Gigahertz-range scanning of near-UV frequency-stabilized probing lasers

David Sieger

Bachelorarbeit in Physik
angefertigt im Institut für Angewandte Physik

vorgelegt der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität
Bonn

Juli 2026

Ich versichere, dass ich diese Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt sowie die Zitate kenntlich gemacht habe.

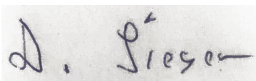

Bonn, 30.07.2026
Datum

Unterschrift

Gutachter und Betreuer: Prof. Dr. Sebastian Hofferberth
Gutachter: Dr. Frank Vewinger

Contents

CHAPTER 1

Introduction

The first laser – light amplification by stimulated emission of radiation – was built in 1960 in California by Theodore Maiman [1]. It quickly became one of the most important inventions of that time period, with uses spreading from something as simple as playing a DVD to applications in the medical field or the cutting of sturdy materials. It is also an important tool in research, in particular for precise control over atoms. In the Nonlinear Quantum Optics (NQO) group in Bonn, lasers are used to trap and cool atoms, and to excite atoms into Rydberg states. This demands precise control over the wavelength of the monochromatic and coherent light emitted by the laser. In the Ytterbium Quantum Optics lab (YQO) [2], which is part of the NQO group, lasers are used to cool and trap a cloud of ytterbium atoms, and then to excite them to the Rydberg state via an intermediate state, creating a three-level system in which the transitions are driven by a probe and a control laser. This allows for various studies to be performed, such as on ultralong-range Rydberg molecules [3] or on interacting photons via the Rydberg nonlinearity [4].

To keep the lasers' frequency stabilized, the lasers have to be locked. For this, a master laser is used, which is locked to an ultrastable cavity; the other lasers are then locked to the master laser using optical phase-locked loops (OPLL), which use beatnotes to stabilize them. At the moment, this only allows the probe laser an 800 MHz range around the $6s6p\,^1P_1$ intermediate state, currently limited by a voltage output limitation and more fundamentally by the decaying gain profile of the Radio Frequency (RF) locking electronics.

The aim of this thesis is to increase this range to multiple gigahertz, without requiring more optical power than the previous iteration of the setup. This enables greater flexibility in measurements and a larger accessible range. For example, measuring the optical depth of the atom cloud consists of scanning the probe laser over a wide frequency range around the resonance. Also when measuring Stark eigenenergies of Rydberg states at finite electric fields, a large frequency range is valuable due to the large polarizabilities leading to quick energy changes with electric field.

This work builds upon existing electronics that are currently used for the probe laser. It aims to improve the design of the phase-locked loop (PLL) box by using different RF components, in order to achieve a larger scan range and to optimize the lock behavior.

Chapter 2 describes the optical setup and the characteristics of the lasers themselves. Chapter 3 then discusses the limitations of the PLL boxes currently used in the experiment, how their design is changed in order to increase the locking range to multiple gigahertz, and verifies the results of this new design using the actual beatnote. In Chapter 4, a closed feedback loop is built, and different settings of the Proportional-Integral-Derivative (PID) feedback electronics – allowing for a slow and a fast feedback path – are optimized in order to obtain a stable lock; the lock quality and its behavior at different lock points are also analyzed. Finally, Chapter 5 tests the phase lock in a long-term scan over several days comparable to its use in the YQO experiment, and Chapter 6 discusses the limitations of the setup as well as possibilities for further improvement. It also covers what would need to be changed in the current setup of the experiment in order to implement the new design.

CHAPTER 2

Optical setup

In order to determine whether the testing setup provides a good simulation of its use in the YQO experiment, and to be able to trace back whether any limitations arise specifically from it, the lasers used as well as the optical setup are characterized and discussed in this chapter.

2.1 Laser system

For testing the phase locking of lasers, two high-power near-ultraviolet (UV) lasers from Toptica were used, more specifically external-cavity diode lasers (ECDLs) in the Littrow configuration. These are based on a diode and an external grating, which reflects the first order back into the diode laser and is thereby able to improve the lasing of the device; the zeroth order is then used for the application. This setup allows the laser to be controlled in two different ways: first, through the laser current, which directly impacts the power of the laser as well as its wavelength; and second, by adjusting the angle of the grating through a piezo, which extends proportionally to the applied voltage, thereby enabling precise movement [5]. In practice, the current and the piezo are not adjusted independently: a setting of the laser, called the feedforward factor, forces a current change proportional to the piezo change, thereby extending the mode-hop-free range (MHFR). Both lasers are of the same model (Toptica DL Pro HP) and can produce light in the near-UV regime, ranging from 398.5 to 400.5 nm, and are specified to have an MHFR of 20 GHz [6]. Since both lasers were built according to YQO specification, no general public data sheet exists for this configuration; the data sheets used here are only accessible with NQO wiki access.

2.2 Measured stability and feed-forward

After careful, iterative tuning of the lasers by adjusting the current as well as the temperature of the diode, an MHFR of just 8 to 10 GHz was achieved. This was done by finding the middle of the MHFR and then applying a large piezo scan until a mode hop occurred, and then readjusting the values to see whether the mode hops could be pushed out further. A specific mode was chosen to maximize the MHFR, which included the frequency of 751.5255 THz, the frequency used in the main experiment by the already existing lasers. This is the frequency needed for the transition from the ytterbium ground state $6s^2\,{}^1S_0$ to the intermediate state $6s6p\,{}^1P_1$, as required for the experiment. By further reducing the current of the laser by 40 mA, to 96 mA, the MHFR was improved to 12 GHz. This is still far off from the claimed 20 GHz, but this may have been achieved in a different mode, which is much more stable. Without the feedforward factor, the MHFR was 5 GHz at best, regardless of which mode was tested. Since these are free-running lasers, some drift is expected; after letting the laser run for a few days, a drift of up to a few hundred MHz was observed, as shown in Fig. 2.1. The jump at around 43 h into the measurement is not a mode hop, and could be caused by something like a drastic change in pressure or temperature, or by vibrations – for example, someone entering the lab. The mode-hop regions, on the other hand, tend not to change as much over the same period of time, roughly staying at the same wavelength and only changing by up to 50 MHz across days of running.

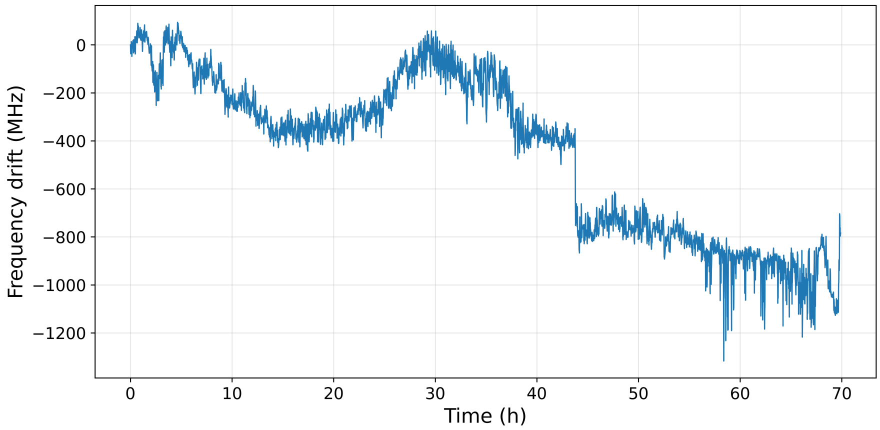

Figure 2.1: Frequency drift of a free-running laser, measured with a wavemeter over roughly 70 h, relative to the frequency at the start of the measurement.

2.3 Lasing threshold

The lasing threshold is the minimum current at which stimulated emission begins to dominate over spontaneous emission and the laser starts producing a coherent beam. As soon as stimulated emission dominates, the current response changes, producing a kink in the photodiode signal versus current curve, which is used to determine the threshold [7].

The lasing threshold was measured by applying a large scan of the laser current from 0 to 150 mA and simultaneously measuring the signal of the laser's internal photodiode. This test was conducted to determine whether a coarse realignment of the laser grating was necessary: if the alignment is poor, the measured threshold current is higher, since stimulated emission then only comes to dominate at a higher current. This was found not to be the case, as shown in Fig. 2.2. The lasing threshold remained consistent with, or better than, the value specified in the laser's datasheet [6], which gives 34 to 35 mA for the slave laser and 35 to 36 mA for the master laser, compared to the measured values of (33.79 ± 0.13) mA for the master laser and (34.44 ± 0.23) mA for the slave laser.

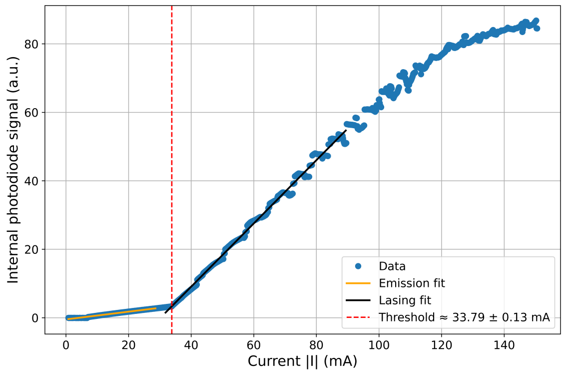

(a) Master laser.

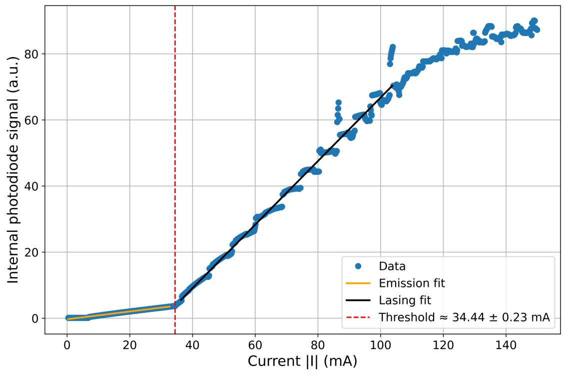

(b) Slave laser.

Figure 2.2: Lasing threshold determined from the laser current scan, for both lasers. Below and above the threshold, the photodiode signal is each fit with a line, since both the spontaneous-emission and the lasing regime are expected to behave linearly with current; the lasing threshold is taken as the intersection of the two fits, giving (33.79 ± 0.13) mA for the master laser and (34.44 ± 0.23) mA for the slave laser.

2.4 Beatnote generation and detection

When two laser beams are overlapped, a beatnote can be observed. Each laser can be described as an oscillating electric field with angular frequency ω; when the intensity of the overlapped beams is measured with a photodiode, whose signal is proportional to E^2, contributions appear at $\omega_1 + \omega_2$ and $\omega_1 - \omega_2$. Of these, only the difference frequency $\omega_1 - \omega_2$ can actually be resolved, since the sum frequency and the frequency itself lie far beyond the detector's bandwidth. By measuring this beatnote frequency, the difference between ω_1 and ω_2 can be determined directly [8].

The setup in which the beatnote is generated can be seen in Fig. 2.3. In order for the laser beams to overlap, a temporary mirror was used to reflect the beam after the BS onto a free part of the optical table; the beams were then checked for overlap at two different distances from the mirror by performing a beam walk using the second PBS cube after the slave laser and the mirror after it. After achieving this, the lens and the photodiode were put in place, with the lens positioned roughly 5 cm away from the photodiode to achieve optimal focusing. The diode was then positioned so as to maximize the beatnote signal on the spectrum analyzer. Later, a repositioning of the lens was done to further optimize the beatnote signal. At

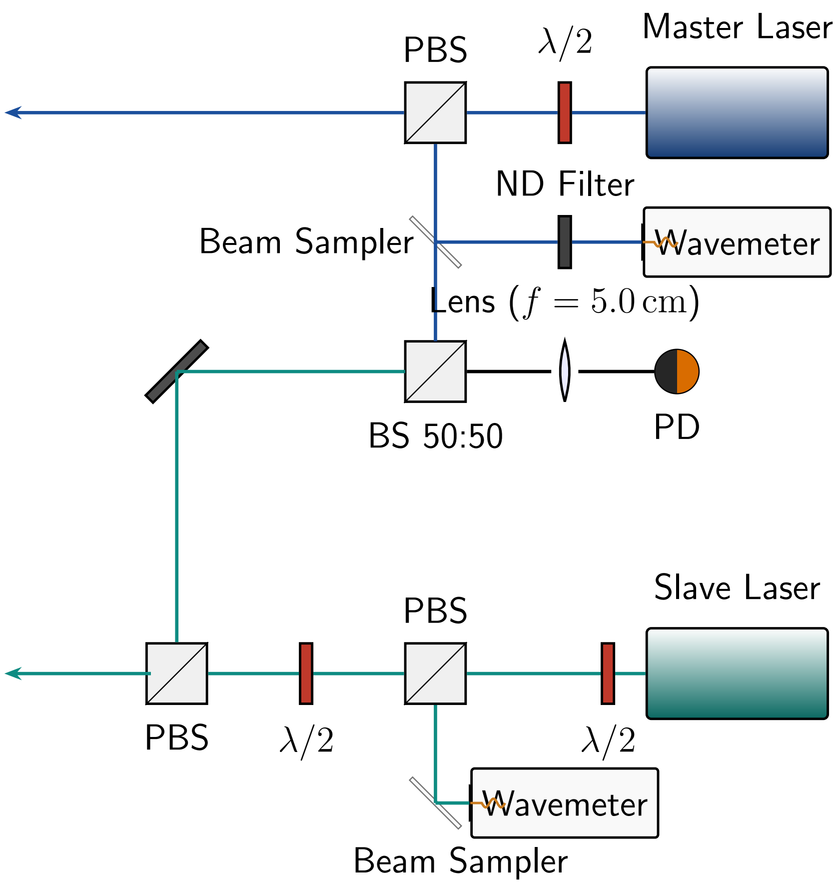

Figure 2.3: Optical setup used to phase-lock the slave laser to the master laser. Each laser passes through a half-wave plate ($\lambda/2$) and a polarizing beamsplitter (PBS) before being sent to the experiment. A beam sampler picks off a small fraction of each beam for wavelength monitoring via a fiber-coupled wavemeter; the master branch additionally passes through a neutral-density (ND) filter. A further pick-off from each laser is combined on a non-polarizing 50:50 beamsplitter (BS 50:50) and focused ($f = 5.0\,\mathrm{cm}$) onto a photodiode (PD) to generate the beatnote signal used by the phase-locked loop.

this point, an old PLL box was used to amplify the signal and relay it to the spectrum analyzer, which is described in more detail in [9].

The slave laser and the master laser were chosen arbitrarily, and both are free-running lasers; in a real experiment, the master laser is most likely locked to a cavity or an atomic transition, which here was not the case. The $\lambda/2$ plates can also be turned, adjusting the polarization and thereby how the PBS splits the laser beam; since the wavemeter requires a specific amount of power, only the second $\lambda/2$ plate was turned to adjust the power, while the power of the master laser remained the same. For the optimizing process of the beatnote, a power of 0.5 mW for the master laser and 1.5 mW for the slave laser was chosen.

CHAPTER 3

PLL electronics and beatnote verification

This chapter focuses on the design of the new PLL box: what was important, where the previously used box reached its limits in being able to scan across the whole range, and how the final version was reached. This is first done using vector network analyzer (VNA) measurements, i.e., the purely electrical response of the PLL box, before verifying the design using the actual optical setup and beatnote.

3.1 PLL and amplifier electronics

The phase-locked loop (PLL) works by comparing the beatnote signal to a reference signal and then generating an error signal, which can be used to eventually close the loop; more detail on this can be found in Chapter 4. For all the PLLs used in the experiment, an EVAL-ADF4007EBZ1 evaluation board [10] is used, containing the ADF4007 chip, which works as a high frequency divider/PLL synthesizer [11]. Some manual changes were made to the wiring of the evaluation board in order to bypass the op-amp following the error signal, since its bandwidth is too low and would otherwise result in worse lock behavior.

In order for the comparison to work, the beatnote signal, which has the form of a sine wave, needs to be digitized; for the digitization to work, the amplitude of the signal needs to be between −10 and 5 dBm for frequencies of 1 to 6 GHz, while for frequencies of 0.5 to 1 GHz and 6 to 7 GHz this range is further reduced to −5 to 5 dBm [11]. In order for the beatnote to reach this amplitude, the signal from the photodiode first needs to be biased and then amplified.

