan generator. Furthermore, we connect the error signal of the blue spectroscopy setup to the oscilloscope. This signal gives us a reliable frequency reference, as we can distinguish the different strontium isotopes in the error signal. The oscilloscope was set to average the spectra over 8 scans, to reduce the effect of noise. An example of the averaged oscilloscope traces can be seen in Figure 7.2.

The measurement is repeated for different temperatures of the strontium oven, which allows us to calculate the different atomic fluxes and determine at which temperature we want to run the oven. To see if there is any hysteresis while heating and cooling, we took the measurements in steps while heating up and also while cooling down. It showed that it takes approximately 15 minutes for the temperature in the oven to reach equilibrium. Furthermore, the nozzle is set at a 50 °C higher temperature, to hopefully avoid the nozzle getting clogged.

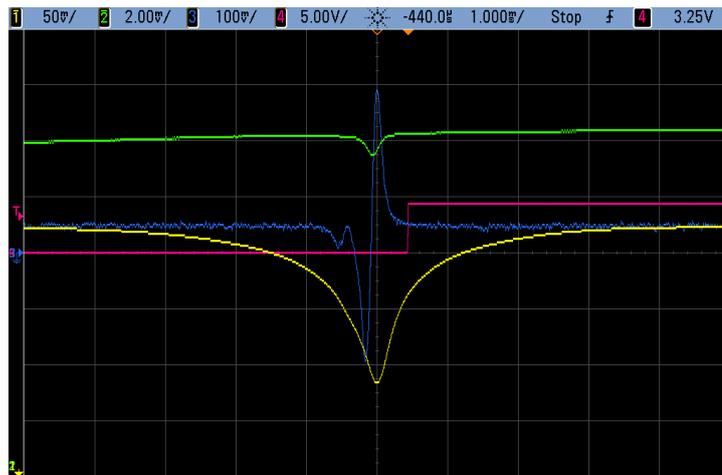


Figure 7.2: The measured atomic beam data on the oscilloscope, measured at 575 °C. Probe beam from the nozzle (yellow), probe beam after the differential tube (green), the error signal from the spectroscopy cell (blue), and the trigger of the frequency scan (pink).

Here I should note, that the first time the strontium oven is heated after a break of vacuum, the flux of atoms at low temperature is much lower than expected. We suspect that the oxide layer on the strontium has to be broken by heating to high temperatures, as the effect disappears after one heating cycle. The temperature of the oven where we observed the increase of flux is approximately 650 °C. A similar effect was reported in Ref. [18].

Unfortunately, we did not have an intensity stabilization system set up for these measurements, so the background intensity of the probe lights changed during the sweep. However, the background can be mostly compensated by proper normalization.

7.3 Data analysis

I analyzed the raw data from the oscilloscope using Python. First of all, the time axis from the oscilloscope had to be converted to a frequency axis. The conversion is done by fitting the error signal from the spectroscopy cell and by including the known information on the isotope shifts and their abundance. The error signal is modeled as a derivative of a Voigt function. A Voigt function is a convolution of a Lorentzian and a Gaussian depending on both their full width half maximums (FWHM). The Lorentzian shape comes from the natural linewidth of the two-level transition and we

