Frequency stabilization of an external cavity diode laser

Cédric Huber
Master Thesis

Supervisors
Oliver Wipfli, Dr. Ilia Sergachev
& Prof. Dr. Jonathan Home
Trapped Ion Quantum Information Group, ETH Zürich

September 29, 2021

Abstract

We report an attempt of passive frequency stabilization of an external cavity diode laser (ECDL) by placing it in a pressure-proof housing. We have performed a series of measurements in order to quantify the stability over a large range of time scales. Additionally, we have written a Python script that determines if a laser is in single- or multimode operation by analyzing the interference patterns generated by a Fizeau wavelength meter.

Contents

1 Introduction

Trapped-ion quantum information experiments and more generally speaking atomic and molecular physics experiments require lasers to drive optical transitions. While the transition frequency is typically on the order of hundreds of terahertz, the linewidth of dipole-allowed transitions in earth-alkali ions such as calcium will be on the order of tens of megahertz and quadrupole transitions can even have linewidths as narrow as 1 Hz. Resonantly driving earth-alkali ions with lasers therefore requires a relative frequency stability on the order of 10^{-8} for dipole-allowed transitions down to 10^{-15} for quadrupole transitions [1]. Over the past decades the passive frequency stability of commercially available lasers has improved to levels that come close to the required stability for dipole-allowed transitions, however, drifts on a timescale of minutes to hours are still large compared to transition linewidths, making active frequency stabilization a necessity.
The frequency of lasers can be stabilized by using the dispersive or absorptive feature of a suitable resonance frequency supplied by either microscopic references such as atoms, molecules, and ions or macroscopic ones provided, e.g., by a Fabry–Pérot resonator [2]. A widely used method is the Pound-Drever-Hall (PDH) technique where the typical frequency reference is an optical cavity. We then say that the laser is locked to a cavity. However, this frequency lock is sometimes accidentally broken owing to a mode hop of the laser source that induces a discrete frequency jump. In that event the laser becomes unusable until parameters are readjusted, which in practice can cause anything from a short measurement interruption up to a complete loss of the ion. Because mode hops are difficult to predict in general, maintaining the frequency locking for a long time remains challenging [4]. For applications that require a laser linewidth of about 1 MHz, however, using a precision wavelength meter (WLM) as a reference suffices and bypasses the technical complexities of the PDH method [5]. Another advantage of this method is that it is possible to stabilize the laser frequency at an arbitrary value such that additional frequency shifts are not needed.
Both of the aforementioned stabilization techniques are active. Active stabilization schemes usually involve an electronic feedback system, where fluctuations of some parameters are converted to an electronic signal, which is then used to act on the laser. Passive schemes do not involve electronics and are based on purely optical effects [6].
A popular commercial external cavity diode laser is Toptica's DL pro which is the one used for the measurements within the framework of this thesis. Our model can be tuned from 840 to 875 nm and its typical applications are Cs cooling or Ca ion repumping. The DL pro is an external cavity diode laser (ECDL) in the Littrow configuration, see figure 1. We have focused on improving its passive stability, rather than active control. The advantages of passive frequency stabilization methods include unmodulated laser output and freedom to stabilize the laser frequency to any desired value. High passive stability also provides a reliable basis for more refined active stabilization methods to maintain single-mode operation of the laser and to prevent mode hops [7].

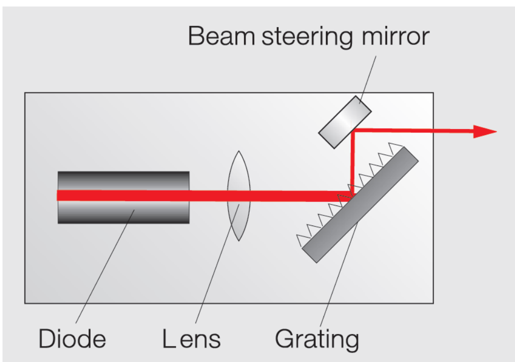

Figure 1: An external cavity diode laser in the Littrow configuration. Figure taken from [12].

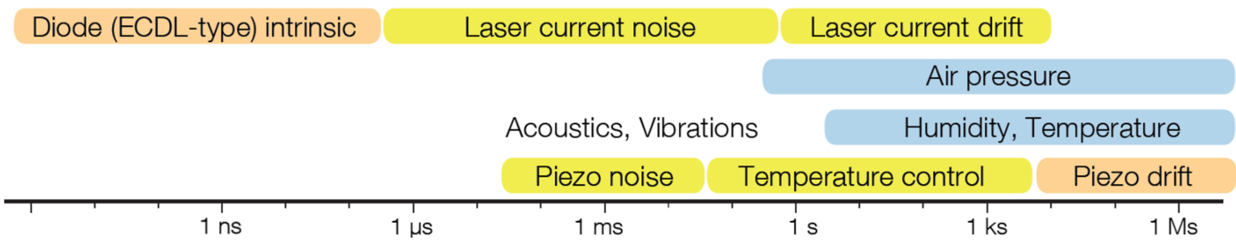

Figure 2: Main causes for frequency variations in ECDLs and their respective time scales. Figure taken from [12].

The manufacturer of the DL pro, Toptica, determined several causes for frequency variations and their respective time scales, see figure 2. The main observation we have made is that the frequency drifts of the DL pro are highly correlated with air pressure. Pressure fluctuations change the optical path length of the laser resonator. However, neither our laboratory environment nor the laser housing has air pressure stabilization, meaning that the free-running laser frequency is effectively tuned by daily weather. Thus we designed and tested a pressure chamber which isolates an entire Toptica DL pro from atmospheric pressure fluctuations. As it turns out, we are not the first ones that have tried this. We know of two research groups that made the same observations with their ECDL [7, 8] and then built a housing for it [3, 4]. They report great frequency stability, although it is not clear how much it was improved due to the housing.

The DL pro also offers an air pressure compensation feature. The laser controller measures the ambient pressure and compensates pressure changes by adjusting the piezo voltage, thereby adjusting the external cavity length in order to keep the optical cavity length constant. While this feature reduces frequency drifts and mode hops of a free-running laser it is not suitable for a locked laser since any change in frequency, regardless of its cause, is compensated right away anyway. Additionally, the described effect is only compensated to first order. Furthermore, the pressure is measured with a sensor inside the laser controller and not inside the actual laser housing. As a consequence, this feature cannot be used in combination with a pressure chamber.

Our goal is to quantify the frequency stability of our laser over a large range of time scales. A useful tool which helps us do this is the Allan deviation. As a reference, we use another 854 nm DL pro which is locked to a cavity with the PDH technique. Thus, we essentially have three different laser states which we can compare: the frequency stability of our laser when the chamber is open and when it is closed and the other laser which is locked to a cavity. So the questions of interest are how much does the stability improve when the laser is in a pressure chamber and how does it compare to a locked laser.

We also combine the light of both lasers on a photodiode which allows us to measure the beat note frequency, the difference of the two optical frequencies, and the combined linewidth of both lasers.

Another vital part of frequency stability is a laser's ability to remain in single-mode operation, which is what the second part of this thesis is about. In order to measure the frequency of a laser we use HighFinesse wavelength meters (WLMs). The input light is transmitted through a number of Fizeau interferometers which creates interference patterns that are recorded with photodiode arrays. A computer program allows to graphically display the patterns. They are then used to obtain the frequency of the light. However, the WLM just determines the frequency regardless of whether the laser is in single- or multi-mode operation. Therefore, the program lacks a feature that does exactly this. Up to now, one has had to look at the displayed interference patterns and assess in what operation state the laser is.

Thus we have decided to write a Python script that determines the spectral purity of a laser by reading in the interference patterns. For that we use the analytic formula for the transmission of a Gaussian beam by a Fizeau interferential wedge [20, 21] and compare the analytic patterns to the measured ones. The script periodically reads in the current pattern and returns the purity of the laser light.

2 Diode lasers

Diode lasers are the most common type of lasers produced, with a wide range of uses that include fiber optic communications, barcode readers, laser pointers, CD/DVD/Blu-ray disc reading/recording, laser printing, laser scanning and light beam illumination [9]. And of course scientific research. Their success story is driven by the fact that they are compact, conveniently operated, cost effective, and highly efficient. We start with explaining how a p-n junction allows for recombination of an electron with a hole which generates radiation in the form of an emitted photon. We continue with how such a junction can be used to build a diode laser. However, the emission spectrum of a solitary diode laser is broad, and the lasing wavelength is not well defined. Fortunately, there are solutions to this problem. One of them is the external cavity diode laser (ECDL), which will be discussed at the end of this section.

2.1 From the p-n junction to the diode laser

A p-n junction is a homojunction between a p-type and an n-type semiconductor. It acts as a diode, which can serve in electronics as a rectifier, logic gate, voltage regulator (Zener diode), or tuner (varactor diode); and in photonics as a light-emitting diode (LED), laser diode (LD), photodetector, or solar cell.

The junction consists of a p-type and and an n-type section of the same semiconductor materials in metallurgical contact. The p-type region has many holes and few mobile electrons, i.e. the holes are the majority carriers and the electrons are the minority carriers. In the n-type region it is the other way around, as one can see in figure 3a. The blue dots are electrons in the conduction band and the white dots are holes in the valence band. The bands are separated by a bandgap of energy E_g. Within this gap lies the Fermi energy E_f ($f(E_f) = 1/2$ where $f(E)$ is the Fermi-Dirac distribution function). For an intrinsic semiconductor we have n = p because electrons and holes are created in pairs, which implies that E_f must lie precisely in the middle of the bandgap. In the case of a p-type (n-type) semiconductor E_f lies below (above) the middle of the bandgap.

When the two regions are brought into contact, see figure 3b, electrons and holes diffuse from areas of high concentration toward areas of low concentration. Thus, electrons from the n-region move over to the p-region where they recombine with holes while leaving behind positively charged ionized donor atoms. Similarly, holes diffuse from the p-region into the n-region where they recombine with electrons leaving behind negatively charged acceptor atoms. As a result, a narrow region on both sides of the junction becomes nearly depleted of mobile carriers. This region is called the depletion layer and it contains only the fixed charges. Positive ions on the n-side and negative ions on the p-side. The fixed charges create a built-in electric field in the depletion layer that points from the n-side to the p-side of the junction. At some point, this field stops the diffusion of further mobile carriers and an equilibrium is established. The result is a potential difference V_0 between the two sides of the depletion layer that causes the energy bands to bend. In thermal equilibrium the entire structure is described by one Fermi-Dirac distribution which means that the Fermi levels in both regions must align.

An externally applied potential will alter the potential difference between the p- and the n-side, see figure 3c. If the junction is forward biased by applying a positive voltage V to the p-region, an electric field is produced in opposite direction of the built-in field. The external bias voltage destroys the equilibrium and this leads to a misalignment of the Fermi levels in the p- and n-regions, as well as in the depletion layer. In this state of quasi-equilibrium there are two Fermi levels in the depletion layer: a valence band quasi-Fermi level E_{f_v} and a conduction band quasi-Fermi level E_{f_c}. Thanks to the forward bias the height of the potential-energy hill is reduced by an amount eV, where e is the elementary charge. The excess majority carrier holes and electrons that enter the n- and p-regions, respectively, become minority carriers and recombine with the local majority carriers. Their concentration therefore decreases with distance from the junction. This process is known as minority carrier injection. A portion of these electron-hole recombinations are radiative which is quantified by the internal quantum efficiency η_i of a semiconductor material. It is defined as the ratio of the radiative

electron-hole recombination to the total recombination (radiative and nonradiative) coefficient.
This means that electron-hole recombination can used to emit light from a semiconductor material. In order to produce discernible radiation sufficient numbers of electron-hole pairs are required which is why we have to forward bias the p-n junction by a large enough voltage, see figure 3d. The external source of energy causes holes from the p-side and electrons from the n-side to be forced into the common junction region by the process of minority carrier injection, where they recombine and emit photons. The recombination radiation is called injection electroluminescence. Electroluminescence is a phenomenon where light is emitted by a material that is subjected to an electric field. The conduction band and valence band quasi-Fermi levels, E_{f_c}, E_{f_v}, lie within the conduction and valence bands and a state of quasi-equilibrium exists within the junction region. The quasi-Fermi levels are sufficiently well separated so that a population inversion is achieved and net gain may be obtained over the bandwidth $E_g \leq h\nu \leq E_{f_c} - E_{f_v}$ within the active region.
A device that is closely related to the diode laser is a light-emitting diode (LED). A LED simply is a forward-biased p-n junction, usually fabricated from a direct-bandgap semiconductor material, that emits light via injection electroluminescence. The light emitted from an LED is generated by spontaneous emission. If the forward voltage is increased beyond a certain point the number of electrons and holes in the junction region becomes sufficiently large such that population inversion is achieved, whereupon stimulated emission (i.e. emission induced by the presence of photons) becomes more prevalent than absorption. Under these conditions and with appropriate feedback, the junction region may be used as a diode laser. The light generated by a diode laser arises from stimulated emission.
In its simplest configuration, the optical feedback is provided by plane mirrors, which are usually implemented by cleaving the semiconductor material along its crystal planes, see figure 4. The sharp refractive index difference between the crystal and the surrounding air causes the cleaved surfaces to act as reflectors. Thus, the semiconductor crystal acts both as a gain medium and as a Fabry-Perot optical resonator. Provided that the gain coefficient is sufficiently large, the feedback converts the junction into an optical oscillator, i.e. a laser.
Semiconductor photon sources such as LEDs and diode lasers are highly efficient electronic-to-photonic transducers, i.e. they can efficiently convert electronic into photonic energy.
If electron-hole pairs are generated at a steady rate R (pairs per unit volume per unit time) by means of an external injection mechanism, then this leads to the generation of a photon flux Φ (photons per second), generated within a volume V of a semiconductor material, of

$$\Phi = \eta_i R V, \tag{1}$$

where η_i is the internal quantum efficiency. Direct-bandgap semiconductors are usually used to make LEDs and diode lasers because η_i is substantially larger than it is for indirect-bandgap semiconductors.
There are three efficiencies associated with diode lasers:

- The internal quantum efficiency η_i, which accounts for the fact that only a fraction of all electron-hole recombinations are radiative.
- The extraction efficiency η_e, which accounts for the fact that only a portion of the light lost from the cavity is useful.
- The power-conversion efficiency η_c, which is the ratio of the emitted optical power to the electrical power supplied to the device.

The spectral intensity of light emitted by a diode laser is characterized by the following three features:

1. The spectral width of the gain coefficient is relatively large because transitions occur between two energy bands rather than between two discrete energy levels.
2. Semiconductors tend to be homogeneously broadened. Nevertheless, spatial hole burning allows the oscillation of many longitudinal modes at the same time. Spatial hole burning is particularly

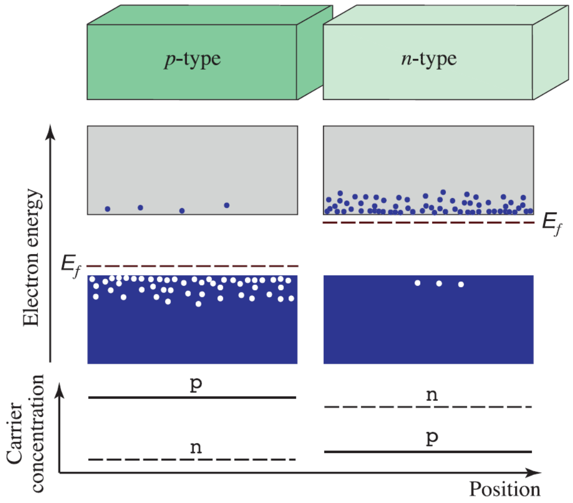

(a) A p-type and n-type semiconductor before contact.

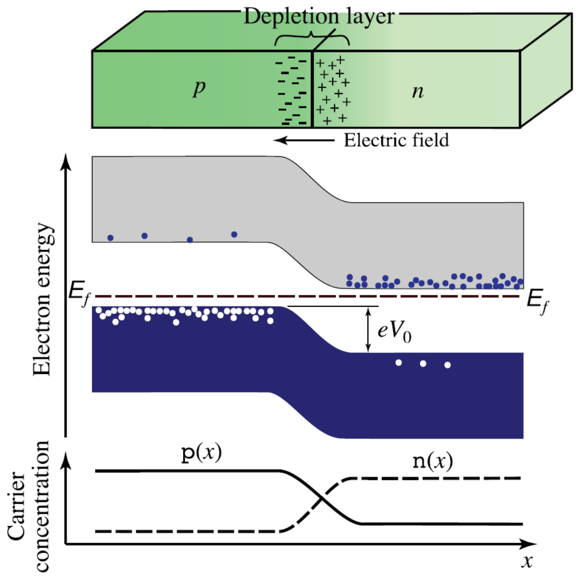

(b) A p-n junction in thermal equilibrium.

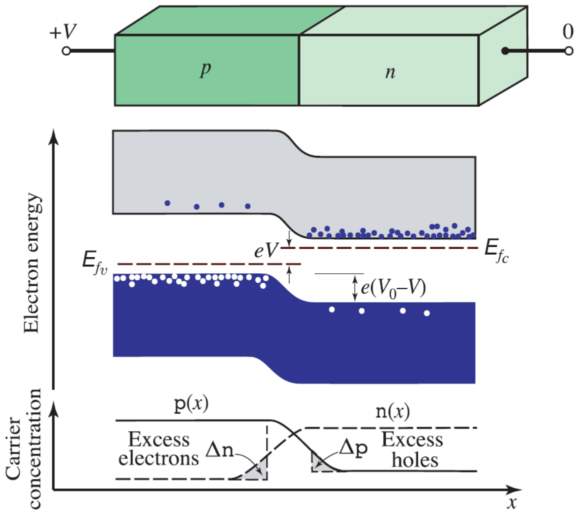

(c) A forward-biased p-n junction.

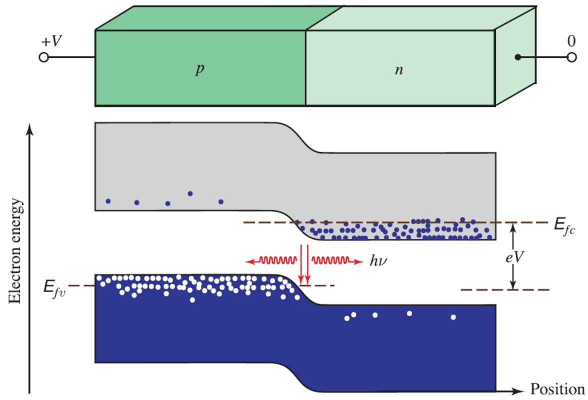

(d) A heavily doped p-n junction that is strongly forward biased.

Figure 3: Energy levels and carrier concentrations for three different scenarios. E_f is the Fermi energy that in (c) splits up in valence and conduction band quasi-Fermi levels E_{f_v} and E_{f_c}. In (b) V_0 is the potential difference between the p-side and n-side of the junction. In (c) this difference is reduced to $V - V_0$ where V is the applied voltage. e is the elementary charge. In (d) large amounts of electrons and holes radiatively recombine within the junction region which creates photons with energy $h\nu$. The concentrations of mobile electrons p(x) and holes n(x) are shown as functions of the position x. Figures taken from [10].

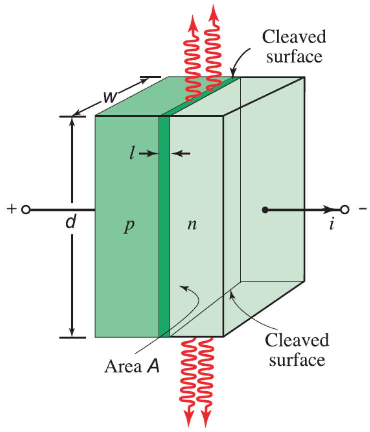

Figure 4: A diode laser. Two surfaces that are perpendicular to the plane of the p-n junction act as reflectors that form an optical resonator of length d and cross-sectional area lw. The other two perpendicular surfaces are often roughened to eliminate feedback. Charge carriers travel perpendicularly to the p-n junction, whereas photons travel in the plane of the junction. The laser beam originates from a small area and has a large divergence with opening angles of several tens of degrees. The opening angle can be reduced by a lens of short focal length close to the front facet of the diode laser. Figure taken from [10].

prevalent in short cavities in which there are few standing-wave cycles. This permits the fields of different longitudinal modes, which are distributed along the resonator axis, to overlap less, thereby allowing the partial spatial hole burning to occur.

3. The semiconductor resonator length d is short compared to most other types of lasers. The frequency spacing is therefore relatively large. Nevertheless, many such modes can generally be supported within the broad bandwidth B over which the gain exceeds the loss. The number of possible laser oscillation modes is $M \approx B/\nu_F$ where ν_F is the spacing between neighbouring modes, i.e. the free spectral range (FSR) [10].