The setup of the multiple boxes that are actively used in the experiment contains three amplifiers (ZX60-6013E-S+ [12]) as well as the bias tee (ZFBT-6G+ [13]) and a 2 m SMA cable connecting the photodiode to the box. For the photodiode, a Hamamatsu G4176-03 photodetector [14] is used, which generally has a bandwidth of 12 GHz, and so it should not have a significant effect on the attenuation and does not need to be considered for now. As can be seen in Fig. 3.1, these amplifiers do not have a flat gain profile, and even the cable has some loss, so the resulting beatnote signal falls off by 20 dB, making it impossible to achieve the −10 to 5 dBm range across the whole range of the desired beatnote frequency, and thus to lock the laser across a range of multiple GHz.

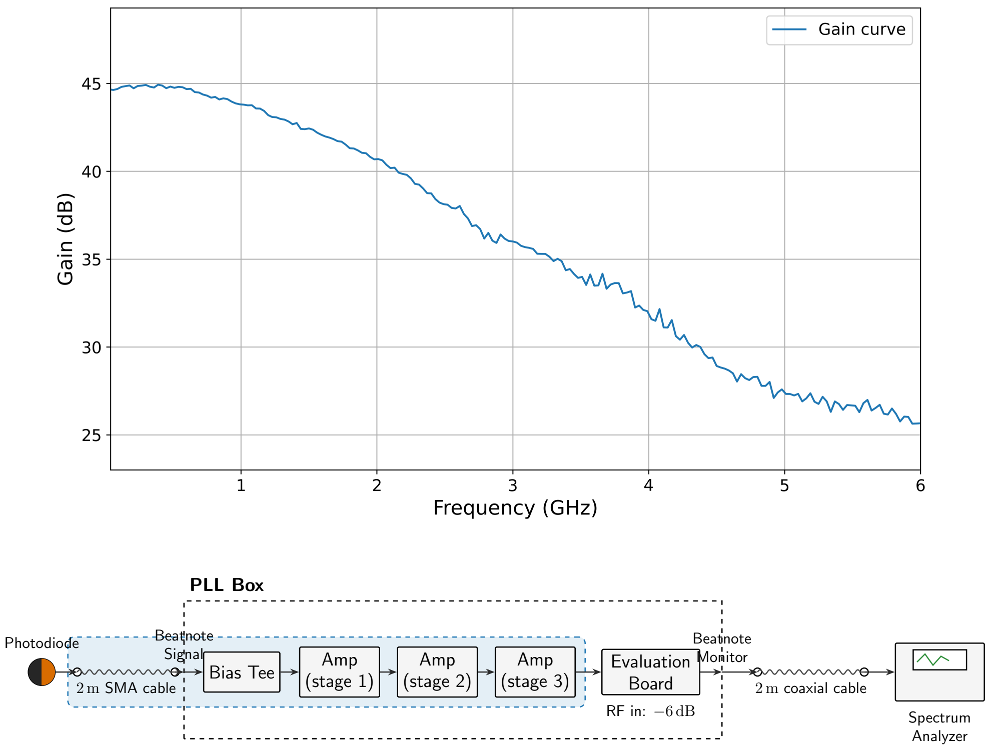

Figure 3.1: Gain of the RF chain (bias tee, three ZX60-6013E-S+ amplifiers, and the 2 m SMA cable from the photodiode) as a function of frequency, measured with a vector network analyzer. The highlighted part of the RF chain below indicates exactly which components this measurement covers.

3.2 Design goals and required changes for multi-GHz scanning

In order to improve this, an RF chain is needed before the evaluation board which ultimately has as flat a gain as possible across the desired beatnote range and relatively low noise. The aim is to change only about 5 dB over the whole range, so that getting the signal into the desired −10 to 5 dBm range has some slack, meaning reoptimization on the optical side has more room and does not need to be done as often. For this, flat-gain, low-noise amplifiers and passive equalizers with an almost linear attenuation across 6 GHz have been ordered. To determine which components to use, the amplification was simulated either by extrapolating the values from the datasheets of different amplifiers or components, or by measuring them directly using a VNA; different setups were then tested to see whether a flat gain curve across that range could be achieved, and the combination seen in Fig. 3.2 was finally chosen and built into a box. However, the voltage distribution board had to be changed compared to previous iterations [9] because of the different voltages needed by the amplifiers.

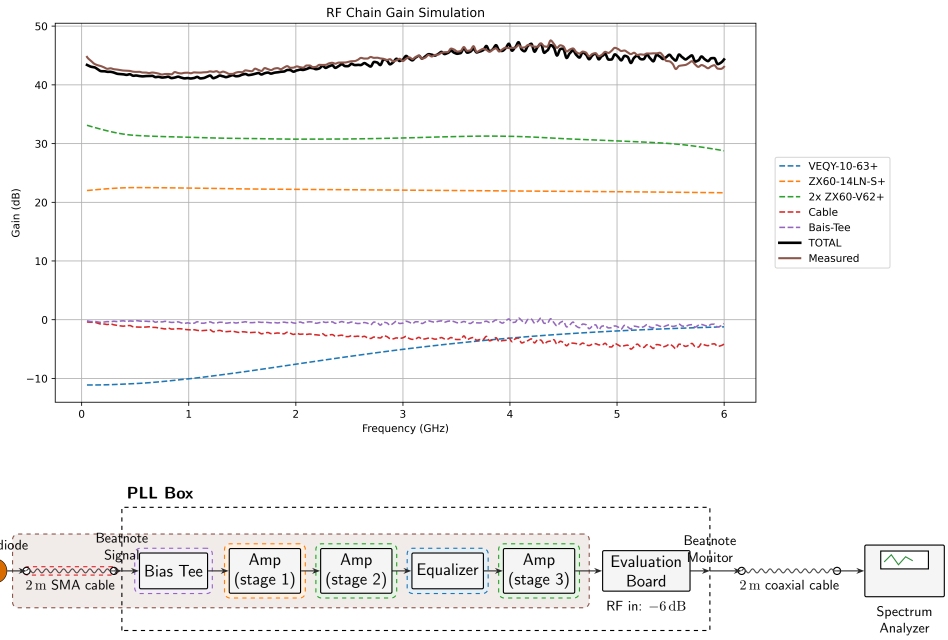

Figure 3.2: Simulated gain of the RF chain (bias tee, first-stage amplifier ZX60-14LN-S+, two later-stage amplifiers ZX60-V62+, an equalizer VEQY-10-63+, and the cable), obtained from datasheet values, compared to the total simulated gain and to the gain of the assembled chain measured with a vector network analyzer, alongside the corresponding stage of the RF chain, with each component colour-matched to its trace above; the measured gain corresponds to the whole highlighted section, from the input cable up to and including the third amplifier stage.

The bias tee used is the ZFBT-6G+ [13], which works up to 6 GHz with an almost flat gain behavior; it takes the photodiode signal and separates out the DC part, which is not needed here. After this, the first-stage amplifier ZX60-14LN-S+ [15] follows; it has a relatively flat gain of about 22 dB and operates in the range of 50 MHz to 10 GHz, with a noise figure of only 2 dB. The later-stage amplifiers are ZX60-V62+ [16], which operate from 50 MHz to 6 GHz and also have a relatively flat gain of about 15 dB and a noise figure of about 5 dB. The equalizer is the VEQY-10-63+ [17], a passive equalizer with an attenuation slope of 10 dB, which works from DC to 6 GHz. Lastly, a standard 2 m SMA cable is used to connect the PLL box to the photodiode. The different contributions of these components can be seen in Fig. 3.2.

After successfully building the design of the PLL box, the measurement obtained with the vector network analyzer is very close to the predicted total, as shown in Fig. 3.2. Only toward the upper end of the range does the measured gain become slightly lower than predicted; this is caused by the parasitic effects of the small SMA connectors inside the box accumulating at high frequencies.

Later, another change was made: the cable leading from the photodiode to the box was replaced by a low-loss cable [18], as can be seen in Fig. 3.3; this reduces the loss before the first amplification stage. The equalizer, which was compensating for the loss of the old cable, could also be removed, so switching out the cable increased the amplification especially at lower frequencies. The overall flatness of the gain stays about the same, with a maximum change of 6 dB, which is acceptable and hard to improve further.

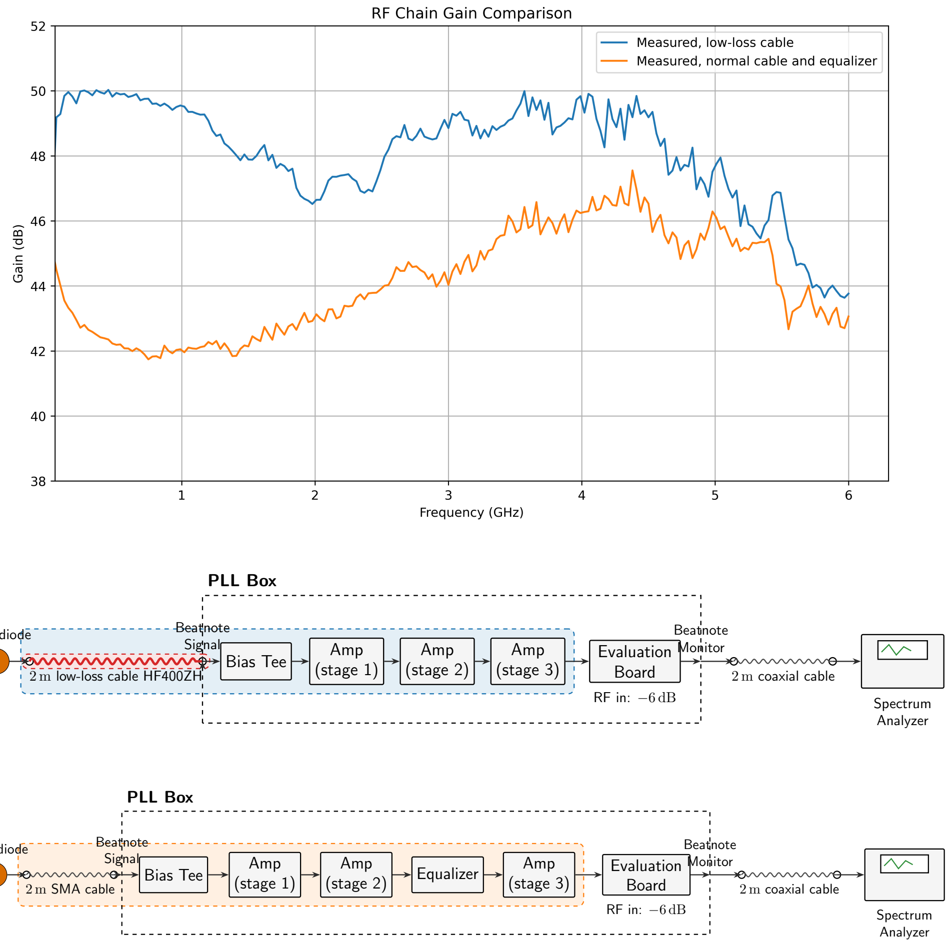

Figure 3.3: Gain of the RF chain, measured with a vector network analyzer, before and after replacing the input cable with a low-loss cable (HF400ZH [18]) and removing the equalizer; below, the corresponding stage of the RF chain after the change, with the measured sub-chain highlighted in blue and the new low-loss cable highlighted in red; and, at the bottom, the stage of the RF chain before the change, with the measured sub-chain highlighted in orange.

This is the final design that was chosen. The schematic of the PLL box, showing all components inside the box and how they are connected, can be seen in Fig. 3.4(a); the layout of the voltage distribution board is shown in Fig. 3.4(b). A full parts list for this final design, with references to the corresponding datasheets, is given in Table A.1. Here, some changes had to be made to the previously used design [9], since the amplifiers required different voltages: every amplifier used 12 V, whereas now the first-stage low-noise amplifier uses 6 V, while the second- and third-stage amplifiers require 5 V. Since the lab power supply provides the box with 15 V, a step-down conversion was chosen so that each voltage regulator only has to perform a small step-down, so that the dissipation is spread out over the different regulators and no single regulator overheats.

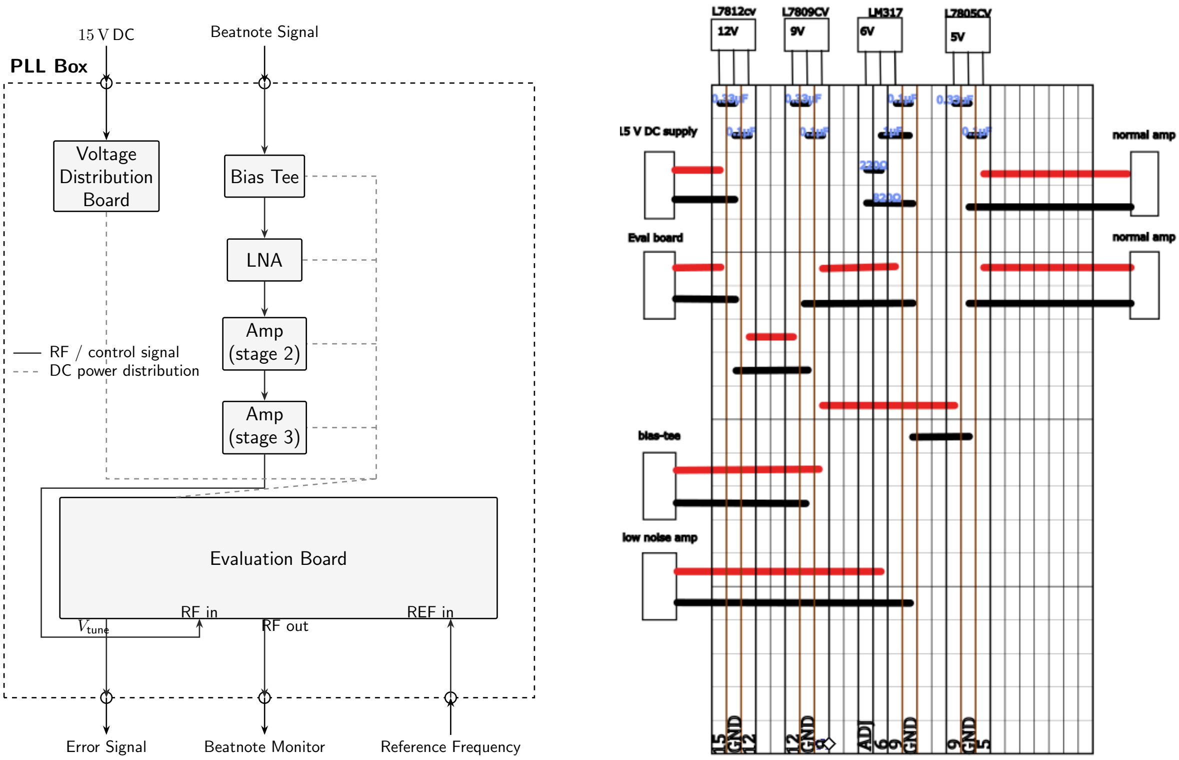

(a) Schematic of the PLL box.

(b) Voltage distribution board.

Figure 3.4: Final PLL box design: schematic and voltage distribution board.

In the voltage distribution board shown in Fig. 3.4(b), the red cables are the ones carrying a voltage, while the black cables are grounded; the black cables with blue numbers are connected to either a capacitor or a resistor. The voltage regulators are located at the top of the board, with the number on each one indicating the voltage it transforms its input into. The components named here are equivalent to those listed in Table A.1.

3.3 Beatnote measurements with the new setup

In order to characterize the whole system, including the effect of the lasers as well as the response of the photodiode, the beatnote was measured using the whole setup. To do so, the master laser was positioned at around 751.5220 THz, so that at a beatnote of 3.5 GHz the slave laser has the frequency needed for the transition from the ytterbium ground state $6s^2\,{}^1S_0$ to the intermediate state $6s6p\,{}^1P_1$, as required for the experiment. The slave laser was manually moved with the piezo controller, and the amplitude of the beatnote observed on the spectrum analyzer was measured. The measurement can be seen in Fig. 3.5.