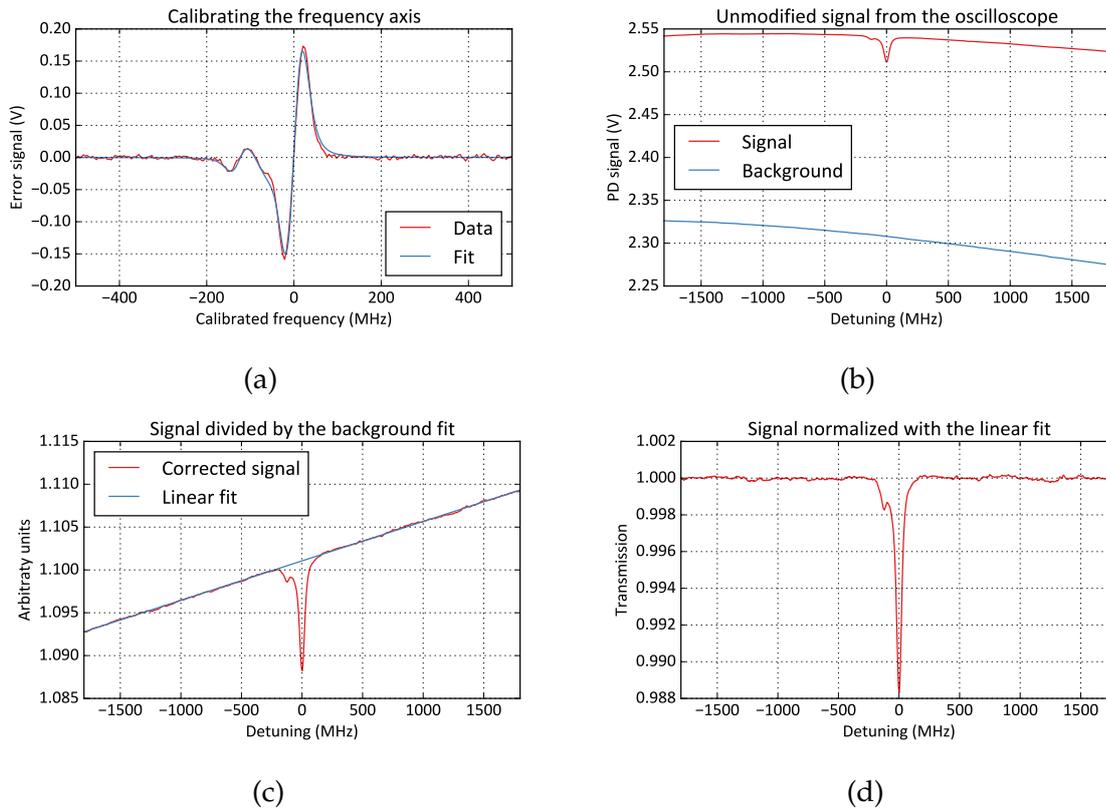


Figure 7.3: An example of processing the raw data from the oscilloscope to get the normalized transmission. Fitting the error signal to calibrate the frequency scale (a). Plot of the absorption after the differential pumping tube and the background taken after blocking the atoms (b). The signal after dividing with the fit of the background and the fit of the residual slope (c). The final processed transmission signal (d).

assume a Gaussian broadening to describe any other broadening effects.

The fit can be seen in Figure 7.3(a). It seems that the fitting function describes the data well. The free parameters are the FWHM of the Gaussian broadening, the amplitude of the error signal, the frequency offset and the frequency scale we wanted to determine.

After the frequency scale is correct, we want to normalize the photodiode signal, so that it shows us the transmission $\frac{I}{I_0}$. It turns out that the intensity modulation during the laser sweep can be fitted well with a quadratic function $ax^2 + bx + c$. We can fit it to a background signal, if we block the atoms, or even to the probe 2 signal, by excluding the region where there is absorption. Next, we can divide our absorption

signal with the fit of the background. If the fit was not made from data of the same channel or recorded at a slightly later time, there was usually an additional linear slope on the corrected signal. In that case, I fitted the slope and divided it out to get the normalized signal. An example of normalization with all the steps can be seen in Figure 7.3.

Unfortunately, it proved that the normalization method is not perfect. By comparing the different phases of the scan, we were able to see that the shape of the transmission signal gets slightly deformed. The deformation is not significant if we want to compare the shape of the velocity distribution, but could pose some problems when fitting a function. Therefore, for future measurements it would be preferable to use intensity stabilized light to assure a flat background.

After the normalization algorithm is properly working, we can look at the transmission spectra for all of the measured temperatures. In Figure 7.4 we can see how the light gets absorbed after the titanium nozzle. The velocity distribution is much broader than what we would expect for an ideal nozzle. After we removed the nozzle from the strontium oven, we realized that the strontium seems to have melted and flowed onto the back of the nozzle. We suspect that the barrier of strontium might be the cause for the unexpected velocity distribution. For the future, we ordered a boat shaped insert for the oven, that will contain the strontium and should prevent blocking of the nozzle.

In Figure 7.5, we show the transmission graphs for the second probe beam, which measures the atomic beam after the differential pumping tube. The atomic beam gets so well collimated by the narrow and long tube, that the width of the absorption is mostly limited by the linewidth of the transmission. However, the small aperture also greatly decreases the number of atoms that can be transmitted. The atomic flux we can get in this section of the vacuum system is crucial for decreasing the loading time of our magneto-optical trap.

