Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

GRAVITATIONAL WAVE DETECTION

Femtometer-precision displacement sensing via heterodyne cavity-tracking for use in gravitational-wave detectors

Shreevathsa Chalathadka Subrahmanya

2025

Femtometer-precision displacement sensing via heterodyne cavity-tracking for use in gravitational-wave detectors

Dissertation
zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Physik
der Universität Hamburg

vorgelegt von

Shreevathsa Chalathadka Subrahmanya

Hamburg
2025

Gutachter/innen der Dissertation:	Prof. Dr. Oliver Gerberding Apl. Prof. Gerhard Heinzel
Zusammensetzung der Prüfungskommission:	Prof. Dr. Oliver Gerberding Apl. Prof. Gerhard Heinzel Prof. Dr. Roman Schnabel Prof. Dr. Peter Schleper Prof. Dr. Ludwig Mathey
Vorsitzende/r der Prüfungskommission:	Prof. Dr. Ludwig Mathey
Datum der Disputation:	07.04.2025
Vorsitzender des Fach-Promotionsausschusses PHYSIK:	Prof. Dr. Wolfgang J. Parak
Leiter des Fachbereichs PHYSIK:	Prof. Dr. Markus Drescher
Dekan der Fakultät MIN:	Prof. Dr.-Ing. Norbert Ritter

Cover Photo: Concept by SCS; painted by Abinaya Omkumar; photographed by Christian Darsow-Fromm.

To my mother, Dr. Umadevi M.

Abstract

Displacement and inertial sensors are used to measure motion in many high-precision experiments. A prime example of such a use case is ground-based gravitational-wave detectors. To enable the detection of gravitational waves, the seismic isolation systems of these detectors employ various inertial and local displacement sensors to reduce the transmission of ground motion to the core optics of the interferometer.

This thesis investigates a displacement sensing technique, named heterodyne cavity-tracking, within the context of current and future ground-based gravitational-wave detectors, as well as future lunar-based concepts. Precise length measurements with femtometer or sub-femtometer readout noise at frequencies below 10 Hz would be extremely beneficial in increasing the overall sensitivity of ground-based gravitational-wave detectors, enabling more astrophysical observations. A specific case of a lunar-based gravitational-wave detection concept involves deploying a femtometer-class inertial sensor array on the Moon's surface to identify gravitational-wave signatures.

This thesis presents a heterodyne-stabilized optical cavity-based interferometer scheme that can serve as a compact, high-sensitivity displacement sensor with a fringe-scale operating range. The dynamic range for displacement readout is aimed at being increased by pushing the precision into the femtometer regime while not completely compromising the achievable operating range.

Specifically, this work studies two aspects of the interferometer scheme that are crucial for achieving the desired displacement readout precision of below 1 fm/$\sqrt{\text{Hz}}$ at around 1 Hz. The first aspect is the development of an ultra-fast, high-bandwidth readout system, the GHz Phasemeter, based on a field-programmable gate-array (FPGA) algorithm. The developed phasemeter is the highest bandwidth, highly stable, and fastest phase-tracking instrument reported to date. This implementation has been extended to include a Digital Heterodyne Controller (DHC), which is required for the laser-locking control loop. The second aspect involves the study of laser frequency references aimed at suppressing frequency noise in the readout. A potential frequency reference made of chalcogenide glass, which is 'athermal' in nature, has been investigated during the course of this work.

Finally, these components are combined to realize a displacement sensing scheme and have been analyzed within the scope of this thesis. An overall displacement readout noise floor of less than 20 fm/$\sqrt{\text{Hz}}$ has been achieved for frequencies above 5 Hz with the current experimental setup. A dynamic range of six orders of magnitude for displacement measurement has been demonstrated, reaching a maximum motion of about 0.15 µm. The work presented here was conducted between 2020 and 2024 at the Institute of Experimental Physics, University of Hamburg, Germany.

Keywords: field-programmable gate-array, frequency measurement, interferometric displacement sensing, laser frequency stabilization, phasemeter.

Kurzfassung

Weg- und Inertialsensoren werden in vielen hochpräzisen Experimenten zur Messung von Bewegungen eingesetzt, ein typisches Beispiel sind erdbasierte Gravitationswellendetektoren. Um die Detektion von Gravitationswellen zu ermöglichen, verwenden die seismischen Isolationssysteme dieser Detektoren verschiedene Inertial- und lokale Wegsensoren, welche die Übertragung von Bodenbewegungen auf die Kernoptik des Interferometers reduzieren.

In dieser Arbeit wird eine Technik zur Messung von Verschiebungen, das sogenannte „heterodyne cavity-tracking“, im Zusammenhang mit aktuellen und zukünftigen erdbasierten Gravitationswellendetektoren, sowie zukünftigen Konzepten auf dem Mond untersucht. Präzise Längenmessungen mit Femtometer- oder Subfemtometer-Ausleserauschen bei Frequenzen unter 10 Hz wären vorteilhaft, um die Sensitivität erdbasierter Gravitationswellendetektoren zu erhöhen und eine größere Anzahl astrophysikalischer Beobachtungen zu ermöglichen. Ein Konzept der Detektion auf dem Mond beinhaltet die Installation eines Systems aus Inertialsensoren mit Femtometergenauigkeit auf der Mondoberfläche, um Gravitationswellensignaturen zu identifizieren.

Diese Arbeit stellt ein, auf einem optischen Resonator basierendes, Interferometer mit Heterodynstabilisierung vor, welches als kompakter, hochempfindlicher Wegsensor mit einem Arbeitsbereich in der Größenordnung einer Wellenlänge eingesetzt werden kann. Das Ziel ist es, den dynamischen Bereich für das Auslesen der Verschiebung zu erhöhen, indem eine Genauigkeit im Femtometerbereich erreicht wird, während ein relativ großer Arbeitsbereich erhalten bleibt.

Insbesondere werden zwei Aspekte des Interferometerschemas untersucht, die entscheidend sind, um die gewünschte Auslesegenauigkeit von weniger als $1\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$ bei etwa 1 Hz zu erreichen. Der erste Aspekt ist die Entwicklung eines ultraschnellen Auslesesystems mit hoher Bandbreite, das GHz Phasenmeter, welches auf einem „field-programmable gate-array (FPGA)“-Algorithmus basiert. Das entwickelte Phasenmeter ist das bisher schnellste, stabilste und breitbandigste Phasenverfolgungsinstrument. Diese Implementierung wurde um einen „Digital Heterodyne Controller (DHC)“ erweitert, der für den Laser-Locking-Regelkreis benötigt wird. Der zweite Aspekt betrifft die Untersuchung von Laserfrequenzreferenzen zur Unterdrückung des Frequenzrauschens beim Auslesen. Eine mögliche Frequenzreferenz aus Chalkogenidglas, die „athermisch“ ist, wurde im Rahmen dieser Arbeit untersucht.

Schließlich werden diese Komponenten kombiniert, um ein Verschiebungsmessverfahren zu realisieren, das im Rahmen dieser Arbeit untersucht wurde. Mit dem aktuellen Versuchsaufbau konnte ein Gesamtausleserauschen unter $20\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$ für Frequenzen oberhalb von 5 Hz nachgewiesen werden. Es wurde ein dynamischer Bereich von sechs Größenordnungen für die Verschiebungsmessung erreicht, mit einer maximalen Bewegung von etwa 0,15 µm. Die hier vorgestellten Arbeiten wurden zwischen 2020 und 2024 am Institut für Experimentalphysik der Universität Hamburg durchgeführt.

Contents

CHAPTER 1

Gravitational-wave detection

More than a billion years ago, two black holes with masses of approximately $36\,M_\odot$ and $29\,M_\odot$ were trapped by gravity and began spiraling around each other. About one billion years ago, far, far away in the vast and endless space, at a luminosity distance of about 410 Mpc, this binary black hole (BBH) merged, releasing approximately $3\,M_\odot$ of energy in the form of gravitational waves. These waves then began their journey through the universe at the speed of light and entered the Milky Way roughly 100 000 years ago. On September 14, 2015, at 09:50:45 UTC, they briefly interacted with the Laser Interferometer Gravitational-Wave Observatory (LIGO) test masses before propagating onward undisturbed [1]. This brief interaction caused the test masses to move by about 10^{-18} m at around 100 Hz, which was loud enough for the LIGO detectors to 'listen' to!

This very first direct detection of gravitational waves, an event named GW150914, occurred just 100 years after Albert Einstein predicted the existence of gravitational waves [2]. The ability of directly detect gravitational waves has opened up gravitational-wave astronomy as a tool to reveal the mysteries of the universe and to better understand fundamental physics by granting us access to signals beyond the electromagnetic spectrum. The field has continued to grow since then, resulting in a total of four observing runs[1] to date, detecting about a hundred gravitational-wave events. We have stepped into a new era of multi-messenger astronomy, expanding our observable area of the universe.

The exceptional aspect of doing science is that once a proof-of-principle experiment is demonstrated, we do not stop. Instead, we push ourselves to improve the experiment and strive to understand the system better. The same happened in the case of gravitational-wave detectors. The first observing run (termed O1) of Advanced LIGO (aLIGO) detected three BBH mergers, whereas the second observing run resulted in eight more detections [3]. One of these events, GW170817, involved gravitational waves from the merging of binary neutron star (BNS). This became the first event subsequently observed in light by dozens of telescopes across the entire electromagnetic spectrum. This joint gravitational and electromagnetic observation probes insights into astrophysics, dense matter, gravitation, and cosmology [4].

The constant effort to increase the sensitivity of the detectors from one observing run to the next made it possible to drastically increase the number of detected events. By the end of the third observing run (O3), the detector network had observed a total of 90 confirmed events since LIGO/Virgo operations began [5].

[1] A period of observation in which gravitational-wave detectors are taking data.

The gravitational-wave strength is expressed using a dimensionless quantity called strain. It is the ratio between how much the distances are stretched or compressed because of the passing by gravitational-wave and the original distance (baseline). Mathematically, it is represented as

$$h = \frac{\Delta L}{L} \propto \frac{1}{r} m^{5/3} f^{2/3}. \tag{1.1}$$

Although strain increases as the distance from the source, r, decreases, it also depends on the effective mass of the system, m, and the frequency of the gravitational waves, f. As a result, objects on Earth do not produce enough strain for us to detect. Only more massive astronomical objects can create a strain significant enough to be detected on Earth.

The strain induced by gravitational waves as they pass through the Earth is typically several orders of magnitude smaller than 10^{-20}. For example, the event GW150914 had a peak strain of 10^{-21} [1]. Comparing this value to the diameter of a single proton, which is on the order of 1 fm, we realize the remarkable sensitivity of these gravitational-wave detectors.

1.1 Gravitational-wave detectors

A passing gravitational wave stretches and compresses space-time. Hence, the most straightforward way of detecting these waves is by dropping two objects, so-called test masses, in free space and then monitoring the distance variation between them. However, this can only be a thought experiment. In reality, various external interactions have influence on the test masses, preventing them from being in a 'free-fall' condition. Earth's gravity, temperature fluctuations in the detector's surroundings, and other noise contributions from the environment are examples of factors that prevent the test masses from being sensitive only to gravitation waves.

Nevertheless, an improved version of the Michelson interferometer has proven capable of directly detecting gravitational waves on the Earth's surface. The end mirrors in the interferometer arms serve as the two test masses from the previously mentioned thought experiment. For a specific signal direction and polarization, the gravitational wave compresses one arm while stretching the other. Generally, when a gravitational wave passes through the interferometer, it induces a non-zero differential arm length between the two interferometric arms. The corresponding differential phase modulation is then converted into an amplitude modulation by the interferometer, enabling the search for signatures of gravitational waves.

The aforementioned laser interferometric approach remains crucial and successful for the detection of signals from astrophysical sources. However, there are other proposals to detect both high-frequency and low-frequency gravitational waves that are not accessible with laser interferometric methods. Pulsar timing arrays currently probe the nanohertz to microhertz frequency band to detect gravitational-wave remnants from past mergers of supermassive black holes [6]. Approaches for high-frequency gravitational-wave detection with resonant microwave cavities are being investigated through small-scale laboratory experiments [7]. These experiments aim to be sensitive to gigahertz-frequency signals. Here, we restrict our discussion to laser interferometer-based detection.

1.1.1 Current detectors

Even though a single detector is sufficient to successfully detect gravitational waves, having a network of detectors spread around the globe makes it easier to pinpoint the origin of these waves. In other words, we improve the sky localization by having more than one detector that detects the same signal.[2] Having a network of detectors ensures that the transient is a real signal and not an artifact of the individual detectors.

At the moment, there are three ground-based detectors that can detect gravitational waves: LIGO Hanford Observatory (LHO) and LIGO Livingston Observatory (LLO) in the United States of America [8], and Virgo in Italy [9]. Because of the direct upgrades made to the initial gravitational-wave detectors, the currently running detectors are classified as second-generation detectors.

In the worldwide network of gravitational-wave detectors, there are two additional detectors currently in operation: GEO600 in Germany [10] and Kamioka Gravitational Wave Detector (KAGRA) in Japan [11]. GEO600 is limited in sensitivity because it has only a 600 m long arm length. This project currently aims to improve the required technology for gravitational-wave detectors and acts as a test-bed. KAGRA is the first underground gravitational-wave detector in the world that also uses cryogenic mirrors. Unfortunately, the current sensitivity of this detector is not high enough to contribute to successful detections.

All these detectors are not replicas of one another. Specific details, such as the length of the interferometer arm and how the test masses are suspended, vary from one detector to another. However, LHO and LLO are sometimes referred to as 'twin' detectors because of their very similar hardware usage. Nevertheless, the overall principle of the current detectors remains the same. Each one is a Michelson interferometer with Fabry-Pérot cavities in the arms [8].[3] Such a cavity helps to increase the laser power in the interferometer arms, enhancing the signal strength. There are two more high-reflective mirrors placed in the symmetric and antisymmetric ports of the Michelson interferometer, forming power recycling and signal recycling cavities, respectively. All these cavities are used to increase the signal-to-noise ratio (SNR) of the detector, making the infinitesimal impact of gravitational waves visible. In the full configuration with power and signal recycling and Fabry-Pérot cavities in the arms, the interferometer scheme is called a dual-recycled Fabry-Pérot Michelson interferometer (DRFPMI) (Fig. 1.1).