2.2 Frequency noise

In contrast to the majority of gas lasers and solid-state lasers the linewidth of a solitary diode laser is determined by the quantum process of spontaneous emission. Each photon spontaneously emitted into the laser mode can be multiplied by stimulated emission and the resulting field amplitude adds to the field in the diode laser. On the other hand, the small line-quality factor Q of the laser resonator does not sustain a very well defined phase of the internal field. As a result, the phase of the combined field shows considerable fluctuations owing to the statistically fluctuating contributions of the spontaneously emitted photons leading to a large Schawlow–Townes linewidth. The Schawlow–Townes linewidth is the fundamental (quantum) limit for the linewidth of a laser [11].
In diode lasers the spontaneous emission leads to an additional fluctuation of the index of refraction

$$\Delta n = \Delta n' + i\Delta n'' \tag{2}$$

within the duration of the relaxation oscillation, i.e. during $\approx 1\,\text{ns}$. Here, n' describes the dispersion and n'' the absorption. In essence, there is a coupling between the phase and the amplitude of the light wave expressed by Henry's coupling parameter

$$\alpha = \frac{\Delta n'}{\Delta n''} \tag{3}$$

which leads to an excess line broadening with respect to the Schawlow–Townes linewidth $\Delta\nu_{\text{ST}}$. α depends on the material of the diode laser. For GaAs and $\lambda \approx 850\,\text{nm}, \alpha \approx 4$ has been determined.

The emission line of a diode laser is found to be a Lorentzian with a full width at half maximum (FWHM) given by the so-called "modified Schawlow–Townes linewidth":

$$\Delta\nu_{\mathrm{LD}} = \frac{\pi h\nu(\Delta\nu)^2}{P}(1+\alpha^2), \tag{4}$$

with photon energy $h\nu$, FWHM $\Delta\nu$ of the resonator and output power P. It holds that $\Delta\nu_{\mathrm{LD}} = \Delta\nu_{\mathrm{ST}}(1+\alpha^2)$ [2].
In general, processes that are faster than the laser frequency measurement itself add to the laser's technical linewidth. Slower processes will "only" lead to a frequency drift between successive measurements. When comparing linewidth values, it is also important to consider the time scale of the measurement [12].

2.3 Frequency stability

In order to find the frequency ν_m of the m-th longitudinal mode the phase shifts occurring at the laser facets have to be taken into account. The facet at the rear of the diode laser is often coated as a mirror of high reflectivity and the laser field can be thought of as a standing wave with a node at the facet. This is obviously not true for the other facet that serves as an output coupler with a typical coefficient of reflection of $R \approx 35\%$. The following relation holds:

$$m\lambda = 2n(\nu)L + \varphi\frac{\lambda}{2\pi} = m\frac{c}{\nu} \qquad \Rightarrow \qquad \nu_m = \frac{mc}{2n(\nu_m)L + \frac{\varphi c}{2\pi\nu_m}}. \tag{5}$$

Hence, the frequency ν_m of the emitted light depends on the mode number m, on the phase shift φ due to reflection, on the length $L = L(T)$ of the laser crystal and on the index of refraction n. The speed of light is denoted by c. In the case of small variations of these parameters one can assume

$$\frac{\Delta\nu}{\nu} = \frac{\Delta m\nu_F}{\nu} - \frac{\Delta\varphi\nu_F}{2\pi\nu} - \frac{\Delta L}{L} - \frac{\Delta n}{n} \tag{6}$$

with the free spectral range ν_F [13]. The first contribution of (6) can lead to mode hops of about 100 GHz. The phase φ in the second term of (6) can be varied in particular by coupling back a part of the light emitted by the diode laser. This effect can be used for frequency stabilization of diode lasers. On the other hand, spurious radiation reflected back, e.g., by the window of the housing of the diode laser, by the collimating lens or from other optical components, can alter the frequency of the laser, in particular if the phase of the back-reflected light fluctuates.
The last two contributions of (6) are influenced by the temperature of the laser. For small temperature variations the length $L(T)$ of the laser diode is expected to vary linearly with the temperature. The index of refraction n influences the laser frequency in a complicated way since it varies with the frequency ν, temperature T, the injection current I and the laser power $P(n = n(\nu, T, I, P))$. Temperature fluctuations affect the index of refraction via different effects. In general, raising the temperature of a laser diode increases its wavelength where the monotonic variation is interrupted by discontinuous jumps.
Besides the ambient temperature, the temperature of the diode laser is also affected by the injection current. With a nearly constant voltage drop across the p-n junction the dissipated power, and hence the temperature increase, is proportional to the injection current. Consequently, a smooth increase of the injection current results in a red detuning of the frequency of the laser due to the associated temperature variation. To achieve good long-term frequency stability these values ask for a low-noise power supply with low ripple. Variation in the current, in general, also affects the index of refraction by the changed number of the free charge carriers [2].

2.4 Linewidth reduction by optical feedback

The linewidth of a solitary diode laser of the Fabry–Pérot type can be as high as tens to hundreds of megahertz. It has been shown that feedback from an external mirror can be used to change the linewidth of a diode laser. Let us consider a solitary diode laser where part of the emitted radiation is coupled back into the diode laser, see figure 5. In this case, the electromagnetic field in the laser is composed of a superposition of the internal field and the back-coupled field. Depending on the phase between both fields the amplitude of the resulting field is increased or reduced. Due to the strong coupling between the phase and the amplitude of the diode laser a sudden power fluctuation in the active medium will lead to a fluctuation of the phase of the field in the diode laser. Consequently, the field coupled back after an external round trip time $\tau_d = 2L_d/c$ also suffers from a phase shift. Depending on the phase difference between the internal and the backcoupled field, the initial perturbation will be increased or reduced.

In the latter case, the influence of the internal field leads to reduced frequency fluctuations. As a consequence, the frequency of the laser is very sensitive to the phase of the back-reflected radiation. A small fraction $\beta = P_R/P_l$ of the power P_l of the laser coupled back into the laser is sufficient to influence the frequency of the diode laser considerably. For a more specific characterization of the feedback process not only the fraction β of the power of the laser diode that is coupled back, but also the ratio of the round-trip times τ_d/τ_D need to be taken into account where τ_D is the round-trip time inside the laser diode. The feedback is characterised by the parameter

$$C = \frac{\tau_d}{\tau_\mathrm{D}} \frac{1 - r_2^2}{r_2} r_3 \sqrt{1 + \alpha^2}. \tag{7}$$

Different regimes corresponding to different values of C have been identified.

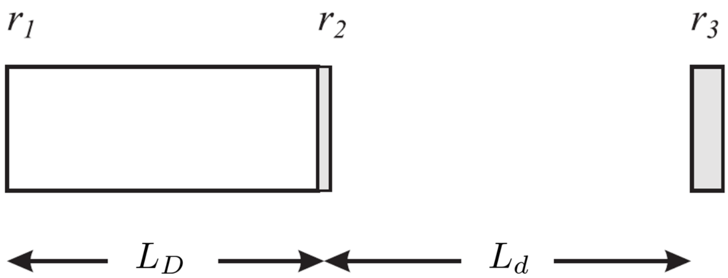

Figure 5: Diode laser with feedback from an external reflector referred to as an external cavity laser. For $r_2 \ll 1$ the device becomes an external cavity diode laser. Figure taken from [2].

Depending on the elements and configurations used to accomplish the feedback, several terms have been used. The laser diode by itself without any extra elements is referred to as a "solitary diode laser". If an external reflective element is added to the solitary diode laser the combined arrangement is termed "External Cavity Laser" (ECL) since an external cavity is formed by the reflective element and the output coupler of the solitary laser. In general, only a fraction of the power reflected by the external reflector is fed back into the active zone of the laser diode. If the output coupler (front facet) has a low reflectivity, e.g., from an anti-reflective (AR) coating ($r_2 \ll 1$), the laser cavity is formed by the rear mirror of the laser diode and the external reflector. This is called the strong feedback regime, where $C \gg 1$ holds. The laser then acts as an "Extended Cavity Diode Laser" (ECDL). Often these laser are also referred to as "External Cavity Diode Laser" (ECDL), which is what we use [2].

2.5 External cavity diode laser

Unfortunately, in ordinary solitary diode lasers, because of the semiconductor's inherent broad gain spectrum, usually more than one mode will operate simultaneously resulting in multiple output wave-

lengths and broad spectral linewidth. As can seen in figure 7, the wide gain profile of the semiconductor supports many internal modes simultaneously, each with a different frequency. Fortunately, the laser system can be forced to operate in a single longitudinal mode by placing the diode laser in an external resonant cavity comprising a wavelength selector. Through insertion of the wavelength selective elements in the system, the external cavity configuration permits operation in a single mode with a linewidth which is considerably less than that of a solitary laser diode. Moreover, selection and tunability of the emission wavelength by external control can be executed without the complications associated with variation of the temperature or pumping level. Wavelength tunability is one of the main important features of these lasers.
Rather than using a plain mirror commercially available ECDLs often comprise a diffraction grating placed at a distance of a few centimetres from the diode laser. There are two principal designs of the ECDLs with diffraction gratings as dispersive wavelength selectors. They are named Littrow and Littman configurations. The laser we are using is in the Littrow configuration. The lasing wavelength in each of these designs is determined by the center of the grating dispersion curve, which is given by the grating equation

$$m\lambda = a(\sin\phi_i + \sin\phi_d). \tag{8}$$

Here $m \in \mathbb{Z}$ is the diffraction order, λ is the center wavelength, a is the grating groove spacing, ϕ_i is the incident angle measured with respect to the grating normal and ϕ_d is the diffracted angle.
In the Littrow configuration, see figure 6, the grating is positioned at the Littrow angle under which the incident and diffracted angles are equal, i.e. $\phi_i = \phi_d = \theta$. Equation (8) reduces then to the following

$$m\lambda = 2a\sin\theta. \tag{9}$$

Most frequently the 1st order diffracted beam ($m = 1$) is directed back into the laser diode to generate optical feedback. The retro-reflected laser frequency ν_g imposed by the grating is given by

$$\nu_g = \frac{c}{2a\sin\theta}. \tag{10}$$

This is the center frequency of the grating feedback band-pass filter whose width $\Delta\nu_g$ can be calculated from equation (10) and is inversely proportional to the number of grating lines N illuminated. That is,

$$\frac{\Delta\nu_g}{\nu_g} = \frac{1}{N}. \tag{11}$$

Tuning of ECDLs in Littrow configuration is accomplished via wavelength selective properties of the grating feedback. According to eq. (10), change of the angle θ by rotating the grating about the axis parallel to the grooves tunes the laser to a unique diffracted wavelength while nearby wavelengths are returned from the grating at differing angles and thus are suffering greater losses. The emitted laser wavelength is determined by a combination of the standing-wave conditions, see figure 7. As the grating rotates the laser tunes by hopping from one external cavity longitudinal mode to the next. To tune between the external cavity longitudinal modes, the cavity length must be adjusted in such a way that the cavity longitudinal mode wavelength matches the wavelength of the grating in the retro-reflection conditions. A mode hop occurs when the difference between the modal and grating frequencies (ν_m and ν_g, respectively) exceeds half of the FSR, i.e., when

$$|\nu_m - \nu_g| \geq \frac{1}{2}\frac{c}{2L_{\mathrm{EC}}}, \tag{12}$$

where L_{EC} is the optical path length of the external cavity. Proper tuning involves therefore changing the diffraction angle as well as the external cavity length so that, the mode structure of the external cavity and the diffraction peak of the grating remain superimposed. This is equivalent to maintaining the same number of waves in the cavity at all wavelengths and places strict requirements on the mechanical tolerances and relative motions associated with the laser/grating/mirror combination. Continuous tuning requires simultaneous cavity-length changing and grating rotation, for example by using mechanical devices that synchronize the two motions. Fortunately, a continuous tuning range, greater than half of the FSR, can be accomplished by rotating the grating about a carefully chosen pivot. It has been shown that the optimum pivot is located at the intersection between the grating plane and the line parallel to the rear facet of the diode [14].

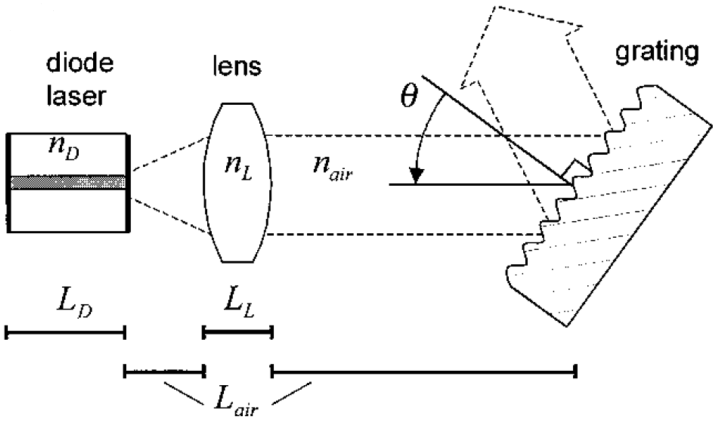

Figure 6: Schematic layout of an ECDL in Littrow configuration. The total optical path length in the cavity is $L_{\mathrm{EC}} = L_D n_D + L_L n_L + L_{\mathrm{air}} n_{\mathrm{air}}$. The rear facet has a highly reflective coating while the front facet features an anti-reflection coating. Figure taken from [3].

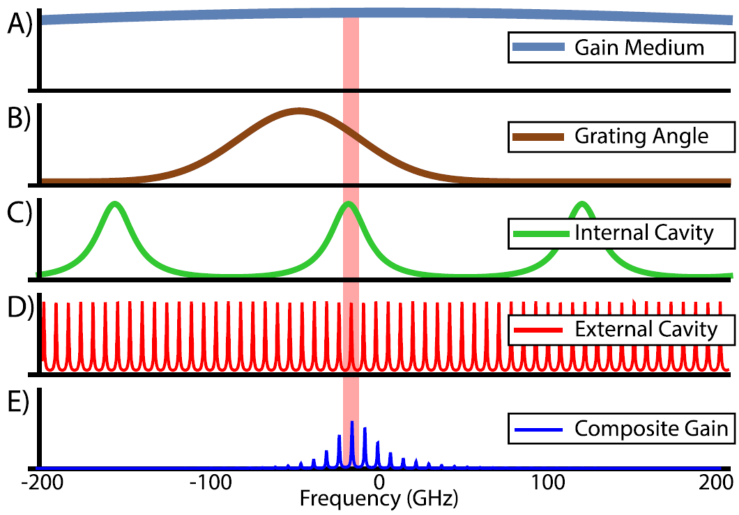

Figure 7: Mode selection for an ECDL. The semiconductor gain profile, the grating filter, the internal diode laser modes and the external cavity modes determine the lasing mode(s). The mode with the largest overall gain is selected. Figure taken from [15].

3 Allan deviation

A measure to quantify frequency stability of clocks, oscillators and amplifiers is the Allan variance $\sigma_y^2(\tau)$, also known as two-sample variance. The Allan deviation is the square root of the Allan variance, $\sigma_y(\tau)$ [16].

Let us assume the instantaneous frequency of the laser that we measure with the wavelength meter is $\nu(t)$ or ν_i for a series of discrete readings. The fractional frequency $y(t)$ is then the normalized difference between the frequency $\nu(t)$ and the nominal frequency ν_0:

$$y(t) = \frac{\Delta\nu(t)}{\nu_0} = \frac{\nu(t) - \nu_0}{\nu_0}. \tag{13}$$

The function $y(t)$ can be reduced to a discrete series of measurements by averaging over the measurement time τ:

$$\bar{y}_i = \frac{1}{\tau}\int_{t_i}^{t_i+\tau} y(t)\,\mathrm{d}t\,. \tag{14}$$

Analogously for discrete readings ν_i.

If one has N samples, the measuring time τ of a single sample and the time T between consecutive measurements which may differ from τ by the dead time $T-\tau$, then one can calculate the N-sample variance as follows

$$\sigma^2(N,T,\tau) = \frac{1}{N-1}\sum_{i=1}^{N}\left(\bar{y}_i - \frac{1}{N}\sum_{j=1}^{N}\bar{y}_j\right)^2. \tag{15}$$

It is now generally agreed to follow a proposition made by Dave Allan which is to select from all possible sample variances the expectation value of the so-called two-sample variance with $T=\tau$. This so-called Allan variance $\left\langle\sigma_y^2(2,\tau,\tau)\right\rangle$, which is also written as $\sigma_y^2(\tau)$, is defined as

$$\sigma_y^2(\tau) = \left\langle\sum_{i=1}^{2}\left(\bar{y}_i - \frac{1}{2}\sum_{j=1}^{2}\bar{y}_j\right)^2\right\rangle = \frac{1}{2}\left\langle(\bar{y}_2 - \bar{y}_1)^2\right\rangle. \tag{16}$$

According to the definition it has to be ensured that there is no dead time between two adjacent measurements. From the squared normalised frequency differences between two adjacent pairs ν_i and ν_{i+1}, the mean value is computed and divided by 2 to give the Allan variance $\sigma_y^2(\tau)$ for the particular measuring time τ:

$$\sigma_y^2(\tau) = \frac{1}{2(N-1)}\sum_{i=1}^{N-1}(\bar{y}_{i+1} - \bar{y}_i)^2. \tag{17}$$

To make a good approximation of the expectation value ($\langle\rangle$) a sufficiently large number of frequency differences has to be used.

To derive the data for longer times, e.g. $\tau = 3\tau_0$, the consecutive values of $\bar{y}_{1,\tau} = (\bar{y}_{1,\tau_0} + \bar{y}_{2,\tau_0} + \bar{y}_{3,\tau_0})/3$, $\bar{y}_{2,\tau} = (\bar{y}_{4,\tau_0} + \bar{y}_{5,\tau_0} + \bar{y}_{6,\tau_0})/3$, $\bar{y}_{3,\tau} = \ldots$ are determined in post processing to estimate the Allan variance for the time $\tau = 3\tau_0$ and accordingly for all other times τ.

To make even better use of the stored data, roughly n times more values of $\bar{y}_{i,\tau=n\tau_0}$ can be obtained if the data processing is done such that $\bar{y}_{1,\tau} = (\bar{y}_{1,\tau_0}+\bar{y}_{2,\tau_0}+\bar{y}_{3,\tau_0})/3$, $\bar{y}_{2,\tau} = (\bar{y}_{2,\tau_0}+\bar{y}_{3,\tau_0}+\bar{y}_{4,\tau_0})/3$, $\bar{y}_{3,\tau} = \ldots$.

Let us take a look at two examples: In case of a linear frequency drift, i.e. the normalized difference $y(t) = at$, where a is the slope of the drift. With $\bar{y}_1 = [at_0 + a(t_0+\tau)]/2$ and $\bar{y}_2 = [a(t_0+\tau)+a(t_0+2\tau)]/2$ it follows that

$$\sigma_y(\tau) = \left\langle a\tau/\sqrt{2}\right\rangle = \frac{a}{\sqrt{2}}\tau. \tag{18}$$

Or for a harmonic modulation: $y(t) = \frac{\Delta\nu_0}{\nu_0}\sin(2\pi f_m t)$:

$$\sigma_y(\tau) = \frac{\Delta\nu_0}{\nu_0}\frac{\sin^2(\pi f_m \tau)}{\pi f_m \tau}. \tag{19}$$

The Allan variance $\sigma_y^2(\tau)$ is a useful time-domain measure of the frequency instability of an oscillator. It allows one to select the ideal oscillator for a particular application [2]. The Allan variance is a function of the sample period τ. A low Allan variance is a characteristic of a clock with good stability over the measured period. The Allan deviation is widely used for plots (conventionally in log–log format) and presentation of numbers. It gives the relative frequency stability, making it easy to compare with other sources of errors.
For example, an Allan deviation of $1.3 \cdot 10^{-9}$ at observation time 1 s (i.e. $\tau = 1$ s) should be interpreted as there being an instability in frequency between two observations 1 second apart with a relative root mean square (RMS) value of $1.3 \cdot 10^{-9}$. For a 10 MHz clock, this would be equivalent to a 13 mHz RMS movement [16].
The Allan deviation is also used to identify types of oscillator and measurement system noise, see figure 8. Usually the stability of an oscillator improves as the averaging period gets longer, since some noise types can be removed by averaging. At some point, however, more averaging no longer improves the results. This point is called the noise floor, or the point where the remaining noise consists of nonstationary processes such as aging or random walk [17].