The atomic density can be extracted from the transmission measurements using numerical integration. First, I took the negative of the logarithm of the transmission to get the optical depth in dependence on detuning. Next, I calculated the area under the function using the trapezoidal integration method. To avoid accumulating errors that show up due to non-perfect normalization, only data close to the absorption dip was integrated. Last of all, I inserted the value of the integral in Eqn. 7.11, getting the density of atoms. The equation should also be valid for the second probe beam, as in this case we can approximate the velocity distribution to be a delta function. The resulting atomic densities are plotted in Figure 7.6. The ratio between the atomic density at the nozzle and after the differential pumping tube is approximately 2200 for every temperature.

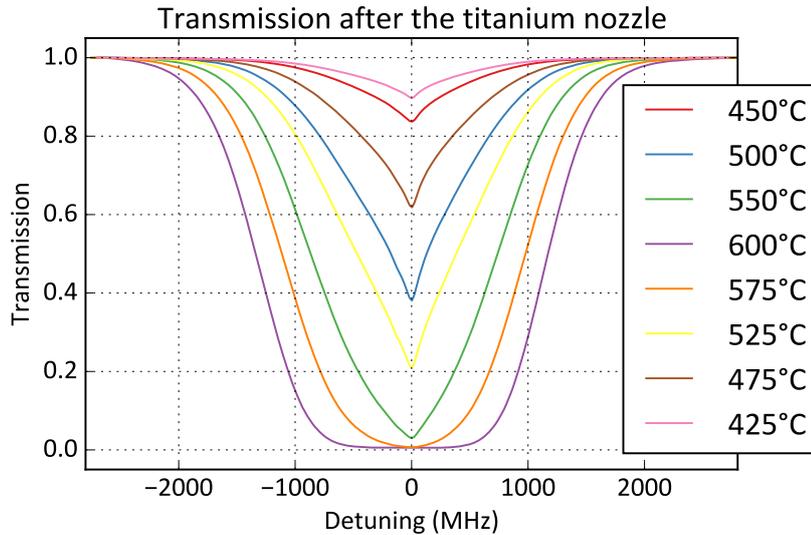


Figure 7.4: The transmission of the first probe beam passing through the atomic beam right after the nozzle. The transmission is plotted for different temperatures of the strontium oven, labeled in the order in which the measurements were taken.

For the atomic flux, I inserted the atomic density into Eqn. 7.12 and calculated the mean longitudinal velocity for every oven temperature. The resulting atomic flux can be seen in Figure 7.7. The ratio between the flux after the nozzle and the flux after the differential pumping tube is approximately 400. We assume that this ratio could be lower if the collimating properties of the nozzle were better. Therefore, we decided to replace the titanium nozzle with the steel nozzle.

After opening the vacuum chamber, we replaced the strontium in the oven and inserted the steel nozzle. Once the vacuum chamber was pumped back down, we started heating up the oven. The atomic flux from the oven was very low until we reached 650 °C in the oven and 700 °C at the nozzle. Since we are not certain exactly where the temperature sensors and heaters are located, the actual temperature could even be slightly higher. The melting point of strontium is close to 770 °C. During the cooling the optical depth was better comparable to the previous measurements with the titanium nozzle. We can see the transmission plots in Figure 7.8. The velocity distribution is much narrower than the one measured for the titanium nozzle. However, it seems that the total flux is decreased. We also checked the absorption after the differential pumping tube, but it seemed to be much lower. Possibly, the pointing of the oven should have been realigned to correctly point the better collimating nozzle.