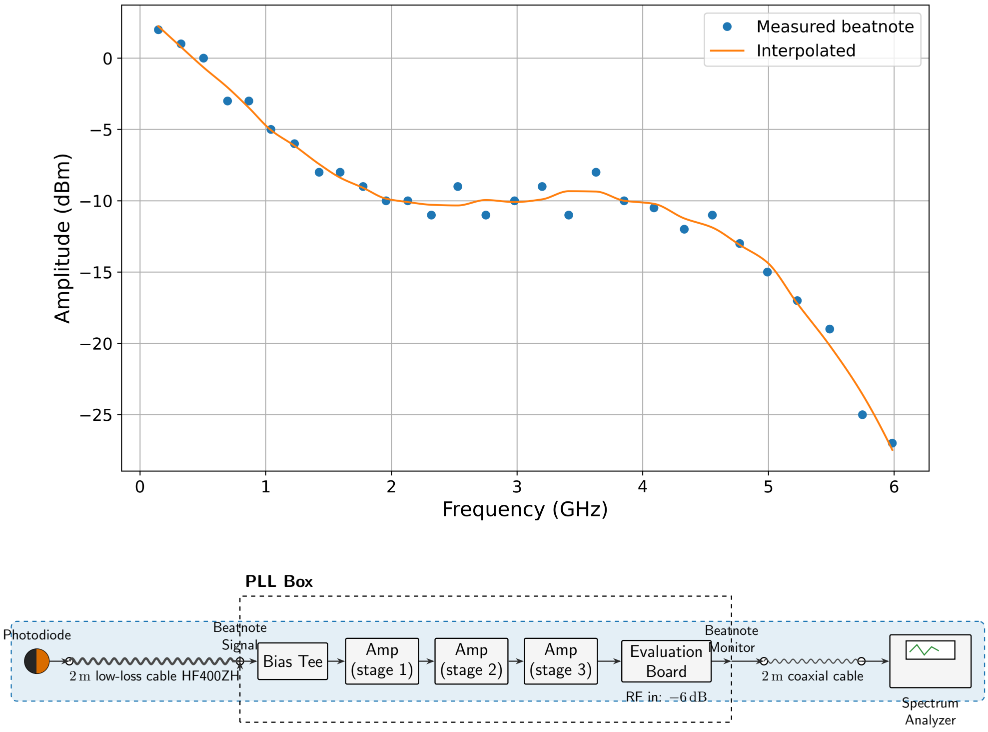

Figure 3.5: Amplitude of the beatnote observed on the spectrum analyzer, as a function of frequency, with the slave laser manually scanned via its piezo controller. Below, the corresponding stage of the RF chain: the raw reading integrates the whole signal path, from the photodiode to the spectrum analyzer itself, highlighted in blue.

However, in order to obtain the actual signal that reaches the evaluation board, the loss of the cable after the monitoring output (Fig. A.1), as well as the 6 dB attenuation caused by the monitoring output itself, needs to be subtracted [10].

The overall flatness decreased by much more than originally expected, to a change of over 12 dB across the whole range, so that the signal barely stays within the 15 dB range required for the evaluation board to work, see Fig. 3.6. However, with the lasers as they currently are, the signal does not stay entirely within the evaluation board's range and falls out of range at 5.7 GHz. To fix this, the overall signal needs to rise by about 3 dB: since the beatnote amplitude scales as $\sqrt{P_1 P_2}$, its power scales as $P_1 P_2$; increasing the laser power of the slave laser from 1.5 mW to 3.5 mW therefore gives an expected increase of $10\log_{10}(3.5/1.5) \approx 3.68$ dB, more than enough to compensate.

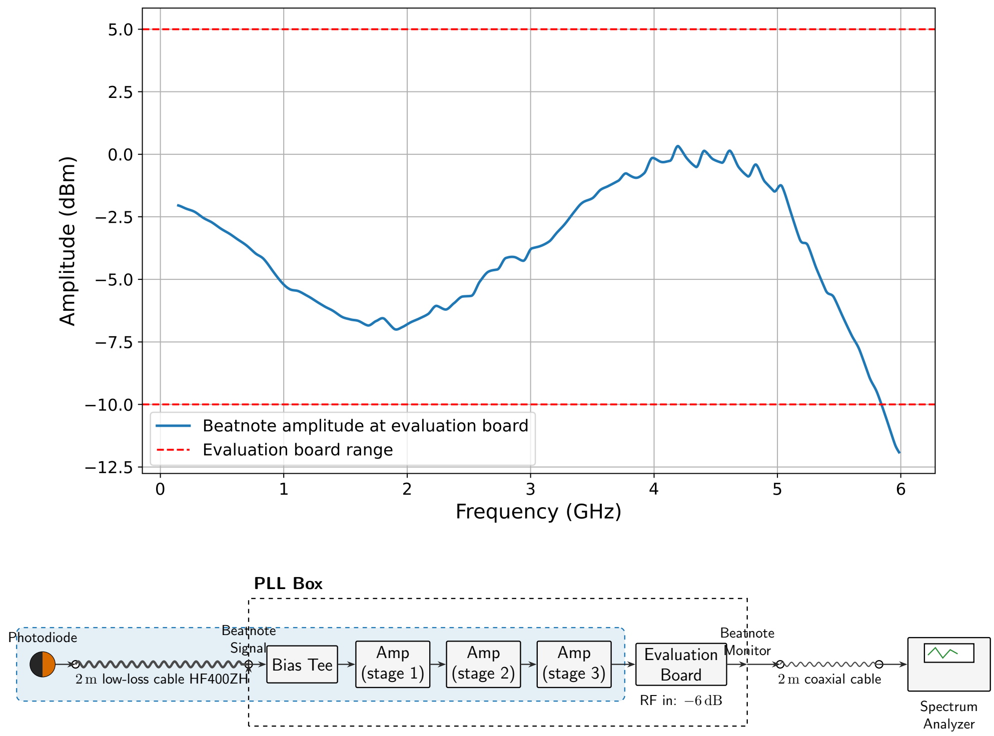

Figure 3.6: Beatnote amplitude at the evaluation board, obtained by correcting the measured beatnote for the monitoring cable loss and the 6 dB attenuation of the monitoring output, compared to the −10 to 5 dBm range required by the evaluation board. Below, the corresponding stage of the RF chain, with everything the corrected signal represents, from the photodiode up to the evaluation board, highlighted in blue.

Nevertheless, more change in the beatnote range was observed than originally believed, so to further investigate this, the electrical RF chain can be compared to the corrected beatnote signal, as shown in Fig. 3.7. Both curves are shown as a normalized version, with the normalization chosen such that both signals are zero at 3 GHz; the reason for this is that the laser was originally optimized at this position, meaning it was checked to be single-mode with a good mode shape at this point. Although the same tendencies in the behavior can be seen, some differences remain, which are caused either by the photodiode response or by the behavior of the laser modes themselves – for example, reducing the overlap by slightly changing shape or position.

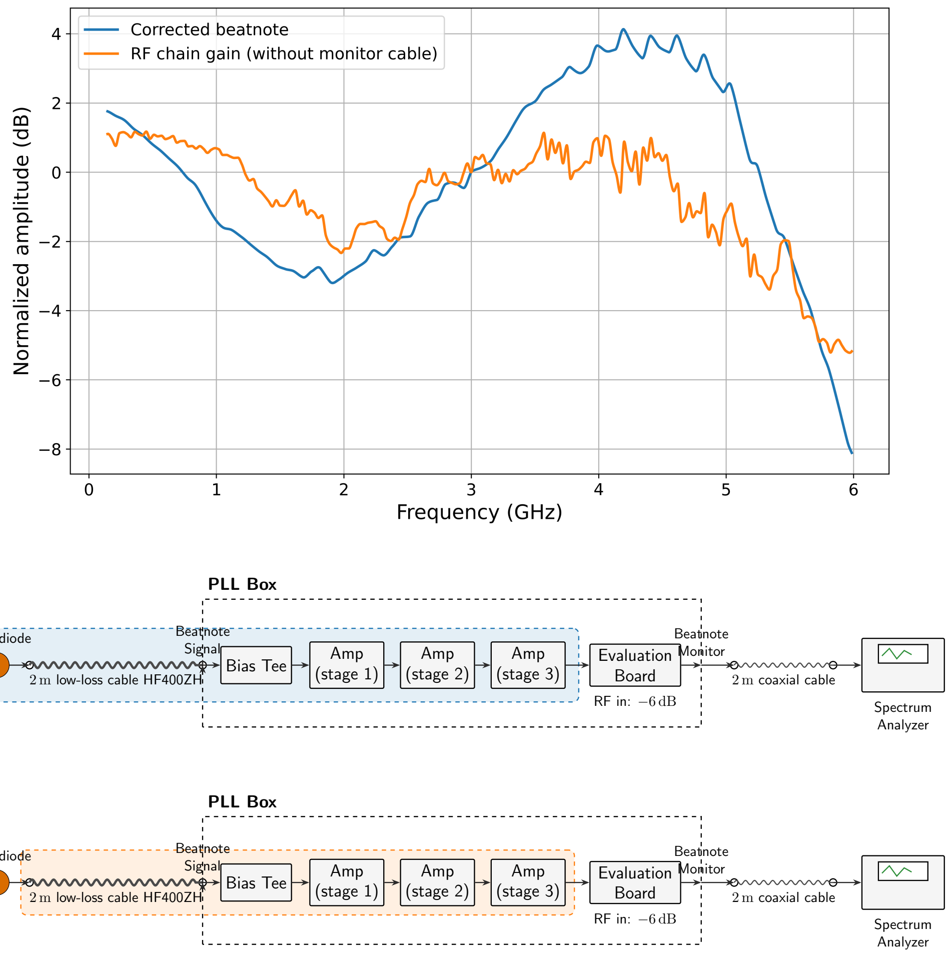

Figure 3.7: Corrected beatnote amplitude compared to the electrical gain of the RF chain (without the monitoring cable), both normalized to be zero at 3 GHz, the frequency at which the laser was originally optimized. Below, the corresponding stage of the RF chain for each trace: the corrected beatnote, from the photodiode up to the evaluation board, highlighted in blue; and the RF chain gain, without the photodiode, highlighted in orange.

To probe this further, the middle optimization point was changed to be at the beginning of the scan rather than in the middle, but an almost similar behavior was observed, so it is most likely that the photodiode responds differently depending on the frequency, which was originally assumed to be negligible. The estimated diode response is plotted in Fig. 3.8. It can be seen to change by 7 dBm over the 6 GHz range, assuming that the laser's own behavior does not influence the system at all, which is debatable, so this is more of an estimate than a precise measurement. Still, this makes it harder to find an optimization point where the whole 6 GHz range falls within the range required for the evaluation board to work.

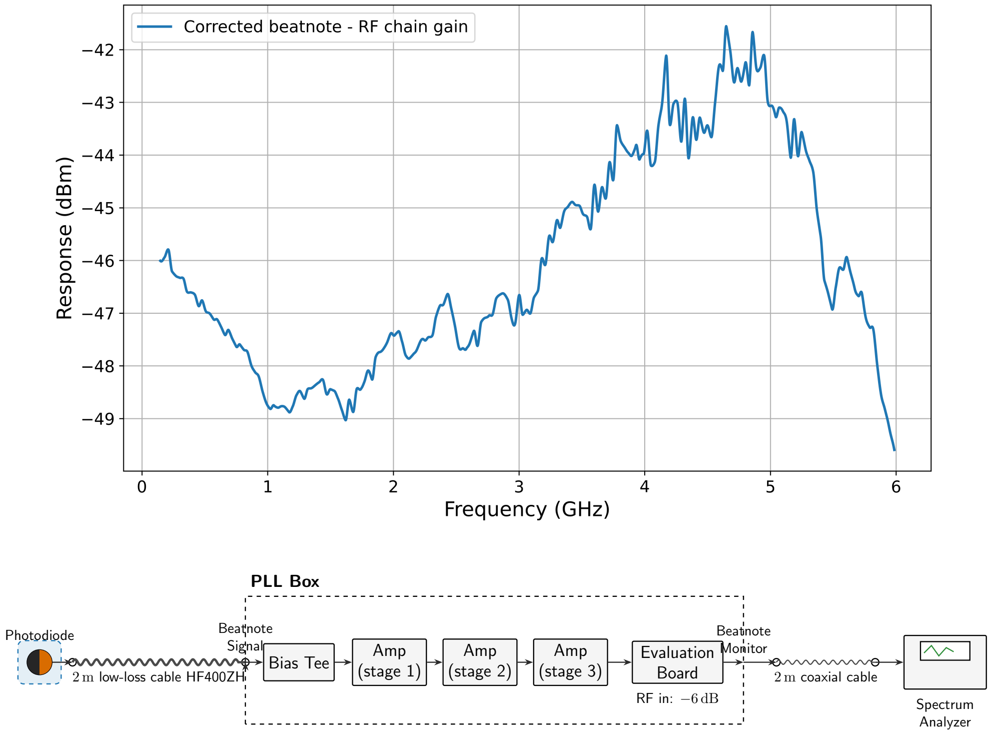

Figure 3.8: Estimated photodiode response, obtained by subtracting the RF chain gain (Fig. 3.7) from the corrected beatnote amplitude at the evaluation board (Fig. 3.6). Below, the corresponding stage of the RF chain, with the photodiode highlighted in blue.

CHAPTER 4

Optimizing lock parameters

This chapter covers building the feedback loop and optimizing the PID and PLL settings in order to obtain a working fast and slow feedback, which together create a lock, and testing the lock's performance under different lock conditions. This is done mostly by analyzing the spectrum around the beatnote and checking the position of the servo bumps, and how they behave at different reference frequencies and optical powers used.

4.1 Setup of the full feedback loop

The full setup, and how the individual components are connected to form the complete feedback loop, can be seen in Fig. 4.2. The loop only closes once the laser light reaches the photodiode, which is not shown here for the sake of optical clarity.

The fully functioning loop works by comparing the beatnote signal coming from the diode and converting it into a digital pulse train. This signal is then divided by a factor N, which can be chosen via the positioning of jumpers on the eval board; a factor of 64 was chosen in order to get the largest possible range of the beatnote. The board is only able to accept a reference frequency between 20 to 240 MHz [10], which is then halved again, giving the relation

$$\frac{f_{\mathrm{ref}}}{2} = \frac{f_{\mathrm{beat}}}{N},$$

so that the beatnote is 32 times the input reference signal chosen for the Direct Digital Synthesizer (DDS) Arduino [19]. Its internal clock, generated from an external reference crystal [20], is specified to 50 ppm, which the phase lock transfers directly onto the beatnote: the beatnote frequency is therefore accurate to 5×10^{-5} of its set value, i.e. 40 kHz at 800 MHz and 290 kHz at 5 750 MHz. A beatnote between 640 MHz and 7.68 GHz can in theory be chosen; with the available settings, this is the largest reference-frequency variability that can practically be achieved within the 0 to 6 GHz domain [11].

The frequencies are then compared by a dual flip-flop circuit and generate an error signal; the functional block diagram can be seen in Fig. 4.1. The error signal gives out a high signal when the reference frequency is higher than the beatnote frequency, and gives out a low signal when the beatnote frequency is higher; if the two are the same, the signal averages out to about the middle of the high and low signal levels. This can be chosen to be 0 to 5 V or 0 to 3 V via another jumper; 0 to 3 V was chosen since it has a shorter rise time [9].

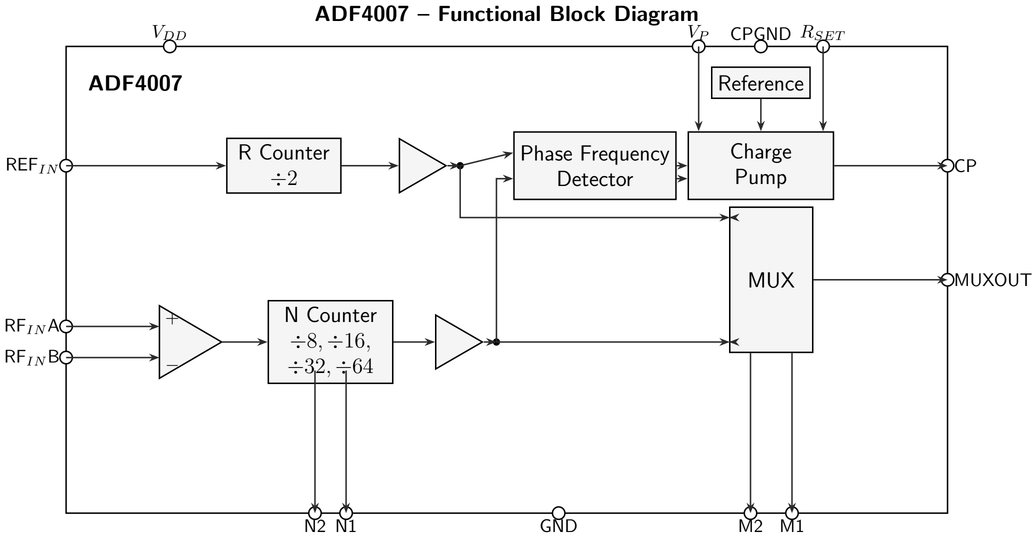

Figure 4.1: Functional block diagram of the ADF4007 PLL synthesizer used on the evaluation board [11]. The reference input is divided by 2, the beatnote input by a selectable factor of 8, 16, 32, or 64 (64 was chosen here), and the phase frequency detector and charge pump then generate the error signal from the two divided signals.