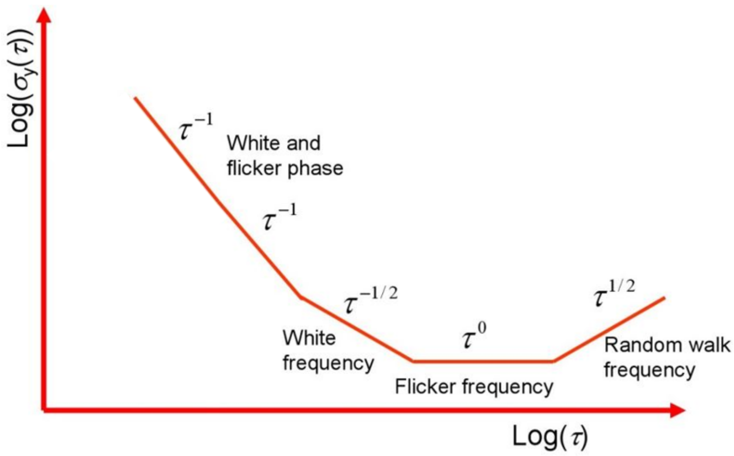

Figure 8: Different types of noise and their respective slope. The Allan deviation cannot distinguish between white phase noise and flicker phase noise. Figure taken from [17].

4 The pressure chamber

In this section we present the design of the pressure chamber[1], see figures 9 and 10. We tried to design it in such a way that we are as flexible as possible to keep doors open for unexpected changes and measurements down the road. At the beginning we thought that two clearance holes at the front and the back would be enough, but thankfully we settled for four holes which simplified connecting all the sensors and working with the vacuum pump.
We wanted to have a chamber which we can open from the top. That way we can make adjustments to the laser and everything else that is inside the chamber fairly quickly. We made the chamber large enough such that a DL pro can easily fit inside and there is still enough space for the sensors and all the cables. Both the chamber and the top are made out of aluminium.
We decided to make the walls 25 mm thick, thus allowing us to evacuate the chamber if desired. On each side it features a through hole with 41.2 mm diameter. Around each of these holes there are six M5 screw threads which in combination with bulkhead clamps and an O-ring allow us to connect other KF40 components, like a feedthrough, to it. Because the material of a screw is much harder than aluminium, a threaded hole drilled into aluminium may wear out fairly quickly. In order to provide a more durable threaded hole, we used threaded inserts, i.e. Helicoils. On the top side, there are eight M6 screw threads, where we also used Helicoils.
Because the chamber was milled out of one block it happens that the material deforms and the chamber wobbles when set on a flat surface. That is why, on the bottom side, we added four M4 screw threads such that one can attach pedestal pillar posts to the chamber. On the top side, there is a 4.1 mm deep and 6.5 mm wide groove for an O-ring with a 5 mm cord diameter. The O-ring we used was made out of nitrile rubber (NBR) with an inner diameter of 335 mm. On the inside bottom, we have six additional M6 screw threads. These can be used to fasten the laser to the chamber. For a few of the edges we added chamfers such that one does not hurt oneself that easily. As it later turned out, we should have also added chamfers to the clearance holes because the sharp edges damage the shielding of the cables when moving them around.
Objects as tall as 95 mm can fit inside our chamber. We chose this height because the DL pro with its protective cover is 90 mm tall. If one takes this cover off the height of the laser decreases by a few centimeters, i.e. it would be possible produce a smaller chamber. With its 460 mm, the chamber is also quite long. With the right cable management and sensor setup, it should also be possible to make the chamber a bit shorter.

[1]Special thanks to Walter Bachmann from the engineering office for the helpful discussions regarding the design of the pressure chamber.

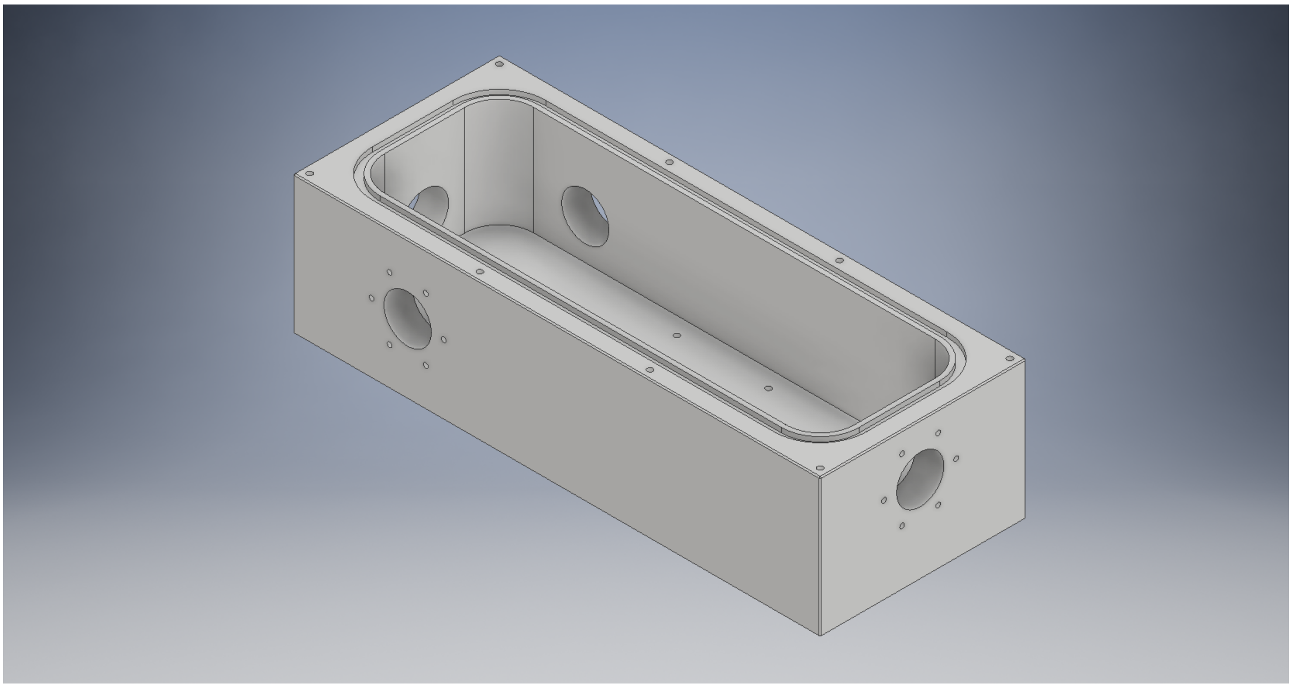

Figure 9: The pressure chamber is 460 mm long, 180 mm wide and 120 mm tall. The inner clearance is large enough to hold a fully assembled Toptica DL pro.

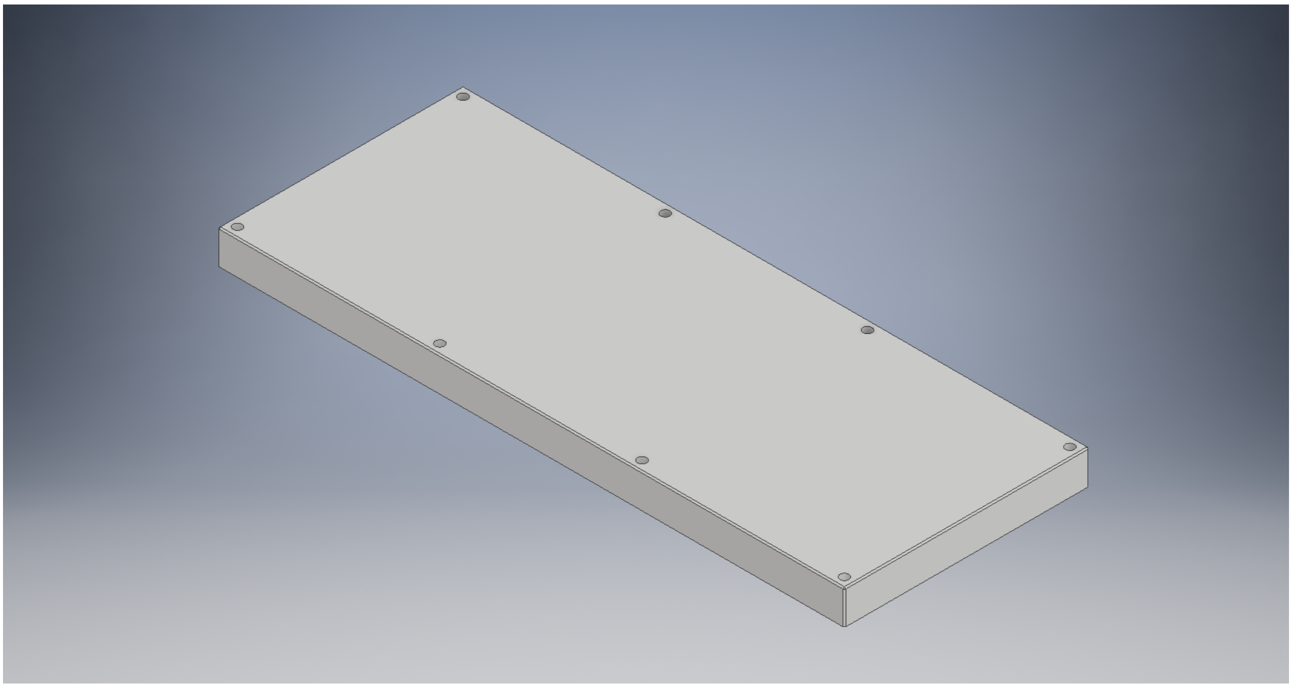

Figure 10: The top of the chamber is 460 mm long, 180 mm wide and 25 mm thick. It features through holes for M6 screws.

5 Setup

A schematic of the setup can seen in figure 11. The sensors are connected to a Raspberry Pi over I^2C and SPI. We designed two printed circuit boards onto which we soldered the sensors, thus guaranteeing a reliable operation of the sensors. The laser is operated with three cables: a D-Sub 9, a D-Sub 15 and a two-pole Lemo cable. For connecting the laser with the feedthroughs, we fabricated our own cables. Because we could not find an appropriate feedthrough for the Lemo cable, we decided to use a regular BNC feedthrough together with our self-made Lemo to BNC cables. Due to the fact that the feedthroughs invert the pins of the D-Sub cables, we soldered two D-Sub plugs back-to-back. This reverses the effect of the feedthroughs.

In order to measure the pressure inside and outside the chamber we bought two Adafruit BMP388 sensors. They have a relative and an absolute accuracy of $\pm 8\,\mathrm{Pa}$ and $\pm 50\,\mathrm{Pa}$, respectively. For determining the air temperature and the relative humidity we decided to use two Adafruit SHT40 sensors which have $\pm 1.8\,\%$ relative humidity accuracy and $\pm 0.2\,^{\circ}\mathrm{C}$ temperature accuracy. Additionally, we installed two resistance temperature detectors (Adafruit PT1000 RTD - MAX31865), where we attached one of the sensors to the housing of the laser and one to the chamber. These two sensors proved to be unreliable and in the end were not able to provide any useful additional information.

We used a HighFinesse WS8-10 wavelength meter for measuring the frequency of the lasers. It has an absolute accuracy of 10 MHz and a measurement resolution of 0.4 MHz. The locked laser which we used as a reference is located in another laboratory. Thanks to fiber splitters we were able to measure the wavelength of both lasers individually as well as combine their light and send it onto a photodiode. The difference frequency between the two lasers, the so-called beat frequency, is measured by a spectrum analyzer. The photodiode we use is an ET-2030A with a bandwidth of 30 kHz to 1.2 GHz.

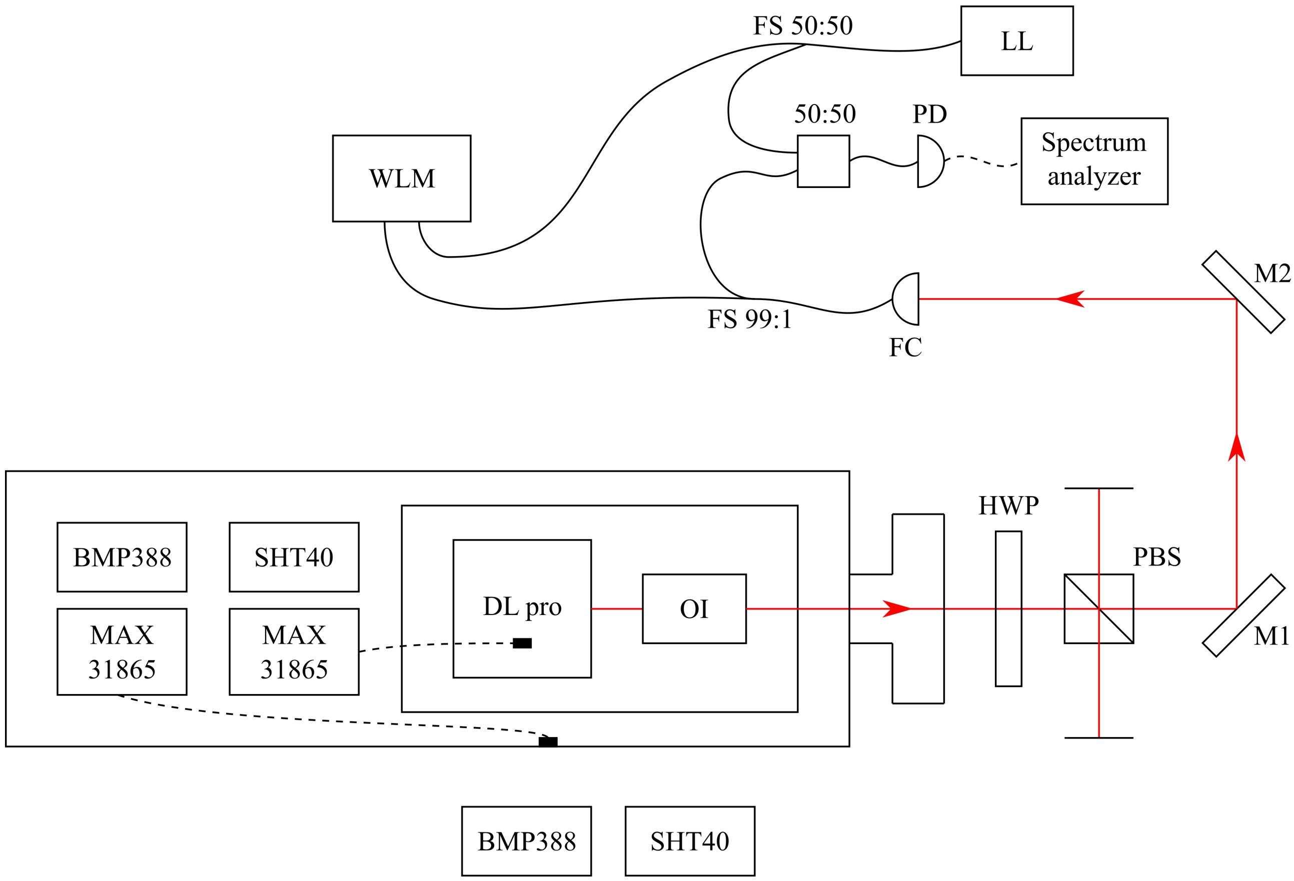

Figure 11: A DL pro operating at 854 nm is in the pressure chamber. Its output first goes through an optical isolator (OI) and leaves the chamber through a viewport. Then the light passes through a half-wave plate (HWP), followed by a polarizing beam splitter (PBS). By rotating the half-wave plate we can control the amount of power that is transmitted by the PBS. Two mirrors (M1, M2) help us couple the light into a fiber. FC is the fiber collimator. A 99:1 fiber splitter (FS) allows us to simultaneously measure the wavelength with a HighFinesse WS8-10 wavelength meter (WLM) and send the larger portion of the power to a 50:50 fiber coupler. The box labelled LL represents the locked laser (LL) located in another laboratory. We split its light and send half of the power to the WLM and the other half to the 50:50 fiber coupler. The combined light of both lasers goes on to an ET-2030A photodiode (PD). The produced electric signal is processed by a spectrum analyzer. Inside the chamber we have placed four sensors: a BMP388 (barometric pressure and temperature sensor), a SHT40 (temperature and humidity sensor) and two MAX31865 for precision temperature sensing. Outside the chamber, there is another BMP388 and SHT40.

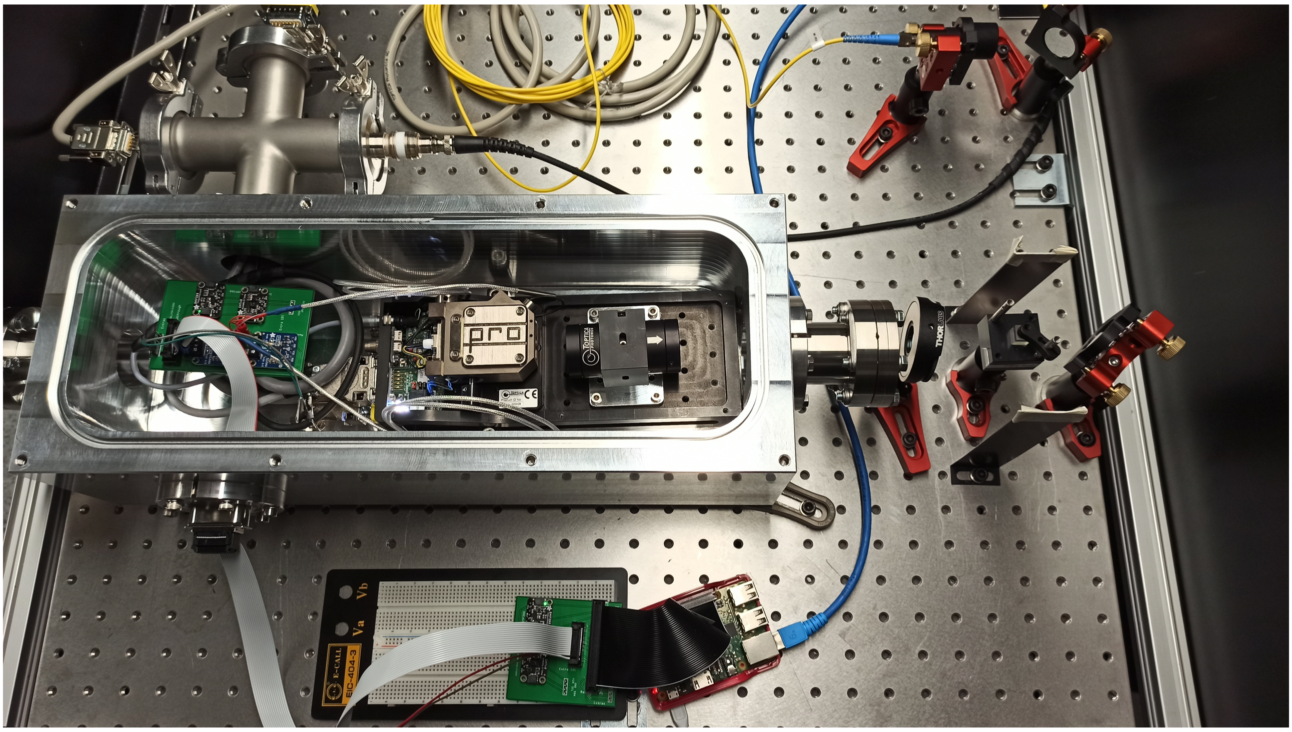

Figure 12: In this picture, part of the setup can be seen. The chamber is open. On the right, attached to the chamber, is a KF-CF adapter with a viewport. Because the viewport reflects some light, we set the laser at an angle with respect to the chamber in order to reduce reflections back into the laser. At one of the long sides of the chamber (above the chamber in the picture), there is a four way cross with three feedthroughs: a BNC, a D-Sub 9 and a D-Sub 15. The three cables connected to these feedthroughs are used for operating the laser. On the other long side, we have another D-Sub 15 feedthrough connected to flat ribbon cables on both sides. Each cable goes to a custom made printed circuit board. The circuit board outside the chamber is connected to a Raspberry Pi. On the left of the chamber, we attached a gate valve which we can connect to a vacuum pump.

6 Experimental data

6.1 Beat note measurement

One thing of interest was how the chamber influences the linewidth of the DL pro and how much the frequency fluctuates over short time scales. Thus we decided to perform a beat note measurement with the locked laser as a reference. We connected an oscilloscope to the photodiode and recorded the voltage over 100 ms with a sampling rate of 100 MS/s (10^7 data points). Then we Fourier transformed the data, where we get a Nyquist frequency of 50 MHz, and fitted a Lorentzian to the peak:

$$L(x, x_0, w_0, a) = \frac{a}{1 + \left(\frac{x - x_0}{2w_0}\right)^2}, \tag{20}$$

where x_0 is the center frequency, w_0 is the full width at half maximum (FWHM) and a is a proportionality constant.
What made this task challenging was the fact that the combined spectrum of both lasers did not simply have one peak. Actually even the light of the locked laser alone was a superposition of several frequencies, see figure 13. The frequency with an offset of +400 MHz is the most dominant one and is even visible on the interference patterns of the WLM, see figure 26. It is generated by a +200 MHz acousto-optic modulator which is passed twice. Therefore we had to be careful that the peak of interest did not overlap with other peaks. An example of a fit to the transformed data can be seen in figure 14.