After another break of vacuum, which oxidized the strontium in the oven, we did

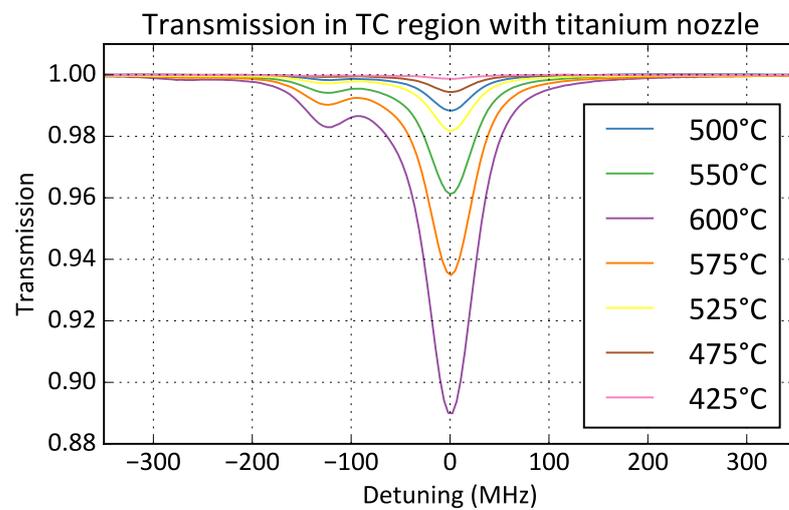


Figure 7.5: The transmission of the second probe beam passing through the atomic beam right after the differential pumping tube. In this section of the vacuum system we can perform transverse cooling (TC) to decrease the transverse velocity of the atoms. The transmission is plotted for different temperatures of the strontium oven, which are labeled in the order in which the measurements were taken.

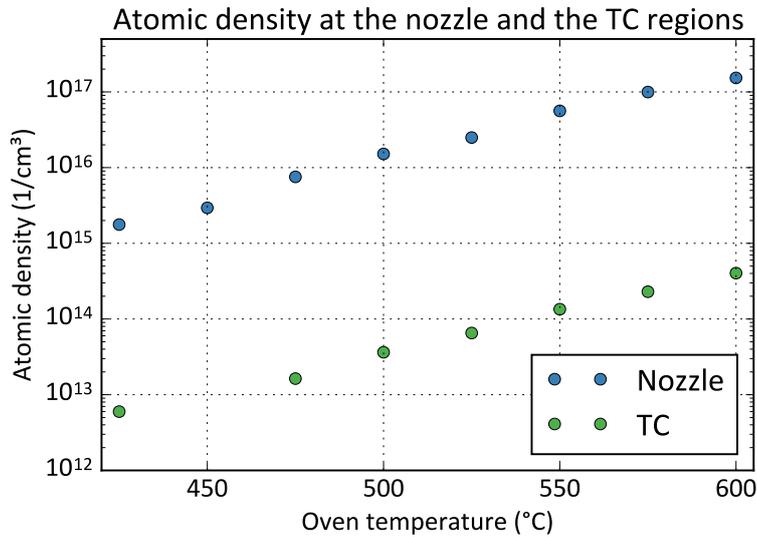


Figure 7.6: Atomic density, calculated for different oven temperatures, just after the titanium nozzle of the oven and in the transverse cooling region, close after the differential pumping tube.

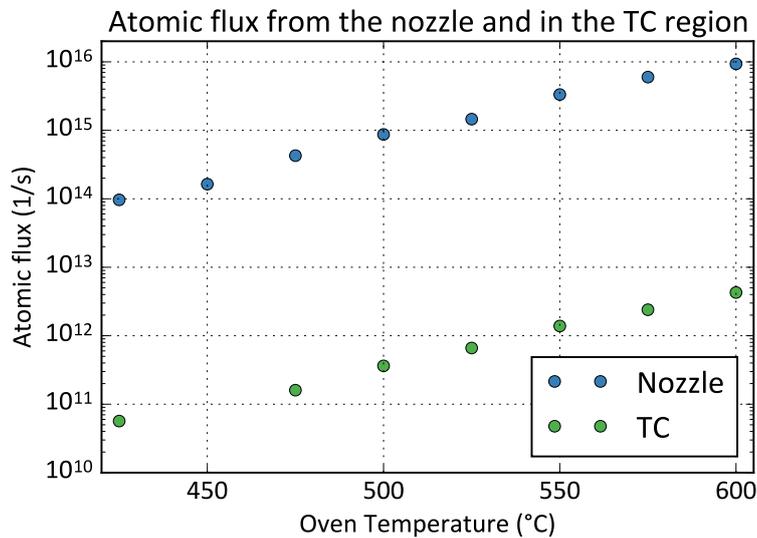


Figure 7.7: Atomic flux at different temperatures of the oven. The flux from the titanium nozzle and the flux in the transverse cooling region from the differential pumping tube are calculated using the atomic density.