LHO and LLO have an arm length of 4 km, whereas Virgo and KAGRA have arms that are 3 km long. The exact length of the arms depends on the detector sites and feasibility. This combination of several kilometers of arm length and the suppression of all noise sources in the detectors to a level below $10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ [4] is what makes the detectors sensitive enough to detect gravitational-wave strains of the 10^{-21} regime in the frequency range of a few tens of hertz to kilohertz. Improved understanding of the noise sources and their mitigation has brought the current detectors to a relative sensitivity of 10^{-23} [5]. The prominent astronomical sources for the current ground-based detectors are the compact BBH inspirals.

[2] The localization of a source in the sky can be estimated using the arrival time differences of the gravitational-wave signal in the detector network.

[3] This statement does not apply to GEO600. It has folded arms instead of Fabry-Pérot cavities.

[4] An explanation of this type of unit is provided in Appendix A.

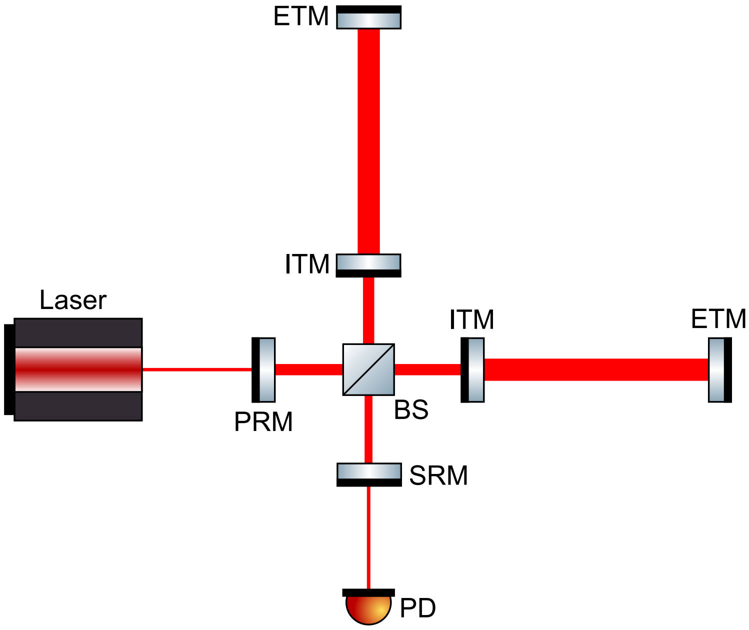

Figure 1.1: Simplified interferometer topology used in second-generation gravitational-wave detectors. This configuration is known as DRFPMI. Input mode cleaners, squeezers, filter cavities, etc., as used in the actual detectors, are not shown here. (PRM: power recycling mirror, SRM: signal recycling mirror, BS: beam-splitter, ITM: input test mass, ETM: end test mass, PD: photodetector. This and other optical setup sketches in this thesis were created using the ComponentLibrary by Alexander Franzen: `https://www.gwoptics.org/ComponentLibrary/`; accessed on January 29, 2025.)

1.1.2 Future ground-based detectors

In the near future, a gravitational-wave detector will be built in India. This LIGO Aundha Observatory (LAO) will be a nearly identical version of the two LIGO detectors. Once operational, this detector will significantly improve the sky localization of sources; and because of its location on Earth relative to other second-generation detectors, it will also help fill in blind spots[5] in the current detector network [12].

There are also proposals to upgrade the aLIGO detectors to further increase sensitivity, resulting in 'LIGO A#.' This upgrade is planned for 2029, after the observing run O5. LIGO A# will include heavier test masses, increased power in the arm cavities, better optical coatings, improved squeezing efficiency, and so on. Another proposal along similar lines is under the name 'LIGO Voyager.' This proposal explores changing to a longer laser wavelength, employing heavier silicon test masses, and other enhancements in this regard [13]. Additionally, there are preliminary plans for 'Virgo_nEXT,' the possible evolution of Virgo beyond the advanced Virgo+ project [14].

While the second-generation ground-based detectors are operating and detecting gravitational waves, several future ground-based detectors are already planned and being worked on. These future detectors will have significantly higher sensitivity, as well as a broader detection band, and are classified as

[5] Each detector has four blind spots, all located in the plane of the arms. Signals passing through these specific points will go undetected.

third-generation detectors. The Einstein Telescope (ET) [15], led by the European community, and Cosmic Explorer (CE) [16, 17], led by the United States, are the two main candidates at present.

Einstein Telescope (ET)

The ET is planned to be built in the next decade. The interferometer layout and detector site are currently under discussion. This detector is expected to exceed the sensitivity of the current detectors by a large margin. For example, at 3 Hz, ET is anticipated to be seven orders of magnitude more sensitive than the aLIGO detectors.

As the ET design is still in the planning stages, nothing can be stated conclusively at this time. However, several features are common among some of the current design ideas. The detector will be located underground, which helps reduce the contribution of seismic noise and isolates it from Newtonian noise sources arising from fluctuations in the local Newtonian gravitational filed. The ET is planned to have two combined interferometers, optimized for different frequency bands. One will be a low-power, cryogenic interferometer for low-frequency gravitational waves, while the other, operating at room temperature with high circulating laser power, will aim for high-frequency signals.

Cosmic Explorer (CE)

The design concept for CE includes two detectors: one with 40 km long arms and the other with 20 km. This increased arm length significantly improves the sensitivity of the detector. The amplitude of the observed signals will increase with effectively no corresponding increase in noise [16]. The interferometer layout, being the traditional Michelson topology on the surface of the Earth, provides the added advantage of utilizing existing expertise with ground-based detectors.

Upgrading the current detectors from one observing run to another and building next-generation detectors are critically important. Such upgrades are intended to enhance both the sensitivity of the detectors and expand the observational frequency band. This will result in an increased number of observable sources, a significant increase in the detection rate, and more frequent multi-messenger astronomy. Fig. 1.2 illustrates this graphically.

1.1.3 Other planned missions

The sensitivity of current and future ground-based detectors is limited to frequencies above 1 Hz. Below this frequency, seismic and Newtonian noise couplings to the test masses make it nearly impossible to maintain the free-fall condition, even with state-of-the-art suspension and control systems. From Fig. 1.2, it is evident that detectors sensitive below the audio-frequency band are of high importance for detecting signals from massive binaries or extreme mass ratio inspirals (EMRIs). Consequently, to achieve this goal, the scientific community is planning for non-terrestrial detectors. Such initiatives include planned missions in space and proposed gravitational-wave detector concepts on the Moon.

Laser Interferometer Space Antenna (LISA)

The LISA is a space-borne gravitational-wave observatory consisting of three spacecraft in heliocentric orbits [19]. The space environment offers direct benefits, including the elimination of terrestrial seismic noise, an excellent vacuum along the interferometer arms, and the freedom to choose sufficiently long

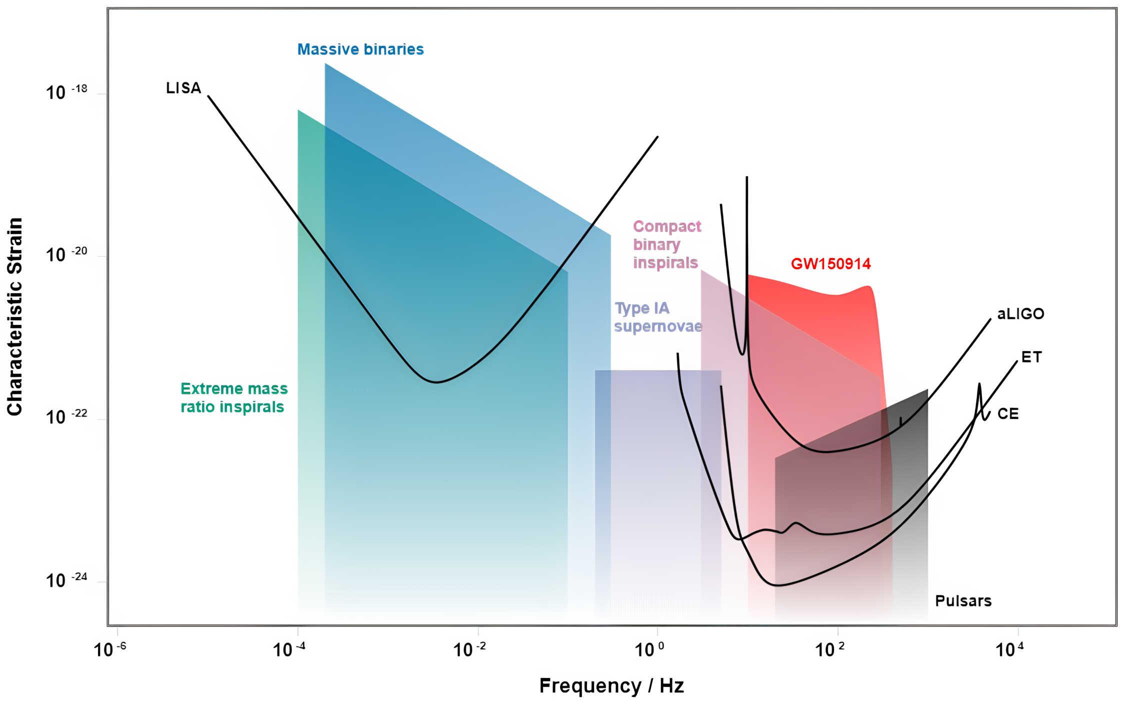

Figure 1.2: Gravitational-wave strain of different astronomical sources and sensitivities of current and future detectors. Strain of the first detection, the event GW150914, is shown as well. (Created using the Gravitational Wave Sensitivity Curve Plotter: `https://gilsay.physics.gla.ac.uk/gwplotter/` [18]; accessed on January 29, 2025.)

arm lengths to be sensitive at low frequencies. After numerous iterations on the design concept, the latest version features spacecraft separated by 2.5×10^6 km in a triangular formation, following Earth at a distance of approximately 50×10^6 km in its orbit around the Sun. The launch is expected in the late 2030s. With LISA, we will be able to observe the entire universe directly using gravitational waves. LISA is anticipated to be sensitive enough to detect the gravitational-wave background emitted at the beginning of the universe. The detection of intermediate to massive black holes will enable us to explore the low-mass end of the massive black hole population, at cosmic times as early as redshifts of about 10 [20].

A dedicated technology validation mission for future space-based gravitational-wave detectors has already been carried out with the LISA Pathfinder mission.[6] The results demonstrated a relative acceleration noise that not only exceeded the requirements of the precursor mission but also met the standards set for the LISA mission [21], confirming a bright future for space-based gravitational-wave detectors.

Lunar Gravitational-Wave Antenna (LGWA)

This proposal aims to bridge the gap between terrestrial and space-based detectors. With the LGWA, an array of high-end seismometers will be deployed on the Moon to monitor its normal modes in the

[6] Launched on December 3, 2015, and concluded on June 30, 2017.

frequency band from 1 mHz to 1 Hz [22]. This would enable the detection of heavier BBH inspirals compared to ground-based detectors and provide early warnings for compact BBH and BNS systems, among other capabilities. The LGWA mission is still in its very early phase, with the technological aspects of the payload currently under study and development [23]. The current plan is to have an array of four seismic stations in a permanently shadowed crater at the Moon's South Pole, each containing two femtometer-class inertial sensors. To test these technologies, a precursor mission called 'SoundCheck' is also under planning [24].

The proof mass of these inertial sensors is suspended from a frame that is rigidly attached to the lunar surface. The frame follows the Moon's elastic response to a passing gravitational wave, producing a differential motion between the proof mass and the frame. Therefore, the gravitational wave signal is recorded in this differential motion. In theory, one such inertial sensor would be enough for the detection of gravitational waves. However, to distinguish between a real signal and local phenomena, such as meteor strikes, several of these inertial sensors are required. Clearly, the detection principle of LGWA is quite different from that of laser interferometer-based detectors. Nonetheless, the requirement ultimately comes down to achieving a sub-femtometer level readout of the proof mass motion of the inertial sensors, which motivates the development of high-precision displacement sensing techniques.

There are a handful of other ideas and proposals for non-terrestrial gravitational-wave detectors. Taiji [25] and TianQin [26] are two additional ongoing programs focused on space-based detection. For lunar gravitational-wave detection, there are proposals such as Gravitational-wave Lunar Observatory for Cosmology (GLOC). The GLOC was proposed as a triangular, laser-interferometric concept with suspended optics on the Moon [27]. As of now, such concepts appear to be technologically distant.

1.2 Low-frequency noise in ground-based detectors

As mentioned previously, freely falling test masses are necessary to detect gravitational waves. Ground-based detectors face the challenge that the ground below them moves by many orders of magnitude more than a gravitational-wave strain. The causes of ground motion can range from ocean waves to local disturbances such as trains or cars. Therefore, care must be taken to isolate the detector from these unwanted disturbances.

Current gravitational-wave detectors use state-of-the-art passive and active isolation systems [28–30] to reduce seismic noise in the detection bandwidth to a level below the anticipated sensitivity and maintain the optics precisely at the desired state. The detectors employ multi-stage pendulums as a passive suspension system, with the optics suspended within these pendulums. Below the resonance frequency of the pendulums, ground displacement is fully transmitted through the suspension to the optics. Above the resonance frequency, motion is suppressed as $1/f^2$. Each suspension stage adds a $1/f^2$ slope above it. Thus, the higher the frequency above resonance, the greater the suppression provided by the passive isolation system. Consequently, the currently used four-stage suspension chain offers a $1/f^8$ suppression slope to the core optics. This slope creates a 'seismic-wall' that determines the detector's sensitivity at low frequencies.