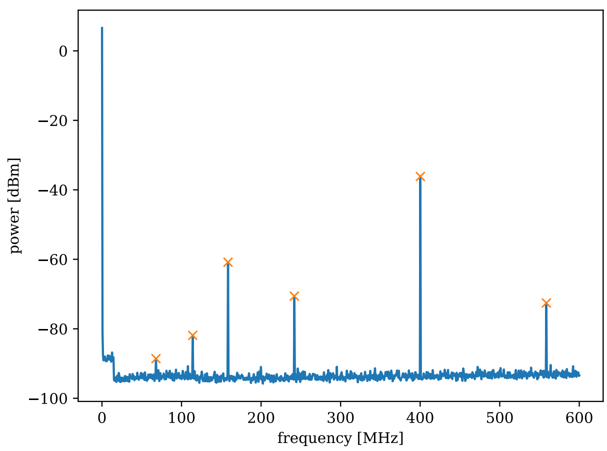

(a) Spectrum of the locked laser alone. The marked peaks have frequencies of 67.9, 114.1 , 158.5, 241.9, 400.2, 558.6 MHz.

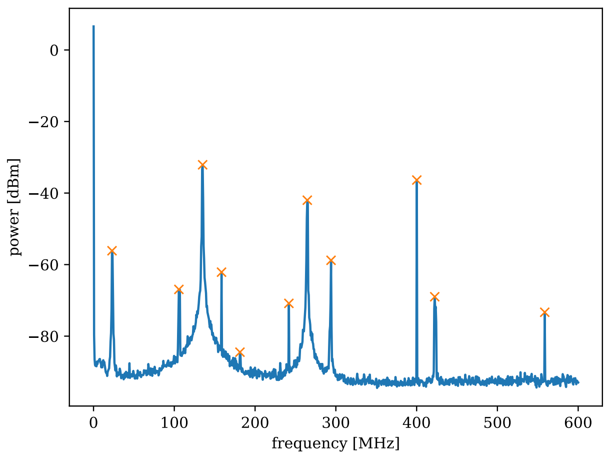

(b) Spectrum of both lasers combined. The marked peaks have frequencies of 22.9, 105.7, 135.1, 158.5, 181.3, 241.9, 264.7, 294.1 , 400.2, 422.4, 558.6 MHz. The broad peak corresponds to the beat note frequency, which is 135.1 MHz is this case.

Figure 13: Spectra measured with a spectrum analyzer.

In order to get the linewidth for different time scales, we divided the sample of 10^7 data points into a number, ranging from 1 to 1000, of smaller samples. We Fourier transformed each smaller sample individually and then we fitted the data. We repeated this process twenty times, half when the chamber was open and the other half when it was closed, and averaged over all measurements. The result can be found in figure 15. It agrees with what we expected. For short time scales there is no difference in linewidth between the two cases and as we go to longer time scales the difference becomes more prominent. We measure a linewidth of 0.78 MHz and 0.63 MHz at a timescale of 0.1 s for when the chamber is open and when it is closed respectively. We want to note that the measured linewidth is the "sum" of both lasers. We can assume that the linewidth of the locked laser is significantly smaller

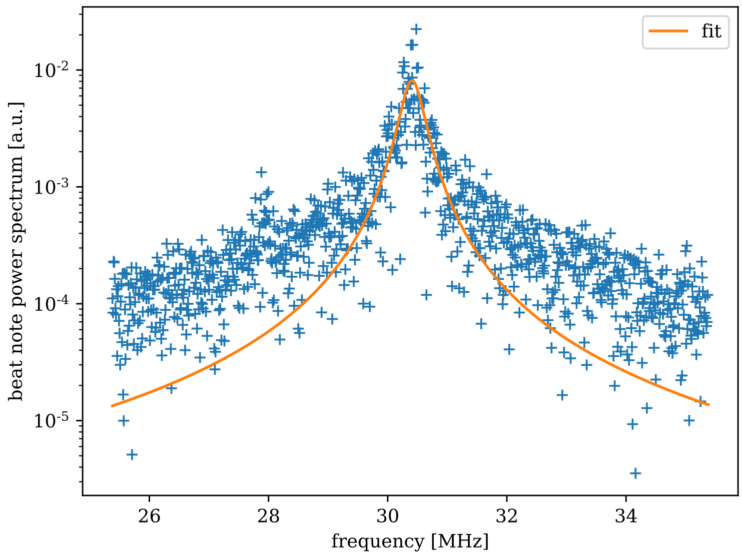

Figure 14: Fourier transform of the measured voltage generated by the photodiode (time scale: 0.1 ms), chamber closed. Lorentzian fit: $x_0 = 30.41$ MHz, $w_0 = 0.41$ MHz

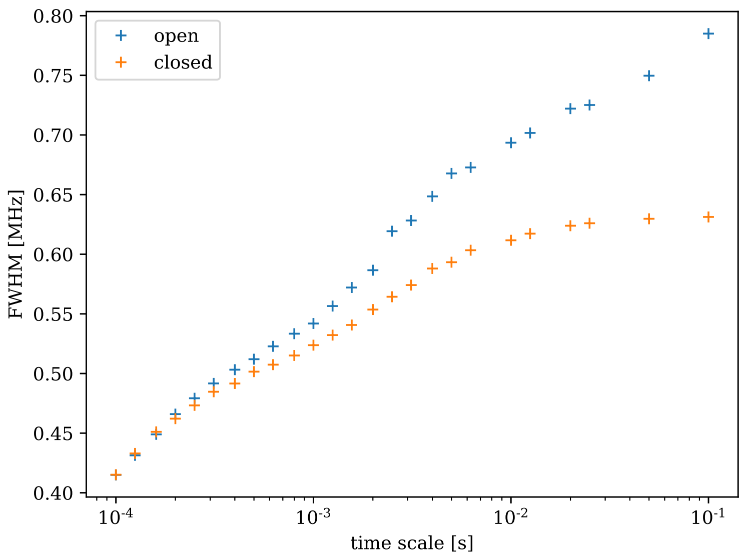

Figure 15: Linewidths w_0 over different time scales, obtained with a beat note measurement.

than the one of the free-running laser. Therefore the measured linewidths should be dominated by the contribution of the free-running laser. Toptica specifies the linewidth of the DL pro to be in the range of 10 kHz and 300 kHz over a time of 5 μs [12].
In order to examine the frequency stability and calculate the Allan deviation for these short time scales, we divide the 100 ms measurements into 10^4 samples. This allows us to obtain a value for the center frequency x_0 in 10 μs intervals and we can calculate the Allan deviation for these very short time ranges, see figure 16. As before the plotted values are averages. One can see the values of the Allan deviations for τ ranging over three orders of magnitude. For that, the values are quite stable. They range from $1.6 \cdot 10^{-10} - 2.9 \cdot 10^{-10}$ for the closed chamber and from $2.18 \cdot 10^{-10} - 3.45 \cdot 10^{-10}$ for the open chamber. However, there are some conspicuous fluctuations which look like some form of sinusoidal noise, see equation (19). However, pure sinusoidal noise would lead to an Allan deviation that decreases with $1/\tau$, which we cannot recognize in our case. The Allan deviations do not really decrease. Therefore, this kind of behaviour is most likely created from multiple different types of noise and noise sources. Similar fluctuations for small τ have been reported in [3].

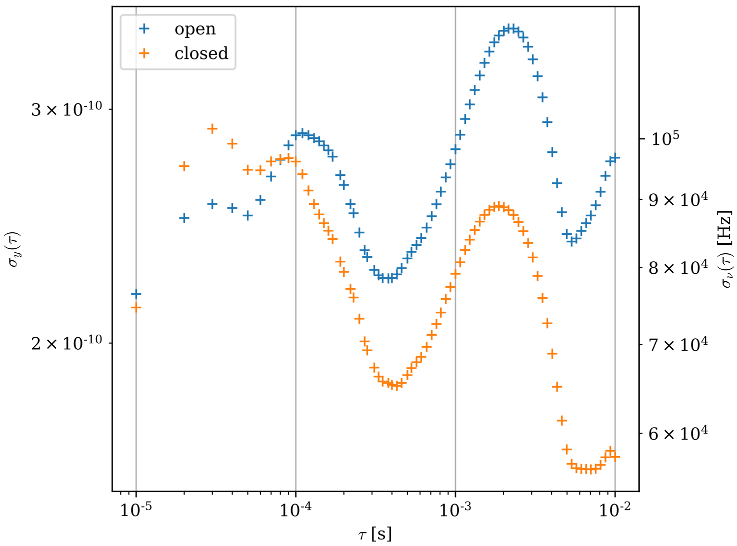

Figure 16: The Allan deviation for the beat note measurement.

6.2 Long term measurements

Measurements over more than one day were what we initially were most interested in. Thus we recorded the frequency of our laser with an approximate sampling rate of 1 Hz over the entire period of this project unless adjustments to the setup or technical difficulties prohibited it. Unfortunately this happened more often than we would have liked. In particular the air conditioning unit of the laboratory was often unstable and temperature fluctuated by several degree Celsius while the unit should in principle be stable up to 0.1 °C. In what follows we present the findings of two long term measurements, one during which the chamber was open and the other when the chamber was sealed.

6.2.1 Chamber open

From earlier measurements, which motivated us to build a pressure chamber in the first place, we knew what we could expect for the case of an open chamber, namely that the frequency fluctuations are dominated by the change in ambient pressure, see figure 17. If we assume that the frequency fluctuations are solely due to pressure changes, i.e.

$$\Delta\nu(\Delta p) = p_c \cdot \Delta p, \tag{21}$$

where p_c is a proportionality constant, we get $p_c = -42.43\,\text{MHz/hPa}$ and a coefficient of determination[2] of 0.936 for the fit, see figure 17a. Therefore the measured data is quite well replicated even by this simple model.

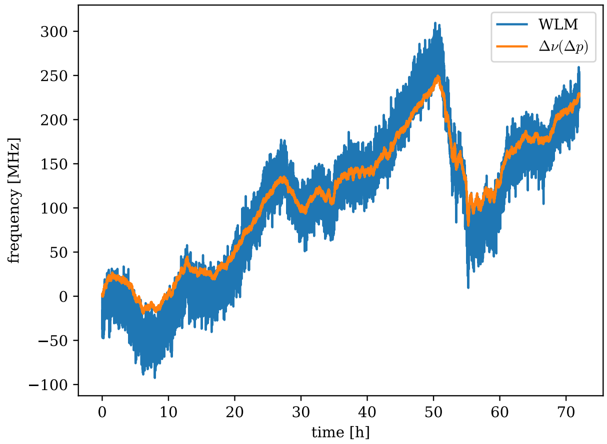

(a) Initial frequency: 350.862481 THz. A linear model of frequency as a function of pressure only, $\Delta\nu(\Delta p)$, is shown in orange. We get a proportionality constant p_c of $-42.43\,\text{MHz/hPa}$ and a coefficient of determination of 0.936.

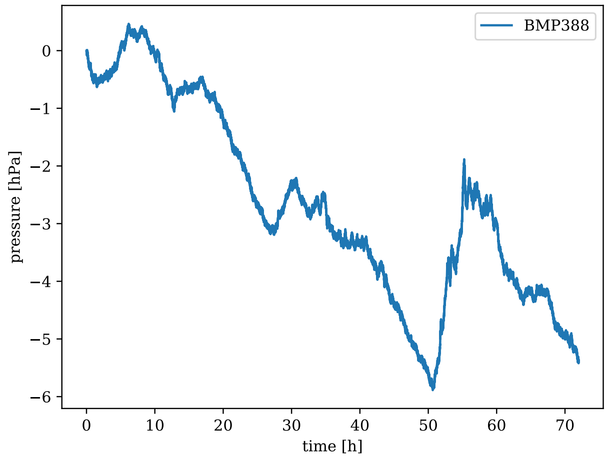

(b) Initial pressure: 959.58 hPa

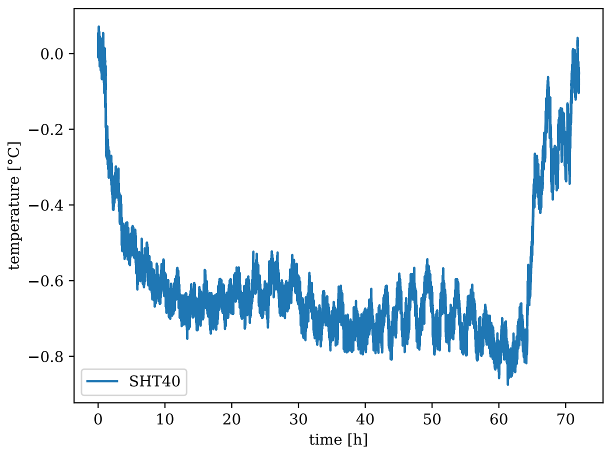

(c) Initial temperature: 21.28 °C

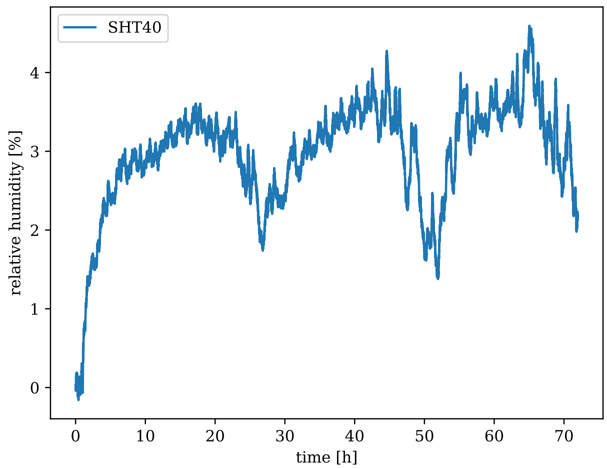

(d) Initial relative humidity: 55.12 %

Figure 17: Example of a measurement over 72 hours when the chamber is open. The initial measurement values are set to zero.

6.2.2 Chamber closed

When the chamber is closed, pressure cannot explain the frequency fluctuations nowhere near as well, see figure 18. Other parameters such as temperature start to have a noticeable impact. The frequency for a closed chamber is significantly less noisy compared to the previous case. It is also obvious that for the pressure and temperature sensor exactly the opposite is the case. One can see how much more stable the pressure is thanks to the chamber. While the pressure outside changes over a range of 9.14 hPa, inside it is only 0.81 hPa. Due to the fact that the volume in the chamber is constant, pressure and temperature are proportional. Leakage of air in the presence of a pressure difference between the

[2]See equation (38) for a definition.

inside and outside of the chamber can still lead to a drift in air mass. If the outside pressure changes such that there is a considerable difference compared to the inside pressure, then we can expect an increased leakage rate. Nevertheless, the more important factor is the temperature. If the inside temperature is not stable, then the pressure will not be stable either. We extend our model to the following:

$$\Delta\nu(\Delta p, \Delta T, t) = p_c \cdot \Delta p + T_c \cdot \Delta T + t_c \cdot t, \tag{22}$$

where the frequency fluctuations depend not only on the change in pressure, Δp, but also the change in temperature, ΔT. Furthermore we introduce a linear time dependency. t_c quantifies the magnitude of the constant drift. We obtain the following values for the coefficients: $p_c = -37.72\,\mathrm{MHz/hPa}$, $T_c = -67.96\,\mathrm{MHz/K}$ and t_c of $-0.83\,\mathrm{MHz/h}$. Toptica claims that the temperature sensitivity of the DL pro is $\ll 100\,\mathrm{MHz/K}$ [12]. The measured pressure and temperature are quite noisy and as a consequence so is $\Delta\nu(\Delta p, \Delta T, t)$. Nevertheless, the coefficient of determination is with a value of 0.984 quite high. If we also incorporate the relative humidity into our model, the coefficient of determination rises to 0.986 and the proportionality constants change accordingly. For example p_c is then -32.78 MHz/hPa. Because the increase in quality of the model is negligible, we refrained from including the relative humidity into our considerations. In addition, the accuracy of the SHT40 sensor that measures it is only $\pm 1.8\,\%$. Yet we note a remarkable correlation between the frequency and the relative humidity. For $\Delta\nu(\Delta h_{\mathrm{rel}}) = h_c \cdot \Delta h_{\mathrm{rel}}$, where Δh_{rel} is the change in relative humidity, we calculate a coefficient of determination of 0.925.
Since the building and the air conditioning unit did not provide the temperature stability that we would have liked to have, we decided to calculate how the frequency would have behaved if it did. First we analyze how the pressure behaves as a function of temperature and time:

$$\Delta p(\Delta T, t) = T_p \cdot \Delta T + t_p \cdot t. \tag{23}$$

For the example at hand, we get $T_p = 3.13\,\mathrm{hPa/K}$ and $t_p = -0.322\,\mathrm{Pa/h}$. We can use this to subtract the change in frequency, which is caused by temperature, from the measured frequency fluctuations $\Delta\nu$. We obtain

$$\begin{aligned} \Delta\nu' &= \Delta\nu - (p_c \cdot \Delta p(\Delta T) + T_c \cdot \Delta T) \qquad (24)\\ &= \Delta\nu - (p_c \cdot T_p \cdot \Delta T + T_c \cdot \Delta T). \qquad (25) \end{aligned}$$

The result $\Delta\nu'$ is plotted in figure 18f. We have effectively calculated the theoretical frequency fluctuations if we managed to perfectly stabilize the temperature. What remains is a near constant drift. A linear fit reveals a drift of -0.71 MHz/h. The exact cause of this drift is unknown to us. Although we were monitoring as many metrics as possible, we did not measure any quantity that correlates with said drift.
We calculated the Allan deviations for these two examples of long term measurements, see figure 19. As is clear, we gain a lot of stability for the short time scales. For measurement times of about ten seconds the Allan deviation is two orders of magnitude smaller for when the chamber is closed. But also for other observation times we gain at least almost one order of magnitude. The temperature corrected frequency variations (violet) show that on longer time scales better temperature stability would allow us to decrease the Allan deviation up to a factor of three. On short time scales, the deviation is actually much worse than the measured values. This is because of the fact that the temperature corrected frequency variations are quite noisy due to the noisy temperature measurement. If one compares figures 18a and 18f, one can see that the measured frequency in the former figure has significantly smaller short time fluctuations. We also added the Allan deviation of the linear fit to figure 19. As we have learnt in section 3, the Allan deviation is then proportional to τ. This illustrates how the Allan deviation would behave if we managed to filter out all types of noise.
To our knowledge, pressure only influences the frequency of the laser by changing the refractive of the air through which the light travels within the external cavity. The temperature, on the other hand,

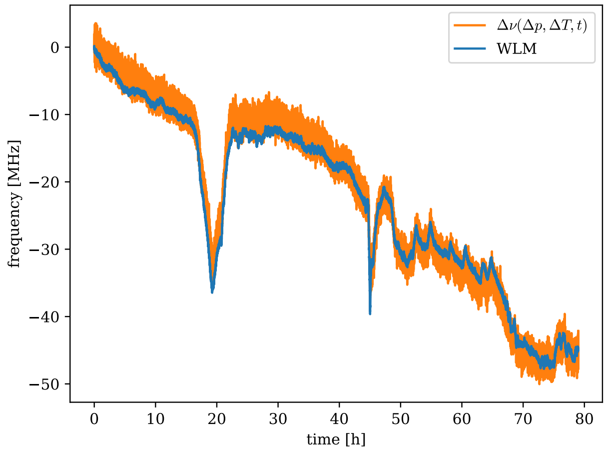

(a) Initial frequency: 350.861881 THz. In orange one can see a linear model with proportionality constants $p_c = -37.72\,\mathrm{MHz/hPa}$, $T_c = -67.96\,\mathrm{MHz/K}$ and a constant drift t_c of $-0.83\,\mathrm{MHz/h}$. The coefficient of determination is 0.984.

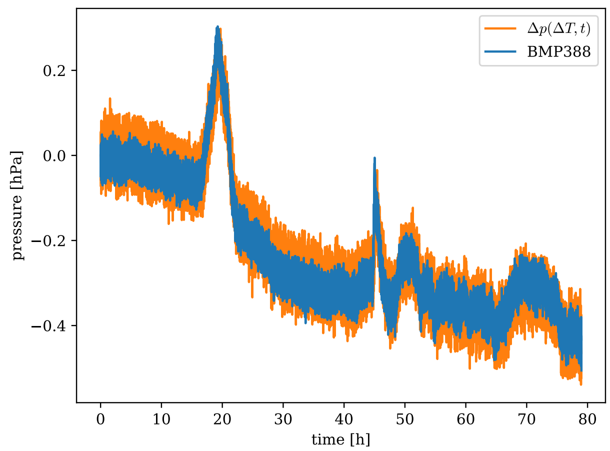

(b) Initial pressure: 960.43 hPa. A linear model that uses proportionality constants $T_p = 3.13\,\mathrm{hPa/K}$ and $t_p = -0.322\,\mathrm{Pa/h}$ is shown in orange. The coefficient of determination is 0.945.