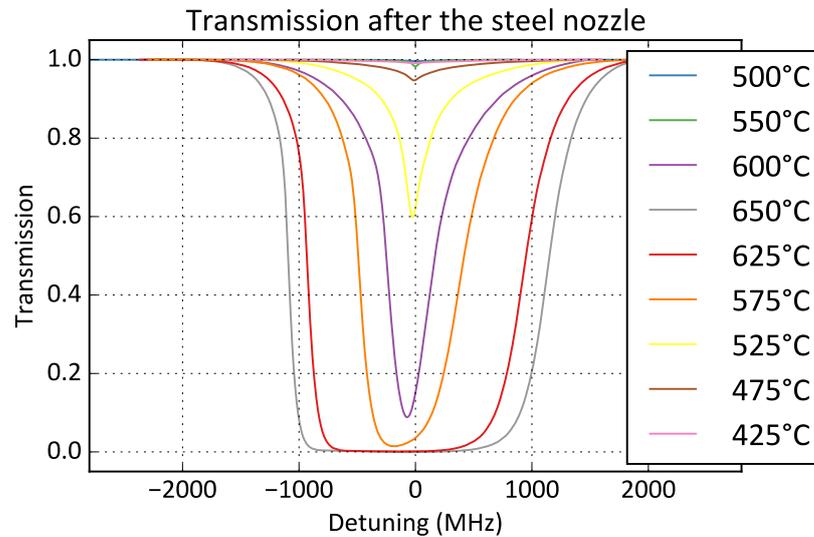


Figure 7.8: The transmission after the steel nozzle for different oven temperatures. This measurement was taken during the first heating of the replaced strontium, therefore there is barely any strontium during the heat-up. The legend is sorted in the order in which the measurements were taken.

another quick measurement to check if the atomic flux can be restored. Surprisingly, it seems that after the initial heat-up the velocity distribution after the steel nozzle looked much more like the one in Figure 7.4. I suspect that the strontium flowed onto the back of the nozzle again. The nozzle should be investigated when the new insert for the strontium oven is put in place.

Chapter 8

Conclusion

I report on the design of a frequency-stable diode laser system and have determined the requirements of the laser system for cooling and trapping of strontium on the 3P_1 transition. I have designed a setup that provides all the optical frequencies required to trap the different strontium isotopes using a magneto-optical trap. A system of acousto-optic modulators allows for precise control of the frequency and power required to optimize the final temperature of the atomic cloud. The necessary acousto-optic modulators were bought and tested to show that the specified diffraction efficiency can be reached, but the final setup is yet to be assembled. We found that sufficient optical power can be provided using the tapered amplifiers we built.

To reference the frequency of the diode laser to the atomic transition, we decided to use a home-built spectroscopy cell. Although the final setup for the modulation transfer spectroscopy is not yet operational, I have managed to assemble the spectroscopy cell and to test that it can be heated to high enough temperatures to observe absorption on the 3P_1 transition. While studying the absorption, I have also learned that the low saturation intensity for this transition means that it is best to use as little optical power as possible to optimize the error signal.

As we did not have any suitable photodetectors available at the time, I decided to learn more about photodetectors and to design boards that can be tailored to such specific needs. I report on the latest version of the photodetector board, which has been tested and shown to be limited by shot noise for most of the interesting range of operation. We found that using a spectrum analyzer is a great way to characterize the output of the photodetectors. After several iterations, the photodetectors seem to be working as expected and the explanation on the principles of operation should be a helpful guide for assembling more in the future.

The core part of the designed laser system is the stabilization of the laser frequency using the Pound-Drever-Hall technique. I report on how this technique works and how it is implemented. My main work here was the design of the housing for the reference cavity. I analyzed the limitations to the stability of an optical cavity and applied this knowledge to construct the vacuum system that houses the reference cavity. The high vacuum environment with two stages of temperature control and

good acoustic isolation should allow us to decrease the linewidth of the cavity to the level where thermal noise in the mirrors is the limiting effect. I aimed for this level of stability, as we hope to later use the same reference cavity for stabilizing the laser used for the 3P_0 clock transition of strontium. The vacuum housing has been assembled and the vacuum levels were tested. The temperature stabilization system and the characterization of the stability of the cavity are yet to be finished. I have achieved a lock of the laser using a test cavity setup, but the setup and parameters will have to be optimized once the high-finesse cavity can be used.