As motion at frequencies below the resonance of the passive isolation system is directly coupled to the test masses, alternative means are needed to keep the optics isolated. This is achieved through a combination of feedforward and feedback controls [28–30]. These controls act on a platform mounted to the ground, and the actuators on it drive the platform in all degrees of freedom. High-performance

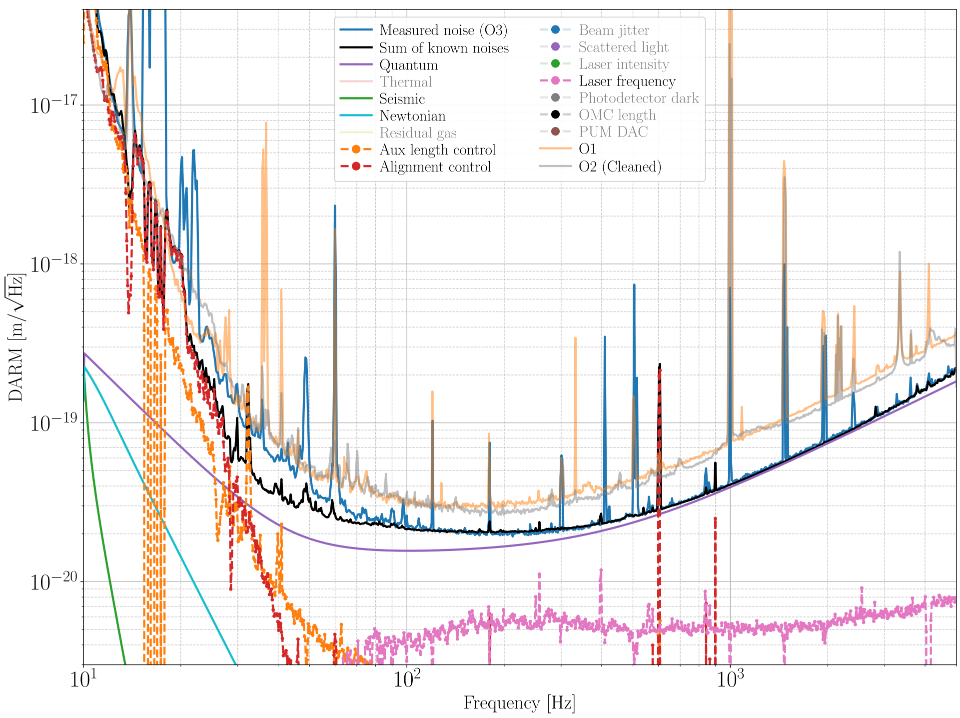

Figure 1.3: Noise budget of the LHO during the observing run O3. The seismic-wall, along with the auxiliary length and alignment control noise, limits the sensitivity at the low-frequency end. For better readability, only some noise sources are kept in the plot. (DARM: differential arm length. Created using Noisebudget Interactive Plotter: `https://ccahilla.github.io/lho_noisebudget.svg`; accessed on January 29, 2025.)

inertial seismometers, geophones, and other displacement sensors determine how these actuators should respond. The current active isolation systems of aLIGO use commercial sensors such as the Trillium 240 force feedback seismometer, the Geotech Model S-13, and the L-4C geophone [28, 31].

Unfortunately, this effort to keep the platform isolated also contributes as unwanted disturbances to the total noise budget of the detector, limiting sensitivity at low frequencies. The noise budget of the LHO during O3 is shown in Fig. 1.3. Seismic and Newtonian noise are major contributors at the low-frequency end. The active isolation system, which aims to eliminate contributions from the former, results in auxiliary length and alignment control noise, limiting sensitivity in that frequency regime. Higher sensitivity at lower frequencies, where the signals are strong and persistent for a longer time, is extremely beneficial. For this purpose, we need to reduce the root-mean-square (RMS) motion of the mirror suspensions with a better pre-isolation system to decrease auxiliary length and alignment control noise. This, in turn, requires improved displacement and inertial sensors.

Along with that, if low-frequency ground motion is not well suppressed, it can result in the mirrors drifting away from their anticipated operating point. Such a loss of lock decreases the duty cycle of the detector, reducing the operational time [32]. Therefore, a better sensing and control scheme is of great importance when aiming to increase the functionality of second-generation detectors or when transitioning to third-generation detectors.

The 1–10 Hz frequency band is of great interest due to the variety of astrophysical sources expected to lie within it [33]. To detect signals from these sources, the control bandwidth should be adapted to provide the required sensitivity in the sub-hertz regime. Active isolation at such low frequencies with sufficient precision demands a new class of highly sensitive inertial sensors, tilt sensors, displacement sensors, and advanced control schemes.

1.3 Overview of displacement and inertial sensors

1.3.1 Displacement sensors

A displacement sensor (also known as a displacement gauge) is, as the name suggests, used to measure the position or change in position of an object. It can determine how far an object has moved or its distance from a reference point. There are various types of displacement sensors, including:

- *Linear variable differential transformers (LVDTs)*: These sensors convert positional changes into electrical signals, thereby measuring linear displacements. They operate based on the principle of electromagnetic induction and consist of a primary coil and two secondary coils wound around a tube, housing a movable ferromagnetic core whose position is measured. Sensors with nanometer resolution and centimeter dynamic range[7] have been developed for low-frequency modal damping of the seismic attenuation system chains at Virgo [34]. However, due to the use of electrostatics, these systems cannot be implemented in all applications.

- *Capacitive sensors*: These sensors retrieve displacement information by measuring the change in capacitance between two plates. For example, sensors with a few tenths of picometer resolution in the 0–1 Hz bandwidth has been reported for space applications [35]. Capacitive sensors are generally smaller and lighter, and are excellent for short-range applications.

Other types include optical displacement sensors, eddy current sensors, linear proximity sensors, and ultrasonic displacement sensors [36–38]. In this thesis, the focus is on laser interferometric sensors, as they are one of the prime candidates for realizing more sensitive displacement and, in turn, inertial measurements.

1.3.2 Inertial sensors

Inertial sensors measure specific forces, rotational rates, and the orientation of an object. A classical inertial sensor can be modeled as a damped harmonic oscillator housed in a box. We then measure how the distance between the edge of the box and the test mass (also referred to as the proof mass) changes. With the knowledge of system's response to external stimuli, we can derive the force acting on the box or the desired ground motion. For the mathematical derivations, one can refer to [31], for example.

There are two main types of inertial sensors: gyroscopes and accelerometers. The latter is widely used in space-based missions like LISA and in active isolation systems. Accelerometers measure the rate of change of velocity of an object and operate on several principles, including piezoelectric, capacitive, and micro-electro-mechanical system (MEMS) technology. Technological developments

[7] The ratio between the maximum displacement that can be measured (or operating range) and the sensitivity (or precision).

over the last few decades have made it possible to create miniaturized high-precision sensors based on MEMS [39].

The advantage of MEMS-based sensors is that they can achieve quality factors[8] well into the millions [40]. However, with typical test masses on the order of micrograms, these sensors exhibit higher noise compared to large, high-performance seismometers. Larger test masses and lower resonance frequencies are essential for good low-frequency performance.

To measure changes in the test mass position of an inertial sensor with very high precision, laser interferometric readout is one of the most widely chosen options.

1.4 Laser interferometric sensors

Laser interferometers are an excellent choice for displacement sensing and the readout of inertial sensors. This is due to their high sensitivity, non-contact measurement, and resistance to magnetic couplings. A prime example is the classic LISA-type heterodyne interferometers, such as those used in the LISA Pathfinder mission. They have proven to be exceptional in measuring the position of test masses, with residual displacement measurement noise in the tens of femtometer regime [41].

Many compact interferometers have been developed in the ground-based gravitational-wave community for two primary applications [42]. Firstly, they serve as local displacement sensors for local damping or on the intra-vacuum seismic isolator (ISI). The other application involves the development of high-resolution inertial sensors, where one mirror is fixed to an inertial mass.

1.4.1 Existing systems

The active isolation system of aLIGO uses various sensing and actuation methods. Local sensing is conducted using optical sensors, as well as commercial seismometers and geophones. Actuation is primarily performed via electromagnetic motors. Two well-known (to the gravitational-wave detection community) and custom-made displacement sensors with optical readout are introduced below.

Birmingham Optical Sensor and Electro-Magnetic actuator (BOSEM)

BOSEMs belong to the category of ‘shadow sensors.’ They are compact, ultra-high vacuum compatible, non-contact, and low-noise position sensors with integrated electromagnetic actuators [43, 44]. BOSEMs are currently used in aLIGO to provide actuation and local sensing of the suspensions. A similar type of shadow sensor is also used in KAGRA [45].

The optical readout of the BOSEM is based on a shadow-sensing scheme. An opaque ‘flag’ is placed between an infrared light emitting diode (LED) and a photodiode. The flag is rigidly mounted to the measurement surface, partially blocking the light coming from the LED. When the measurement surface or the platform moves, the signal on the photodiode varies accordingly, indicating a change in the photodiode output related to the flag’s position. Modifications, such as adding a lens after the LED, improve both the linearity and sensitivity of the BOSEM.

The best BOSEMs with ‘enhanced’ specifications have a resolution of $45\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ down to $0.1\,\mathrm{Hz}$. Standard BOSEMs exhibit $0.3\,\mathrm{nm}/\sqrt{\mathrm{Hz}}$ sensor noise in the 1–10 Hz band [44].

[8] The quality factor (or Q factor) is a measure of energy loss in the resonator. A higher Q factor means lower energy dissipation, leading to more stable and accurate sensor performance.

Homodyne Quadrature Interferometer (HoQI)

With the aim of improving local sensing in gravitational-wave detectors, a compact laser interferometer-based sensor called HoQI has been developed at the University of Birmingham [46]. It has the potential to replace the existing BOSEMs for local position sensing in the suspensions. HoQIs can also be used for inertial sensor readout, for example, instead of the coil-magnet readout of GS-13 geophones.

A standard two-beam interferometer typically has a path length difference of less than a quarter of a wavelength within its operating range. To increase the dynamic range, HoQI uses two nearly orthogonal quadratures of the interferometer output and employs a polarization scheme to generate the required differential phase shift [47].

The interferometric layout and readout schemes are detailed in [46]. HoQI has achieved a peak sensitivity of $20\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$ at 70 Hz and a sensitivity of $70\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ at 10 mHz. This performance is orders of magnitude better than that of the existing displacement sensors used in aLIGO. Consequently, it is now scheduled to include HoQIs as relative displacement sensors in the big beam-splitter suspension (BBSS) during the LIGO A+ upgrade[9]. During O5, the first stage, called the top mass (M1), of BBSS will feature BOSEM-sensing and BOSEM-actuation. The next stage, called the intermediate mass (M2), will incorporate HoQI-sensing and BOSEM-actuation [48]. This setup will allow for the investigation of how improved sensors affect the performance of the suspended optics.

1.4.2 Ongoing developments

There are many displacement and inertial sensors under development within the LIGO scientific community. All of these aim to improve sensitivity at frequencies below a few tens of hertz. A non-exhaustive list of the ongoing developments is presented here.

- *HoQI*: HoQI is the most mature design among the ongoing research efforts. The current package size of the HoQI makes it challenging to incorporate into, e.g., the ISI. Therefore, a new compact version is being prototyped at the University of Birmingham. Attempts to implement the HoQI-readout for the commercial L-4C and GS-13 instruments have already been carried out [49], and further improvement efforts are ongoing. It is then expected to further reduce the low-frequency noise of these inertial sensors.

- *SmarAct sensor* (customized): This compact displacement sensor utilizes commercially available ‘SmarAct C01’ sensing heads. The optical setup consists of two Michelson interferometers that use deep-frequency modulation (DFM) techniques [50] to obtain a linear, relative displacement readout over multiple interference fringes [51]. A sensitivity of $0.3\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ above 0.1 Hz has been demonstrated.

- *Compact Balanced Readout Interferometer (COBRI)*: This displacement sensor is under development at the University of Hamburg. An unequal arm length Michelson interferometer is realized using an 11 mm long quasi-monolithic component, forming the on-axis optical sensor head [52]. This compact sensor also utilizes the DFM techniques, and the sensor output will be read out using an efficient and optimal algorithm implemented on a field-programmable

[9] Planned to take place after the observing run O4.

gate-array (FPGA) [53]. The considered balanced readout option helps eliminate some of the correlated noises from the displacement readout.

There are many other ongoing efforts as well. To further increase the sensitivity of DFM-based multi-fringe sensors, optical resonators can be integrated into the setup, resulting in Resonantly enhanced Deep-Frequency Modulation Interferometry (ReDFMI) [54]. Early testing of interferometric sensing using Digital Interferometry (DI) is underway at the Australian National University. This technique employs pseudo-random modulated signals at different time-of-flight delays to result in differential phase [55]. A 6-axis seismometer comprising a mass suspended by a single fused-silica fiber is being designed and analyzed at the University of Birmingham [56]. To measure local gravitational forces with high precision, a dual torsion pendulum-based sensor is under investigation at the Australian National University [57, 58]. University of Louvain (UCLouvain), in collaboration with the Dutch National Institute for Subatomic Physics (Nikhef), is further developing the Rasnik, a 3-point alignment system, as a displacement sensor [59, 60]. The goal is to reach sub-picometer sensitivity with such sensors as well.

1.4.3 Limitations of current sensors

All the sensors mentioned above have their sensitivity limited to the picometer or sub-picometer regime in the 1–10 Hz band. As pointed out earlier in Sec. 1.2, improving the sensitivity of the active isolation system is crucial for reducing associated noise sources in current and future detectors that limit their sensitivity at low frequencies. For this purpose, a displacement sensor with higher sensitivity that is compatible with the detector's environment is needed. If such a sensor also comes with a high operating range, it would provide an additional benefit. This is because all of these sensors are fundamentally limited by a trade-off between resolution and dynamic range [42].

Regarding planned missions like the LGWA, the core requirement boils down to achieving sub-femtometer precision in displacement readout. So far, to the best of our knowledge, no promising candidates have been explored in this regard. Developing a sensing mechanism for such a high-precision readout of inertial sensors is essential.

At the same time, there is a need to improve the readout schemes of the existing commercial inertial sensors employed by aLIGO. Currently, these sensors have internal noises that are not fundamentally limited by the suspension thermal noise of the proof mass [46]. Capacitance-based readouts can achieve high precision but have a limited operating range. They must be very close to the target object, which may be problematic for the suspension systems of the detector. Inductive readouts suffer from sensitivity to electromagnetic interference [28]. Some other readout schemes employ superconducting inductance measurements to achieve very high sensitivity [23]. However, the requirement for a cryogenic environment limits their applications.

To address these requirements, this thesis investigates a compact laser interferometric scheme for high-precision displacement sensing. This scheme, referred to as *heterodyne cavity-tracking*, represents the most precise readout that can be realized with laser interferometry. It is expected to achieve sub-femtometer noise levels at frequencies around 1 Hz and to provide a fringe-scale[10] operating range. It has the potential to improve the readout of existing inertial sensors or compact opto-mechanical sensors that are under development [61, 62].