The error signal is given to the oscilloscope and then to the FALC PRO [21], a PID controller that specializes in ultra-fast response and low-noise regulation. It has a setting that allows the offset of the incoming error signal to be chosen, which must be set so that the signal is symmetric around 0. The PID has different settings to optimize the slow and the fast feedback – discussed further in Sections 4.2 and 4.4 – and it produces its own error signal, which is also monitored on the oscilloscope.

Both the fast and the slow feedback are observed on the oscilloscope and sent to the laser controller, which in turn sends the signal to the laser. This either adjusts the current – and therefore the wavelength – on a relatively small scale, at a frequency on the order of MHz, or it is the slow signal, which generally fixes the long-term drifts as well as the changes in the reference frequency once the laser is locked.

The oscilloscope triggers on an external trigger given out by the laser controller, at the frequency of the scan the laser applies; this is generally a piezo scan, but in theory it could also be changed to a current scan, as has been done in Fig. 2.2.

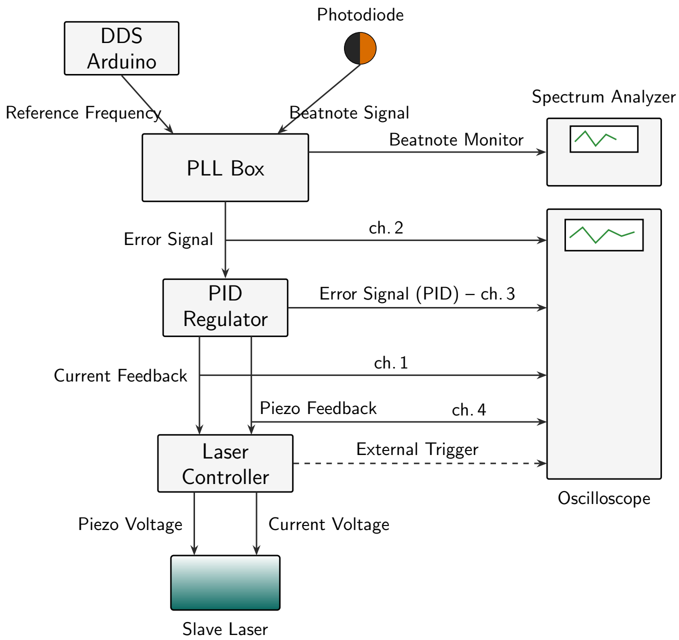

Figure 4.2: Setup of the full feedback loop. The photodiode's beatnote signal and a reference frequency from a DDS Arduino enter the PLL box, which outputs a beatnote monitor signal (to a spectrum analyzer) and an error signal. The error signal is fed to a PID regulator, which produces a current feedback and a piezo feedback signal for the laser controller; the laser controller in turn drives the slave laser via a piezo voltage and a current voltage. The oscilloscope monitors the error signal, the PID's own error signal, and both feedback signals, and is triggered externally from the laser controller.

4.2 Fast feedback

For optimizing the fast feedback, the functions of the FALC PRO as a PID controller need to be considered. It has an overall gain, equivalent to P. It then has three different integrators, I_1, I_2, and I_3, and two derivative stages, D_1 and D_2.

First, one needs to determine whether the error signal input has the correct polarity. Using only the gain, it is increased while observing whether the beatnote moves toward or away from the lock point; if it moves away, the inverting input is used instead. Next, I_1 and I_2 are adjusted iteratively together with the gain, looking at a small frequency range around the beatnote to check for servo bumps. These are moved as far away from the beatnote as possible while keeping them small and avoiding ringing (i.e., second-order servo bumps), improving both their position and height. D_1 and D_2 are then adjusted iteratively in the same way to find an optimum. The effect the different settings have on one another is quite convoluted. The difference the settings can make is still visible, for example when optimizing D_1, shown in Fig. 4.3; the optimal settings found for the phase lock are shown in Table 4.1, with the resulting lock shown in Figure 4.5.

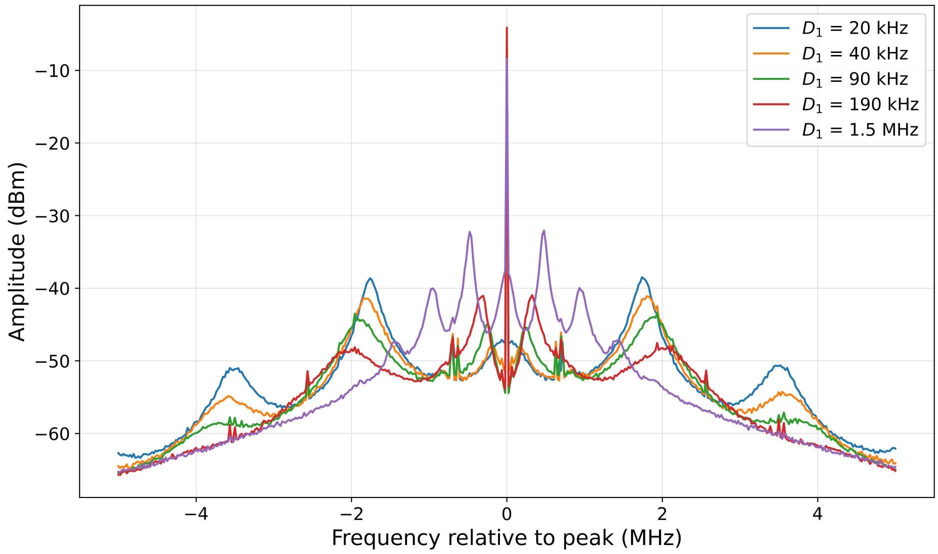

Figure 4.3: Peak-aligned beatnote spectra for different values of D_1. Lower values of D_1 push the servo bumps further out and reduce their height. With incorrectly chosen settings, third-order servo bumps can also appear, as seen here for $D_1 = 1.5$ MHz.

Parameter	Value
Gain (P)	−19 dB
I_1	7 MHz
I_2	10 kHz
D_1	10 kHz
D_2	700 kHz

Table 4.1: Optimal FALC PRO settings found for the phase lock.

As a side note, the lock performance was optimized with the master laser at roughly 751.5220 THz and the slave laser at 751.5255 THz, resulting in a beatnote of 3.5 GHz. This was done so that the range of values the beatnote could later achieve would be symmetric around 751.5255 THz, which is the frequency used for the probe laser in the main experiment. For this optimization, a power of 3.5 mW for the slave laser and 0.5 mW for the master laser was chosen.

4.3 Lock quality

To analyze the quality of the lock, the spectrum analyzer in the setup shown in Fig. 4.2 can be used directly, looking at the width of the beatnote peak itself, the position of the servo bumps, and sidebands in the lock. Sidebands can be introduced through specific frequencies inside and outside the lock, which then end up modulating the laser itself.

The first quantity to consider is the width of the beatnote when optimizing for what is observed on the spectrum analyzer. A few settings need to be considered: the bandwidth can be changed, as can the scan range, which then indirectly changes the scan duration; the amount of averaging across the scan can also be changed, to get rid of some noise when observing the beatnote. The most important of these is the bandwidth of the spectrum analyzer, which was minimized at 10 Hz, meaning peaks narrower than 10 Hz can no longer be correctly resolved. This can be seen in the beatnote width shown in Fig. 4.4, which is already a strong indicator of a good lock.

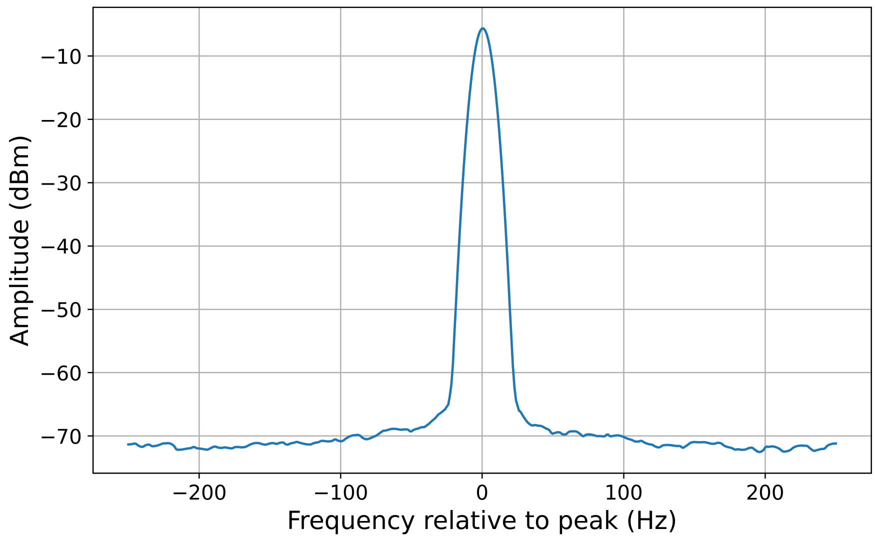

Figure 4.4: Beatnote peak of the locked laser at a beatnote position of 1 GHz, with the frequency axis centered on the peak. Measured with a 500 Hz scan span, a resolution bandwidth of 10 Hz, and 100-fold trace averaging.

Next, the height and position of the servo bump peaks are considered, for which a medium scan width is needed. Enough averaging is required to remove most of the noise so that the shape can be seen clearly, so a medium bandwidth for the spectrum analyzer was chosen. Figure 4.5 shows the final, optimized lock, with the servo bumps positioned at ±1.43 MHz away from the carrier; this represents a substantially better height and position of the servo bumps than any previously achieved in the optimization process.

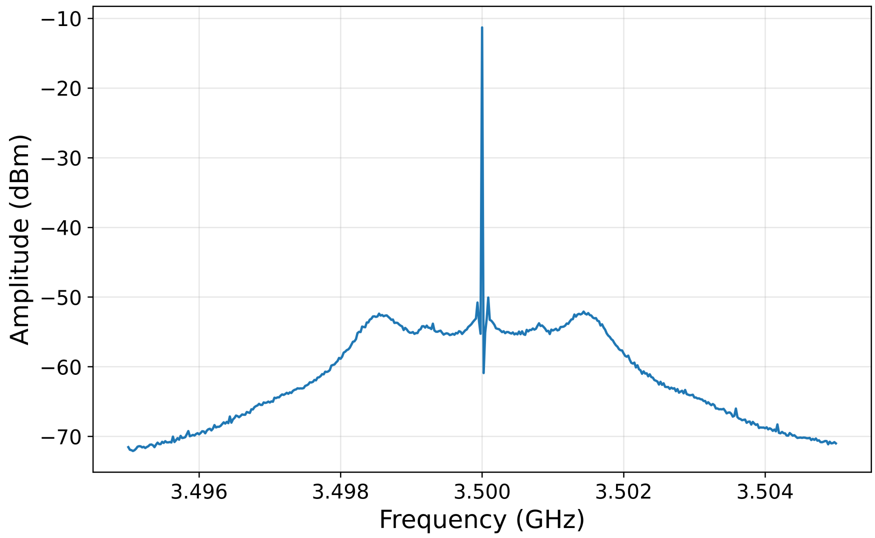

Figure 4.5: Beatnote peak of the optimized lock at 3.5 GHz, with the servo bumps visible on both sides of the carrier. Measured with a 10 MHz scan span, a resolution bandwidth of 100 Hz, and 20-fold trace averaging.

The residual phase error was extracted from the beat-note spectrum of Fig. 4.5 using

$$\frac{P_{\text{carrier}}}{P_{\text{total}}} = \exp\left(-\sigma_\varphi^2\right), \tag{4.1}$$

calculated similarly to Appel, MacRae and Lvovsky [22]. Integrating the phase-noise pedestal over both sidebands, from 22 kHz to 5 MHz around the carrier, gives $\sigma_\varphi^2 = 1.56\,\text{rad}^2$ and a root-mean-square residual phase error of $\sigma_\varphi = (1.25 \pm 0.22)\,\text{rad}$. The uncertainty is dominated by the detector and averaging mode of the spectrum analyzer. The measurement was taken at a beatnote frequency of 3.5 GHz with a divider ratio $N = 64$. When comparing this to the values given in [22], it must be taken into account that they used a different evaluation board; although similar to the one used, their ADF4107 can change N and R digitally, where R represents the factor by which the reference frequency is divided, whereas the ADF4007 used here essentially has a fixed R of 2.

Reference	N	R	$(\Delta\varphi)^2$
Mini-Circuits JTOS-400, phase locked to 216 MHz [22]	96	3	$0.08\,\text{rad}^2$
This work, Arduino DDS reference at 3.5 GHz/32	64	2	$1.56\,\text{rad}^2$

Table 4.2: Comparison of the residual phase noise $(\Delta\varphi)^2$ between the reference used in [22] and the one used in this thesis; both share the same divider ratio $N/R = 32$.

Both references share the same divider ratio $N/R = 32$, yet the residual phase noise achieved here is considerably larger, showing that the limit for achieving a lower $(\Delta\varphi)^2$ is not reached. Essentially, this means that only $\exp(-\sigma_\varphi^2) \approx 0.21$, i.e. about 21 %, of the power is in the coherent carrier, so this is more of a frequency lock that works by comparing the phase to the reference frequency, rather than an actual phase lock. In the YQO experiment, however, this is exactly what is needed, and phase coherence is not important. The loss of phase coherence is most likely caused by the delay in the circuit itself and the large N that was chosen, and limits how far the servo bumps can be pushed out; this issue could be partly improved by reducing the cable length of the setup. However, for the application in the main experiment this is not necessary, since the natural linewidth of the $6s^2\,{}^1S_0 \rightarrow 6s6p\,{}^1P_1$ transition, (29.13 ± 0.06) MHz [23], is large in comparison to the residual phase noise at the pedestal and the servo bumps.

Another good and more straightforward way to tell whether the servo bumps are too large is to check for second-order servo bumps appearing in the spectrum-analyzed beatnote signal. Their presence suggests that the servo bump amplitude is large enough that, at integer multiples of the frequency offset at which the servo bumps are positioned from the carrier (the beatnote peak), another, smaller servo bump can be observed. This is created because, at that frequency, the noise of the beatnote is suppressed very poorly, and in some cases is even larger than that of the unlocked beatnote. This was not the case for the optimal lock found here, but one example of this behavior is shown in Fig. 4.6.

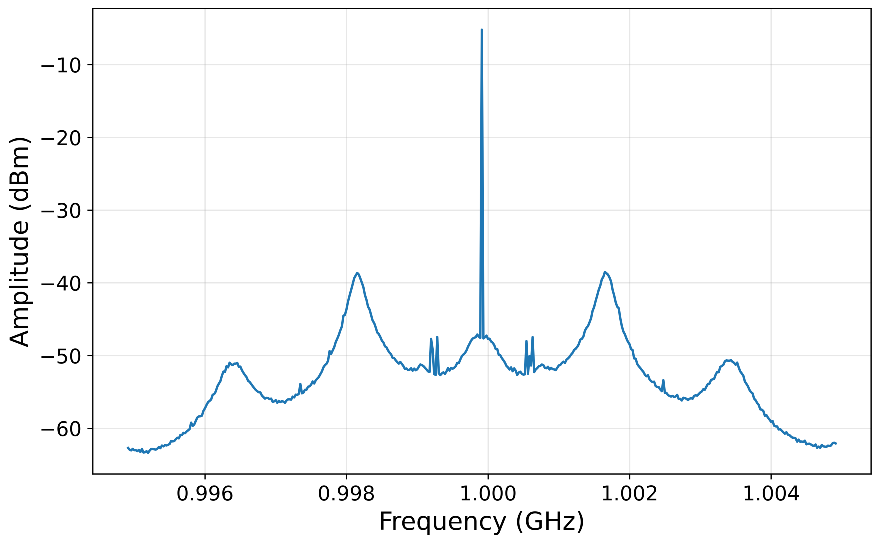

Figure 4.6: Beatnote peak of a lock exhibiting second-order servo bumps: in addition to the first-order servo bumps close to the carrier, a second, pair of bumps appears further out. Measured with the same settings as Fig. 4.5: a 10 MHz scan span, a resolution bandwidth of 100 Hz, and 20-fold trace averaging.