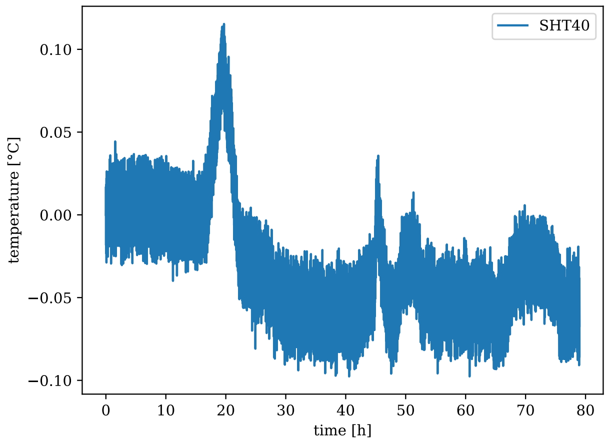

(c) Initial temperature: 23.62 °C

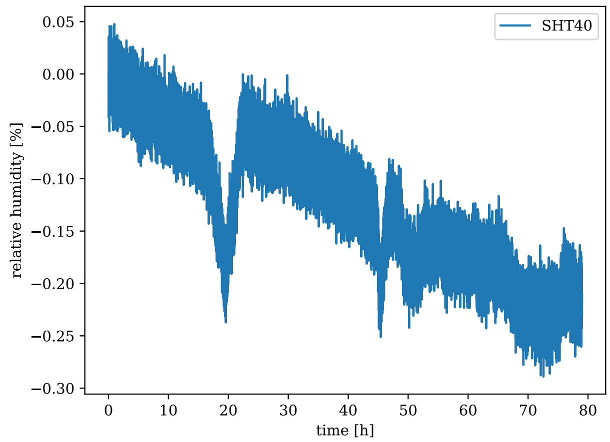

(d) Initial relative humidity: 49.57 %

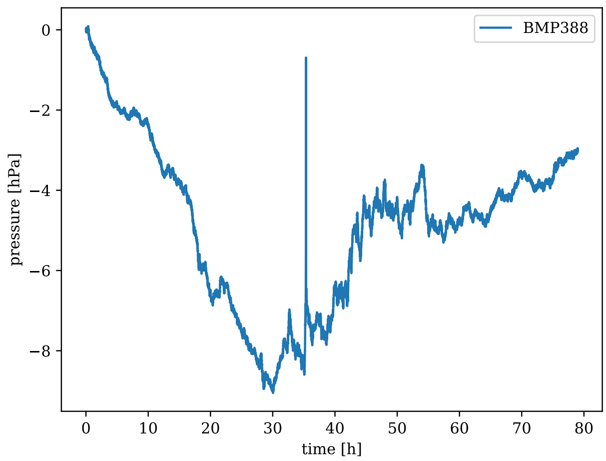

(e) Pressure measured outside of the chamber. The large increase in pressure after 35 hours is no mistake, the pressure sensors of the laser controller (DLC) and the wavelength meter measured a similar spike. Initial pressure: 958.34 hPa

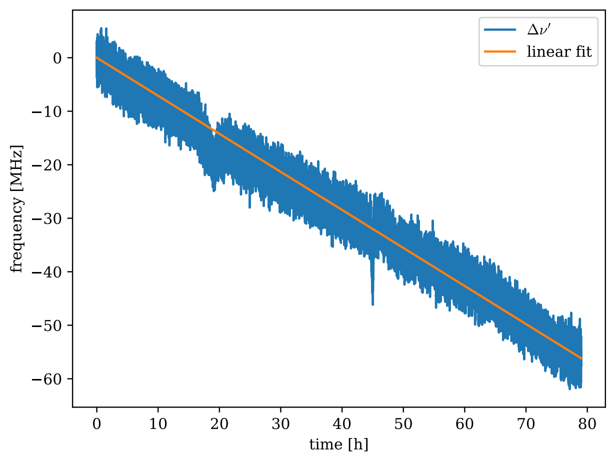

(f) In blue one can see the frequency fluctuations deduced by subracting the temperature effect. A linear fit reveals a drift of -0.71 MHz/h.

Figure 18: Example of a measurement over 79 hours when the chamber is closed. The chamber is set on a water-cooled breadboard, which heats up to 24 °C. The initial measurement values are set to zero.

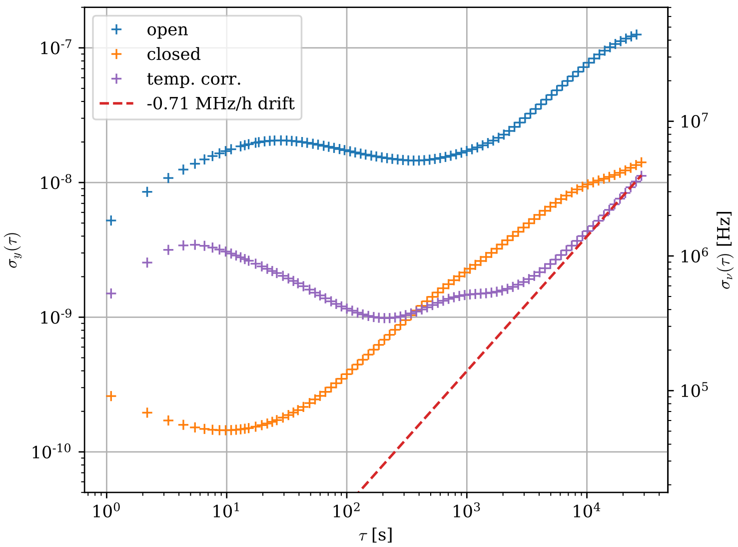

Figure 19: The Allan deviation for the two examples when the chamber is open and when it is closed. In order to simulate perfect temperature stability in the laboratory we corrected the closed chamber curve for temperature fluctuations to get the violet curve, see figure 18f.

influences quite a few components of an ECDL, for example the optical path length of any involved lens or the line spacing of the grating and of course also the refractive index of air.

In [3], the authors have studied the frequency fluctuations caused by different effects on the components of their ECDL in detail. Unfortunately, we were not able to perform a similar analysis for the DL pro because Toptica does not release any specific information about their laser's design. For example it would have been interesting to know the exact optical path length of the laser diode, L_D, or the external cavity, L_{EC}, see figure 6. With L_{EC} and L_{air}, we could have calculated a theoretical air temperature and pressure dependency of the frequency due to the change of the refractive index of air, n_{air}:

$$\frac{\mathrm{d}\nu}{\mathrm{d}p} = -\nu \frac{L_{\mathrm{air}}}{L_{\mathrm{EC}}} \frac{\mathrm{d}n_{\mathrm{air}}}{\mathrm{d}p} \tag{26}$$

and

$$\frac{\mathrm{d}\nu}{\mathrm{d}T_{\mathrm{air}}} = -\nu \frac{L_{\mathrm{air}}}{L_{\mathrm{EC}}} \frac{\mathrm{d}n_{\mathrm{air}}}{\mathrm{d}T_{\mathrm{air}}}. \tag{27}$$

In ideal conditions with a pressure-proof enclosure, both of these effects are eliminated, because $(n_{\mathrm{air}} - 1)$ is proportional to the density of air [3]. As we have already seen, in reality there is most likely leakage to some degree and thus said effects are not completely eradicated.

The only thing we can do is estimate L_{EC}. If we assume that during a mode hop the laser jumps from one resonator mode of the external cavity to a neighbouring mode, i.e. the jump in frequency is exactly the free spectral range of the external cavity, then L_{EC} must be $c/2\nu_F = 4.4 - 4.8\,\mathrm{cm}$ for $\nu_F = 3.1 - 3.4\,\mathrm{GHz}$ (the magnitude of the mode hops we measured lied between 3.1 and 3.4 GHz, see section 6.4).

6.2.3 Additional measurements

In figure 20, we summarize the Allan deviations of several long term measurements. For small τ, the spread of the deviations is relatively small for each of the three examined cases, especially for a closed chamber. We can see that for five out of seven closed chamber measurements we observe a linear drift for τ being larger than approximately $3 \cdot 10^2\,$s. There are two exceptions marked with letters A and B in figure 20. In the first case (A), the smaller frequency stability is due to insufficient temperature stabilization. The air conditioning system in our laboratory caused temperature fluctuations of up to 0.8 K within time ranges of 25 minutes. The temperature inside the chamber did of course not vary by that much but it was still enough that the frequency of the emitted light followed these fluctuations which is reflected in an increased Allan deviation. For τ larger than $2 \cdot 10^3\,$s, these temperature fluctuations average out and the deviation gets inline with the other measurements. In the second case (B), we simply did not measure any clear long term frequency drift in upward or downward direction. The frequency fluctuations we measured do not correlate with either temperature or pressure inside the chamber. Special about this measurement is that while the outside temperature was between 20.56 and 20.66 °C, which is quite stable by our standards, the air inside the chamber was heated to more than 24 °C by means of a resistor that we placed inside the chamber. The inside temperature remained within a range of 0.04 °C. This was one of the approaches we tried over the course of this project which we hoped would allow us to increase temperature stability. The idea was that if something inside the chamber gives off heat continuously, the air temperature inside would be less prone to temperature fluctuations outside the chamber. The other approaches were wrapping the chamber in styrofoam or setting the chamber on a water-cooled breadboard connected to a chiller which allowed us to set the temperature of the breadboard to a range of different values. None of these methods have clearly increased temperature stability. Until the end of the project, we were predominantly exposed to the arbitrariness of the air conditioning system.

Of course, we wondered what is the cause of the measured linear drifts. As can be seen in figure 2, piezo drifts can lead to long term frequency variations. What is meant by this is a phenomenon called creep which is the drift of the displacement of a piezoelectric translator when a constant electric field is applied [18]. In our case, creeping changes the orientation of the grating of the ECDL and the length of the external cavity. The displacement of the piezoelectric translator follows a logarithmic shape over time. We have observed this logarithmic shape ourselves, see figure 21. At the start of the frequency measurement we set the piezo voltage to a new value and then left it unchanged for the remainder of the measurement. Over the first ten hours the change in frequency does indeed follow said shape, however at later times the drift becomes almost linear again. This is because a frequency drift due to a linear drift in temperature drowns out the piezo drift. We suspect that, if the piezo voltage is kept constant for a long time, the piezo drift which should give a logarithmic frequency change, over a short window of time, cannot be distinguished from a linear drift. Therefore we would say that piezo drifts at least contribute to the linear long term drifts we measured.

Another contributing effect could be diode aging. It refers to the phenomenon that one needs to decrease the diode laser current over time in order to keep the laser frequency constant [19]. Therefore if you operate the diode laser at a constant current, the frequency will increase over time. In the cited paper, they report diode aging (they call it frequency aging) of 0.14 MHz/hour for a distributed-feedback diode laser. Others have reported drifts up to 2.78 MHz/hour due to aging. Unfortunately we have no idea how large diode aging is for a DL pro. It is noteworthy that diode aging is temperature dependent. At higher diode temperatures the aging process is faster.

On a side note, the Allan deviations for the locked laser are somewhat idealized because we neglect that the laser jumps out of lock from time to time. Sometimes already after less than half a day. In that case, the deviations are much higher.

Thanks to the chamber, the frequency stability is improved the most around time scales of about 10 s. The chamber reduces not only air temperature and pressure fluctuations but also the temperature fluctuations of the diode itself. The housing of the whole ECDL is placed on a thermoelectric cooler which stabilizes the temperature of the diode and all other components of the external cavity. The laser controller allows us to set the target temperature of the cooler and read out the temperature of the diode. In figure 22, the target temperature was set to 20 °C. The chamber was initially closed. We opened it after about 3.75 hours and consequently the temperature fluctuations measured at

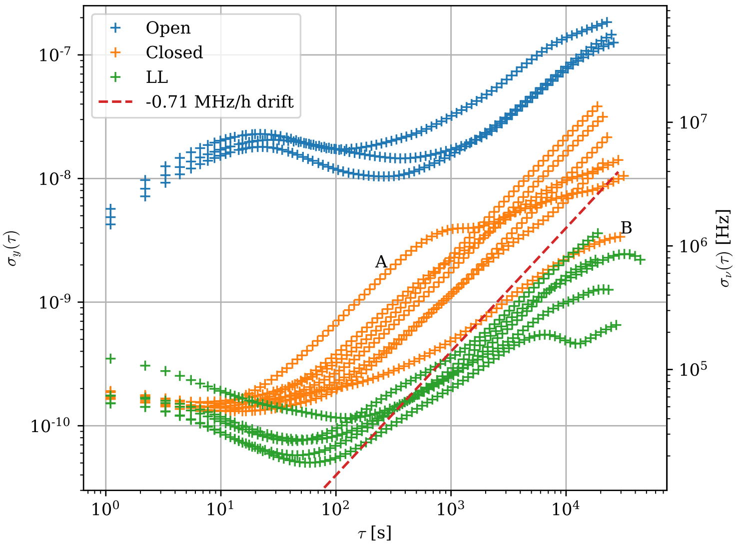

Figure 20: The Allan deviation for several long term measurements. We differentiate between three cases: open chamber, closed chamber and the locked laser (LL). Again we added the constant drift of figure 18f.

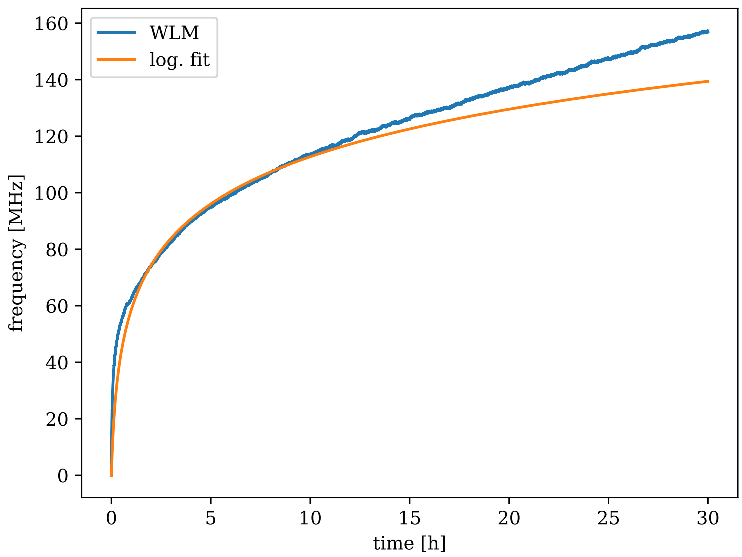

Figure 21: The piezo voltage was altered up to the point where this laser frequency measurement starts and then left unchanged. We fitted a logarithm to the first 10 hours of the measured data. The drift in frequency for the second half of the measurement is close to being linear. The chamber was sealed during the measurement.

the diode increased significantly. If we consider 4000 temperature measurements, then we obtain a standard deviation of 0.09 mK for when the chamber is closed and a standard deviation of 0.87 mK for when the chamber is open. This clearly shows how the sealed chamber supports the thermoelectric cooler at its job which leads to increased stability. As can be seen in figure 2, temperature control leads to frequency variations on time scales ranging from approximately 40 ms to more than an hour.

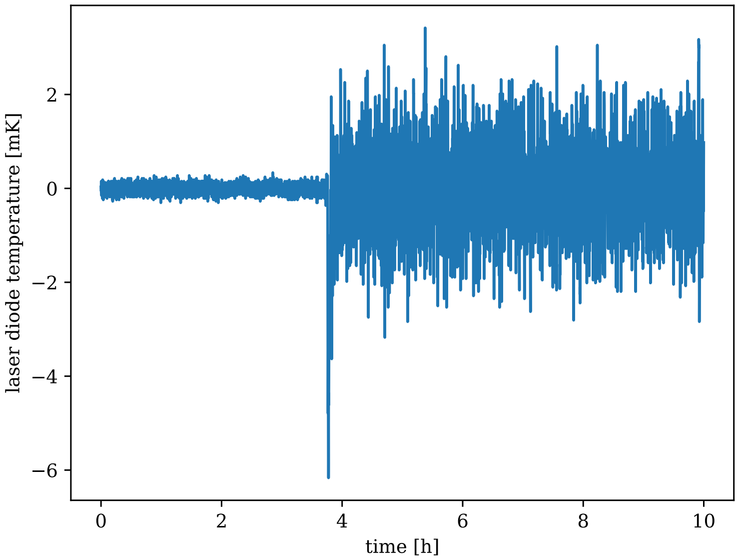

Figure 22: Laser diode temperature fluctuations. 0 mK corresponds to 20 °C. After about 3.75 hours, the chamber was opened and the magnitude of the fluctuations increased from 0.09 mK to 0.87 mK.

6.3 Opening door

In the following we shortly present two measurements where we opened door to the hallway as furiously as we could. In the first case the chamber was open and in the second it was closed. We recorded the frequency of the laser with the WLM and a sampling rate of 50 Hz, see figure 23. Over the time range of a little over 20 seconds, the frequency varies quite significantly such that sometimes it is not immediately discernible when the door was opened. That is why we added a plot of the difference quotient. With the help of the moving average over five consecutive data points we can quickly detect when the door was opened. When the chamber is open, the opening of the door leads to an increase of about 3 MHz over approximately 0.3 s. In the case of a closed chamber there is no indication of someone opening a door whatsoever.

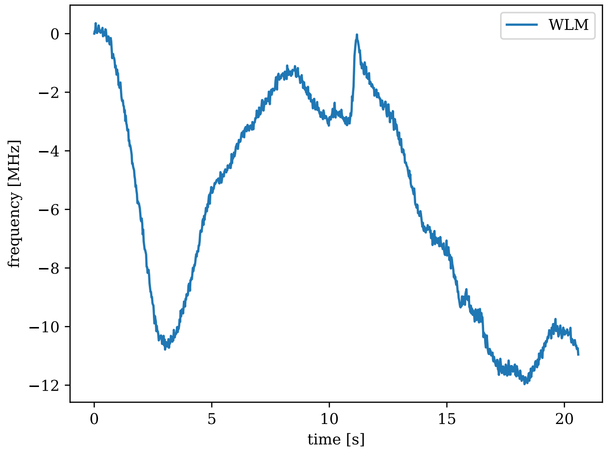

(a) Chamber open. Door was opened at $t \approx 11\,\mathrm{s}$.

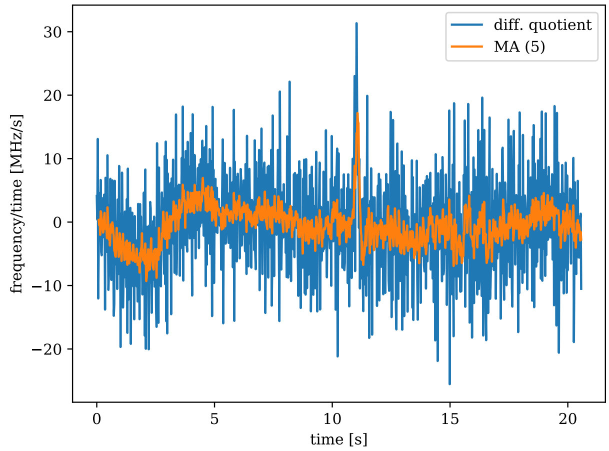

(b) The difference quotient of the measurement to the left. A moving average (MA) over five data points was added. The moving average clearly increases when the door was opened.

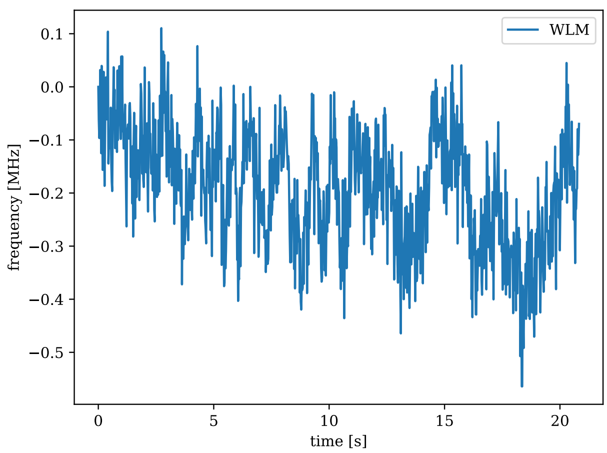

(c) Chamber closed. Door was also opened at $t \approx 11\,\mathrm{s}$.

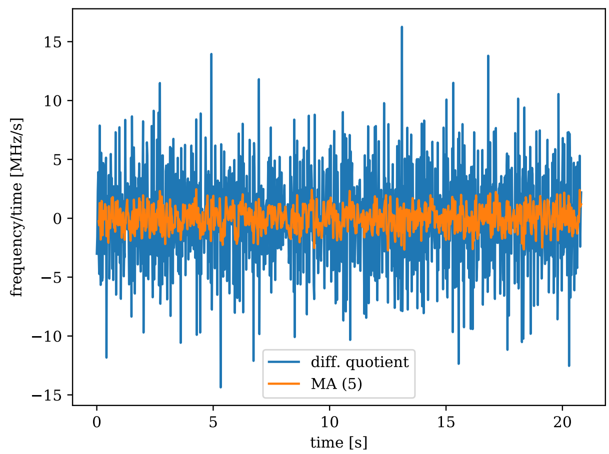

(d) The difference quotient of the measurement to the left. A moving average (MA) over five data points was added. The moving average shows no sign of a fast increase.