Apart from the work on the main project of my master thesis, I was also involved in the planning and building of the main experimental setup. I got to learn a lot about how a laboratory is organized and how it grows from an empty room to a place where state-of-the-art science can be performed. Apart from shorter tasks like ordering different parts and helping with the assembly of the vacuum system, I took on additional projects to help with the progress in the main experiment.

I report on the design of an injection-locking setup. For the 461 nm 1P_1 transition tapered amplifiers are not available and we found using injection-locked lasers a good alternative to obtain the necessary optical power. To improve the stability of these lasers, I designed a new breadboard where optics parts are referenced using dowel pins. A prototype of the laser module was produced and I assembled and tested it to see what further changes had to be made. I include a guide on how the parts for the injection-lock module were designed and how one can assemble them. The new mechanical design is working well and we are already planning to use a similar setup for the alignment of the magneto-optical traps. Furthermore, I have adapted the electronic circuit used to control the laser diodes. I describe the latest version of the board and show some of the measurement data that helped characterize its operation.

Last of all, I report on the measurements I did on the strontium atomic beam and on how I analyzed the data. The velocity distribution in the atomic beam can be probed using absorption spectroscopy. Using these measurements, we were able to determine the atomic flux in the beam and were able to estimate how well the oven nozzles are performing. As we found that strontium seems to be sticking to the nozzle, we have ordered an insert for the oven and hopefully the quality of the atomic beam will be improved.

Unfortunately, I did not get to completely finish all of my projects or to see the arrival of the experiment to the point where they can be implemented and put to use. Therefore, one of the main goals of this master thesis is to gather what I have learned about these projects and to act as a guide on how they were done.

Acknowledgments

First of all, I would like to thank Sebastian Blatt for the excellent supervision during my stay at the MPQ, for being an inexhaustible source of knowledge and never leaving my questions unanswered, and for motivating me to learn about every aspect of experimental physics by offering the wide variety of projects. Furthermore, I would like to thank him for all the advice given regarding the writing of this thesis and for all the effort spent when checking it for mistakes.

I would like to thank Immanuel Bloch, for giving me the opportunity to do a master thesis in his group. I have learned so much by working among such experts in the field and by being able to use the state-of-the-art equipment.

I would like to thank the other members of our team; Andre, Rodrigo, Annie and Stepan for all the successful collaboration in the building of our experiment and for making the whole journey in the lab and also aside of it a pleasant and fun experience. Especially, I would like to thank Andre for being a great office mate, for sharing his knowledge on optical cavities and for the companionship that extended beyond the work hours.

I would like to thank our interns Qing, Zhichao, and Yunpeng for the work done on their projects and for the help on the early stages of the experiment.

I would like to thank all the members of the Bloch's group for the interesting discussions around lunch time, the good pieces of advice, and all the help I received whenever I found I am missing a piece of equipment or do not understand how it works. Thank you for the friendly atmosphere at the group retreat and at the institute and for all the fun games of kicker played.

I would like to thank our well prepared technical staff for all the help with the more technical aspects of our projects. I would like to thank Anton for teaching me everything I now know about mechanical design, for having the patience to help turn my ideas into actual pieces of equipment and for always helping me find the correct tool and technique for the job. I would like to thank Karsten and Ollie for all the help with the electronics, for teaching me how to create printed circuit boards, how to understand the complicated pieces of equipment and how to test the devices.

Last but not the least, I would like to thank my family for the moral and financial support and for standing behind me no matter what path I decide to take. I would like

to thank my friends for always making me feel welcome and giving me the feeling I never left, when I did get the chance to visit home.