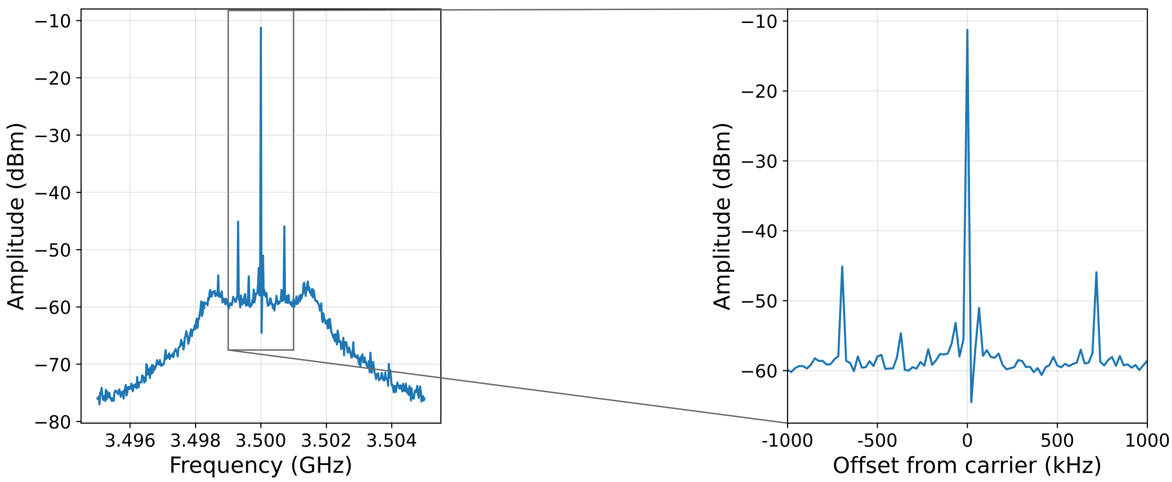

Figure 4.7: Full beatnote spectrum (left) showing a small sideband pair offset by roughly 711 kHz from the carrier, with a zoomed view of the marked region (right) showing the sidebands more clearly. Measured with a 10 MHz scan span, a resolution bandwidth of 30 Hz, and no trace averaging.

When looking at the spectrum of the locked beatnote, small sidebands around 711 kHz could be observed which did not move when the PID settings were changed, and which always remained at the same position.

By further investigating using noise-hunting antennas, it was found that they were caused by a PID rack located next to the setup, which was radiating a noise forest of narrow peaks ranging from 400 kHz to 40 MHz. However, only some of these peaks radiate into the system and end up as an actual modulation of the laser, and therefore as a sideband, as shown in Fig. 4.7. Since the peaks are very narrow, they have little area under the sideband, so the stability of the lock is not affected heavily. After turning off the PID rack, these sidebands also disappeared, confirming the cause.

4.4 Slow feedback

For optimizing the slow feedback, the frequency at which the slow feedback is effective was investigated. For this, a scan was applied to the locked laser, and it can be seen that if the slow feedback is not enabled, the fast feedback has to compensate for the entire scan. The slow feedback was then enabled with different slew rates, and it was observed how much of the scan is compensated by the fast feedback and how much by the slow feedback. This was repeated across the different scan frequencies that were available, mapping out the response, as shown in Fig. 4.8. The measurement was quite limited at low frequencies below 0.02 Hz, which are in principle the most important for long-term drifts, but as can be seen, for all slew rates at 0.02 Hz the slow feedback almost exclusively corrects for the applied scan, and it can be seen that they fail at different frequencies depending on the slew rate. In order to determine which slew rate is best suited, different slew rates were tried while changing the reference frequency of the beatnote over several steps, from a few MHz to multiple GHz, observing whether the lock successfully survived the jump. The slew rate of 258.98 V/V/s performed best. This unit, rather than the more typical V/s, arises from normalizing the output signal instead of using volts directly, so that the controller output is given by $u = V_{\text{out}}/V_{\text{max}}$. Using an interpolation of the points and finding where the signal decreases by 3 dB, the bandwidth of the slow feedback for this slew rate was estimated to be (2.20 ± 0.25) Hz, showing that it mostly acts on the timescale of subseconds to hours, correcting for long-term drifts, without affecting the performance of the fast feedback aspect of the lock.

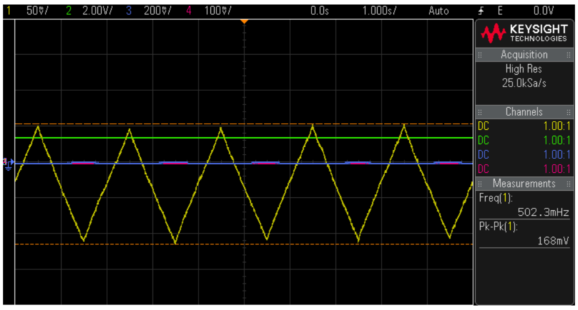

(a) Oscilloscope trace without slow feedback enabled, showing the fast feedback (yellow, 50 mV per division), the slow feedback (red, 100 mV per division), the FALC PRO error signal (blue, 200 mV per division), and the PLL box error signal (green, 2.00 V per division).

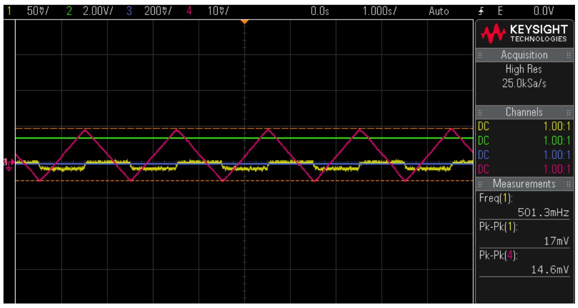

(b) Oscilloscope trace with slow feedback enabled, showing the same four signals as in (a), at the same scales.

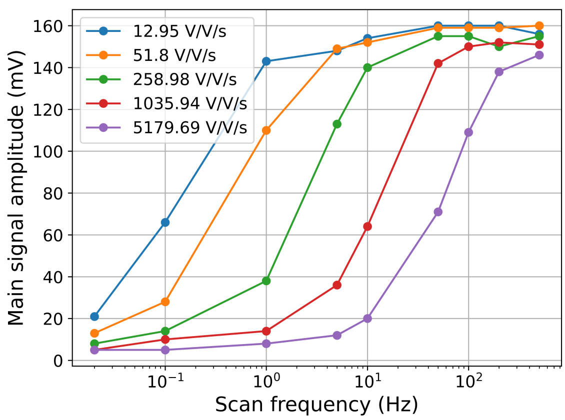

(c) Main signal amplitude.

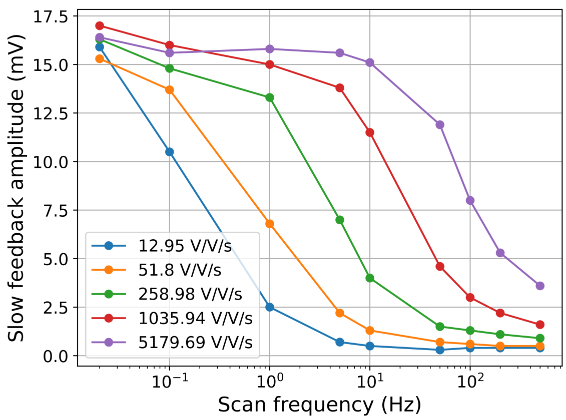

(d) Slow feedback output amplitude.

Figure 4.8: Response of the feedback loop to a scan applied to the locked laser, across different slew rates of the slow feedback. Top: oscilloscope traces of the scan without (left) and with (right) slow feedback enabled. Bottom: amplitude of the slow feedback output (right) and of the main signal (left), compensated by the fast feedback, as a function of scan frequency, for different slew rates.

4.5 Lock behavior at different lock points

When actually scanning the laser, it is not enough to check the lock behavior at a single point; it must also be ensured that the lock does not fail, or worsen significantly, at different positions throughout the scan. Additionally, the different jumps performed by the system were tested, in order to determine whether there is a maximum jump size that is possible, or whether the jump range has no impact on the lock at all.

To test the position dependence of the lock, the laser was first optimized so that no second beatnotes would appear when scanning it with the laser controller across the full 6 GHz range. It was then locked at a beatnote frequency of 3.5 GHz, and finally, using the reference frequency, the beatnote was changed to different values, resulting in beatnote frequencies between 800 MHz and 5.5 GHz.

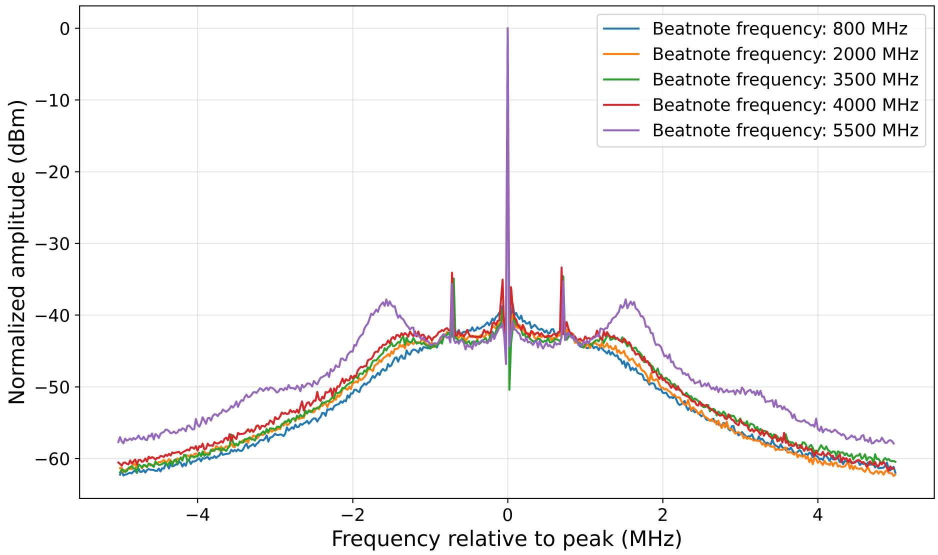

Figure 4.9: Cable-loss-corrected beatnote spectra at different lock points, normalized to 0 dBm at the peak and centered on the peak. The noise pedestal and the servo bumps grow noticeably at higher beatnote frequencies, most visibly at 5.5 GHz.

As can be seen in Fig. 4.9, the lock quality changes depending on position. This is because the response of the laser-current modulation is not independent of the position within the MHFR; this effect is stronger especially near the end of the MHFR, as can be seen at a beatnote frequency of 5.5 GHz. It is therefore possible that the changes in current have more or less of an effect on the laser frequency, which can be viewed as similar to changing the settings of the PID, resulting in a different lock quality. This effect was minimized by using an MHFR larger than the 6 GHz scanning range, but as can be seen, it is still noticeable, and the lock quality worsens especially at high beatnote frequencies.

In order to check whether the jump range of the beatnote frequency has any effect, or whether there is some limit to the range different jumps can span, different jumps were tested to see whether the lock remains locked, or whether its behavior changes. First, the maximum distance was tested: a jump from 800 MHz to 5.5 GHz, and the lock remained locked in either direction.

It was then tested whether the starting position from which the lock jumps has an impact on the lock quality itself. Different lock positions were therefore tested by going from the starting position at 3.5 GHz to some intermediate point in the lock, and then, by changing the reference frequency once again, jumping to a beatnote frequency of 2 000 MHz. As can be seen in Fig. 4.10, the lock is not impacted by the jump size or starting position and always locks the same way. This is exactly what was expected, and makes the lock very resilient once the scan is later actually applied.

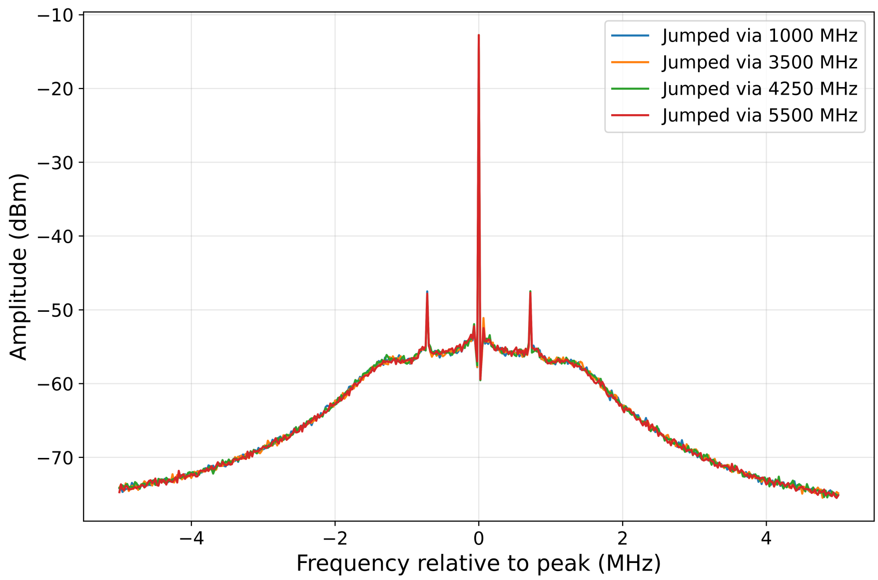

Figure 4.10: Beatnote spectra at a beatnote frequency of 2 000 MHz, reached via different intermediate jump points starting from 3.5 GHz. The traces are indistinguishable, showing that the lock is unaffected by the size or starting position of the jump.

4.5.1 Sidebands appearing at different lock points

It was also observed that sidebands appear when locking the laser in the low-GHz range, and after further investigation it was found that these sidebands are positioned exactly at half of the input reference frequency. Recalling from the beginning of this chapter that the reference frequency is 32 times smaller than the beatnote frequency, at a beatnote frequency of 800 MHz this comes out to a reference frequency of 25 MHz; in Fig. 4.11, the sidebands can indeed be seen at 12.5 MHz (half the reference frequency) and 25 MHz (the reference frequency itself) from the carrier.

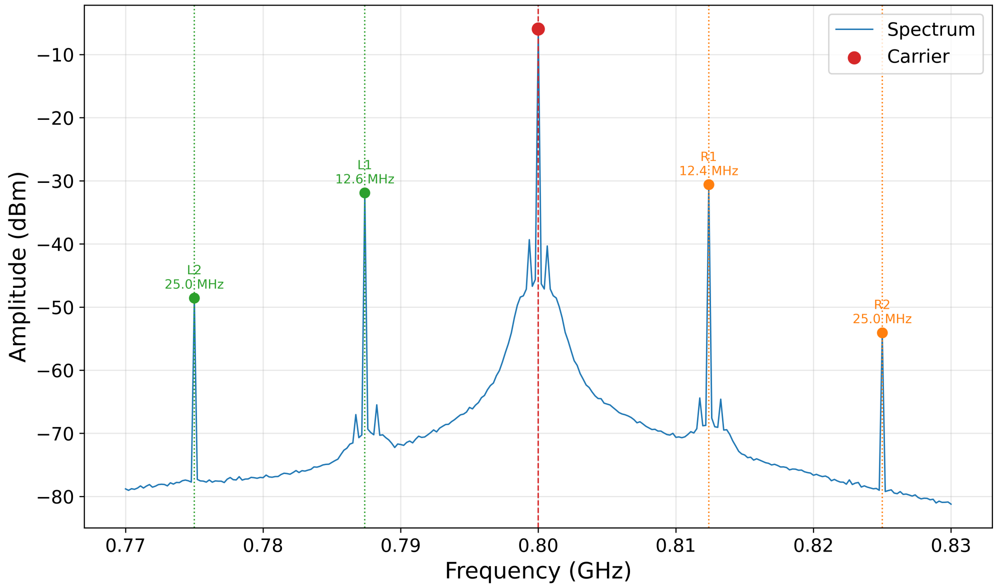

Figure 4.11: Carrier and sidebands of a beatnote locked at 800 MHz, with the two closest sidebands on each side labeled by their offset from the carrier. Measured with a resolution bandwidth of 100 Hz and 10-fold trace averaging.