Figure 23: Comparison of a frequency measurement when the chamber is open and when it is closed. During each measurement the door to the hallway was opened once and then slowly closed.

6.4 Mode hop analysis

A phenomenon that can stop laser frequency stabilization from working altogether is a mode hop. According to Toptica, the mode-hop-free tuning-range of a DL pro lies between 20 and 50 GHz, which is in good agreement with what we measured. This is achieved by simultaneously varying grating angle, length of the external cavity and laser diode current [12]. For that a feature called feed forward has to be turned on, which means that as one changes the piezo voltage, the diode current is adjusted by means of a correction factor as well. If feed forward is deactivated, then the mode-hop-free tuning-range shrinks dramatically to about 1.5 and 3.5 GHz. When changing the piezo voltage, the resonator modes of the external cavity move in frequency space, see figure 7. At a certain point, another mode suddenly takes over all the optical power and the frequency of the output light changes rapidly, i.e. a mode hop occurs. The goal of this section is to examine when these mode hops occur and what role the pressure plays. We present two measurements. Both of these measurements took place when the chamber was closed and the frequency was determined with the WLM. Feed forward was turned off since we wanted the laser to hop.

For the first measurement, see figure 24a, we wanted to see by how much we can tune the frequency until a mode hop occurs. So the idea was that if the laser starts in mode number one, we increase the frequency by changing the piezo voltage until the laser hops to mode number two. Then we decrease it such that the laser hops back to mode number 1 and keep decreasing until we reach mode 0. Finally, we return back to the initial mode number 1. So there are four mode hops. The question then is if shortly thereafter we repeat this procedure, do the mode hops occur at the same frequencies? Since not much time passes between these two courses of action, the environmental influences should be nearly identical. Therefore, the expectation was that the mode hops occur at nearly the same frequencies.

During the second measurement, see figure 24b, we decreased the pressure inside the chamber by opening the gate valve such that the pressure inside adjusts to the pressure outside. The pressure dropped by 6.4 hPa, which lead to an increase in frequency of 386 MHz. At the same time, we observed that the relative humidity dropped by 0.25 %. Assuming that the change in frequency was only due to the change in pressure, we get a proportionality constant of -60.3 MHz/hPa. The pressure inside was higher because when you close the chamber and you fasten the eight screws at the top, then this reduces the volume while the exchange with outside is already limited. This leads to an increased pressure compared to outside.

Table 1 summarises all relevant values. For the case of no pressure change the first mode hop happened at 2525.7 MHz with a magnitude of 3301.5 MHz, while during the second run it happened at 2515.7 MHz with a magnitude of 3305.7 MHz. As we can see the differences in magnitude of the corresponding mode hops were as small as 1.3 MHz but never larger than 56.4 MHz. Similarly for at what frequency the mode hops occurred, we have differences between 1.2 and 28.8 MHz with one exception of 651.5 MHz for the third mode hop in the case of a pressure change.

This simple measurement shows that the frequencies at which mode hops occur stay approximately at the same values, even when the ambient pressure changes. This means that a change in pressure alone can push the laser closer to a mode hop or even make it hop if the change is large enough. Of course you want to stay away from mode hops as far as possible. With a pressure chamber such as ours we reduce the changes in pressure. If we assume that our laser drifts with 1 MHz/h and pressure is kept constant, then a mode hop will maybe occur after about 2500 hours (3.4 months), i.e. after an increase of 2.5 GHz, but if pressure changes lead to an additional increase in frequency of say 400 MHz, then the laser will already mode hop after 2100 hours. This analysis is of course an oversimplification but we think the main point is valid, namely that a laser inside the pressure chamber will mode hop less often.

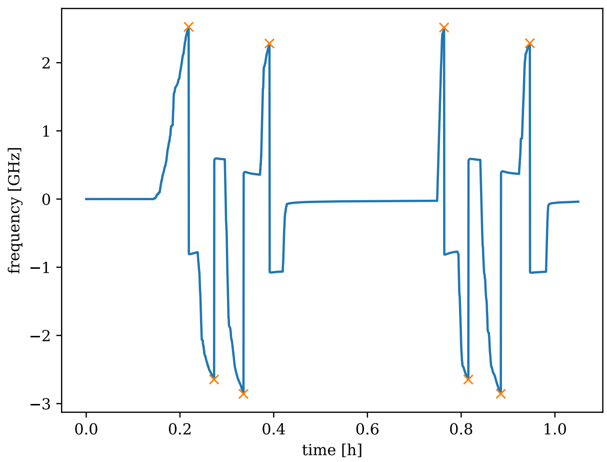

(a) No deliberate pressure change during the measurement.

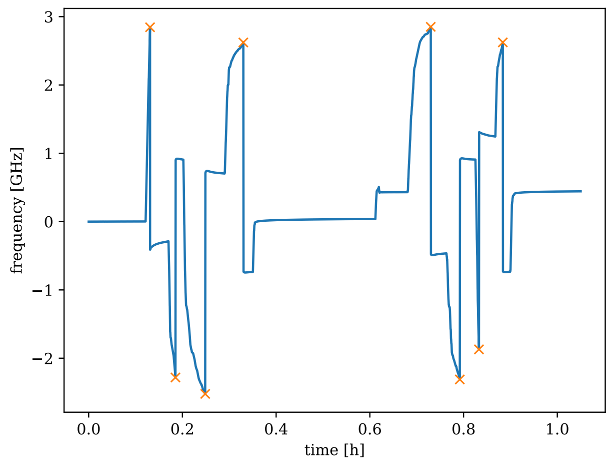

(b) After about 0.6 hours, the pressure inside the chamber was reduced by about 6.4 hPa, which lead to an increase in frequency of 386 MHz.

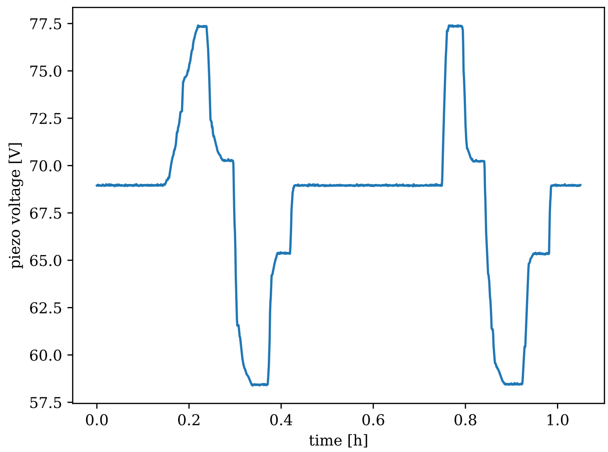

(c) The measured applied piezo voltage for the top left figure. For the top right figure it looks very similar.

Figure 24: The two upper plots show the frequency measured by the WLM while tuning the laser by changing the piezo voltage in order to force the laser to mode hop. The orange crosses mark when the mode hops occurred. The initial frequency was set to zero.

Magnitude of mode hop	1	2	3	4
No pressure change	3301.5	3182.1	3177.1	3338.4
	3305.0	3173.4	3199.8	3331.0
Pressure change	3256.0	3114.1	3167.1	3306.9
	3311.8	3170.5	3168.4	3333.9
Frequency of mode hop	1	2	3	4
No pressure change	2525.7	-2645.8	-2857.5	2281.7
	2515.7	-2647.1	-2855.9	2285.7
Pressure change	2845.0	-2281.4	-2521.6	2622.3
	2850.7	-2310.2	-1870.1	2623.5

Table 1: Summary of measured values which all are in units of MHz. The upper half indicates the magnitude of the mode hops while the lower half shows at what frequency relative to the initial frequency the mode hops occurred. Per run there were four mode hops. Two runs were executed in succession. While in the first case there was no pressure change between the two runs, in the second case there was one.

6.5 Summary & conclusion

Our measurements have shown that a pressure-proof housing is able to improve the frequency stability of a DL pro. In figure 25, we put the Allan deviations from the beat note measurement and the long term measurements together. We averaged the Allan deviations of the long term measurements. In order to fill in the gap between these two types of measurements, we used the WLM to record the frequency of the lasers with a sampling rate of 50 Hz over a time range of 10^3 s. Again we averaged the deviations over several measurements.
As is obvious from said figure, the frequency stability is increased for almost all examined observation times when using a closed pressure chamber compared to the open chamber situation. The largest improvement was identified for observation times around 20 seconds. The laser in the sealed chamber is even able keep up with the locked laser for τ being smaller than 10 s. The long term stability of the DL pro is predominantly dependent on the ambient pressure fluctuations and therefore at the mercy of the weather. If one is lucky and pressure is relatively stable, then the stability can be quite good. Nevertheless, the chamber essentially allows us to eradicate this kind of randomness. The chamber has allowed us to create a more controlled environment where good temperature stability becomes crucial.
Even though we logged as much data as possible to get as much information as possible, it is not always clear why the frequency behaved as it did. As of now we cannot satisfyingly explain all these near linear long term drifts and it seems that it is not possible to get rid of them with only passive stabilization schemes. We believe that they are the sum of different effects such as piezo drift and diode aging. Possibly also long term temperature and pressure drifts that we cannot detect. Because the piezo drift decreases logarithmically, it is recommendable to change the piezo voltage as rarely as possible. Unfortunately we often had to change the piezo voltage in order to get the laser to the desired frequency, for the beat note measurements for example.
In general it is challenging to isolate the different influences on the laser frequency. What would make it easier would be less noisy temperature and pressure sensors with better resolution. Often when the temperature inside the chamber was quite stable, we could recognize a small temperature drift over several hours, but if one zoomed in to about one hour, the temperature fluctuations just look like white noise and no drift is discernible.
Since the refractive index of air also depends on the relative humidity and the carbon dioxide content, it might be worth to consider the influence of these two quantities for future projects. Furthermore, it would be interesting to create an active stabilization technique where instead of adjusting the piezo

voltage in order to compensate frequency variations one would adjust the pressure or the temperature inside the chamber. While it would be hard to compensate fluctuations on short time scales, it should be possible to reduce the long term drifts.

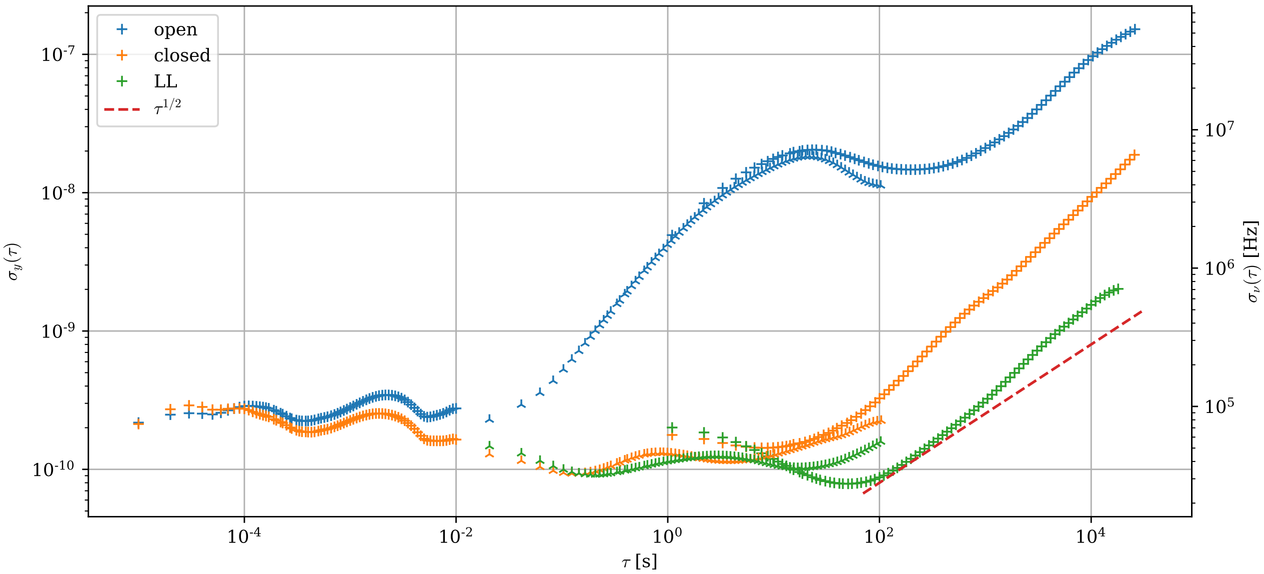

Figure 25: All Allan deviations averages for the whole time range. Data points at timescales shorter than 10^{-2} s were measured with the beat note method, while points at longer timescales were recorded with the WLM. In the case where we measured the frequency with the WLM with 50 Hz we used a different type of marker than the plus. In red one can see the slope of random walk of frequency noise.

7 Single- and multi-mode

In our laboratories, there are quite a few HighFinesse wavelength meters (WLM). They are designed for wavelength measurements of continuous wave and pulsed lasers. They comprise an optical unit which is designed to create interference patterns, obtained by a certain number of Fizeau interferometers. For the WS8-10, which we use, it is six.
The signal from two photodiode arrays is transmitted via the USB communication cable to a computer. A computer program then graphically displays the interference patterns and calculates the wavelength based on calibration performed with a reference light source [23].
Because most use cases of ECDLs rely on single-mode operation, we basically always want the lasers to be in single-mode operation. When talking about lasers, single-mode operation can mean one of two things. The first is single-transverse-mode operation, where a laser operates on a single kind of transverse resonator mode. Laser oscillation may then still occur on multiple axial (longitudinal) modes, which have essentially the same transverse shape but differing optical frequencies, separated by the free spectral range. In the second case, the term "single-mode" really indicates operation on a single resonator mode, which usually is an axial mode. This is more precisely called single-longitudinal-mode operation or single-frequency operation. In that case, the laser is a single-frequency laser, and the laser linewidth is fairly small. Single-frequency operation is usually more difficult to achieve than just single-transverse-mode operation, because it is not sufficient to introduce spatially varying loss or gain. Factors which make it more challenging are all those reducing the mode competition, e.g. inhomogeneous saturation via spatial hole burning [24].
When we talk about single-mode operation in this thesis, we mean single-frequency operation. Everything else will be called multi-mode operation. The WLM measures two interference patterns, see figure 26. The upper interference pattern represents the recording of five Fizeau interferometers and the lower one of an additional long Fizeau interferometer. Only the lower one is going to be relevant to us, so henceforth when we say interference pattern we mean the lower one. By looking at the interference pattern one can easily see if the laser is in pronounced multi-mode operation or approximately in single-mode one. In single-mode the interference patterns consist of roughly equally spaced fringes, see lower pattern in figure 26. They look similar to a weighted sum of Lorentzians. As one can see, the fringes are not symmetric. On the right side of each peak there are bumps but these are just a consequence of the nature of a Fizeau wedge as we will later see.
In order to get rid of having to regularly check the interference patterns to see if a laser is still in single-mode operation, we decided to write a Python script that continuously reads in these patterns and evaluates the current state of the laser, i.e. quantifies the spectral purity. The idea then is that as soon as the purity of the laser light decreases below a predefined threshold, it triggers an alarm and one can improve the purity by adjusting certain laser operating parameters, either manually or, even better, by using another script.

7.1 Theory

As already mentioned, the optical unit of a HighFinesse WLM consists of Fizeau interferometers. So the first step towards our goal is to understand what a Fizeau interferometer is and how they work. How are these transmitted interference patterns formed and what information do they convey? Fortuntely, Elena Stoykova has worked intensively on this topic and in [20] she developed a procedure for calculating the transmitted interference pattern for a power-normalized Gaussian beam by using its plane-wave expansion. In [21], she and Marin Nenchev extended the work on Fizeau wedges examining the transmitted as well as the reflected intensity distributions for positive and negative incidence. In the following, we will summarize their results.
A Fizeau wedge or Fizeau interferometer consists of two reflecting plane surfaces separated by a gap with linearly increasing thickness, see figure 28. Depending on the reflectivity of these surfaces, the gap thickness, and the apex angle, the Fizeau wedge may find various applications in metrology, spectroscopy, and laser resonator design.

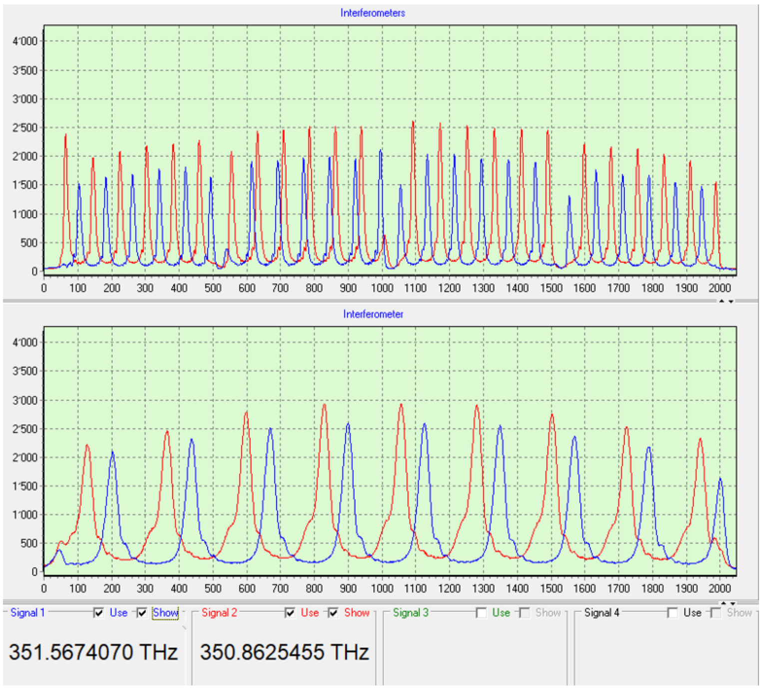

Figure 26: Screenshot of the WLM software interface. The upper interference pattern represents the recording of 5 Fizeau interferometers and the lower one of an additional long Fizeau interferometer which together result in high measurement precision [23]. Both lasers are in single-mode operation, as can be seen from the lower interference patterns. At the bottom their respective frequency in THz can be read. Laser 2 (red) is the locked laser. Its pattern has an additional small side lobe on the left of each peak because of an acousto-optic modulator that creates a +400 MHz sideband.

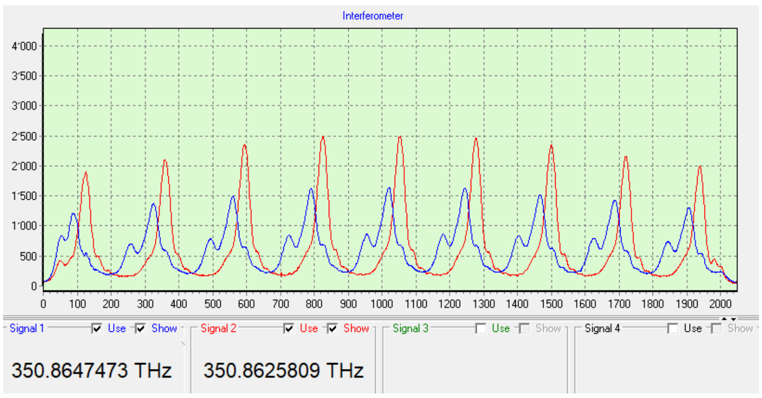

Figure 27: Screenshot of the WLM software interface. Laser 1 (blue) is in multi-mode operation, i.e. there is at least one additional mode lasing. Laser 2 (red) is in single-mode operation.

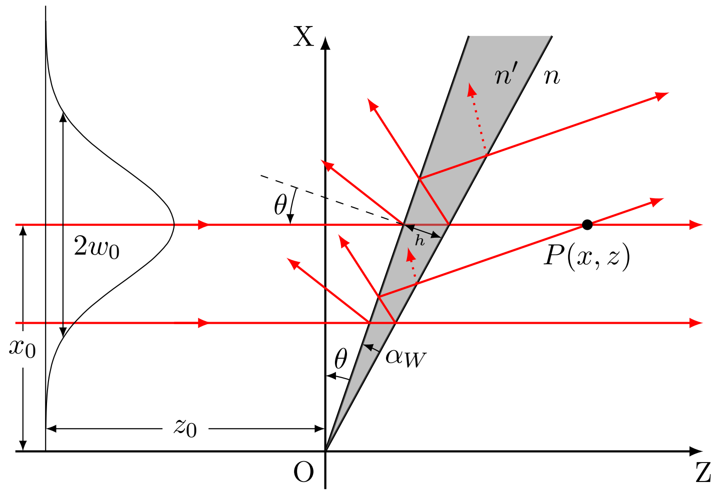

Figure 28: A Fizeau wedge with opening angle α_W whose apex is at the origin O of the coordinate system XOZ. The wedge and surroundings consist of materials with refractive index n' and n respectively. The optical axis of the incoming Gaussian beam is parallel to the Z-axis and crosses the X-axis at $x = x_0$, its minimum waist w_0 is attained at $z = z_0$. The angle between the optical axis and the normal to the front wedge surface is called θ. The wedge thickness h is measured along the normal to the front wedge surface at the point at which the center of the incident beam enters the wedge. In red, there are two parallel light beams that go through the wedge and meet in point $P(x = x_0, z)$. For a plane-wave one also has to consider light beams with a larger number of reflections. Once one knows the intensity at P for a plane-wave, one can calculate the intensity for a Gaussian beam by making use of a plane-wave expansion.