[10] In the scale of a laser wavelength or an optical fringe.

Although, the main focus here is on current and future gravitational-wave detectors, there is a wide range of other possible applications. Such compact interferometer sensors are needed in fields such as atom interferometry [63], particle accelerators [64], and satellite controls [65]. In general, the development of high-precision, fringe-scale displacement sensors will significantly contribute to the field of metrology.

1.5 Thesis structure

This thesis presents an interferometric displacement sensor suitable for use in current and future ground-based gravitational-wave detectors. This sensing scheme is also highly suitable for non-terrestrial detectors, such as LGWA, where sub-femtometer displacement readout is a prime requirement [23].

Part I of the thesis primarily focuses on the laser interferometric scheme used for heterodyne cavity-tracking. This part begins with a general introduction to laser interferometry and also discusses the lasers used, along with attempts at its frequency noise characterization. From this part, it becomes clear that a digital signal processing (DSP) system and a stable laser frequency reference are the two main aspects of this interferometer scheme. These topics are then discussed separately in the next two parts of the thesis.

Part II addresses the digital implementation of the readout and control systems based on a radio-frequency system-on-chip (RFSoC). Steps to get started with such an evaluation board are explained in Chapter 4. For the readout, a gigahertz-bandwidth instrument was developed during the course of this work. The development and characterization of this instrument are described in Chapter 5 and published in IEEE Transactions on Instrumentation and Measurement [66]. For the laser frequency control, a Digital Heterodyne Controller (DHC) was realized on the same evaluation board, as discussed in Chapter 6.

Part III presents the study of laser frequency references. A chalcogenide glass material that has the potential to serve as a compact, portable frequency reference is proposed in this part, and preliminary tests conducted with these samples are detailed in Chapter 7. The in-house constructed Ultra-Low Expansion (ULE) glass-based stable optical cavities are presented in Chapter 8, including the construction procedure and characterization.

Finally, the experimental demonstration of the heterodyne cavity-tracking is presented in Part IV. Heterodyne cavity-locking is validated first in Chapter 9 using cavities constructed with commercial mirror mounts. Later, the in-house constructed stable cavities were used to determine the displacement readout noise floor. To probe the operating range of our sensor, intentional actuation on the cavity length was performed, and the tracking performance is presented in Chapter 10. Initial results of this work are published in Optics Express [67].

The final chapter of the thesis, Chapter 11, provides a summary and discusses the future scope of this work.

Part I

Interferometric scheme

CHAPTER 2

Basics of laser interferometry

This chapter aims to provide the basics of laser interferometry techniques. Before delving into that, we will briefly discuss the lasers used in precision experiments, focusing primarily on the type of laser utilized in this work. We will also cover the frequency noise of the lasers and various frequency stabilization techniques involving optical cavities that are used to minimize it.

2.1 Laser as a measuring tool

Laser interferometry is the tool of choice for high-precision measurements [42, 68]. The monochromatic nature of lasers and their short wavelengths make them extremely useful in various fields of science. We can use the wavelength of the laser as a ruler for metrology applications.

Different types of lasers are commercially available. Solid-state lasers, such as neodymium-doped yttrium aluminum garnet (Nd:YAG) in a non-planar ring oscillator (NPRO) configuration [69], are among the most commonly used in experiments worldwide [70, 71]. Semiconductor-based lasers, known as laser diodes, represent another category. However, the emission spectrum of these laser diodes is broad, and the lasing wavelength is not well-defined. The wide gain profile of the semiconductor supports many modes simultaneously, each with a different frequency. Therefore, additional frequency-selective feedback mechanisms are generally introduced into the laser cavity to achieve precise wavelength selection and a narrow emission line-width.

2.1.1 External cavity diode lasers

The lasers used in the work presented in this thesis are Continuously Tunable Lasers (CTLs) from Toptica Photonics AG [72]. This type of laser is a grating-stabilized external cavity diode laser (ECDL). They are tunable from 1510 nm to 1630 nm, but they were mainly operated around 1550 nm. Such tunable single-frequency diode lasers utilize a laser diode and a frequency-selective element, such as a grating, for laser frequency selection and tuning. An external motor attached to the grating allows for wider tuning ranges for the CTL.

An ECDL incorporates an optical grating mounted in front of the laser diode, while a second resonator is formed 'externally' between the laser diode's back facet and the feedback element. Schematic of such an ECDL, shown in Fig. 2.1(a), is known as a Littrow-type configuration. The

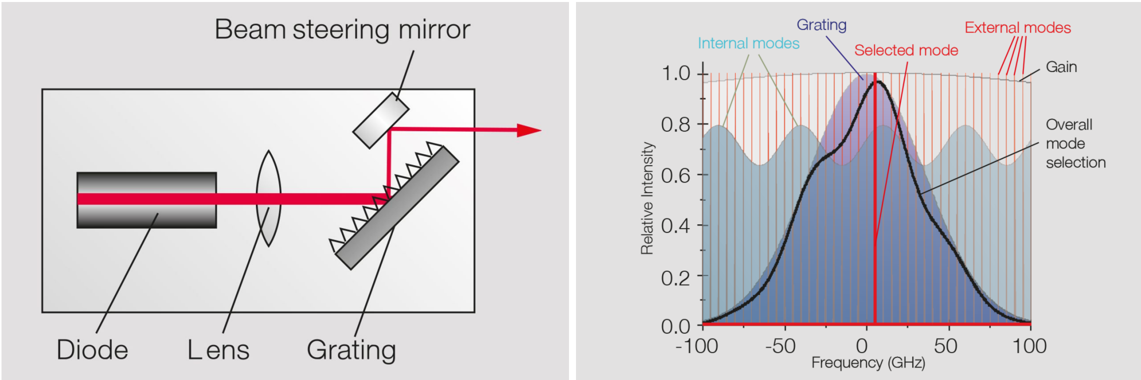

(a) (b)

Figure 2.1: (a) Schematic of an ECDL. By altering the angle of incidence on the grating, the wavelength of the laser is changed. (b) Mode selection mechanism in an ECDL. The mode with the largest overall gain is chosen for lasing. (Image credit: Toptica Photonics AG.)

grating filter, the semiconductor gain profile, and the modes of the internal laser diode and external cavity determine the lasing mode.

To change the wavelength of an ECDL, the spectral response of the filter is tuned. This can be done, e.g., by altering the angle of incidence on the grating. As the laser naturally operates on the mode with the largest overall gain, it will hop to another longitudinal mode and start emitting at a new wavelength. The overall mode selection mechanism can also be understood using Fig. 2.1(b).

When a laser can operate in multi-modes, there may also be transitions between different sets of modes, leading to what is termed mode-hopping. The CTLs achieve wide, mode-hop-free tuning through accurate synchronization of various parameters. This includes the simultaneous variation of the grating angle, the length of the external cavity, and the laser diode current. To achieve a tuning range of 120 nm without a single mode-hop, CTLs have active feedback mechanisms to keep the aforementioned tuning elements synchronized.

Adjusting the length of the external cavity allows for fine-tuning of the laser wavelength. It is accomplished using piezoelectric actuators that hold the grating. Additionally, the laser diode current can be changed independently, allowing for additional fine-tuning of the wavelength. However, these capabilities for actuating the laser frequency over a wide range unfortunately come at the cost of making the laser system susceptible to external mechanical and acoustic disturbances.

2.2 Interferometry

Interferometry is a technique used to extract information from electromagnetic waves by measuring the intensity or frequency of their interfered electric fields. Interferometers can track the position of a target mirror with precision much smaller than a wavelength. The operating range of conventional interferometers is limited to small-range (sub-wavelength or sub-fringe) measurements. However, there are interferometer techniques that enable an operating range significantly larger than a wavelength (multi-fringe).

Interferometers can be realized in different optical configurations. Notable examples include two-beam interferometers, such as Michelson, Mach-Zehnder, and Sagnac interferometers. It is also

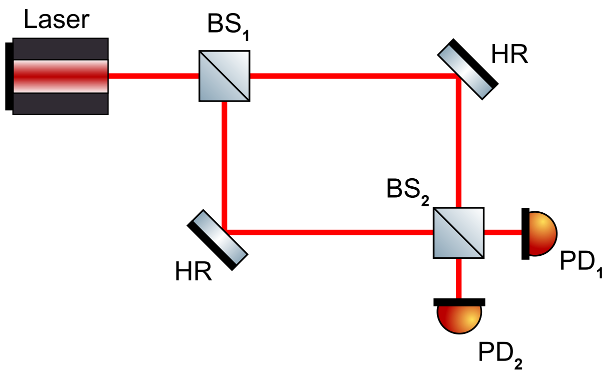

Figure 2.2: Homodyne Mach-Zehnder interferometer. (BS: beam-splitter, HR: high-reflective mirror, PD: photodetector.)

possible to employ optical cavities (also known as optical resonators) to realize resonant interferometers.

Before discussing the particular type of interferometer used in this thesis, two broad classifications of laser interferometers are briefly reviewed below. Apart from those mentioned here, there are many other types of interferometers, such as those involving DFM techniques [50] or using polarization-based techniques [46]. For an extensive overview of different types of compact interferometers, readers can refer to [42].

2.2.1 Homodyne interferometer

When the interfering electric fields have the same frequency, we classify such an interferometer as a homodyne interferometer. The DRFPMI topology used in current gravitational-wave detectors is a homodyne interferometer, as the two interfering fields originate from the same laser and therefore oscillate nominally at the same frequency.

A simple homodyne Mach-Zehnder interferometer is sketched in Fig. 2.2. The laser beam is split into two paths at the first beam-splitter (BS_1) and is then recombined after some time at the second beam-splitter (BS_2). The recombined fields will interfere constructively or destructively, depending on the relative path lengths of the two interferometer arms. It is then possible to determine the relative path length difference with sub-wavelength precision by measuring the power of the interfered beams.

The electric fields accumulate phase as they propagate along each arm of the interferometer, which can be represented by

$$\phi_i = \left(\frac{2\pi}{\lambda}\right) L_i = kL_i \qquad \text{with} \quad i = 1, 2; \tag{2.1}$$

where $k = 2\pi/\lambda$ is the wave number and λ is the wavelength of the electric field. L_i represents the path length of the corresponding interferometer arm. The relative optical path length difference is then given by $\Delta L = |L_1 - L_2|$.

The optical power at the two outputs of BS_2 can then be expressed as

$$P_i(t) = \frac{P_{in}}{2}(1 \mp \cos(k\Delta L)), \tag{2.2}$$

with P_{in} as the input laser power. A detailed derivation is available in [73].

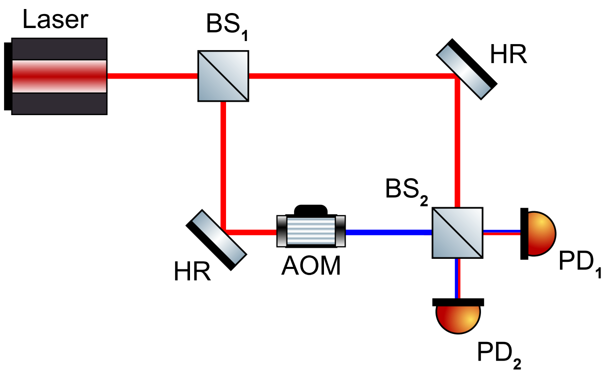

Figure 2.3: Heterodyne Mach-Zehnder interferometer. Laser frequency in one of the interferometer arms is shifted using an AOM. (BS: beam-splitter, HR: high-reflective mirror, PD: photodetector.)

From Eq. 2.2, it is clear that the interferometer output is sensitive to the changes in the differential optical path length. We can express the differential phase between the interferometer arms as

$$\Delta\phi = k\Delta L = \left(\frac{2\pi}{\lambda}\right)\Delta L. \tag{2.3}$$

Thus, considering a laser wavelength on the order of a micrometer, an interferometer amplifies the differential path length by roughly seven orders of magnitude. The output power is independent of the common arm length of the interferometer. This dependency of power on the differential phase generally applies to other types of two-beam interferometers as well.

2.2.2 Heterodyne interferometer

This class of interferometers uses two frequencies to make displacement measurements. Because of the involvement of two frequencies in the interference, they produce a beat note at their difference frequency.

Fig. 2.3 illustrates a simple way of realizing a heterodyne interferometer using a Mach-Zehnder topology. An acousto-optic modulator (AOM) is used in one of the interferometric arms to shift the laser frequency by Ω. In contrast to the homodyne interferometers described in Sec. 2.2.1, the beams also pick up an additional phase corresponding to the difference in laser frequency.

The optical power at the two outputs of BS_2, for a perfect contrast scenario, can be written as

$$P_i(t) = \frac{P_{in}}{2}(1 \mp \cos(2\pi\,\Omega\,t + \Delta\phi)) \qquad \text{with} \quad i = 1, 2; \tag{2.4}$$

where $\Delta\phi = k\Delta L$. Thus, changes in relative optical path length ΔL are captured in the phase of the heterodyne beat note oscillating at Ω. The extraction of this phase from the beat note can be achieved in various ways; one method, which is central to this thesis, is described in Chapter 5.

The advantage of the heterodyne technique over the homodyne technique comes from shifting the parameter of interest from the DC[1] regime to the heterodyne frequency. This significantly mitigates

[1] The term originates in electronics, where 'DC' refers to direct current. In this context, it is used colloquially to describe a constant voltage value.

technical noise in the phase readout process. The phase of a heterodyne beat note can be detected with higher dynamic ranges compared to reading out the power detected by a photodetector. This heterodyne technique is also relevant when the phase information is present in a single, weak carrier field. In such cases, it needs to be recovered by interference with a local strong laser field, as planned for the LISA mission.

2.3 Laser frequency noise

Even though lasers are considered to be monochromatic, the frequency of a laser actually fluctuates around a certain value. This intrinsic fluctuation of the laser frequency is often a significant source of noise in precision interferometric experiments. Narrow line-width lasers are less susceptible to laser frequency noise in the presence of macroscopic optical path length differences [73]. However, this does not eliminate the presence of laser frequency noise.

To understand how these frequency fluctuations couple to the optical phase and thereby limit the sensing of relative optical path length, we start with the optical phase difference between two interferometric arms

$$\Delta\phi = k\Delta L = \left(\frac{2\pi f}{c}\right)\Delta L, \tag{2.5}$$

where c is the speed of light.