Bibliography

- [1] I. Bloch, J. Dalibard, and S. Nascimbene, “Quantum simulations with ultracold quantum gases,” *Nature Physics*, vol. 8, pp. 267–276, Apr 2012. (Cited on page [1](#))
- [2] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, “Optical atomic clocks,” *Rev. Mod. Phys.*, vol. 87, pp. 637–701, Jun 2015. (Cited on page [2](#))
- [3] A. Omran, M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch, and C. Gross, “Microscopic observation of pauli blocking in degenerate fermionic lattice gases,” *Phys. Rev. Lett.*, vol. 115, p. 263001, Dec 2015. (Cited on page [2](#))
- [4] M. F. Parsons, F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt, and M. Greiner, “Site-resolved imaging of fermionic ${}^6\text{Li}$ in an optical lattice,” *Phys. Rev. Lett.*, vol. 114, p. 213002, May 2015. (Cited on page [2](#))
- [5] S. Stellmer, *Degenerate quantum gases of strontium*. Ph.d., University of Innsbruck, 2012. (Cited on pages [5](#), [7](#), [8](#), and [14](#))
- [6] M. M. Boyd, *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time*. Ph.d., University of Colorado, 2007. (Cited on pages [6](#) and [14](#))
- [7] A. D. Ludlow, *The Strontium Optical Lattice Clock: Optical Spectroscopy with Sub-Hertz Accuracy*. Ph.d., University of Colorado, 2008. (Cited on page [5](#))
- [8] NIST, “The national institute of standards and technology (NIST) database,” 2016. [Online; accessed August-2016]. (Cited on page [5](#))
- [9] M. M. Boyd, T. Zelevinsky, A. D. Ludlow, S. Blatt, T. Zanon-Willette, S. M. Foreman, and J. Ye, “Nuclear spin effects in optical lattice clocks,” *Phys. Rev. A*, vol. 76, p. 022510, Aug 2007. (Cited on page [6](#))
- [10] J. E. Sansonetti and G. Nave, “Wavelengths, transition probabilities, and energy levels for the spectrum of neutral strontium (sri),” *Journal of Physical and Chemical Reference Data*, vol. 39, no. 3, 2010. (Cited on page [6](#))
- [11] T. Ido and H. Katori, “Recoil-free spectroscopy of neutral sr atoms in the lambda-dicke regime,” *Phys. Rev. Lett.*, vol. 91, p. 053001, Jul 2003. (Cited on page [6](#))

- [12] A. V. Taichenachev, V. I. Yudin, C. W. Oates, C. W. Hoyt, Z. W. Barber, and L. Hollberg, "Magnetic field-induced spectroscopy of forbidden optical transitions with application to lattice-based optical atomic clocks," *Phys. Rev. Lett.*, vol. 96, p. 083001, Mar 2006. (Not cited.)
- [13] T. Nicholson, S. Campbell, R. Hutson, G. Marti, B. Bloom, R. McNally, R. McNally, W. Zhang, M. Barrett, M. Safronova, G. Strouse, W. Tew, and J. Ye, "Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty," *Nat Commun*, vol. 6, Apr 2015. (Cited on page 6)
- [14] W. Forsythe, ed., *Smithsonian Physical Tables*. Knovel, 9th ed., online version ed., 1954. (Cited on pages 7 and 8)
- [15] A. Nescmeyanov, *Vapor pressure of the chemical elements*. Elsevier Pub. Co., 1963. (Cited on page 8)
- [16] D. Lide, *CRC Handbook of Chemistry and Physics, 84th Edition*. Taylor & Francis, 2003. (Cited on page 8)
- [17] N. Poli, *Raffreddamento ed intrappolamento di atomi di stronzio: verso un nuovo standard di frequenza nella regione ottica*. Ph.d., Universita degli studi di Firenze, Florence, Italy, 2005. (Cited on page 8)
- [18] K. R. Vogel, *Laser Cooling on a narrow atomic transition and measurement of the two-body cold collision loss rate in a strontium magneto-optical trap*. Ph.d., University of Colorado, Boulder, CO 80309-0440, 2005. (Cited on pages 8 and 87)
- [19] J. Vanier and C. Audoin, *The Quantum Physics of Atomic Frequency Standards*. No. v. 1 in *The Quantum Physics of Atomic Frequency Standards*, A. Hilger, 1989. (Cited on page 9)
- [20] I. Courtillot, A. Quessada, R. P. Kovacich, J.-J. Zondy, A. Landragin, A. Clairon, and P. Lemonde, "Efficient cooling and trapping of strontium atoms," *Opt. Lett.*, vol. 28, pp. 468–470, Mar 2003. (Cited on page 10)
- [21] M. Schioppo, N. Poli, M. Prevedelli, S. Falke, C. Lisdat, U. Sterr, and G. M. Tino, "A compact and efficient strontium oven for laser-cooling experiments," *Review of Scientific Instruments*, vol. 83, no. 10, 2012. (Cited on page 9)
- [22] H. Metcalf and P. van der Straten, *Laser Cooling and Trapping*. Graduate Texts in Contemporary Physics, Springer New York, 2001. (Cited on pages 9, 12, and 13)
- [23] S. Stellmer, R. Grimm, and F. Schreck, "Production of quantum-degenerate strontium gases," *Phys. Rev. A*, vol. 87, p. 013611, Jan 2013. (Cited on pages 13 and 14)