The Fizeau wedge transforms a plane wave falling on it into an infinite number of plane waves in both transmission and reflection. The propagation directions of the generated plane waves differ successively by twice the angle α_W at the wedge apex, whereas their amplitudes decrease successively by a factor corresponding to two additional reflections.
The two ways that the beam falls on the front wedge surface have been called negative and positive incidence. At negative incidence the incident ray undergoes multiple reflections within the wedge in the direction of its apex. At positive incidence the multiple reflections occur in the direction of increasing wedge thickness.
The aforementioned approach for calculating the interference pattern for a power-normalized Gaussian beam leads to an analytical solution in the case of an air-gap Fizeau wedge. The transmitted intensity distribution can be calculated from

$$I_T(x, z, \lambda) = a^2 T^2 \left(\left(\sum_{p=1}^{\infty} R^{p-1} \Omega_p \cos \varepsilon_p \right)^2 + \left(\sum_{p=1}^{\infty} R^{p-1} \Omega_p \sin \varepsilon_p \right)^2 \right), \tag{28}$$

where

$$\Omega_p = \frac{\sqrt{\pi}}{\sqrt[4]{b^2 + p_p^2}} \cdot \exp\left[-\frac{b q_p^2}{4(b^2 + p_p^2)} \right], \tag{29}$$

$$\varepsilon_p = \frac{1}{2} \arctan \frac{p_p}{b} - \frac{p_p(q_p^2 - 4 p_p r_p) - 4 b^2 r_p}{4(b^2 + p_p^2)}, \tag{30}$$

$$a = \sqrt[4]{2\pi} \sqrt{w_0}, \qquad b = \pi^2 w_0^2, \qquad p_p = -\pi(z - z_0 + x\nu'_p + z\nu_p)\lambda, \tag{31}$$

$$q_p = 2\pi(x + x\nu_p - z\nu'_p + x_0), \qquad r_p = \frac{2\pi}{\lambda}(x\nu'_p + z\nu_p), \tag{32}$$

$$\nu_p = \cos\xi_p - \cos\xi_1, \qquad \nu'_p = \sin\xi_p - \sin\xi_1, \qquad \xi_p = 2(p-1)\alpha_W. \tag{33}$$

λ	wavelength of the light
α_W	angle at the wedge apex
R	reflectivity of the coatings
T	transmittance of the coatings
w_0	minimum waist of Gaussian beam
x_0	position of the center of Gaussian beam along the X-axis
z_0	position along the Z-axis where the minimum beam waist w_0 is attained.

Table 2: Overview of the parameters for formula (28).

The assumptions that go into the calculation are that refractive indices for the wedge and the surroundings are one ($n = n' = 1$)[3] and that there is no phase shift for external reflections and a π-phase shift for all internal reflections. Furthermore, it is assumed that the field reflection coefficients r and r' of the front and the rear wedge surfaces are identical. Likewise for the field transmission coefficients t and t'. Also the paraxial approximation was used.
Essentially there are seven parameters, see table 2: the wavelength λ, α_W is the angle at the wedge apex, $R = rr' = r^2$ is the reflectivity of the coatings and $T = tt' = t^2$ is the transmittance. The Gaussian beam whose optical axis is parallel to the Z-axis crosses the X-axis at $x = x_0$, its minimum waist w_0 is attained at $z = z_0$.
It is remarkable that there is no dependence of the interference pattern $I_T(x, z, \lambda)$ on the angle, θ, between the incident beam axis and the normal to the front wedge surface in the coordinate system XOZ, see figure 28. In other words, rotation around the wedge apex does not change the transmitted intensity distribution.
We truncate the infinite sums in (28) at certain value $p = p_{\max}$, which depends on the reflectivity R. $p_{\max}$ should be appropriately large for high reflectivities.
One can show that the fringe spacing can be calculated with

$$\delta = \frac{\lambda}{2\tan\alpha_W}, \tag{34}$$

see [26]. Furthermore, simple geometry yields

$$\tan\alpha_W = \frac{h\cos\theta}{x_0}, \tag{35}$$

where h is the wedge thickness along the normal to the front wedge surface at the point at which the center of the incident beam enters the wedge. And we also make use of

$$h \approx \frac{c}{2\cdot\nu_F}, \tag{36}$$

where c is the speed of light and ν_F is the free spectral range (FSR) of the WLM.
Even though the light beam of a diode laser usually has an elliptical shape, once the light is coupled into a single-mode fiber, the output at the other end of the fiber has a Gaussian shape, i.e. the light entering the WLM has a Gaussian shape and the use of formula (28) is justified. Unfortunately, unlike the manufacturer of the WLM, we do not know the values of the WLM-dependent parameters, which are the angle α_W, the reflectivity R, the transmittance T, the angle θ, the exact position of the sensor (photodiode array) relative to the wedge apex and the length of the sensor ℓ. The only parameters that we then would have to estimate/fit are x_0, z_0 and w_0.
But with the help of the formulas above we can estimate a few parameters given an interference pattern

[3] The calculation for non-unity refractive indices has also been done, see [22].

and its corresponding wavelength. For example we can measure the spacing δ and then by making a reasonable guess about the length ℓ of the photodiode array, we can use formula (34) to calculate the angle α_W. For the remaining parameters we have to make guesses. Consequently, it took some time at the beginning until we managed to produce fringe profiles similar to the measured ones. When reading in the interference pattern from the WLM, one gets an array of N values where only their relative magnitude has any significance, which is why we normalize the pattern magnitude to one. Likewise, for the the calculated transmitted intensity pattern, i.e. we divide the values $I_T(x_i, z, \lambda)$ for $i \in M = \{0, 1, 2, \ldots, N-1\}$ by $\max_{i\in M}\{I_T(x_i, z, \lambda)\}$. This allows us to get rid of the parameter T. The idea is that, as long as the laser is in single-mode operation, we should be able to produce a fringe profile with the analytic formula that is very similar to the measured interference pattern. As additional modes start lasing the similarity between the patterns should decrease. We quantify the degree of similarity by calculating the coefficient of determination R^2. Consider a data set that consists of N measured data points $y_1, \ldots, y_N$, each associated with a fitted value $f_1, \ldots, f_N$. If $\bar{y}$ is the mean of the observed data:

$$\bar{y} = \frac{1}{N}\sum_{i=1}^{N} y_i, \tag{37}$$

then the coefficient of determination is

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}, \tag{38}$$

with

$$SS_{\text{res}} = \sum_i (y_i - f_i)^2, \tag{39}$$

$$SS_{\text{tot}} = \sum_i (y_i - \bar{y})^2, \tag{40}$$

where SS_{res} is the residual sum of squares. In the best case, the modelled values exactly match the observed values, which results in $SS_{\text{res}} = 0$ and $R^2 = 1$. A baseline model, which always predicts $\bar{y}$, will have $R^2 = 0$. Models that have worse predictions than this baseline will have a negative R^2 [25].

7.2 Algorithm

In the following we will be doing part of the calculations and steps of the algorithm in the "data point space". What we mean with this is the space where quantities are expressed in units of data points that the photodiode array of the WLM measures. For example if the spacing between two fringes is 10 data points, then in order to get the real distance between them in meters one has to multiply 10 with ℓ/N, where ℓ is the length of the sensor and N is the total number of data points the WLM can measure. For that reason there are no units for plots in the data point space. For the analytic formula we of course need values of all parameters in SI-units that are as close to reality as possible.

1. Guess the values of z, z_0, R and ℓ and put in the measured value of the FSR ν_F of the WLM.
2. Read in pattern and wavelength λ from the WLM, see figure 29a.
3. Estimate parameters $x_0, w_0, \alpha_W, x_{0,\text{G}}$, see figure 29b:
   - Find the peaks of the pattern and save their locations x_i in an array. Cut off one peak on each side of the pattern.
   - Fit a Gaussian $f(x)$ to the remaining peaks which gives us $x_{0,\text{G}}$ and w_0. $x_{0,\text{G}}$ is center of the Gaussian fit to the measured pattern. Scale them to the right size by multiplying them with ℓ/N.

$$f(x) = \exp\left(-2\frac{(x - x_{0,\text{G}})^2}{w_0^2}\right) \tag{41}$$

- Calculate the average spacing between the peaks and scale it to the right size by multiplying it with ℓ/N. Calculate α_W with (34) using the average spacing and the wavelength λ measured by the WLM.
- Calculate the wedge thickness h using equation (36) and then calculate x_0 with (35) where we assume that $\theta = 0$.

4. Make peaks equidistant, see figure 29c:
   - Use formula (34) to calculate the spacing between peaks using the estimate of α_W and the measured wavelength λ. Create an array with the values y_i where the peaks are supposed to be such that they are equidistant.
   - Cubically interpolate the function $g(x_i) = y_i$.
   - Apply this interpolation function to the array with the cut x-values and transform the new array by multiplying it with ℓ/N and adding $x_0 - x_{0,\mathrm{G}}$. In figures 29c and 29d, one can see the stretched data, where the peaks are equidistant, before and then after the array of x-values has been transformed.
5. Maximize R^2 by shifting the x-values in order to get the peaks of the calculated and stretched patterns to overlap as well as possible.
6. Optimize parameters z, z_0, w_0 and R.
7. Read in new pattern and wavelength. Repeat steps 4 and 5. Return R^2. Repeat this step.

7.3 Comments on algorithm

When one starts the algorithm, one has to ensure that the laser is in single-mode operation because the first pattern that we read in is used to estimate and optimize the parameters of formula (28).
We have to make sure that we use reasonable values for z, z_0, R and ℓ. This can be done by plotting the analytic pattern with different values until one receives a pattern that is similar to a measured single-mode pattern. This is necessary because otherwise the algorithm (optimization of parameters) may not work properly. We have guessed ℓ to be 2 cm. Since we are using a fixed value for z for formula (28), we are calculating the patterns in a plane parallel to the X-axis, i.e. we assume the photodiode array is parallel to the X-axis. The sensor of the WLM could of course be tilted with respect to the X-axis, however this should only lead to stretching of the pattern in X-direction.
The FSR ν_F of the WLM can be measured, in our case it is 1.8787 GHz. The simplest method of doing that is to tune the laser frequency until the measured interference pattern overlaps with the initial pattern. ν_F is then the difference between the initial and the final frequency. We are also working on a more sophisticated method where we read in a couple of patterns with their respective frequency and then make use of Fourier transformations in order to extract ν_F.
One of the first things we do is cut off one peak on each side of the interference pattern to get rid of sensor edge effects. Afterwards, we can get a good Gaussian fit to the remaining peaks.
In our example, the peaks we find are at [170, 405, 636, 869, 1096, 1318, 1540, 1762, 1979] and the distances between them are: [235, 231, 233, 227, 222, 222, 222, 217]. Since they are not equidistant, we have to make them so. We calculate the spacing between peaks based on the estimated parameter α_W and the measured wavelength λ. For the first pattern this seems counter-intuitive since we just reverse what we did earlier, but for the following patterns, when the wavelength changes, this makes sense because we make sure that the spacing of the analytic pattern exactly matches the spacing of the stretched pattern. We transform the array with the x-values in such a way that the distance between neighbouring peaks is this spacing. We do not know for sure why the measured peaks are not equidistant. It could be due to optical aberrations or some other effect which we have to compensate in order to make sure that the algorithm's sensitivity is not compromised. Without this compensation

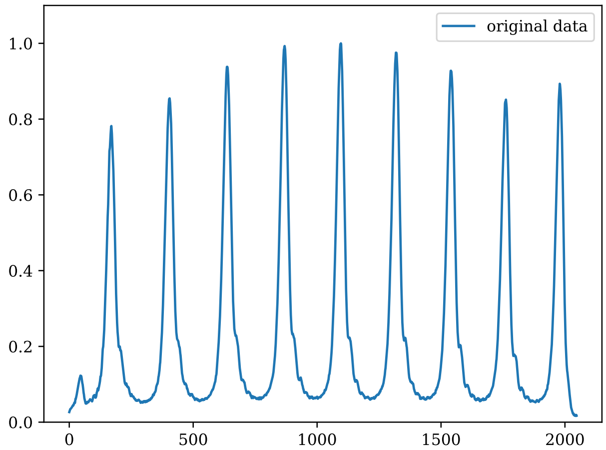

(a) Original data as read from the WLM. The pattern consists of $N = 2048$ data points.

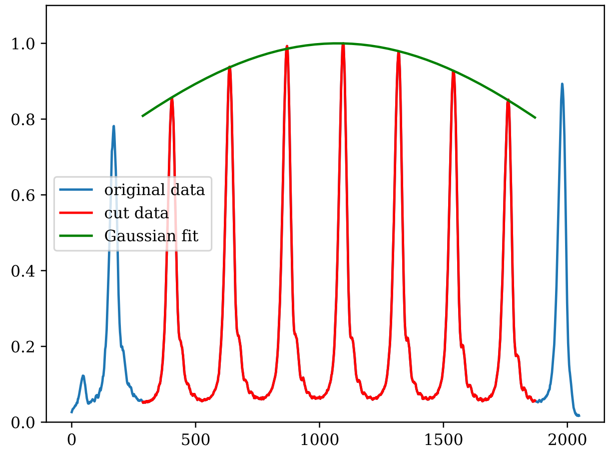

(b) Original data and cut data, where one peak on each side is cut off. A Gaussian is fit to the remaining peaks. The values of $x_{0,\mathrm{G}}$ and w_0 are 1072.99 and 2411.63.

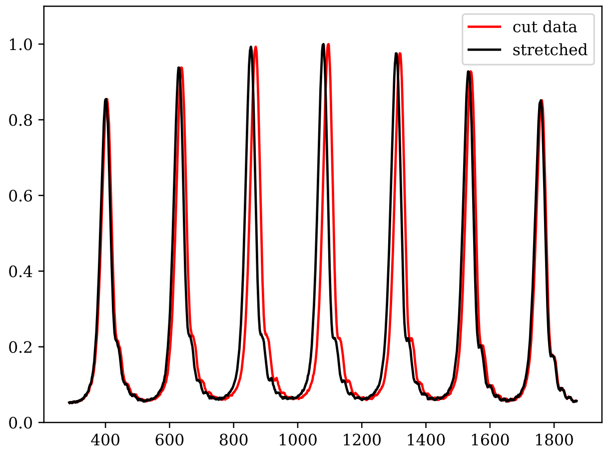

(c) Cut data and stretched data, where the peaks are equidistant.

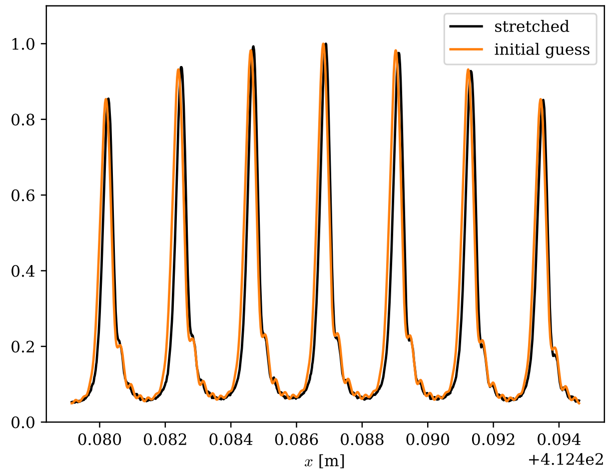

(d) Stretched data and the analytic pattern calculated with formula (28) after they have been transformed. We call it initial guess because the values of the used parameters are the initially guessed ones (z, z_0, R) and the estimated ones (x_0, w_0, α_W). The fringes are not optimally overlapping yet, which is why $R^2 = 0.8454$ is relatively low.

Figure 29: Plots of the different processing steps. The wavelength $\lambda = 854.44\,\mathrm{nm}$.

the coefficient of determination would be very low, even when the laser is in single-mode operation.
Once we have our two arrays with the cut data and our transformed x-values where the peaks are equidistant, we are ready to calculate the coefficient of determination. We put the mentioned array of x-values and the estimated parameters into equation (28), which gives us what we call "initial guess", see figure 29d. As one can see, we obtain the right spacing but the analytic pattern is shifted to the left compared to the stretched data. The coefficient of determination is therefore only 0.8454. As a consequence, we had to introduce a function where we calculate the shift that maximises the coefficient of determination. This shift is added to x_0 and the array of the x-values. The shifted pattern then has a coefficient of determination of 0.9833.
The last step is to numerically optimize the other parameters, namely z, z_0, w_0 and R, using a sequential quadratic programming method. This is done in one step. The parameters change only by a few percent, which, in our example, increases the coefficient to 0.9918.
Now that we have reasonable values for all six parameters, see table 3, we can read in the next pattern together with its wavelength. We repeat step 4 and 5 and return the coefficient of determination.

λ	$854.44\,\mathrm{nm}$
α_W	$193.43\,\mu\mathrm{rad}$
R	0.6327
w_0	$0.02191\,\mathrm{m}$
x_0	$-412.487\,\mathrm{m}$
z_0	$0.09963\,\mathrm{m}$
z	$-0.38221\,\mathrm{m}$

Table 3: Values of the parameters for formula (28) that we obtain for the example in figure 29. z is listed as a parameter because it is also needed for the analytic formula (28).

The idea was that we calculate the parameters for formula (28) and then calculate the analytic pattern. Ideally, this pattern matches the measured one and no shift along the X-axis is necessary. As is clear from the example above, this does not work with our current code. The initial guess is clearly shifted. A very small mismatch already leads to a significantly smaller coefficient of determination. The shift we apply in order to make up for that is $-73.06\,\mu$m. So essentially we are not able to calculate the right array of x-values, i.e. the location of the sensor along the X-axis. And because the sensor does not move, you should only have to calculate the array of x-values once and then use for all following patterns. But the fact that we have to cut off data points on each side of the measured pattern and we also have to make the peaks equidistant makes things a little more complicated. One can easily observe on the live data of the WLM that the interference patterns shift to the left when the frequency of the light increases and shift to the right when the frequency decreases. Therefore the points where we cut off data shift. In our example, see figure 29b, these cut off points are 287 and 1869. The data points out of this range are discarded and not used for the calculation of the coefficient of determination. This is why for each new pattern, as the frequency changes, we need to calculate a new array of x-values. A possible solution might be cutting off a fixed number of data points on each side. However, the transformation we apply in order to make the peaks equidistant might then still cause problems. Part of the mismatch can likely be attributed to uncertainties in our calculated parameters. Their values all rely on guessed (e.g. the length of the sensor) or measured parameters (e.g. the FSR of the WLM or $x_{0,\mathrm{G}}$). For example we have observed that the value of $x_{0,\mathrm{G}}$ fluctuates. This is probably due to noise on the measured interference patterns. Thus we think it is probably not possible to get around a shift for the first pattern, but once one has a good match for the first pattern, it should be possible to read in the new wavelength, put it into the analytic formula and get a good match for the successive patterns, provided that the laser is still in single-mode operation. So no additional shift along the X-axis for each new pattern should be necessary.
When we were testing our script without the shift for each new pattern, the coefficient of determination remained stable above 0.975 for more than 2 hours of measuring, but as soon as we started to tune the

laser by changing the piezo voltage, the coefficient of determination decreased rapidly even though the laser was still in single-mode operation. This shows that our script without the shift cannot handle large changes in frequency. As you can see in the next section, the additional shift we introduced as a consequence solves this problem.
Since we have a relatively low reflectivity R, it turns out that a value of 16 for $p_{\max}$ is already sufficient, see figure 30. Of course one does not want to have an unnecessarily large $p_{\max}$ because this will make the optimization calculations much slower.