Next, we split the laser frequency f into an average quantity f_0 and a fluctuating quantity δf. The path length difference ΔL is similarly split into L_0 and δL. Hence, Eq. 2.5 becomes

$$\Delta\phi = \frac{2\pi}{c}(f_0 + \delta f)(L_0 + \delta L) \tag{2.6}$$

$$\approx \frac{2\pi}{c}(f_0 L_0 + f_0 \delta L + L_0 \delta f). \tag{2.7}$$

The first term in the above equation is referred to as the 'operating point' of an interferometer. We want to maintain the interferometer at that static offset and observe the length signal, which corresponds to the second term of the equation. The third quantity, which contains the laser frequency noise, also couples to the differential phase. Since the photodetector outputs of the interferometer are sensitive only to $\Delta\phi$, the contribution from the laser frequency noise is indistinguishable from the actual differential path length signal. In principle, it is possible to make the interferometer insensitive to frequency noise by well-aligning it with equivalent arm lengths.

For a commercial Nd:YAG laser at 1064 nm, the frequency noise is roughly $10\,\mathrm{kHz}/\sqrt{\mathrm{Hz}}$ at 1 Hz, with a characteristic $1/f$ slope [74]. Therefore, it is often necessary to stabilize the laser frequency using a stable reference, such as an optical cavity. We will discuss more about laser frequency stabilization techniques in Sec. 2.5.

2.4 Optical cavities

An optical cavity allows a beam of light to circulate in a closed path. In its simplest and standard form, it consists of two mirrors facing each other to form a multiple 'bounce' system. An optical cavity constructed using two parallel reflecting surfaces is known as a Fabry-Pérot interferometer or etalon.

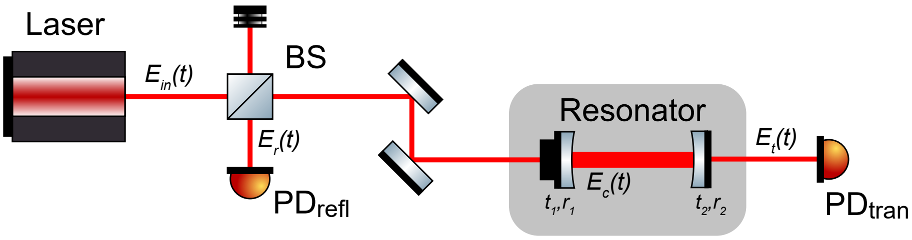

Figure 2.4: Sketch of an optical resonator constructed using two mirrors. The multiple bounces of the electric field inside the resonator, which leads to power build-up, is represented by a thick line. Laser fields at different points in the optical layout are labeled in accordance with the discussion in Sec. 2.4.1. (BS: beam-splitter, PD: photodetector.)

Resonators can be constructed in various configurations. Linear (or standing-wave) resonators and ring resonators are two commonly known forms.

Fig. 2.4 shows a sketch of a two-mirror system forming a linear optical resonator, where the light field bounces back and forth between the two mirrors. A portion of the incoming light from the laser is transmitted through the input mirror of the cavity. This light then circulates around the cavity and interferes with the time-delayed incoming field. The resonator's quality can be quantified by the effective number of such bounces required to reach a steady state. An optical resonator also has an associated dimensionless parameter that characterizes its ability to store energy relative to the energy it loses per cycle. This quantity, known as the quality factor (or Q factor), is an approximate measure of the number of bounces before the light dissipates from the cavity.

For a cavity of a given length, there are only certain optical frequencies for which the round-trip phase shift is an integer multiple of 2π. In terms of laser wavelengths, these are the half-wavelengths that are integer divisors of the cavity length. These frequencies self-consistently reproduce after each round trip and are called the mode frequencies or resonance frequencies.

The frequency spacing of these resonator modes is called the free spectral range (FSR). The ratio of the FSR and the line-width (or full-width at half-maximum (FWHM)) of the resonances is known as finesse, which is an indicator of the resonator quality as well. It is expressed as

$$\mathcal{F} = \frac{\mathrm{FSR}}{\mathrm{FWHM}} \approx \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2}. \tag{2.8}$$

Here, r_1 is the coefficient of reflection of the input mirror, and r_2 is the coefficient of reflection of the end mirror of the cavity. The reflectivity of these mirrors can be calculated as $R_i = r_i^2$.

The approximation in Eq. 2.8 is only valid for two-mirror resonators with very high reflectivity mirrors [75]. It is interesting to note that even though the FSR depends solely on the length of the cavity (the longer the cavity length, the smaller the FSR), cavity finesse is independent of the length. It is entirely determined by the transmission losses or reflectivities of the resonator.

The highest internal circulating power is achieved when the input laser frequency matches one of the resonator modes and the spatial shapes are also aligned. Under these conditions, the reflection from the resonator is minimized, and the transmission through it becomes maximized. The precise spatial matching of the electric field distributions of the input laser beam and the resonator modes is known as mode-matching.

The plot in Fig. 2.5 shows the transmitted and reflected power for high- and low-finesse cavities

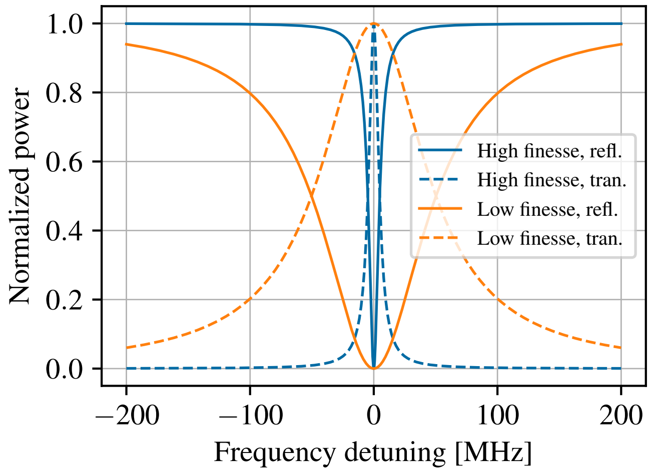

Figure 2.5: The reflected and transmitted power for high- and low-finesse cavities when the input laser frequency is tuned around the resonance frequency. Lossless, impedance-matched cavities with a physical length of 5 cm are considered for the simulation. The high-finesse cavity has mirrors with $R = 99\%$, while the low-finesse cavity mirrors have $R = 90\%$.

as the input laser frequency is tuned around the resonance frequency. The high-finesse cavity has mirrors with a reflectivity $R = 99\%$, whereas the low-finesse cavity has $R = 90\%$. Cavity responses are simulated using 'Finesse3'[2] [76, 77]. At resonance, the reflected power is at a minimum, and the transmitted power is at a maximum.

As the transmittance of both mirrors forming the cavities are equal (i.e., $T_1 = T_2$) in the above example, they are classified as impedance-matched cavities. If $T_1 < T_2$, the cavity is called an under-coupled cavity. When $T_1 > T_2$, we have an over-coupled cavity, where the circulating power is at its largest. Such an over-coupled configuration is used in the LIGO. For any mirror, the sum of reflectivity, transmittance, and losses equals unity, as the energy must be conserved.

Fig. 2.6 shows the normalized amplitude and phase of the reflected field as a function of the tuning of the laser frequency for the aforementioned three cases. The phase of the reflected field jumps by π and 2π as we pass through the resonance of an impedance-matched cavity and an over-coupled cavity, respectively. Such a phase shift of the reflected field does not occur for an under-coupled cavity.

Another important parameter when it comes to alignment or construction of optical cavities is the fringe visibility (also known as contrast or coupling efficiency). For an impedance-matched cavity, it can be defined based on the off-resonance reflected power P_{max} and the reflected power at resonance P_{min} as

$$\text{fringe visibility} = \frac{P_{\text{max}} - P_{\text{min}}}{P_{\text{max}} + P_{\text{min}}}. \tag{2.9}$$

This value ranges between 0 and 1. In an ideal scenario, no power is reflected from the impedance-matched cavity at resonance, resulting in a fringe visibility of 100%.

Similar to the two-beam interferometers discussed in Sec. 2.2, the outgoing fields (either reflected

[2] Frequency domain interferometer simulation software.

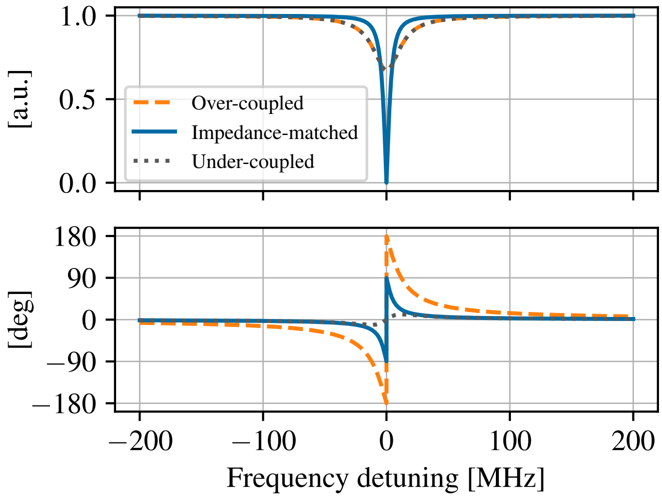

Figure 2.6: Amplitude and phase response of the reflected field for over-coupled, impedance-matched, and under-coupled loss-less cavities. The cavity responses are simulated using Finesse3. The impedance-matched cavity has $T_1 = T_2 = 1\%$. In the case of the over-coupled cavity, $T_1 = 5\%$ and $T_2 = 1\%$. For the under-coupled case, $T_1 = 1\%$ and $T_2 = 5\%$.

or transmitted) of the impedance-matched (or over-coupled) cavity carry information about the optical phase of the resonator, which can be converted into the length of the resonator. This makes these optical cavities promising candidates for displacement sensing as well. The displacement of one optic changes the path length inside the resonator, and multiple bounces amplify the optical phase shift [42].

When using optical cavities, the relative length fluctuations of the cavity can simply be equated to the relative frequency (or wavelength) fluctuations of the laser as [78]

$$\frac{\delta L}{L} = \frac{\delta f}{f} = \frac{\delta \lambda}{\lambda}, \tag{2.10}$$

under the assumption that $\delta L \ll L$. However, the increase in sensitivity compared to two-beam interferometers comes at the cost of a reduced operating range.

2.4.1 Optical cavity transfer functions

In this section, we aim to mathematically understand the influence of the cavity (as in Fig. 2.4) on an incident field $E_{in}(t)$. For simplicity, we assume a loss-less cavity of length L with mirrors having reflectivity $R_i = r_i^2$ and transmittance $T_i = t_i^2$. We also use the convention that a phase shift of $i = \sqrt{-1}$ is gathered after transmission through an interface.

The reflected field from the input mirror of the cavity can be expressed as

$$E_r(t) = r_1 E_{in}(t) + i t_1 r_2 E_c(t - 2T), \tag{2.11}$$

where $E_c(t)$ is the circulating field right behind the input mirror of the cavity. To arrive at the superposition of two fields constructing the cavity reflected field, we consider the delay by twice the cavity single-pass time T. Similarly, for the transmitted field, we can write

$$E_t(t) = it_2 E_c(t - T). \tag{2.12}$$

Following the same analogy, the circulating field of the cavity becomes

$$E_c(t) = it_1 E_{in}(t) + r_1 r_2 E_c(t - 2T). \tag{2.13}$$

In the Laplace domain (with s as the Laplace variable), the above three equations get transformed into[3]

$$E_r(s) = r_1 E_{in}(s) + it_1 r_2 E_c(s) e^{-2sT}, \tag{2.14}$$
$$E_t(s) = it_2 E_c(s) e^{-sT}, \quad \text{and} \tag{2.15}$$
$$E_c(s) = it_1 E_{in}(s) + r_1 r_2 E_c(s) e^{-2sT}. \tag{2.16}$$

Rearranging this set of equations leads to the cavity transfer functions for reflection, transmission, and circulation, respectively, as follows:

$$R(s) = \frac{E_r(s)}{E_{in}(s)} = \frac{r_1 - r_2 e^{-2sT}}{1 - r_1 r_2 e^{-2sT}} \tag{2.17}$$
$$T(s) = \frac{E_t(s)}{E_{in}(s)} = \frac{-t_1 t_2 e^{-sT}}{1 - r_1 r_2 e^{-2sT}} \tag{2.18}$$
$$C(s) = \frac{E_c(s)}{E_{in}(s)} = \frac{it_1}{1 - r_1 r_2 e^{-2sT}} \tag{2.19}$$

In the case of laser interferometry, the incident field $E_{in}(t)$ oscillates at a fixed angular frequency ($\omega = \Im\{s\} = 2\pi f$) and a constant amplitude ($\sigma = \Re\{s\} = 0$). The cavity single-pass time T can be expressed using cavity length L as $T = L/c$. This leads to the steady-state cavity transfer functions as

$$R(f) = \frac{r_1 - r_2 e^{-i2\pi f/\mathrm{FSR}}}{1 - r_1 r_2 e^{-i2\pi f/\mathrm{FSR}}}, \tag{2.20}$$
$$T(f) = \frac{-t_1 t_2 e^{-i\pi f/\mathrm{FSR}}}{1 - r_1 r_2 e^{-i2\pi f/\mathrm{FSR}}}, \quad \text{and} \tag{2.21}$$
$$C(f) = \frac{t_1}{1 - r_1 r_2 e^{-i2\pi f/\mathrm{FSR}}}. \tag{2.22}$$

It is clear that the cavity transfer functions are periodic in f with a periodicity given by

$$\mathrm{FSR} = \frac{c}{2L}. \tag{2.23}$$

This is another approach to arrive at the FSR of the cavity. The magnitudes of $T(f)$ and $C(f)$ are maximized when f/FSR is an integer, at which point the laser field is 'resonating' in the cavity.

[3] A translation of $-2T$ in time results in a factor e^{-2sT} in the Laplace domain.

2.5 Laser frequency stabilization

When the laser used has intrinsic frequency noise that exceeds acceptable levels, it is possible to use a stable reference to make the laser frequency follow that reference. The free-running line-width, or short-term stability of the laser, is often inadequate for many applications without active stabilization of the laser frequency.