- [24] D. McCarron, "A guide to acousto-optic modulators," tech. rep., Technical report, Durham University, 2007. (Cited on page 18)
- [25] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts, "Double-pass acousto-optic modulator system," *Review of Scientific Instruments*, vol. 76, no. 6, 2005. (Cited on page 18)
- [26] E. D. Black, "An introduction to pound–drever–hall laser frequency stabilization," *American Journal of Physics*, vol. 69, no. 1, 2001. (Cited on pages 20, 21, and 22)
- [27] M. J. Martin, *Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock*. Ph.d., University of Colorado, Boulder, CO 80309-0440, 2013-04 2013. (Cited on pages 24 and 44)
- [28] D. J. McCarron, S. A. King, and S. L. Cornish, "Modulation transfer spectroscopy in atomic rubidium," *Measurement Science and Technology*, vol. 19, no. 10, p. 105601, 2008. (Cited on page 25)
- [29] R. K. Raj, D. Bloch, J. J. Snyder, G. Camy, and M. Ducloy, "High-frequency optically heterodyned saturation spectroscopy via resonant degenerate four-wave mixing," *Phys. Rev. Lett.*, vol. 44, pp. 1251–1254, May 1980. (Cited on page 25)
- [30] J. Graeme, *Photodiode Amplifiers: OP AMP Solutions*. Gain technology, McGraw-Hill Education, 1996. (Cited on pages 27 and 29)
- [31] R. P. Scott, C. Langrock, and B. H. Kolner, "High-dynamic-range laser amplitude and phase noise measurement techniques," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 7, pp. 641–655, Jul 2001. (Cited on page 30)
- [32] H. Nyquist, "Thermal agitation of electric charge in conductors," *Phys. Rev.*, vol. 32, pp. 110–113, Jul 1928. (Cited on page 30)
- [33] H. B. Callen and R. F. Greene, "On a theorem of irreversible thermodynamics," *Phys. Rev.*, vol. 86, pp. 702–710, Jun 1952. (Cited on page 41)
- [34] K. Numata, A. Kemery, and J. Camp, "Thermal-noise limit in the frequency stabilization of lasers with rigid cavities," *Phys. Rev. Lett.*, vol. 93, p. 250602, Dec 2004. (Cited on page 41)
- [35] G. Harry, T. Bodiya, and R. DeSalvo, *Optical Coatings and Thermal Noise in Precision Measurement*. Cambridge University Press, 2012. (Cited on page 42)

- [36] T. Legero, T. Kessler, and U. Sterr, "Tuning the thermal expansion properties of optical reference cavities with fused silica mirrors," *J. Opt. Soc. Am. B*, vol. 27, pp. 914–919, May 2010. (Cited on page 42)
- [37] S. A. Webster, M. Oxborrow, and P. Gill, "Vibration insensitive optical cavity," *Phys. Rev. A*, vol. 75, p. 011801, Jan 2007. (Cited on pages 43 and 47)
- [38] J. O'Hanlon, *A User's Guide to Vacuum Technology*. John Wiley Sons, 2003. (Cited on pages 52 and 53)
- [39] J. Sanjuan, N. Guerlebeck, and C. Braxmaier, "Mathematical model of thermal shields for long-term stability optical resonators," *Opt. Express*, vol. 23, pp. 17892–17908, Jul 2015. (Cited on pages 54 and 55)
- [40] E. Kirilov, M. J. Mark, M. Segl, and H.-C. Nägerl, "Compact, robust, and spectrally pure diode-laser system with a filtered output and a tunable copy for absolute referencing," *Applied Physics B*, vol. 119, no. 2, pp. 233–240, 2015. (Cited on page 68)
- [41] M. Inguscio and L. Fallani, *Atomic Physics: Precise Measurements and Ultracold Matter*. OUP Oxford, 2013. (Cited on page 83)

Declaration

I hereby declare that this thesis is my own work, and that I have not used any sources and aids other than those stated in the thesis.

date of submission

author's signature