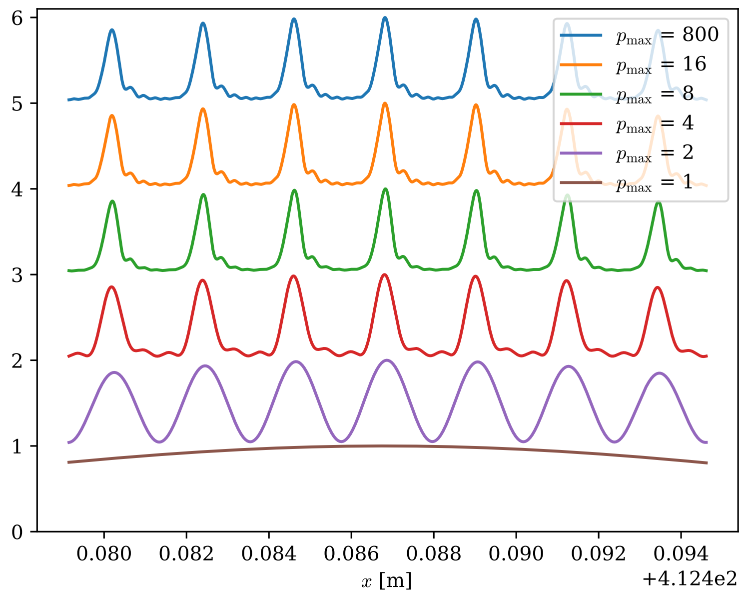

Figure 30: Varying $p_{\max}$. We used the same values for all parameters as in table 3. For $p_{\max} = 1$, we get a Gaussian, which makes sense since we do not take any reflections into account. If we calculate the coefficient of determination where we assume that the pattern with $p_{\max} = 800$ corresponds to the measured data y_i and the patterns with a lower $p_{\max}$ correspond to the model with data points f_i, then we obtain 0.99999929 for $p_{\max} = 16$. Therefore increasing $p_{\max}$ even further does not make much sense. For $R = 0.99$ on the other hand, we get a coefficient of only 0.4428. Then one has to increase $p_{\max}$ to about 50 in order to get a similarly large coefficient as before.

Looking at table 3, you can see that we have two negative values, namely x_0 and z. In the definition of q_p, see equation (32), there is a $+x_0$ which we think should be $-x_0$. What we have now does not make much sense. The center of the incident Gaussian beam is at $x_0 = -412.487\,\mathrm{m}$, but the center of the directly transmitted beam ($p = 1$) is supposed to be at $x = 412.487$, see figure 30.
The parameter z could be negative because we are dealing with negative incidence, i.e. $\theta < 0$. If we had positive incidence, then we would not really be calculating the transmission. We would calculate a pattern originating from a Gaussian that has not yet passed through the wedge.

7.4 Performance of algorithm

In this section, we show the performance of our algorithm on live data, see figures 31 and 32. We recorded 108 patterns together with the wavelength over a time range of 365 seconds, i.e. on average we read in a new pattern approximately all 3.4 s. The wavelength of the laser was tuned by changing the piezo voltage. After pattern 49 the coefficient of determination falls below 0.98 and the laser starts going towards multi-mode operation. On the WLM software you would see that the heights of the peaks decrease as the main mode is loosing power and other modes start lasing. Because we normalise the patterns to 1, the height of the largest peak remains unchanged, but the valleys between the peaks become less deep. The lowest coefficient with a value of 0.1313 is reached for pattern 79. Then we increase the frequency. Between pattern 88 and 89 there is a mode hop after which the laser returns to single-mode operation and the coefficient of determination reflects this by attaining values over 0.99 again.

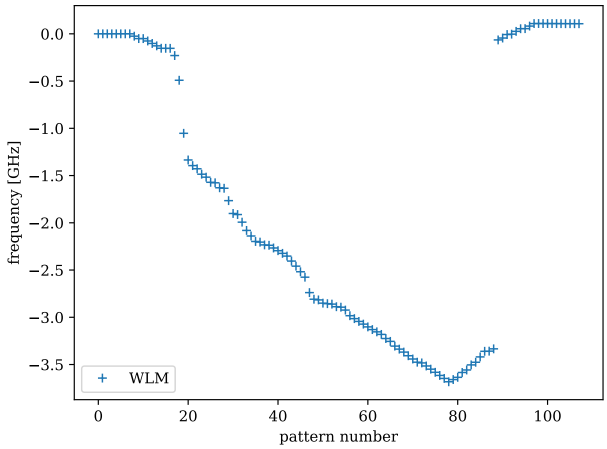

(a) Change in frequency measured with WLM. The frequency was changed by adjusting the piezo voltage.

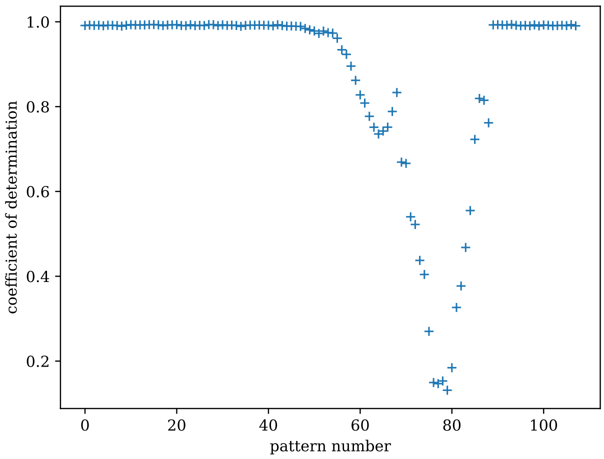

(b) Evolution of the coefficient of determination.

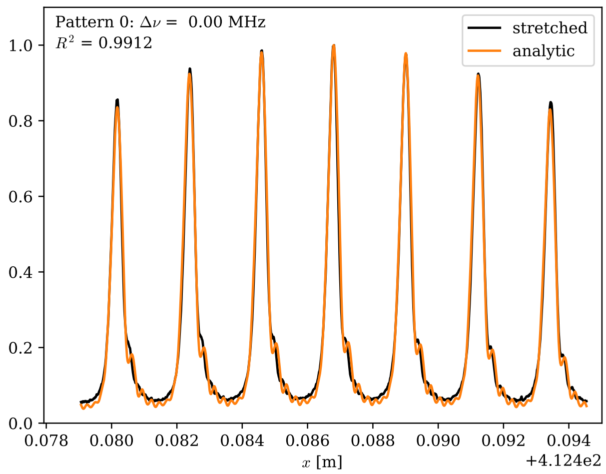

(c)

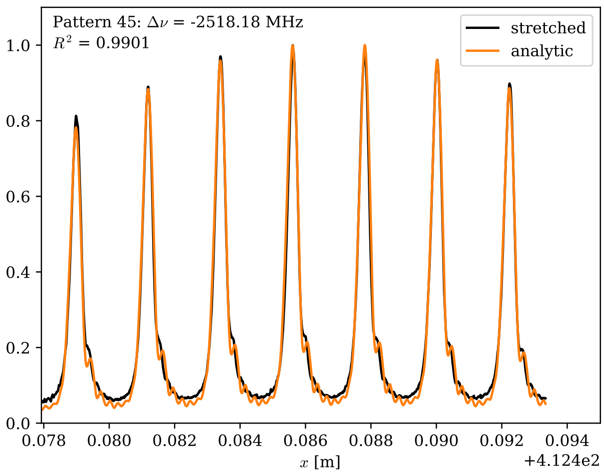

(d)

Figure 31: Performance of algorithm on live data. R^2 is the coefficient of determination and $\Delta\nu$ is the change in frequency relative to the first pattern. The wavelength $\lambda = 854.44\,\text{nm}$.

We have also done long-term tests of about 48 hours of our algorithm where we observed that the coefficient of determination stays above or close to 0.99 when the laser is in single-mode operation. Therefore, in our experience, we would say that 0.98 is a good threshold. A coefficient below that value indicates that other modes are lasing.

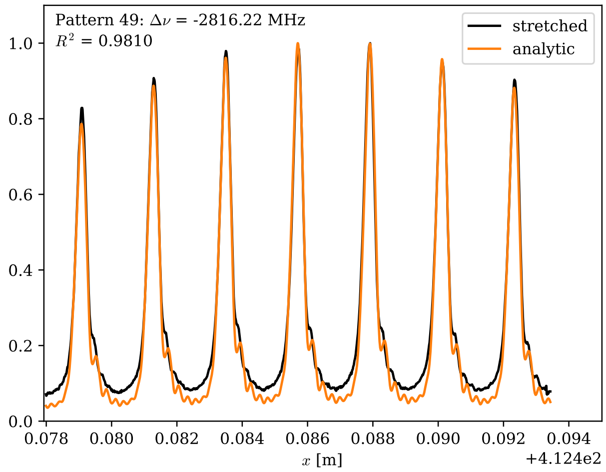

(a)

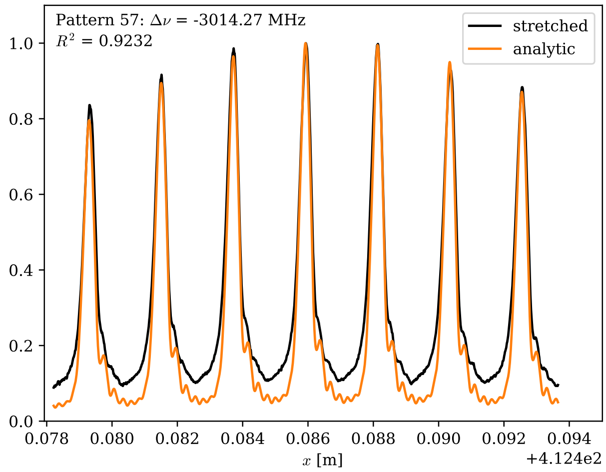

(b)

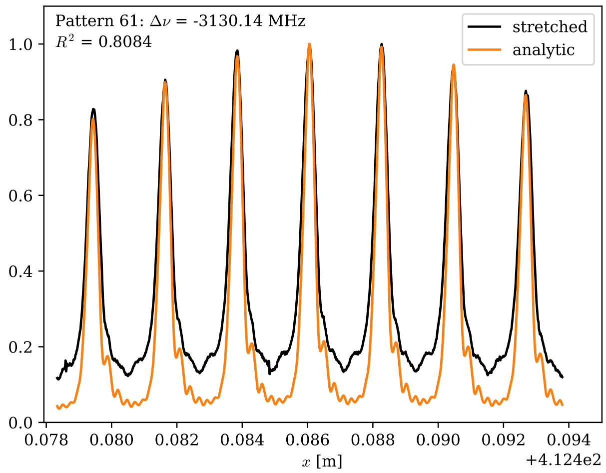

(c)

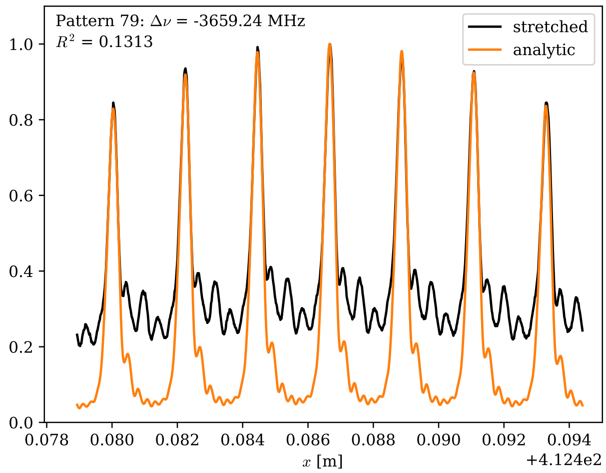

(d)

Figure 32: Performance of algorithm on live data. R^2 is the coefficient of determination and $\Delta\nu$ is the change in frequency relative to the first pattern. The wavelength $\lambda = 854.44\,\mathrm{nm}$.

7.5 Modulated spectrum

As we have already mentioned, because of an acousto-optic modulator and other optical elements, the locked laser has several additional frequencies in its spectrum, the strongest of which is a +400 MHz sideband. This leads to a smaller coefficient of determination, even though the laser is in single-mode operation. One can find an example of this issue in figure 33, where the coefficient of determination is only 0.9539. If we subtract the analytic pattern from the stretched one and normalize it, then we get a pattern with peaks that are shifted relative to the peaks of the stretched pattern, see figure 33b. Subsequently, we calculate the analytic pattern with a new wavelength

$$\lambda = \frac{c}{c/\lambda_0 + 400\,\mathrm{MHz}}, \tag{42}$$

where λ_0 is the wavelength the WLM initially determined. The peaks of this new pattern overlap very well with the ones of the difference pattern. Therefore it should be possible to model the measured pattern of the locked laser as a sum of two analytic patterns:

$$I_T^{\text{measured}} = a \cdot I_T(x, z, \lambda_0) + b \cdot I_T(x, z, \lambda). \tag{43}$$

where $0 \leqslant a, b \leqslant 1$. We have to make sure that the right side is still normalized to 1. As one can see in figure 33c, this approach allows us to obtain an analytic pattern which again has a coefficient of determination over 0.98. We just have to introduce an additional step where we optimize the parameters a and b. As long as the relative power of the two frequencies stays constant, so should the values of a and b. The algorithm should then work as it has before.

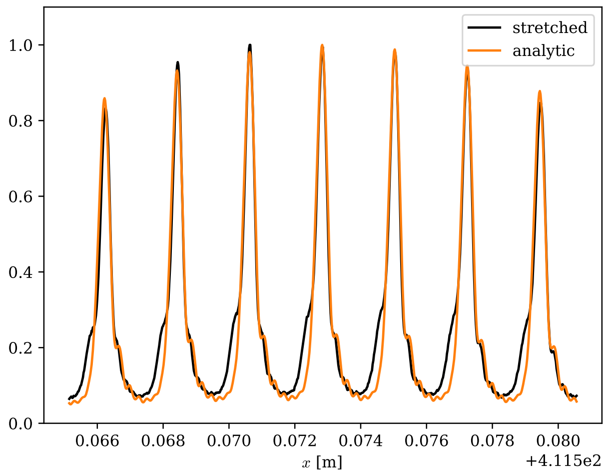

(a) Coefficient of determination of 0.9539.

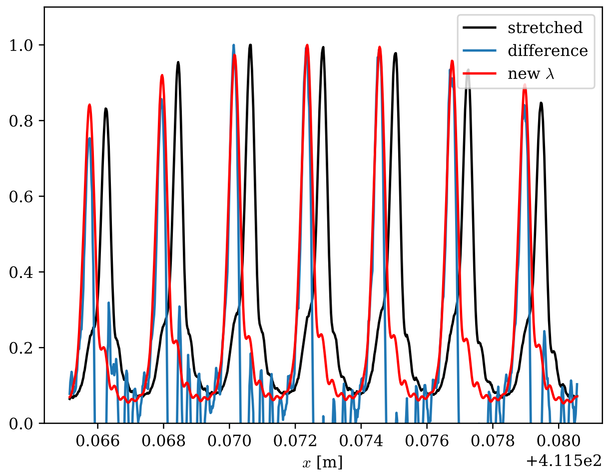

(b) The blue line is the difference between the stretched and the analytic pattern, see figure 33a. The red line is the analytic pattern calculated with a new wavelength which is exactly 400 MHz higher in frequency.

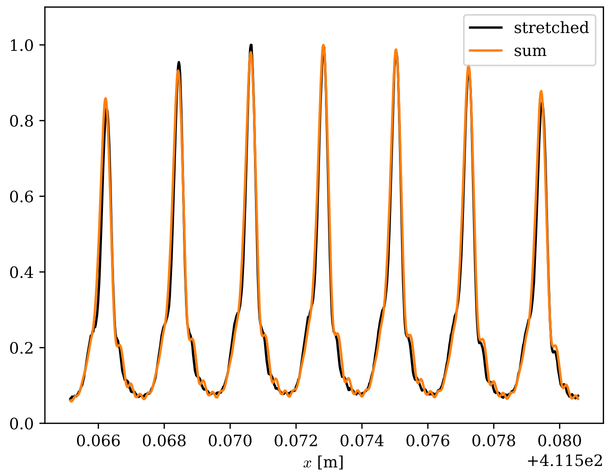

(c) Modelling the measured pattern as a sum of two analytic patterns with $a = 0.972$ and $b = 0.122$, see equation (43). The coefficient of determination is 0.9821. With optimization of the other four parameters (z, z_0, w_0, R) we can even increase it to 0.9899.

Figure 33: Performance of algorithm on a pattern of the locked laser which has a +400 MHz sidband. The wavelength $\lambda = 854.44\,\text{nm}$.

7.6 Conclusion & outlook

In summary, we designed and implemented an algorithm that keeps track of the spectral purity of a laser. It relieves experimentalists from the burden of manually checking the displayed pattern in the WLM software and is a key ingredient for automation. Even though the algorithm works sufficiently well for all practical purposes, there is room for improvement. For example, we are not satisfied with the fact that we have to optimize the shift along the X-axis for each new pattern.

One of the challenges with this project was that we are dealing with a black box, the Fizeau wedge inside the WLM, about which we have no information. As a consequence, there are four parameters whose values one has to guess before starting the script. The values we provide in this thesis are a good starting point but depending on the type of WLM one is using, one will have to adjust them accordingly. The script was tested extensively with a Toptica DL pro and the WS8-10 WLM. We have started using it with other WLMs and have had good experiences with it so far. We hope to see a feature similar to this incorporated into the WLM software someday, which should not be very hard to achieve for the manufacturer of the WLMs.

References

[1] A. Kramida, Y. Ralchenko, J. Reader, and NIST ASD Team (2020), *NIST Atomic Spectra Database*, National Institute of Standards and Technology, Gaithersburg, MD, version 5.8 [Online], 14.09.2021, `https://physics.nist.gov/asd`

[2] F. Riehle, *Frequency standards: basics and applications*, Wiley-VCH Verlage GmbH & Co. KGaA (2004)

[3] H. Talvitie et al., *Passive frequency and intensity stabilization of extended–cavity diode lasers*, Review of Scientific Instruments 68, 1 (1997)

[4] A. Takamizawa et al., *External cavity diode laser with very-low frequency drift*, Applied Physics Express 9, 032704 (2016)

[5] J. Kim et al., *Locking multi-laser frequencies to a precision wavelength meter*, arXiv:2108.02411 [quant-ph] (2021)

[6] R. Paschotta, *Stabilization of lasers*, RP Photonics Encyclopedia [Online], 01.09.2021, `https://www.rp-photonics.com/stabilization_of_lasers.html`

[7] H. Talvitie et al., *Improved frequency stability of an external cavity diode laser by eliminating temperature and pressure effects*, Applied Optics, Vol. 35, No. 21 (1996)

[8] A. Takamizawa et al., *External cavity diode laser with frequency drift following natural variation in air pressure*, Applied Optics, Vol. 54, No. 18 (2015)

[9] Wikipedia, *Laser diode* [Online], 31.08.2021, `https://en.wikipedia.org/wiki/Laser_diode`

[10] B.E.A. Saleh and M.C. Teich, *Fundamentals of photonics*, Wiley Series in Pure and Applied Optics, 3rd edition (2019)

[11] A.L. Schawlow and C.H. Townes, *Infrared and optical masers*, Physical Review 112 (6) (1940)

[12] Toptica Photonics AG, *Tunable Diode Lasers*, product brochure (2018)

[13] H.R. Telle, *Stabilization and modulation schemes of laser diodes for applied spectroscopy*, Spectrochimica Acta Rev. Vol. 15, No. 5 (1993)

[14] B. Mroziewicz, *External cavity wavelength tunable semiconductor lasers - a review*, Optoelectronics review, 16 (4), 347-366 (2008)

[15] S.C. Doret, *Simple, low-noise piezo driver with feed-forward for broad tuning of external cavity diode lasers*, Review of Scientific Instruments 89, 023102 (2018)

[16] Wikipedia, *Allan variance* [Online], 05.09.2021, `https://en.wikipedia.org/wiki/Allan_variance`

[17] National Institute of Standards and Technology (NIST), *Time and Frequency from A to Z* [Online], 05.09.2021, `https://www.nist.gov/pml/time-and-frequency-division/popular-links/time-frequency-z`

[18] H. Jung and D.G. Gweon, *Creep characteristics of piezoelectric actuators*, Review of Scientific Instruments 71, 1896 (2000)

[19] R. Matthey et al., *Methods and evaluation of frequency aging in distributed-feedback laser diodes for rubidium atomic clocks*, Optics Letters 36, issue 17 (2011)

[20] E. Stoykova, *Transmission of a Gaussian beam by a Fizeau interferential wedge*, J. Opt. Soc. Am. A, Vol. 22, No. 12 (2005)

[21] E. Stoykova and M. Nenchev, *Gaussian beam interaction with an air-gap Fizeau interferential wedge*, J. Opt. Soc. Am. A, Vol. 27, No. 1 (2010)

[22] E. Stoykova, M. Deneva and M. Nenchev, *Analysis of Fizeau wedge with a non-air gap by plane wave expansion*, Proc. of SPIE Vol. 11207 (2019)

[23] High Finesse, *Wavelength Meter Angstrom WS 8*, user manual

[24] R. Paschotta, *Single-mode operation*, RP Photonics Encyclopedia [Online], 18.08.2021, `https://www.rp-photonics.com/single_mode_operation.html`

[25] Wikipedia, *Coefficient of determination* [Online], 04.09.2021, `https://en.wikipedia.org/wiki/Coefficient_of_determination`

[26] O. Novák et al., *Fizeau interferometer system for fast high resolution studies of spectral line shapes*, Review of Scientific Instruments 82, 023105 (2011)