There are various methods by which laser frequency stabilization can be achieved. Stable references may include an atomic transition, an interferometer with unequal arm lengths, another stable laser, or an optical cavity. Part III of the thesis is dedicated to such frequency references. Using a stable optical cavity is one of the most common and straightforward options for frequency stabilization [79, 80]. A stable optical resonator will have well-defined resonance frequencies, allowing us to 'lock' the laser frequency to one of the resonances, thereby stabilizing the laser frequency.

Another way to understand this is by treating the optical cavity as a filter for the wavelength of the incident light field. The cavity filters out all wavelengths except for the one that satisfies the resonance condition. This wavelength selectivity enables the cavity to serve as a reference for laser frequency stabilization.

2.5.1 Resonator locking methods

There are several ways to make the laser frequency follow the resonator modes. In this section, we will briefly discuss some common techniques used for this purpose.

Side-of-fringe locking

Side-of-fringe or DC locking is one of the simplest locking techniques. This can be achieved using either the reflected signal from the cavity or the transmitted signal. Fig. 2.7 shows the signals required for achieving side-of-fringe locking. A properly chosen reference voltage is subtracted from the photodetector output at the reflection port of the cavity. The resulting signal has zero-crossings around the reference voltage and exhibits a different sign on either side. This serves as the necessary error signal for the closed-loop control.

What is shown in Fig. 2.7 is also known as mid-fringe locking, as the reference lies exactly between the maximum and minimum of the cavity response. While it is possible to shift the reference point along the side of the fringe, mid-fringe locking provides maximum linearity and sensitivity. At the top of the fringe (at the resonance point), the direction of motion becomes ambiguous, making it nearly impossible to use that point as a locking reference for this method.

Since we operate in the DC regime, noise coupling is higher with this method. Any fluctuations in the laser output power or the relative intensity noise (RIN) of the laser will directly couple into the system and limit its performance.

Dither locking

This locking mechanism allows one to lock to the top of the fringe and is therefore also known as top-of-fringe locking. In this method, the laser frequency is intentionally modulated (or dithered) at various points across the cavity resonance and the response of transmitted or reflected power from the cavity is observed.

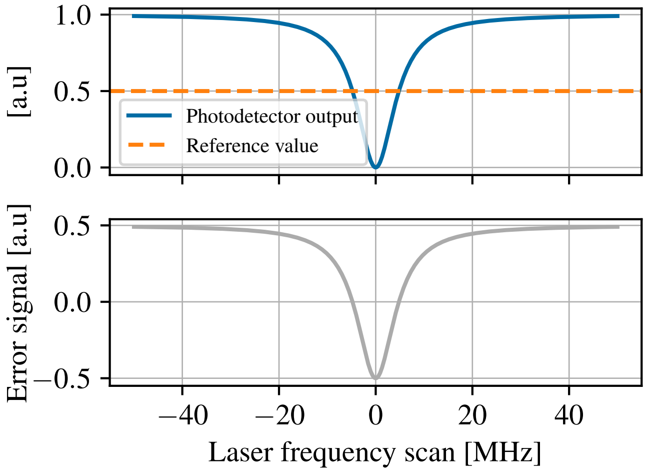

Figure 2.7: Signals for locking a laser to an optical cavity using side-of-fringe locking scheme. The top subplot shows the reflected signal from the cavity along with a chosen reference value. The subtraction of the latter from the former generates an error signal for the control loop, which is displayed in the bottom subplot.

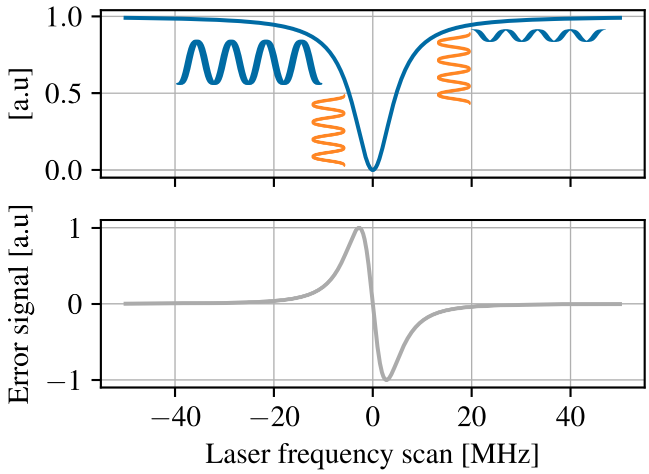

Figure 2.8: Signals for locking a laser to an optical cavity using dither locking scheme. The top subplot shows the reflected signal from the cavity, along with the variations in the photodetector output at two different dither points. The bottom subplot represents the error signal that can be derived from such a cavity response.

Far from the resonance, the variation in the reflected power for a given amount of dithering is nearly zero. As we get closer to the resonance, the same amount of dithering will cause a large variation in the reflected signal. At resonance, the response is flat, and the behavior is mirrored on the other side of the resonance but out of phase. The top subplot of Fig. 2.8 depicts this response.

It is then possible to employ phase-sensitive detection, such as a lock-in amplifier, to derive an error signal, as shown in the bottom subplot of Fig. 2.8. The error signal is essentially the derivative of the cavity response and can be interpreted as the projection of the amplitude and phase response of the

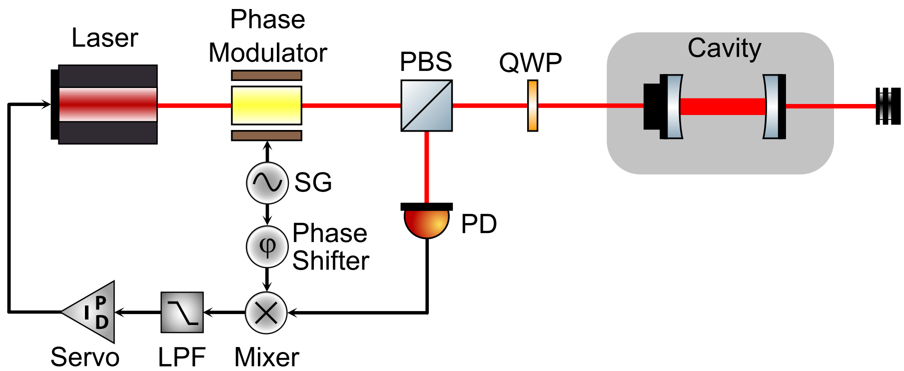

Figure 2.9: Optical layout of a PDH locking setup. (PBS: polarizing beam-splitter, QWP: quarter-wave plate, PD: photodetector, SG: signal generator, LPF: low-pass filter.)

cavity reflection at different dither points.

Pound-Drever-Hall (PDH) locking

The PDH locking technique is one of the most commonly used methods for locking a laser to an optical cavity. Only the basic idea of the PDH locking technique is covered here. A detailed introduction to PDH laser frequency stabilization can be found in [81].

The optical layout of the classical PDH setup is shown in Fig. 2.9. The carrier field from the laser is phase modulated, imprinting phase side-bands onto the carrier. Assuming that the carrier frequency is in the vicinity of the resonance frequency of the cavity, the complex response of the cavity is encoded only in the carrier field. The side-bands are reflected directly from the cavity onto the photodetector. The polarization-based optics used (a combination of polarizing beam-splitter (PBS) and quarter-wave plate (QWP)) distinguish cavity reflection from the incoming laser beam without losing any of the reflected power. If power loss is not a concern, one could simply replace this combination with a beam-splitter. The photodetector output contains a beat note at the phase modulation frequency, along with the cavity phase interaction. Demodulating[4] the photodetector output at the modulation frequency (with the appropriate demodulation phase) yields the expected PDH error signal.

When the modulation frequency is less than the line-width of the cavity used, we obtain an error signal very similar to the one showed in the bottom subplot of Fig. 2.8. If the modulation frequency spans multiple line-widths, we get additional zero-crossings associated with the side-bands.

The main advantage of PDH locking over the two previously discussed methods lies in noise coupling. Since the modulation frequency is typically on the order of tens of megahertz, we generally only have to be concerned with noise in that frequency range. The relevant noise sources typically narrow down to shot noise and electronic noise. This decouples the frequency stabilization from the laser's intensity fluctuations. An additional benefit is that it is possible to suppress frequency fluctuations that occur faster than the cavity response time [81].

The interferometric scheme used in this thesis belongs to the heterodyne category and employs an optical cavity for laser frequency stabilization. The resonator locking method, however, is slightly different from those discussed above. We will discuss it in detail in the following chapter.

[4] In this context, demodulation is the process of mixing two specified signals followed by low-pass filtering to remove higher harmonics.

CHAPTER 3

Heterodyne laser frequency stabilization

The interferometer scheme chosen for displacement sensing presented in this thesis is based on a technique called heterodyne stabilization (HS). This stabilization method is the primary choice for most of the experiments conducted during this work. It is a slightly different approach to laser frequency stabilization compared to those mentioned in Sec. 2.5 and is discussed in detail in this chapter. This chapter also includes a noise analysis of the HS scheme when used for displacement sensing, characterization of the frequency noise of the CTL, and efforts to demonstrate HS using analog electronics.

3.1 Motivation and our approach

The main motivation of this thesis is to develop a displacement sensor with a high dynamic range. The goal is to achieve sensitivity in the sub-femtometer regime at frequencies around 1 Hz, which is about 1000 times higher than that of most currently available compact laser interferometric sensors. For example, homodyne quadrature-based interferometers have demonstrated high precision (with a sensitivity of about 100 fm/$\sqrt{\text{Hz}}$ at 1 Hz) and high-range or multi-fringe readout capabilities [46, 82]. Heterodyne laser interferometers in compact setups with a sensitivity of 11 pm/$\sqrt{\text{Hz}}$ at 0.1 Hz are also reported [83]. While increasing the sensitivity by 2 to 3 orders of magnitude, it is also important not to completely compromise the operating range of the sensor.

In general, when laser interferometry is chosen for displacement sensing, photon shot noise becomes the fundamental limitation on sensitivity [84]. In the classical two-beam homodyne interferometers, as discussed in Sec. 2.2.1, the differential phase information is retrieved from the power of the interferometer output. Interferometric schemes that operate over one or multiple fringes are also limited by other readout noise, such as analog-to-digital converter (ADC) quantization noise, effectively limiting these techniques to sub-picometer or, at most, femtometer regimes [54].

To address these limitations and push the sensitivity into the sub-femtometer regime, this thesis studies the integration of an optical cavity for displacement sensing. This interferometric scheme employs an optical resonator to which a laser is locked. This approach helps to reduce most of the associated readout noise contributions and shot noise to negligible levels. Nevertheless, when cavity-locking schemes are involved, the interferometer operates over a small displacement range, i.e., within the resonance of the optical cavity. In our implementation, the overall operating range is extended by letting the laser follow the cavity motion and reading out the resulting large laser frequency changes. This way, even though multi-fringe readout is not feasible, we are not restricted to narrowband displacement sensing.

Another advantage arises from the readout method of the interferometer. In this technique, the displacement information is encoded in the heterodyne frequency (i.e., the beat note between the two frequencies used). Hence, we focus on the frequency of the photodetector output instead of its power. A frequency-tracking instrument can be used to measure the heterodyne beat note with very high dynamic range.

3.2 Heterodyne stabilization

For the locking of the laser to an optical cavity, we investigate a method called heterodyne stabilization (HS). This novel laser frequency stabilization method, which can be integrated into the LISA mission, was proposed by Eichholz et al. in 2015 [85]. In the context of LISA, HS was an effort to integrate the laser frequency control scheme into the LISA interferometric readout system, rather than relying on a separate laser stabilization unit. Here, by using the same scheme, we aim to increase displacement sensitivity at the cost of using two separate lasers to create the heterodyne beat note.

The basic working principle of HS is very similar to that of the classical PDH technique [86] for stabilizing the laser frequency. The laser frequency is stabilized to one of the resonances of an optical cavity using the cavity interaction phase shift. The key difference is that the PDH scheme uses phase modulators to create the radio-frequency (RF) side-bands (as described in Sec. 2.5.1), whereas HS uses two laser sources resulting in heterodyne interferometry.

Compared to the PDH scheme, here, the frequency of one of the lasers acts as a ‘carrier’ and the other as a ‘side-band.’ Assuming that the carrier frequency is in the vicinity of the cavity resonance, the complex response of the cavity is imprinted only in the carrier field. The side-band field is directly reflected onto the photodetector. Thus, the output of the photodetector contains a beat note at the heterodyne frequency, along with the cavity phase interaction. Hence, demodulating the output signal against the heterodyne beat note allows us to extract the cavity interaction phase, which can be used as an error signal for the laser frequency stabilization loop.

In this thesis, we study two aspects simultaneously: HS as a resonator locking technique for laser frequency stabilization, and a high-bandwidth heterodyne frequency readout when a laser follows a cavity with a movable end mirror, referred to as a dynamic cavity. The latter can also be studied with other locking techniques; however, using HS integrates both readout and control architectures. Throughout this thesis, the term ‘HS’ refers to the scheme used to lock the lasers to the cavities. When this scheme is applied for displacement sensing, we refer to it as heterodyne cavity-tracking.

3.2.1 Analytical understanding

A basic heterodyne setup is sketched in Fig. 3.1. Let $A_i = A_i(x, y)$ be the amplitude profile of the laser field $E_i(t)$, traveling in the z direction with angular frequency ω_i and phase ϕ_i. Then, the superposition of two such laser fields can be represented as

$$E(x, y, z, t) = E_1(t) + E_2(t) = A_1 e^{-i(\omega_1 t - k_1 z + \phi_1)} + A_2 e^{-i(\omega_2 t - k_2 z + \phi_2)}. \tag{3.1}$$

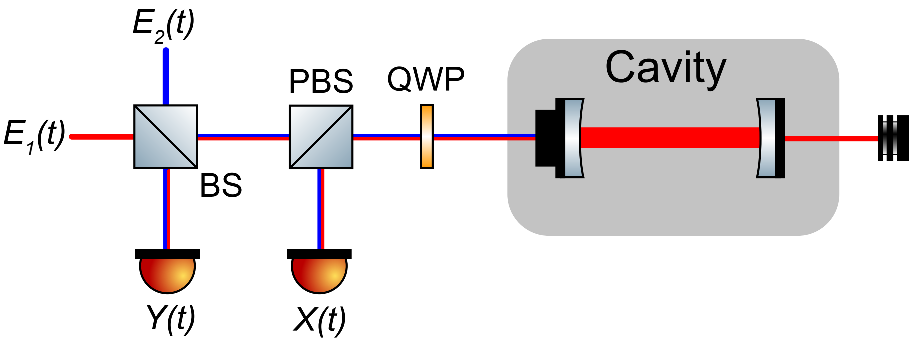

Figure 3.1: Basic heterodyne setup. Two laser fields $E_1(t)$ and $E_2(t)$ superpose and generate a reference beat note $Y(t)$. The superposition is also sent to an optical cavity, where only one laser (red) is in resonance. The reflection from the cavity $X(t)$ can be demodulated with $Y(t)$ to obtain the cavity interaction phase. (BS: beam-splitter, PBS: polarizing beam-splitter, QWP: quarter-wave plate.)