This effect remains present when choosing larger beatnote frequencies, but as the bandwidth of the lock decreases, it is not nearly as pronounced as seen here. Figure 4.12 shows that the sideband amplitude, relative to the carrier peak, decreases fairly consistently with increasing beatnote frequency, with something of an outlier at a beatnote frequency of 1.5 GHz, which may coincidentally be a frequency that is especially strongly suppressed, or may be due to other unforeseen effects.

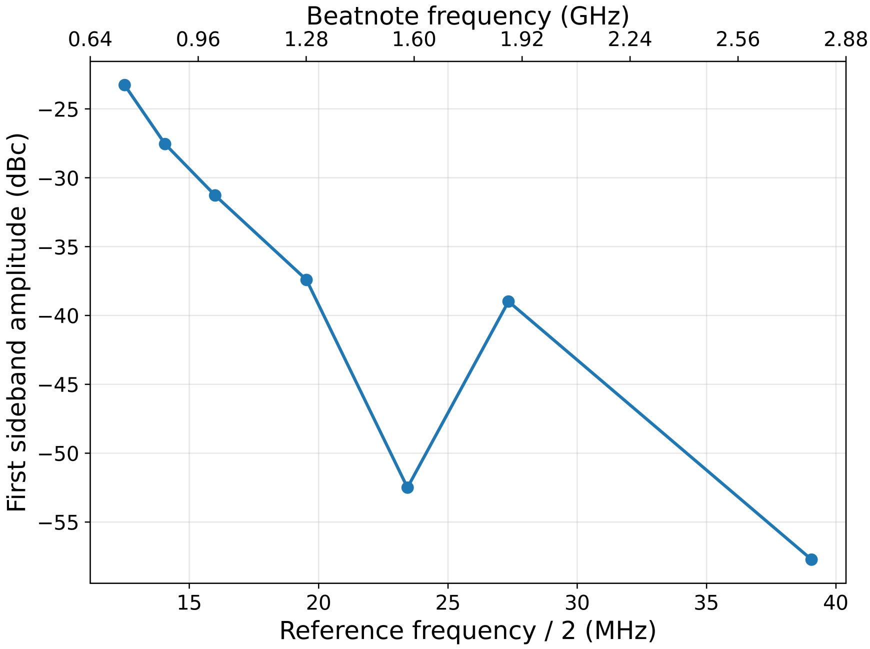

Figure 4.12: Amplitude of the first sideband relative to the carrier, as a function of the reference frequency (bottom axis) and the corresponding beatnote frequency (top axis).

This effect is due to the design of the evaluation board itself, which uses a charge pump to generate the error signal, shown in Fig. 4.1 [11]; the charge pump outputs the signal as a pulse train containing components at half the reference frequency, as well as multiple harmonics of that frequency, as can be seen in Fig. 4.13.

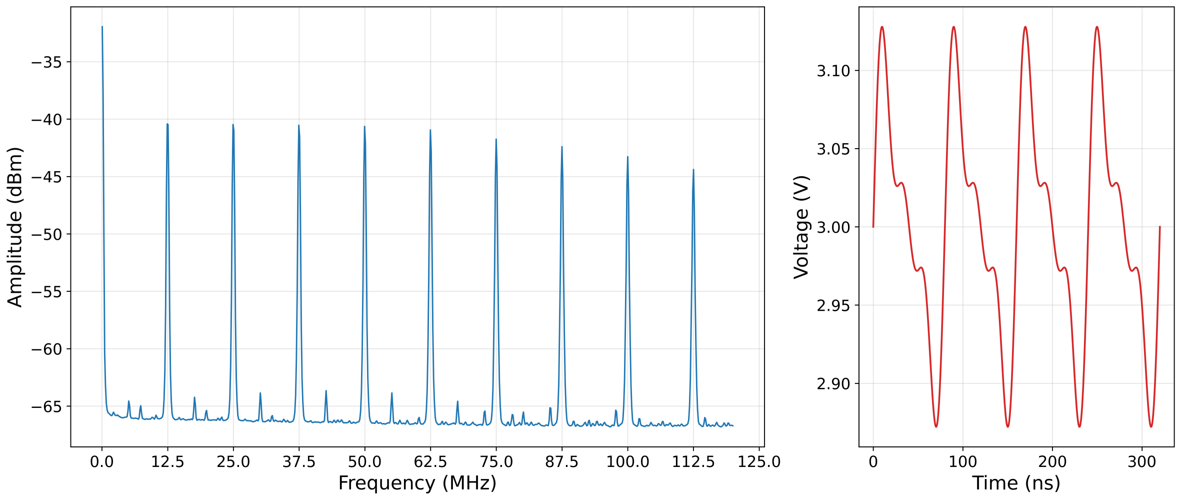

Figure 4.13: Left: spectrum of the PLL error signal while unlocked, at a beatnote frequency of 800 MHz, showing harmonics spaced by 12.5 MHz – half the reference frequency. Right: illustrative sketch (not measured) of the corresponding error-signal output in the time domain: a constant 3 V level with ripple at these harmonics superimposed.

This issue with the sidebands was not observed in previous iterations of the setup because, after the error signal output, an op-amp was used so that it could be connected via a 50 Ω impedance to most PID controllers. This op-amp also acted as a low-pass filter, which suppressed these harmonics; additionally, a lower value of N was chosen, resulting in an overall larger reference frequency, at which the issue is not as pronounced. The FALC PRO does not need a 50 Ω impedance connection, so the op-amp is not needed. A possible solution will be discussed in Section 6.1, but it was not implemented within the time frame of this thesis.

4.6 Lock failure at low optical power

What could not be fully tested is whether the optical power could be reduced without compromising the function of the locking system itself. Some tests were conducted, which can be seen in Fig. 4.14, showing the beatnote amplitude of the locked laser and whether it was possible to lock the laser using a similar approach to the one described in Section 4.5; the same test was then repeated with lower power. The total optical power was split between the slave and master laser, with the master laser fixed at 0.5 mW; the 4 mW, 2 mW, and 1 mW labels therefore correspond to 3.5 mW, 1.5 mW, and 0.5 mW for the slave laser, respectively.

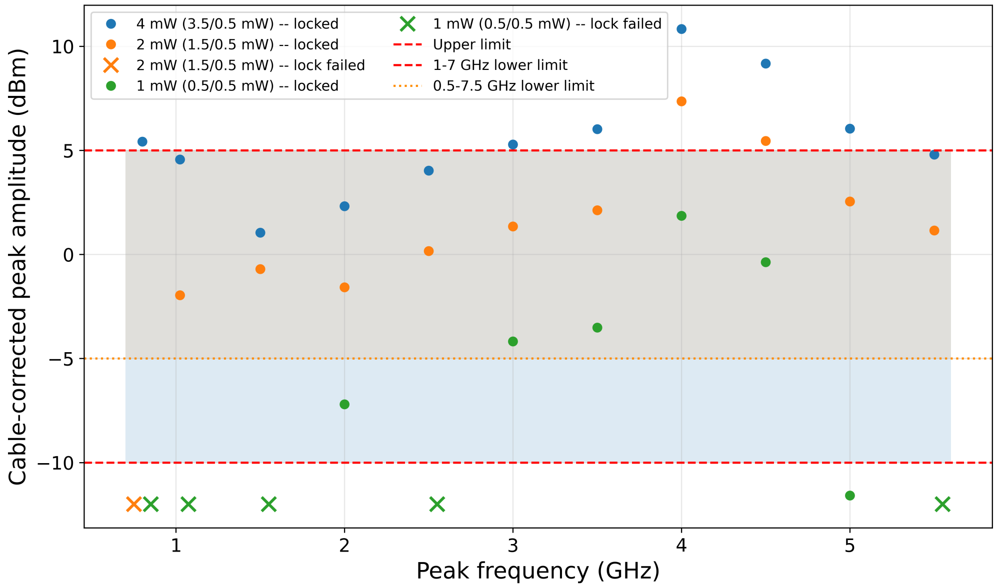

Figure 4.14: Cable-corrected peak amplitude of the beatnote as a function of frequency, for the three optical power splits tested. Filled circles mark frequencies where the laser could be locked; crosses mark frequencies where locking failed, plotted at a fixed amplitude for clarity. The dashed and dotted lines show the evaluation board's upper and lower input-level limits [11].

What is missing from this investigation is the reason for the lock failure, and whether it was a signal-to-noise issue or whether it could be resolved by further amplifying the photodiode signal with an additional later-stage amplifier.

For now, the optical power cannot be significantly reduced without compromising the achievable scan range. However, the power level used in this setup is already low enough to be implemented in the main experiment, leaving sufficient optical power for the master laser to be able to lock at least four lasers in the near-UV. Further details on this are given in Section 6.2.

CHAPTER 5

Performance of large frequency range scan

5.1 Scan parameters

Before looking at the scan itself, it is necessary to understand which parameters can be chosen for a scan, and what kind of scan is actually being performed. As parameters for the scan, the start point as well as the end point and the step size can be chosen; here, the given value for the scan is the beatnote frequency, which in turn is used to calculate the needed reference frequency. The scan that is performed is a step-by-step scan; if the end point would be exceeded, the scan starts again from the beginning.

The time that each step takes cannot be directly chosen; it is instead given by the control system, and heavily depends on whether other experimental parameters are being scanned alongside, but in general the steps take at least a few seconds. Depending on the measurement, the whole scan can take anywhere from several minutes up to hours.

As an example, Fig. 5.1 shows what a scan looks like when observing the signals monitored by the oscilloscope in the setup shown in Fig. 4.2: the fast feedback, the slow feedback, and the FALC PRO's own error signal. It can be seen that, for the different lock points along the scan, there are different but stable constellations between the settings of the fast and slow feedback, and that these appear to be reproducible, as can be seen when the scan repeats itself. For the larger jumps, it takes somewhat longer for the lock to recover.

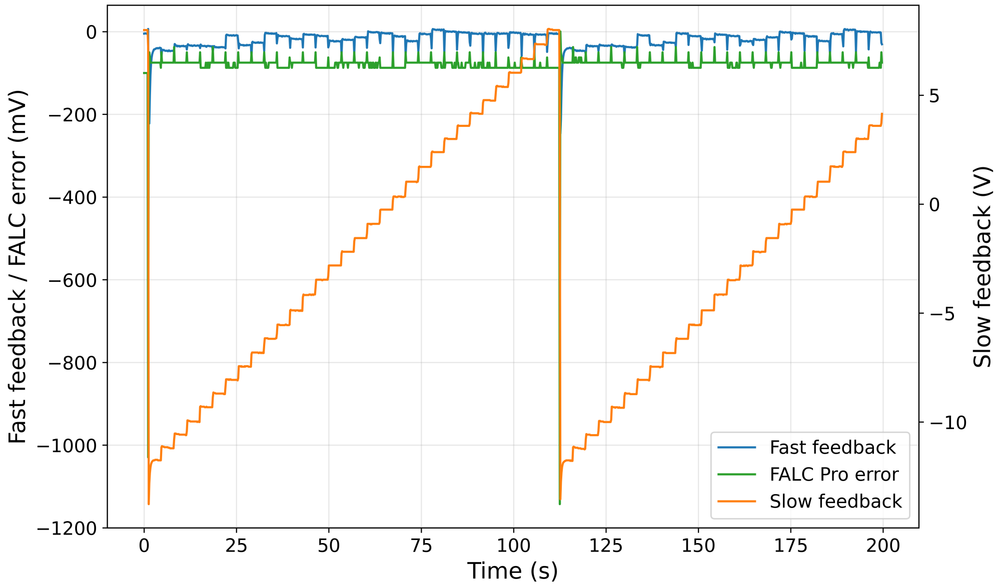

Figure 5.1: Oscilloscope traces of the fast feedback and the FALC PRO error signal (left axis, in millivolts) and the slow feedback (right axis, in volts) during a scan from 800 MHz to 5.5 GHz.

5.2 Demonstration of a 5 GHz scan

To demonstrate the performance of the scan across its full range, the beatnote frequency was repeatedly scanned from 800 MHz to 5 750 MHz in steps of 150 MHz, resulting in 34 iterations per scan, which can be seen on the bottom axis of Fig. 5.2. The top axis shows the time within a single scan, with one scan typically taking about 225 s. The right axis shows the number of completed scans, together with the corresponding elapsed time; over the course of this measurement, 1096 individual scans were completed over roughly 70 h.

The wavelength data used to compute the frequency difference $\Delta\nu$ between the master and slave lasers was measured with a wavemeter, whose accuracy is only 0.4 pm, corresponding to approximately 750 MHz at the 399 nm operating wavelength used here. Both the master and the slave laser were recorded, by continuous channel switching of the wavemeter, so that this absolute difference could be obtained, which is the quantity shown in Fig. 5.2.

As can be seen, the measured scan closely follows the ideal, linear staircase pattern shown at the bottom, and this behavior remains highly consistent and reproducible across the entire measurement.

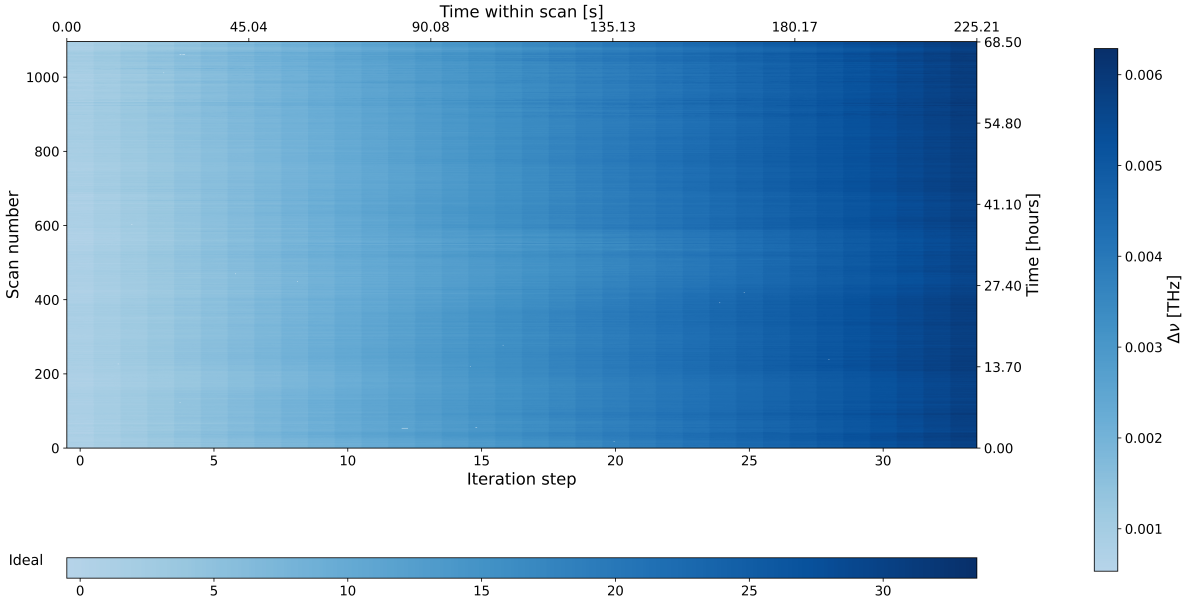

Figure 5.2: Reproducibility of the 800 MHz to 5 750 MHz beatnote scan, shown as a heatmap of the frequency difference $\Delta\nu$ between the master and slave lasers across 1096 repeated scans (left and right axes), as a function of the iteration step within each scan (bottom axis) and the corresponding time within a scan (top axis). Below, the ideal, linear staircase pattern of the intended scan.

To investigate the origin of the largest deviations from the ideal scan seen in Fig. 5.2, the individual scans were ranked by their root-mean-square deviation from the intended step pattern. Figure 5.3 compares one of the best-performing scans to two of the worst-performing scans, alongside the behavior of the free-running master laser during the same time window. It can be seen that the master laser remained comparatively stable during the good scan, while it drifted by several hundred megahertz during the two worst scans; this indicates that the largest deviations from the ideal scan are caused by drifts of the free-running master laser, rather than by a failure of the lock itself.