Assuming that the two laser fields have matching spatial profiles and propagation directions, we can neglect the z-dependency. The above equation is then simplified to

$$E(t) = A_1 e^{-i\omega_1 t} + A_2 e^{-i\omega_2 t}. \tag{3.2}$$

When detected by a photodetector, $E(t)$ generates a voltage $Y(t)$ that is proportional to its intensity

$$Y(t) \propto |E(t)|^2 = A_1^2 + A_2^2 + 2A_1A_2\cos(\Delta\omega t), \tag{3.3}$$

where $\Delta\omega = \omega_1 - \omega_2$ is the heterodyne beat frequency.

If $E(t)$ is reflected off a cavity, each of the two fields $E_i(t)$ experiences the cavity transfer function for reflection $R(f)$ (Eq. 2.20), based on its frequency. We consider that the laser field $E_1(t)$ is near-resonant with the cavity, while the other field is far off from the resonant modes. For frequencies far off resonance, we have $R(f) \approx 1$. Hence, without loss of generality, it can be assumed that $E_2(t)$ does not interact with the cavity and is directly back-reflected at the input mirror. The reflected amplitude and phase of $E_1(t)$ depend upon its frequency mismatch with the cavity resonance, $\delta f = f - f_{\text{res}}$.

The reflected field from the cavity can then be expressed as

$$E_r(t) = R(\delta f)E_1(t) + E_2(t) = R(\delta f)A_1 e^{-i\omega_1 t} + A_2 e^{-i\omega_2 t}. \tag{3.4}$$

This field, when detected using a photodetector, yields the signal

$$\begin{aligned} X(t) \propto |E_r(t)|^2 = |R(\delta f)|^2 A_1^2 + A_2^2 + 2A_1A_2 \times \\ [\Re\{R(\delta f)\}\cos(\Delta\omega t) + \Im\{R(\delta f)\}\sin(\Delta\omega t)]\,. \end{aligned} \tag{3.5}$$

The information from the cavity interaction is contained in $\Im\{R(\delta f)\}$. This can be recovered by demodulating $X(t)$ against the reference $Y(t)$.

We first apply a phase shift of $\pi/2$ to the reference beat note. Eq. 3.3 then becomes

$$Y_{ac}(t) = 2A_1A_2\sin(\Delta\omega t). \tag{3.6}$$

This is then multiplied with the oscillating part of the cavity reflection $X(t)$. That is,

$$X_{ac}(t) \times Y_{ac}(t) = 4A_1^2A_2^2 \cdot \sin(\Delta\omega t) \cdot [\Re\{R(\delta f)\}\cos(\Delta\omega t) + \Im\{R(\delta f)\}\sin(\Delta\omega t)] \,. \tag{3.7}$$

This can be further simplified as

$$\begin{aligned} X_{ac}(t) \times Y_{ac}(t) = 2A_1^2A_2^2 \cdot [\Re\{R(\delta f)\}\,(\sin(2\Delta\omega t) + \sin(0))] + \\ 2A_1^2A_2^2 \cdot [\Im\{R(\delta f)\}\,(\cos(0) - \cos(2\Delta\omega t))] \,. \end{aligned} \tag{3.8}$$

For complete demodulation, the multiplication is followed by a low-pass filter to remove the higher harmonic terms that are generated in the mixing process. Hence,

$$X_{ac}(t) \times Y_{ac}(t) \overset{\text{LPF}}{=} 2A_1^2A_2^2\Im\{R(\delta f)\}. \tag{3.9}$$

Now, let us take a closer look at the information hidden in $\Im\{R(\delta f)\}$. When the laser field $E_1(t)$ is near-resonant with the cavity, i.e., when its frequency mismatch δf is much smaller than the line-width of the cavity (denoted as Δf), we can approximate Eq. 2.20 as[1]

$$R(\delta f) \approx \frac{r_1 - r_2(1 - i2\pi\frac{\delta f}{\text{FSR}})}{1 - r_1r_2(1 - i2\pi\frac{\delta f}{\text{FSR}})}. \tag{3.10}$$

By multiplying both the numerator and the denominator by the complex conjugate of the denominator, and discarding the terms containing squared $\delta f/\text{FSR}$, we arrive at

$$R(\delta f) \approx \frac{r_1 - r_2}{1 - r_1r_2} + i2\pi\frac{r_2(1 - r_1^2)}{(1 - r_1r_2)^2}\frac{\delta f}{\text{FSR}}. \tag{3.11}$$

Using Eq. 2.8, the imaginary part of the reflected field transfer function can then be expressed as

$$\Im\{R(\delta f)\} \approx \frac{2\mathcal{F}}{\pi}\left(\frac{1}{r_1} - r_1\right)\frac{\delta f}{\Delta f} = \mathcal{G}\frac{\delta f}{\Delta f}. \tag{3.12}$$

It is evident that Eq. 3.12 is proportional to δf and thus responds linearly to the differential changes between the laser frequency and the cavity resonance. This differential change is amplified by an optical gain $\mathcal{G}$, which increases as the finesse $\mathcal{F}$ of the cavity increases.

By substituting Eq. 3.12 into Eq. 3.9, the demodulated signal of the heterodyne stabilization scheme becomes

$$X(t) \times Y(t) = 2A_1^2A_2^2\mathcal{G}\frac{\delta f}{\Delta f} = e(\delta f). \tag{3.13}$$

This expression depends solely on δf for a given setup. Consequently, driving $e(\delta f)$ to zero is equivalent to making the laser frequency follow the cavity resonance. In other words, this establishes a lock between the laser frequency and the cavity length.

[1] For $x \ll 1$, $e^x \approx 1 + x$.

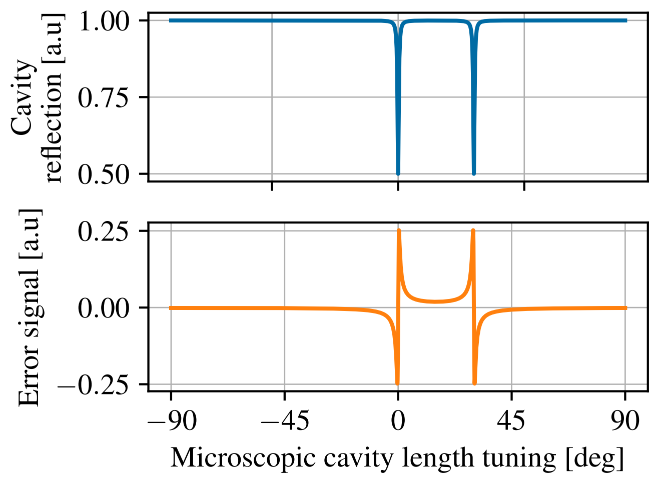

Figure 3.2: Cavity response in the case of heterodyne setup. The top subplot shows the reflected signal from the cavity, as detected by a photodetector, when the cavity length is tuned microscopically over one FSR of the cavity. The bottom subplot illustrates the demodulation of the cavity reflection against the offset frequency.

3.2.2 Simulating HS scheme

So far, we have analytically examined the reflected signal from the cavity in a heterodyne setup, as sketched in Fig. 3.1. Now, we will perform a simulation of such an optical setup using Finesse3 [77], comparing the results with the analytical predictions.

The simulation setup has two laser sources operating at 1550 nm with a frequency offset of 500 MHz between them. These two laser fields are then interfered using a beam-splitter and directed into a 5 cm long impedance-matched cavity with a finesse of about 312. Unlike previous simulations that involved scanning the laser frequency, we now tune the microscopic length of the cavity itself. Microscopic tuning is performed in units of degrees, such that a tuning of 180° changes the cavity length by half the wavelength of the laser (equivalent to one FSR of the optical cavity). As far as the cavity response is concerned, it is important to note that changing the laser frequency and tuning the cavity length are identical, and yield equivalent results.

The reflected signal from the cavity in this setup is plotted in the top subplot of Fig. 3.2. As the simulator interprets the cavity length as an integer multiple of the laser wavelength, one of the lasers is resonant in the cavity when no cavity length tuning is applied (microscopic cavity length tuning = 0°). The other laser remains off resonance at this point. Consequently, unlike the cavity response shown in Fig. 2.5, the reflected power from the cavity drops to only half of its maximum value. When the cavity length is detuned by 30°, the other laser, which is offset by 500 MHz, becomes resonant in the cavity.

If the cavity reflection is then demodulated at 500 MHz (which corresponds to the heterodyne beat frequency), we obtain the bottom subplot of Fig. 3.2. This demodulated signal exhibits zero-crossings at the resonance points, making it suitable for use in a closed-loop control system to establish a lock between the laser and the cavity.

The behavior of the cavity response in the current heterodyne setup is confirmed using the analytical calculations presented in Sec. 3.2.1. It follows that the real part of the cavity reflection is given by

$$\Re\{R(d\phi)\} = \frac{(r_1 - r_2)^2 + 4r_1 r_2 \sin^2(kL + d\phi)}{(1 - r_1 r_2)^2 + 4r_1 r_2 \sin^2(kL + d\phi)}, \tag{3.14}$$

for a loss-less cavity of length L with coefficient of reflectivities r_1 and r_2 for the input and end mirrors, respectively. $k = 2\pi/\lambda$ is the wave number, and $d\phi$ represents the microscopic cavity length tuning in units of radians. This quantity is what is detected by a DC photodetector. Thus, the top subplot of Fig. 3.2 can be reconstructed using the above equation while considering contributions from both laser fields.

Similarly, the imaginary part of the cavity reflection is given by

$$\Im\{R(d\phi)\} = \frac{(r_1^2 - 1)r_2 \sin(2(kL + d\phi))}{1 + r_1^2 r_2^2 - 2r_1 r_2 \cos(2(kL + d\phi))}. \tag{3.15}$$

The bottom subplot of Fig. 3.2 can also be reproduced using this equation by taking into account the contributions from both laser fields. In reality, direct detection of this quantity on a photodetector is not possible; it must be obtained through proper demodulation.

3.2.3 Displacement sensing using the HS technique

From the discussions in Sec. 3.2.1 and Sec. 3.2.2, it is clear that when a lock is established between a laser and a cavity, the laser frequency follows the cavity resonance to which it is locked. The cavity resonance condition, in turn, depends on the length of the cavity. Therefore, it is possible to retrieve information about the cavity length by monitoring the frequency of the laser.

However, direct detection of the laser frequency is impossible, as the bandwidth of currently available photodetectors is limited to a few gigahertz, while the laser frequency is generally in the hundreds of terahertz. Hence, a comparison of the laser frequency with some reference is required. Such a comparison is intrinsic to the HS technique. To access the cavity interaction phase, we demodulate the cavity reflection against the heterodyne beat note created between two lasers (Fig. 3.1). Thus, by monitoring the generated beat note, we can obtain information about the cavity length.

In order to utilize the HS scheme for displacement sensing, an optical cavity must be constructed with one of the mirrors serving as the test mass for which we want to measure displacement. By locking a laser to that cavity and monitoring the heterodyne beat note, we are able to read out the displacement of the test mass. This thought experiment assumes an ideal second laser with a fixed frequency. However, since this cannot be the reality, the displacement sensing becomes relative to the stability of the second laser or the frequency reference used. The second laser could be stabilized using any of the available methods. It is also an option to use another stable reference cavity and lock the second laser to it employing the HS technique. A simple and compact way of realizing such an interferometric scheme for displacement sensing, called heterodyne cavity-tracking, is sketched in Fig. 3.3.

When the length of the dynamic cavity changes, the frequency that resonates within that cavity also changes. The control loop then adjusts the corresponding laser frequency to follow this change. To keep the interferometer compact, short cavities are employed in our implementation. This means a high-bandwidth readout is required to track the changes of the dynamic cavity length. Following

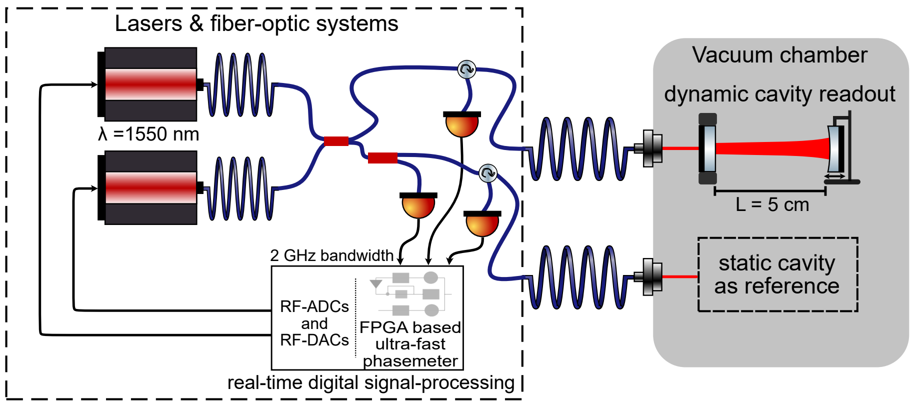

Figure 3.3: Interferometric scheme of heterodyne cavity-tracking. The beat note between the lasers and the signals reflected from the cavities are detected using photodetectors, and are then sent to a digital signal processing (DSP) unit (more details on this will be discussed in Part II). One laser is locked to the resonance of a static cavity, which serves as the reference, while the other laser follows the resonance of the dynamic cavity. The beat note frequency is our signal of interest for displacement sensing.

Eq. 2.10, the one-way motion of the end mirror of the dynamic cavity can be expressed as

$$\Delta L = \frac{\Delta f}{f} \cdot L, \tag{3.16}$$

where Δf is the heterodyne beat note, f is the absolute laser frequency, and L is the absolute physical length of the dynamic cavity.