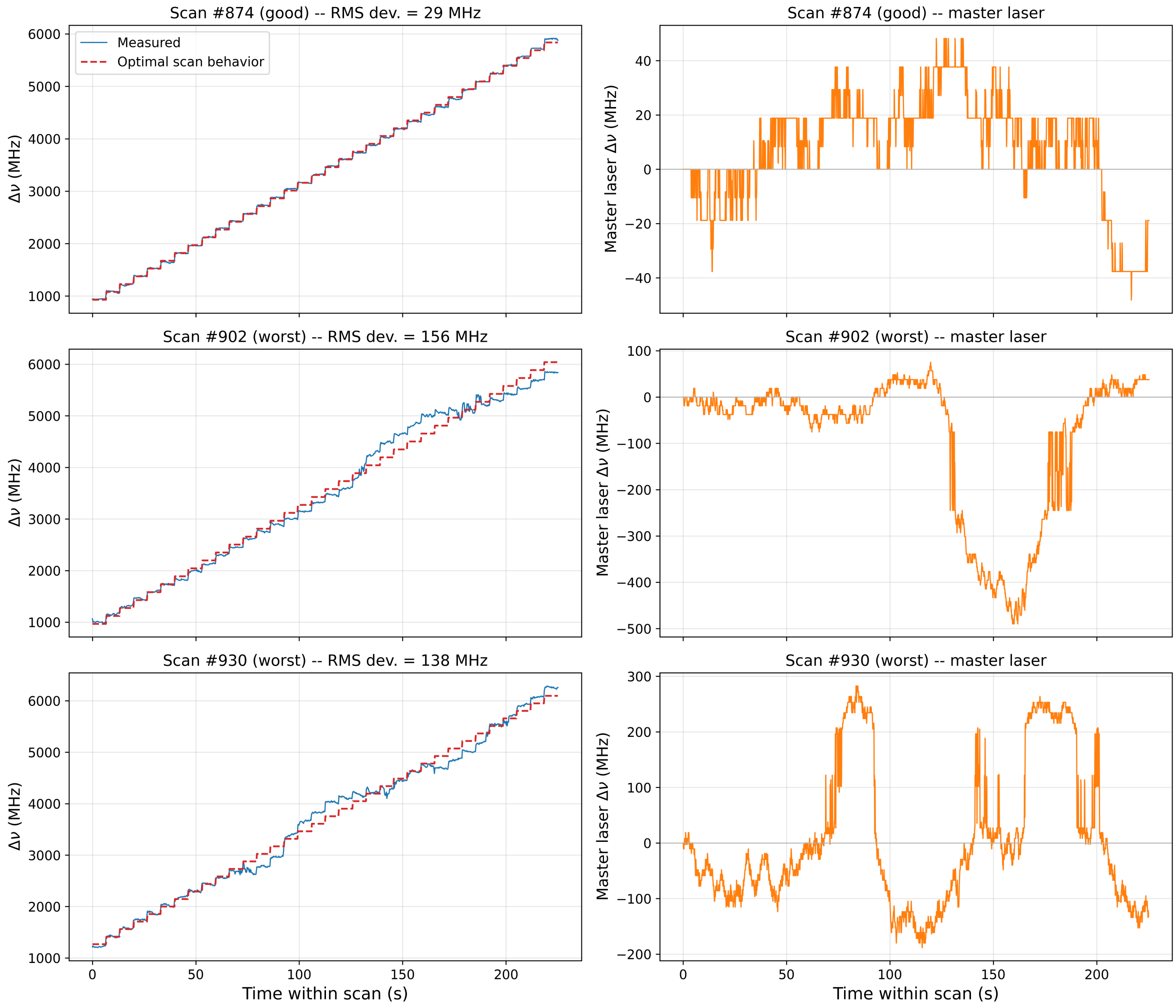

Figure 5.3: Left: measured beatnote frequency difference $\Delta\nu$ (blue) compared to the optimal, step-wise scan behavior (red dashed), for one of the best-performing scans (top, #874) and two of the worst-performing scans (#902 and #930), ranked by their RMS deviation from the optimal scan. Right: the corresponding behavior of the free-running master laser during the same scans, relative to its value at the start of each scan.

After careful consideration, this is most likely, at least in part, a measurement artifact caused by the channel switching between the different lasers on the wavemeter and its integration time: any change in wavelength on a faster timescale than the integration time ends up appearing as an artifact in these measurements. Since the wavemeter alternates between the two lasers, each individual laser is only resampled every 93.43 ms on average, so any change faster than 10.70 Hz ends up appearing as an artifact. This general effect is also exclusive to the test setup, since the master laser in the main experiment is itself stabilized and not free-running. This is still a good indication of the lock working, but it should be seen more as a proof of concept than as a precise measurement, since the wavemeter is far too imprecise compared to the lock itself; the laser not falling out of lock during this time is nevertheless a good indication of a stable lock.

CHAPTER 6

Conclusion

This chapter first gives a short summary of the current limitations of the setup and where they originate from, followed by a short discussion of how they could be addressed and improved. This is followed by a discussion of how the main experiment would need to change in order to optimally use the new frequency-stabilization range of the probe laser, once the PLL box is integrated into the main experiment. Finally, a short summary of this thesis is given.

6.1 Limitations and possible improvements

For now, the maximum achieved scan spans 4 950 MHz, from a minimum beatnote frequency of 800 MHz up to a maximum of 5 750 MHz. The reason the lock fails is most likely different at these two boundaries. At the lower boundary, the sidebands discussed in Section 4.5 appear to be the biggest issue, and are most likely the reason for the instability below 800 MHz. At the upper boundary, the limitation is most likely one of amplitude, since the gain of the RF chain drops significantly toward this end of the range, as shown in Fig. 3.7; this is caused by parasitic effects of small components that become significant at the connections between the individual components, as well as by the amplifiers not having a completely flat gain response.

Increasing the upper end of the scan range further is not easily done, and would require a completely new RF chain design that better compensates for these parasitic effects, by aiming for a positive gain toward the end of the chain beyond 4.5 GHz, which is currently the limit up to which this is achieved. The equalizers purchased for the current design have an almost linear response curve, which is not suitable for achieving this behavior, since it would reduce the gain in the lower-frequency part too strongly while only slightly improving the difference in gain between 4.5 GHz and higher frequencies, overall worsening the signal-to-noise ratio. This might therefore represent a soft cap for the current design, unless major changes are made by using different amplifiers and other RF chain components.

For the sideband issue, the easiest fix would be a low-pass filter that suppresses the harmonics of half the reference frequency originating from the evaluation board's design (Section 4.5). When choosing a suitable reference frequency for this filter, the relevant case is the lowest reference frequency still accepted by the evaluation board, 20 MHz (Section 4.1), since this produces the lowest-frequency harmonic, at 10 MHz. A low-pass filter should therefore be chosen that already strongly suppresses this frequency, while still not affecting the lock itself.

The easiest way to build such a low-pass filter is to use one designed for a low-impedance (50 Ω) signal, which can easily be added after an op-amp such as the one shown in Fig. 6.1. This op-amp converts the high-impedance error signal of the evaluation board into a 50 Ω signal, onto which an LC low-pass filter with a 5 MHz cutoff could be added.

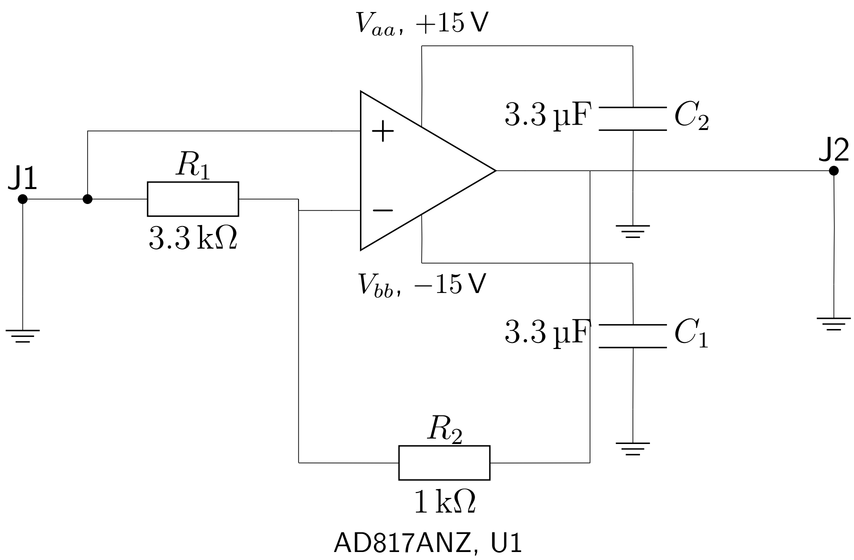

Figure 6.1: Fast output buffer used to convert the evaluation board's high-impedance error signal into a 50 Ω signal, provided by the YQO team [24]. This circuit was previously used with an older model of PID controller that required a low-impedance input; it is fast enough not to affect the lock quality, but its suppression alone is not strong enough to suppress the harmonics discussed above, since its bandwidth depends heavily on the inverting amplifier stage.

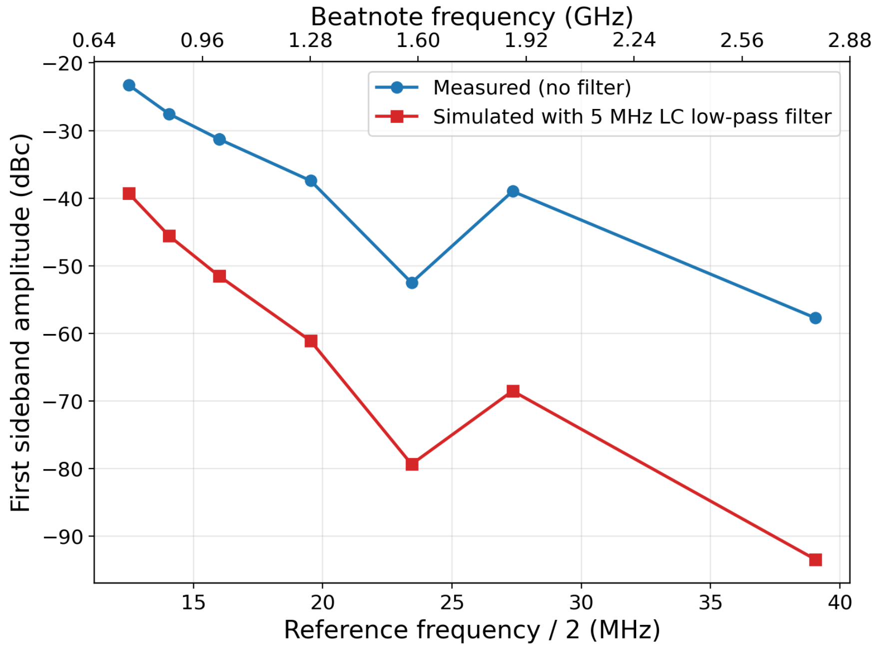

Figure 6.2: Simulated effect of adding a 5 MHz LC low-pass filter after the buffer shown in Fig. 6.1, based on the measured sideband amplitudes from Fig. 4.12: the filter's calculated attenuation at each harmonic frequency is added on top of the measured curve.

The resulting height of the sidebands can be seen in Fig. 6.2; although still present, their height would be reduced drastically, without affecting the lock at lower frequencies and therefore the lock performance itself. Further testing would, however, be necessary to confirm this. A good way to see whether they still affect the lock quality is to compare their height to that of the servo bumps, which lie at around 45 dB below the carrier; anything below that level should not have a significant impact.

The FALC PRO can accept both high-impedance and low-impedance input signals, selectable via a jumper behind the front plate. Suppressing these harmonics might also enable locking at even lower beatnote frequencies, down to the evaluation board's hard cap of 640 MHz (Section 4.1). This modification should therefore be implemented before the setup is used in the main experiment, since accidental excitation of unwanted atomic transitions by the sidebands should be avoided at all costs.

Reducing N should at least be discussed, since it is the underlying cause of several of the issues the lock has: with a lower N, $(\Delta\varphi)^2$ would become smaller, and the sideband issue would also be less pronounced, as mentioned in Section 4.5.1. However, the accuracy of the Arduino DDS would also become worse, since larger reference frequencies would then be needed, and ultimately the scan range would be limited to 3.84 GHz when choosing the next-smallest value, $N = 32$, which would severely reduce the scan range from roughly 5 GHz to only about 3 GHz. Since the lower $(\Delta\varphi)^2$ does not improve the features of the lock needed for the YQO experiment, and the sideband issue can instead be fixed using the methods already described above, reducing N – although worth considering – does not seem to be worth it.

6.2 Impact on the experiment

In the main experiment, a separate master laser operates at 798 nm and is locked to an ultrastable cavity; part of this light, together with its frequency-doubled output at 399 nm, is used to lock several of the other lasers in the experiment, shown in Fig. 6.3. Of relevance for the Probe laser are the $6s^2\,{}^1S_0$ ground state and the $6s6p\,{}^1P_1$ state, which acts as the intermediate step between the ground state and the targeted Rydberg state.

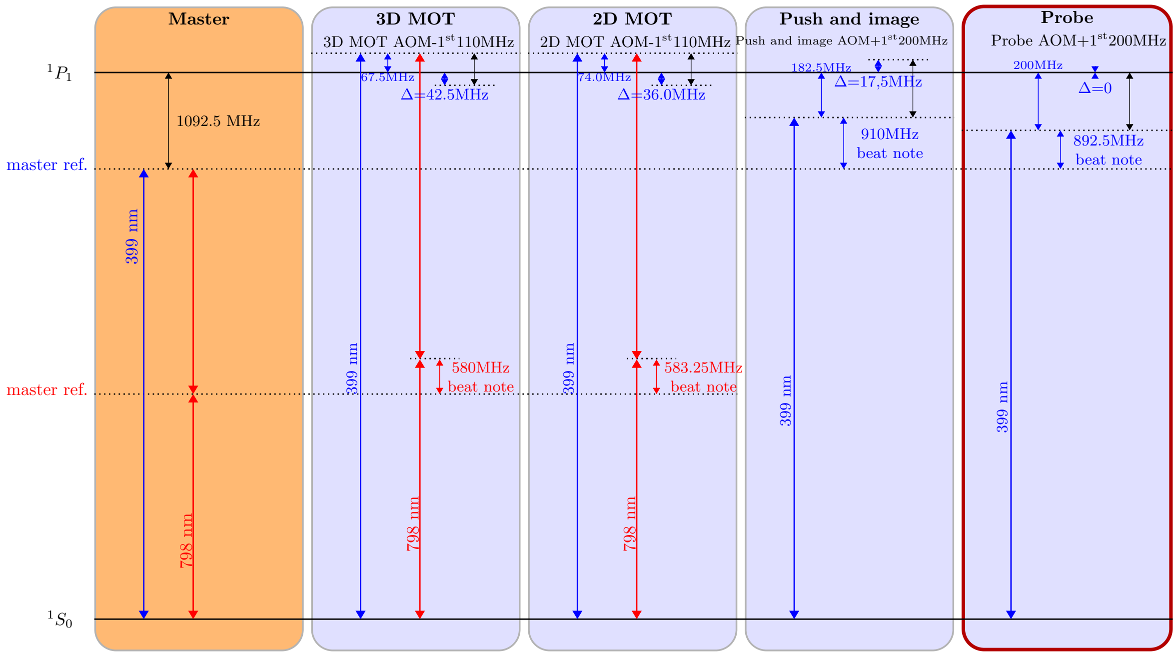

Figure 6.3: Recreation of the laser locking table of the YQO team [24], showing the current laser locking scheme used in the main experiment. The Probe panel, outlined in red, is the one changed by the new design developed in this thesis, shown in Fig. 6.4.

The acousto-optic modulators (AOMs) shown in Fig. 6.3 are used to precisely control the lasers, allowing quick switching between different frequencies or turning the laser off entirely, and therefore between different functions of a given laser. As long as the beatnote positions remain the same relative to one another, the AOMs themselves do not need to be changed; the diffraction order used by each AOM determines whether their shift frequency is added to or subtracted from the signal.

Currently, the Probe laser used for this transition is limited to a scan range of about 800 MHz, symmetrically around a beatnote of 892.5 MHz; this limit is caused by the PID's voltage limitation of the slow feedback already reaching its maximum, thereby limiting the piezo range. With the scan range achieved in this thesis increased to 4 950 MHz, the scan would no longer be symmetric around the $6s6p\,{}^1P_1$ transition; to correct this, the position of the master laser must be adjusted.

To achieve a symmetric scan, the middle of the new scan range must be found: this lies at 3 275 MHz, meaning the reference laser must be moved by 2 382.5 MHz. This is done by shifting the reference laser down by this frequency, to a new position of roughly 3 475 MHz below the $6s6p\,^1P_1$ transition. Consequently, the locks in the blue must also be shifted by 2 382.5 MHz; due to the frequency-halving from the wavelength-doubling stage, the corresponding locks in the infrared chain only need to shift by half of this amount, 1 191.25 MHz. This results in the new laser setup shown in Fig. 6.4.

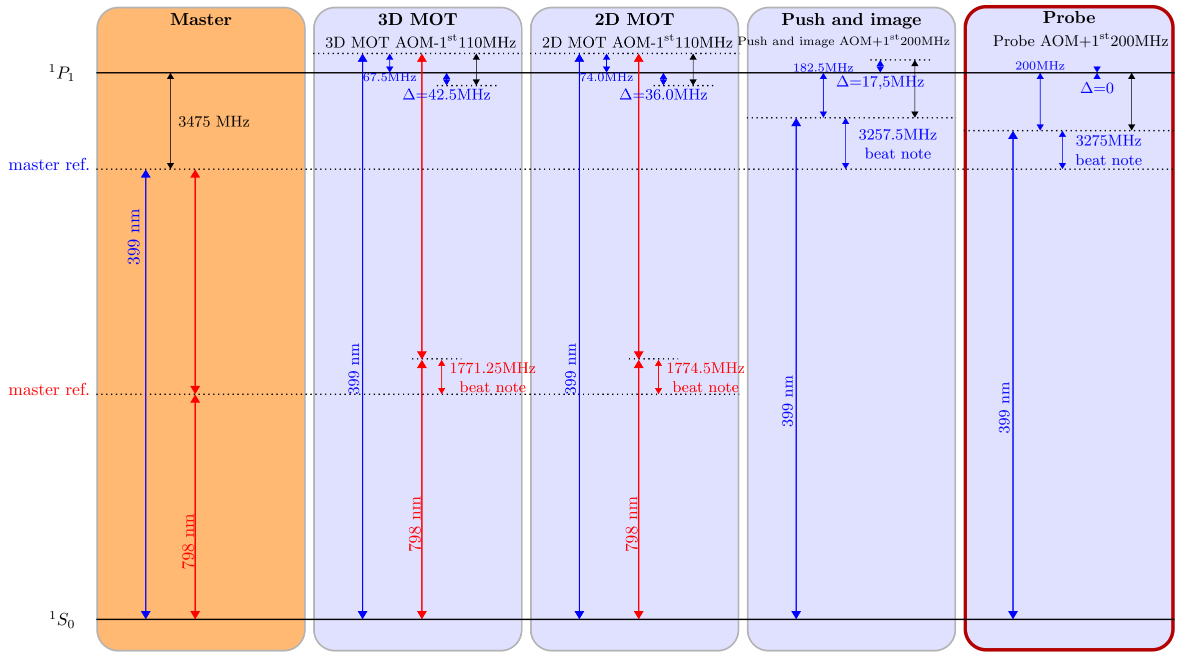

Figure 6.4: The same laser locking table as Fig. 6.3, with the small changes needed to introduce the new Probe laser developed in this thesis, again outlined in red.

These new beatnotes in the 798 nm light, landing at 1 771.25 MHz and 1 774.5 MHz, should still be no problem for the existing infrared setup, especially since that system is not power-limited; this allows the intensity of the infrared laser to be increased, and therefore the beatnote amplitude to a sufficient level, without changing the amplification of its RF chain.

For the Push and image (blue, 399 nm) laser, a locking setup capable of locking in the 3 257.5 to 3 275 MHz range is instead needed. The current setup developed in this thesis is sufficient as it is to maintain a reliable lock. However, since the beatnote of this laser does not need the same large variability, this might not be necessary, and adding a low-noise amplifier to the previous design, which is currently in use, might be enough. Simply increasing the optical power in order to get a higher beatnote amplitude is not an option, since the power of the master laser is limited to 10 mW. With the setup proposed here, roughly 1 mW would be used; in reality, planning for 1 mW per lock for the master laser seems more reasonable. Compared to the previously used 1.5 to 2.5 mW per lock, this reduces the optical power needed from the master laser, allowing the two lasers used in this thesis to be integrated into the main experiment without reaching the limit of the given power budget.

6.3 Summary

In conclusion, the goal of developing a multi-GHz scanning range for the near-UV, frequency-stabilized probe laser could be achieved by first analyzing the previous design [9] and finding that the limitation was caused by the reduced gain across the frequency range from 0 to 6 GHz, resulting from the amplifiers' behavior and the loss of the cable at higher frequencies. This issue was then resolved by using different amplifiers and cables, which flattened the overall gain curve across that range to within 12 dB. The next part of the thesis was spent on setting up the feedback loop and optimizing the parameters of the FALC PRO. The lock behavior was then systematically analyzed at different lock points, and the reasons for this behavior were investigated. Lastly, the ability to scan over 4 950 MHz was tested in long-term use and shown to work over 68 h, using only a fraction of the master laser's available optical power, demonstrating that the goal of the thesis was achieved. Some aspects of the lock, its quality, and its variability – especially regarding the optical power – remain open for further investigation, with the results obtained here already providing a rough outline of further possibilities. Finally, the integration and the effect of the upgraded lock capabilities on the main experiment were formulated.

APPENDIX A

Additional information

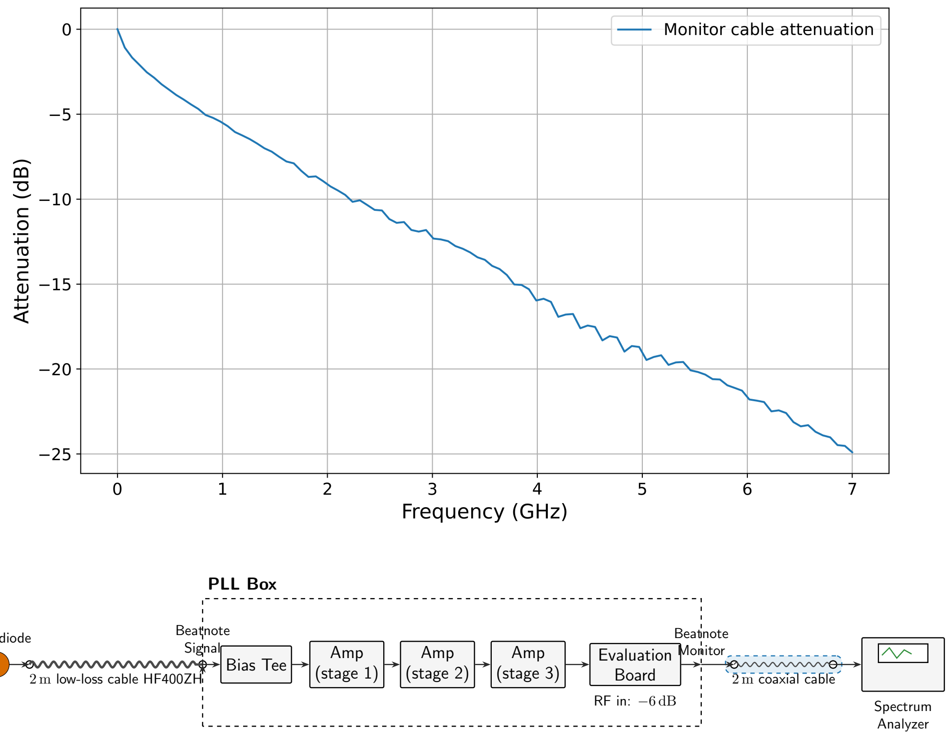

Figure A.1: Attenuation of the monitoring output cable, measured with a vector network analyzer. Below, the corresponding stage of the RF chain, with the monitoring output cable highlighted in blue.

Component	Part number	Reference
Evaluation board	EVAL-ADF4007EBZ1	[10]
PLL/divider chip	ADF4007	[11]
Photodiode	Hamamatsu G4176-03	[14]
Input cable (low-loss)	HF400ZH	[18]
Bias tee	ZFBT-6G+	[13]
First-stage amplifier	ZX60-14LN-S+	[15]
Second-stage amplifier	ZX60-V62+	[16]
Third-stage amplifier	ZX60-V62+	[16]

Table A.1: Parts list for the final PLL box design (Fig. 3.4).

Device	Manufacturer / Model
Spectrum analyzer	Keysight N9322C
Wavemeter	HighFinesse WS6/200
Oscilloscope	Keysight DSO-X 2014A
Vector network analyzer	Rohde & Schwarz ZNB20 (20 GHz)

Table A.2: Measurement devices used throughout this thesis.

Component	Manufacturer / Model
Master laser	Toptica DL Pro HP [6]
Slave laser	Toptica DL Pro HP [6]
Photodiode	Hamamatsu G4176-03 [14]

Table A.3: Optics used throughout this thesis.

Tool	Purpose
Claude	Spelling and grammar correction, as well as some LaTeX formatting
Python (Jupyter Notebook)	Data analysis and visualization
Inkscape and LaTeX	Creation of images
ChatGPT	Coding and debugging of code

Table A.4: Software tools used throughout this thesis.

Bibliography

[1] T. H. Maiman, *Stimulated Optical Radiation in Ruby*, Nature **187** (1960) 493 (cit. on p. 1).

[2] *Ytterbium Rydberg Quantum Optics*, Nonlinear Quantum Optics group, University of Bonn, URL: https://www.iap.uni-bonn.de/nqo/en/research/ytterbium-rydberg-nqo (visited on 27/07/2026) (cit. on p. 1).

[3] T. Legrand et al., *Revealing electron-ytterbium interactions through Rydberg molecular spectroscopy*, Physical Review X (2026), In press, arXiv: 2512.20609 (cit. on p. 1).

[4] O. Firstenberg, C. S. Adams and S. Hofferberth, *Nonlinear quantum optics mediated by Rydberg interactions*, Journal of Physics B: Atomic, Molecular and Optical Physics **49** (2016) 152003 (cit. on p. 1).

[5] L. Ricci et al., *A compact grating-stabilized diode laser system for atomic physics*, Optics Communications **117** (1995) 541 (cit. on p. 3).

[6] *DL Pro HP: Data Sheets for Units 1 (Slave) and 2 (Master)*, Internal data sheet for the slave laser (Unit 1); the master laser (Unit 2) data sheet is available at https://wiki.nqo.uni-bonn.de/images/DL_Pro_HP_2_Data_sheet.pdf. Since both lasers were built according to YQO specification, no general public data sheet exists for this configuration; these documents are only accessible with NQO wiki access., Nonlinear Quantum Optics group, University of Bonn, URL: https://wiki.nqo.uni-bonn.de/images/DL_Pro_HP_1_Data_sheet.pdf (visited on 29/07/2026) (cit. on pp. 3, 5, 43).

[7] C. E. Wieman and L. Hollberg, *Using diode lasers for atomic physics*, Review of Scientific Instruments **62** (1991) 1 (cit. on p. 4).

[8] W. Demtröder, *Laser Spectroscopy 2: Experimental Techniques*, 5th ed., Chapter 7, “Heterodyne Spectroscopy”., Berlin, Heidelberg: Springer, 2015, ISBN: 978-3-662-44640-9 (cit. on p. 5).

[9] C. M. Witte,
Phase-offset frequency stabilization of a diode laser for ytterbium quantum optics experiments, Bachelor's thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2025,
URL: https://www.iap.uni-bonn.de/nqo/en/publications/publications-1/2025-clara-witte-bachelor.pdf (visited on 19/07/2026) (cit. on pp. 7, 10, 11, 18, 41).

[10] *UG-158: Evaluation Board for the ADF4007 PLL Frequency Synthesizer*,
Evaluation Board User Guide., Analog Devices, Inc., 2010,
URL: https://www.analog.com/media/en/technical-documentation/user-guides/UG-158.pdf (visited on 21/07/2026) (cit. on pp. 8, 13, 17, 43).

[11] *ADF4007: High Frequency Divider/PLL Synthesizer*, Datasheet, Rev. 0,
Analog Devices, Inc., 2004, URL: https://www.analog.com/media/en/technical-documentation/data-sheets/ADF4007.pdf (visited on 21/07/2026)
(cit. on pp. 8, 18, 29, 31, 43).

[12] *ZX60-6013E-S+: Coaxial Amplifier, 20 MHz to 6 GHz*, Mini-Circuits,
URL: https://www.minicircuits.com/pdfs/ZX60-6013E-S+.pdf (visited on 21/07/2026)
(cit. on p. 8).

[13] *ZFBT-6G+: Coaxial Bias-Tee, 10 to 6000 MHz*, Mini-Circuits,
URL: https://www.minicircuits.com/pdfs/ZFBT-6G+.pdf (visited on 14/07/2026)
(cit. on pp. 8, 10, 43).

[14] *G4176-03: Photodiode*, Hamamatsu Photonics K.K., URL:
https://www.datasheets360.com/pdf/-5073308062623184073 (visited on 21/07/2026)
(cit. on pp. 9, 43).

[15] *ZX60-14LN-S+: Wideband Low Noise Amplifier, 50 MHz to 10 GHz*, Mini-Circuits,
URL: https://www.minicircuits.com/pdfs/ZX60-14LN-S+.pdf (visited on 14/07/2026)
(cit. on pp. 10, 43).

[16] *ZX60-V62+: Coaxial Wideband Amplifier, 50 MHz to 6 GHz*, Mini-Circuits,
URL: https://www.minicircuits.com/pdfs/ZX60-V62+.pdf (visited on 14/07/2026)
(cit. on pp. 10, 43).

[17] *VEQY-10-63+: Microwave Gain Equalizer, DC to 6 GHz*, Mini-Circuits,
URL: https://www.minicircuits.com/pdfs/VEQY-10-63+.pdf (visited on 14/07/2026)
(cit. on p. 10).

[18] *HF 400 ZH: High Performance Broadband Low Loss 50 Ω Coaxial Cable, with SMA Connectors (SSB-Electronic Art. No. 7362)*, Cable datasheet; connector datasheet:
https://www.koax24.de/storage/datasheet/de/020451_Datenblatt_7362.pdf,
koax24 / SSB-Electronic GmbH, URL:
https://www.koax24.de/storage/datasheet/de/050229_Datenblatt_HF400ZH.pdf
(visited on 21/07/2026) (cit. on pp. 11, 43).

[19] *AD9959: 4-Channel, 500 MSPS DDS with 10-Bit DACs*, Datasheet, Rev. C, Analog Devices, Inc.,
URL: https://www.analog.com/media/en/technical-documentation/data-sheets/ad9959.pdf (visited on 28/07/2026) (cit. on p. 17).

[20] *ABLS Series: HC/49US SMD Low Profile Crystal*, Datasheet, Rev. 04.07.11, Abracon Corporation, URL: https://docs.rs-online.com/0861/0900766b80fb5b7a.pdf (visited on 29/07/2026) (cit. on p. 17).

[21] *FALC PRO: Fast Analog Linewidth Control*, TOPTICA Photonics AG, URL: https://www.toptica.com/products/locking-driving-electronics/falc-pro (visited on 21/07/2026) (cit. on p. 18).

[22] J. Appel, A. MacRae and A. I. Lvovsky, *A versatile digital GHz phase lock for external cavity diode lasers*, Measurement Science and Technology **20** (2009) 055302 (cit. on p. 22).

[23] R. M. Kroeze, S. L. Kristensen and S. Pucher, *^{171}Yb Reference Data*, Table 8, 2026, arXiv: 2509.04416 [physics.atom-ph] (cit. on p. 23).

[24] YQO team, *Internal Designs and Data from the YQO Team*, Internal, unpublished material provided by the YQO team, including the laser locking table of the main experiment and the fast output buffer circuit., Ytterbium Quantum Optics lab, Nonlinear Quantum Optics group, University of Bonn, 2026 (cit. on pp. 37, 39).