It is important to keep in mind that Eq. 3.16 is an approximation, assuming plane waves in a flat-mirror configuration. In reality, for a Gaussian beam, the phase shift deviates from that of a plane wave by the so-called Gouy phase shift. Hence, the effective overall length of the cavity is slightly different from the stable mode FSR due to the Gouy phase accumulated during each round-trip. The Gouy phase correction to the FSR depends on the stable mode but is at most on the order of π (or one wavelength). As the employed optical cavities have a macroscopic length of a few centimeters (e.g., 5 cm), the required correction term is at most on the order of $1.55\,\mu\mathrm{m}/5\,\mathrm{cm}$, which is less than 0.1%.

It is also assumed that the coupling to the cavity mode is not significantly degraded by the moderate motion of the end mirror of the dynamic cavity. In other words, the microscopic mirror displacements have no significant effects on the cavity mode or coupling.

3.3 Noise analysis

In this section, we will analyze different noise sources that affect heterodyne cavity-tracking and explore the fundamental readout noise floor. This approach will provide insights into the maximum sensitivity achievable with this interferometric scheme for displacement sensing. As a reminder, our goal is to reach a sub-femtometer level noise floor at frequencies around 1 Hz.

In this scheme, there are two points where noise can couple to the system: while locking the laser to the dynamic cavity and while measuring the beat frequency. We assume that the frequency reference, or the secondary laser used to create the heterodyne beat, is an ideal system and, hence, does not contribute to the current noise budget.

The optical setup considered for estimating the noise budget includes laser sources operating at 1550 nm with an output power of 10 mW. A fixed frequency offset of 500 MHz is introduced between the two lasers. For the dynamic cavity, we consider a 5 cm long impedance-matched cavity in a flat-curved configuration, with the end mirror having a radius of curvature of 300 mm. The cavity mirrors are 99% reflective, resulting in a finesse of 312.

For the various noise sources, we first estimate their contribution to frequency noise, denoted as $\tilde{S}_f$. It has the unit Hz/$\sqrt{\text{Hz}}$. We can then arrive at the displacement noise, $\tilde{S}_L$ (in units of m/$\sqrt{\text{Hz}}$), using the relationship

$$\frac{\tilde{S}_f}{f} = \frac{\tilde{S}_L}{L}, \tag{3.17}$$

with f as the absolute frequency of the laser and L as the physical length of the cavity.

3.3.1 Noise sources influencing the laser-locking

Shot noise-induced phase noise

Shot noise has historically been described as the noise arising from the statistical distribution of electrons in photodetectors [68]. According to the Schottky formula, the shot noise is given by

$$\tilde{i}_{\text{SN}}(f) = \sqrt{2e\overline{I}} = \sqrt{2eR_{\text{PD}}\overline{P}} \qquad [\text{A}/\sqrt{\text{Hz}}], \tag{3.18}$$

with e as the electron charge ($\approx 1.602 \times 10^{-19}$C), $\overline{I}$ as the mean photo-current, R_{PD} as the responsivity of the photodetector (in units of A/W), and $\overline{P}$ as the mean laser power. Since the above equation is independent of frequency, shot noise appears as 'white noise'[2] in the spectrum.

To arrive at the shot noise-induced phase noise in locking the laser to the dynamic cavity, we estimate the signal-to-noise ratio (SNR) of the scheme and express phase noise as the inverse of it. In HS, the laser is locked to one of the cavity resonances. Let $\phi = 0$ (in radian) be the locking point. The reflected power from the cavity is not zero at that point due to the contribution from the other laser, which is off-resonance. This reflected power will be used to calculate the shot noise. Since the cavity response is taken from the simulation, we have an ideal scenario in terms of fringe visibility.

Our 'signal', where the cavity phase interaction is hidden, is the demodulated signal. Mathematically, it is the derivative of the reflected power from the cavity (see Fig. 3.2). To obtain the phase information, we differentiate the demodulated signal with respect to $d\phi$.

Putting these components together, we express noise (in amplitude spectral density (ASD)) as

$$\tilde{i}_{\text{SN}} = \sqrt{2eR_{\text{PD}}P_{\text{dc}}(\phi = 0)} \qquad [\text{A}/\sqrt{\text{Hz}}] \tag{3.19}$$

and the signal as

$$\frac{di}{d\phi} = R_{\text{PD}}\frac{dP_{\text{demod}}}{d\phi}|_{\phi=0} \qquad [\text{A/rad}]. \tag{3.20}$$

[2] Power is equally distributed over all frequencies.

Therefore,

$$\mathrm{SNR} = \frac{1}{\tilde{\phi}_{\mathrm{SN}}^2} = \frac{|\frac{di}{d\phi}|^2}{\tilde{i}_{\mathrm{SN}}^2}. \tag{3.21}$$

Notice that to calculate the SNR, we squared both the signal and noise quantities and arrived at the power spectral density (PSD) of shot noise.

The quantity of interest, shot noise-induced phase noise, can then be obtained as

$$\tilde{\phi}_{\mathrm{SN}} = \sqrt{2} \cdot \sqrt{2\frac{e}{R_{\mathrm{PD}}}} \cdot \left(\frac{\sqrt{P_{\mathrm{dc}}}}{\frac{dP_{\mathrm{demod}}}{d\phi}} \right)_{\phi=0} \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}]. \tag{3.22}$$

An additional factor of $\sqrt{2}$ is included because of the two uncorrelated white noise sources folding from positive and negative frequencies in the mixing process of the heterodyne scheme. This factor can also be accounted for by using the RMS amplitude of the input signal [87].

To estimate the contribution of $\tilde{\phi}_{\mathrm{SN}}$ to the total noise budget, the photodetector output at resonance and the derivative of the demodulated signal were simulated (similar to Fig. 3.2). The photodetector responsivity was assumed to be 1 A/W, which holds for commercial photodiodes at 1550 nm. The obtained $\tilde{\phi}_{\mathrm{SN}}$ in units of frequency noise has a value of about 54 mHz/$\sqrt{\mathrm{Hz}}$ for all frequencies.[3] This is equivalent to a fundamental displacement noise floor of about 1.4×10^{-17} m/$\sqrt{\mathrm{Hz}}$, which is well below the femtometer regime.

Temperature coupling noise

When the temperature surrounding the dynamic cavity changes, the cavity length is also affected. The displacement noise induced by this change is highly dependent on the spacer material used to construct the cavity and the ambient environment in which the dynamic cavity is placed. At very low frequencies, below a millihertz, the temperature noise exhibits a characteristic $1/f$ slope. Above this frequency, the temperature noise is usually flat, and a value of about 1 µK/$\sqrt{\mathrm{Hz}}$ is achievable [88].

This temperature noise $\tilde{S}_T(f)$ couples to the displacement noise of the dynamic cavity, depending on the linear coefficient of thermal expansion (CTE), denoted as α, of the spacer material. Mathematically,

$$\tilde{S}_L(f) = \alpha \times \tilde{S}_T(f) \qquad [1/\sqrt{\mathrm{Hz}}]. \tag{3.23}$$

Glass materials with near-zero expansion characteristics are commercially available. ULE glass from Corning and Zerodur from Schott are two examples of such materials. Corning ULE Glass 7972 has a mean linear CTE of $0 \pm 30 \times 10^{-9}$ /K [89].

Here, to estimate the temperature coupling noise, a CTE of 0.1 ppm/K and a temperature noise of 1 µK/$\sqrt{\mathrm{Hz}}$ at 1 Hz, with a $1/f$ slope throughout the considered frequency range, are considered. This approach ensures that we do not underestimate the temperature coupling noise in the sub-Hz regime. Under these conditions, at 1 Hz, the displacement readout will have a noise contribution of 5 fm/$\sqrt{\mathrm{Hz}}$ due to temperature coupling. This emphasizes the requirement for stable experimental environments and careful dynamic cavity designs to achieve sub-femtometer displacement sensing at frequencies

[3] To convert phase noise to frequency noise, we used the relation that a phase change of π corresponds to a frequency change of one FSR.

around 1 Hz.

Coating thermal noise

Coating thermal noise is predicted using the Fluctuation-Dissipation theorem. The PSD of the thermally induced fluctuations of the cavity mirror surface, as read by the interferometer, can be expressed as [90]

$$\tilde{S}_{xx}(f) = \frac{2k_B T}{\pi^{3/2} f} \frac{1}{wY} \left[\Phi_{\text{substrate}} + \frac{d}{\sqrt{\pi} w} \left(\frac{Y'}{Y} \Phi_{\parallel} + \frac{Y}{Y'} \Phi_{\perp} \right) \right] \quad [1/\text{Hz}], \tag{3.24}$$

where k_B is Boltzmann constant ($\approx 1.38 \times 10^{-23}$ J/K), T is temperature, w is the laser beam waist radius, d is the thickness of the coating, Φ is the loss angle[4], and Y and Y' are the Young's moduli of the substrate and coating, respectively. $\Phi_{\parallel}$ and $\Phi_{\perp}$ are, respectively, the loss angles associated with the energy density in parallel and perpendicular coating strains. Assuming a homogeneous coating strain, both of these quantities are considered to be equal.

To estimate the coating thermal noise, several assumptions were made:

- Standard lab temperature of 293.15 K.
- For the considered optical cavity parameters, the mode-matching criterion requires a laser beam wait radius of 234.87 µm at the input mirror.
- Assuming that the cavity mirrors are made of N-BK7 material, $Y = 82 \times 10^9\ \text{N/m}^2$ [91].
- In general, high-reflective coatings for 1550 nm are made with alternating layers of tantala (Ta_2O_5) and silica (SiO_2). About 15 layers would suffice to achieve a reflectivity of 99%. This results in a total coating thickness of $d = 3.5$ µm.
- The Young's modulus of tantala and silica are 140 GPa and 73 GPa, respectively [92, 93]. Here, as an average approximation, Y' is considered to be 100×10^9 N/m^2.
- The loss angle of tantala is about 4.5×10^{-4} for a layer thickness of 3.13 µm, while that of silica is about 0.6×10^{-4} for a layer thickness of 3.07 µm [93]. We consider the loss angle to be 10^{-4} for both the substrate and the coating.

Even though these assumed values depend heavily on the experimental setup under consideration, they do not differ by an order of magnitude from each other. Ultimately, the estimated $\tilde{S}_{xx}(f)$ is multiplied by 2 to account for the total contribution of the coating thermal noise from both mirrors of the dynamic cavity.

At 10 Hz, the dynamic cavity setup has a contribution of approximately $3.92 \times 10^{-17}\ 1/\sqrt{\text{Hz}}$ from the coating thermal noise.[5] At 1 Hz, the coating thermal noise-induced displacement noise is about $6.19 \times 10^{-18}\ \text{m}/\sqrt{\text{Hz}}$.

[4] The loss angle is related to the complex refractive index of a material and is used to quantify the amount of energy losses due to scattering and absorption.

[5] As a comparison, at 10 Hz, aLIGO coating thermal noise is on the order $10^{-23}\ 1/\sqrt{\text{Hz}}$ [92].

Quantization noise-induced phase noise

As sketched in Fig. 3.3, in a digital implementation of the laser-locking, the reflected signal from the dynamic cavity detected by the photodetector is sent to a real-time DSP unit. The digitization process of the analog signal is performed using an analog-to-digital converter (ADC). Consequently, certain parameters of the ADC affect the displacement readout performance. The most crucial parameters are the sampling frequency f_s, the number of generated bits N, the additive input noise, and the sampling jitter [87].

When an analog signal is converted into a digital signal by a process called quantization, the corresponding analog value at any particular instant is represented using N bits. The smallest distinguishable value is given by 2^{-N}, referred to as the value of the least significant bit (LSB). ADCs have their own intrinsic noise floor, determined by a voltage noise floor and the digitization noise. The voltage noise at the ADC input is typically at a similar level as the digitization noise, or slightly higher [87].

To arrive at the quantization noise-induced phase noise in locking the laser to the dynamic cavity, we follow a similar approach to that used for the estimation of shot noise-induced phase noise. Specifically, we calculate the SNR and express the phase noise as its inverse. The voltage noise due to the quantization noise from the ADC that digitizes the dynamic cavity reflection is given by[6]

$$\tilde{V}_{\text{quant}} = V_{pp} \frac{2^{-N}}{\sqrt{6 \cdot f_s}} \qquad [\text{V}/\sqrt{\text{Hz}}], \tag{3.25}$$

where V_{pp} is the maximum input range of the ADC. Our ‘signal’ can be expressed as

$$\frac{dV}{d\phi} = (50\,\Omega) \cdot R_{\text{PD}} \frac{dP_{\text{demod}}}{d\phi}|_{\phi=0} \qquad [\text{V/rad}], \tag{3.26}$$

assuming that a $50\,\Omega$ load resistance is used to convert the generated photo-current into a voltage. Hence, the quantization noise-induced phase noise can be calculated as

$$\tilde{\phi}_{\text{quant}} = \sqrt{2} \cdot \frac{1}{\sqrt{\text{SNR}}} = \sqrt{2} \cdot \frac{\tilde{V}_{\text{quant}}}{\left|\frac{dV}{d\phi}\right|} \qquad [\text{rad}/\sqrt{\text{Hz}}]. \tag{3.27}$$

For the current noise budget, a 12-bit ADC sampling the input at a frequency of 4.096 GHz is considered. The maximum input range of the ADC is taken as 1.8 V.[7] The calculation of $dV/d\phi$ is carried out utilizing the simulation used for the estimation of shot noise-induced phase noise. The obtained $\tilde{\phi}_{\text{quant}}$ in units of frequency noise has a value of approximately 76 mHz/$\sqrt{\text{Hz}}$ for all Fourier frequencies, which is equivalent to a displacement noise floor of about 1.97×10^{-17} m/$\sqrt{\text{Hz}}$. Similar to the quantization noise, the electronic noise also couples to the laser-locking loop. However, the noise analysis presented here does not explicitly include its contribution, as it is expected to be of the same order as the shot noise and quantization noise.

Fig. 3.4 shows the contributions from all these noise sources that couple to the laser-locking in the HS scheme. According to the current assumptions and this noise budget, the temperature coupling

[6] More details on quantization are available in Sec. 5.3.2.

[7] These parameters agree with the DSP unit discussed in Part II.

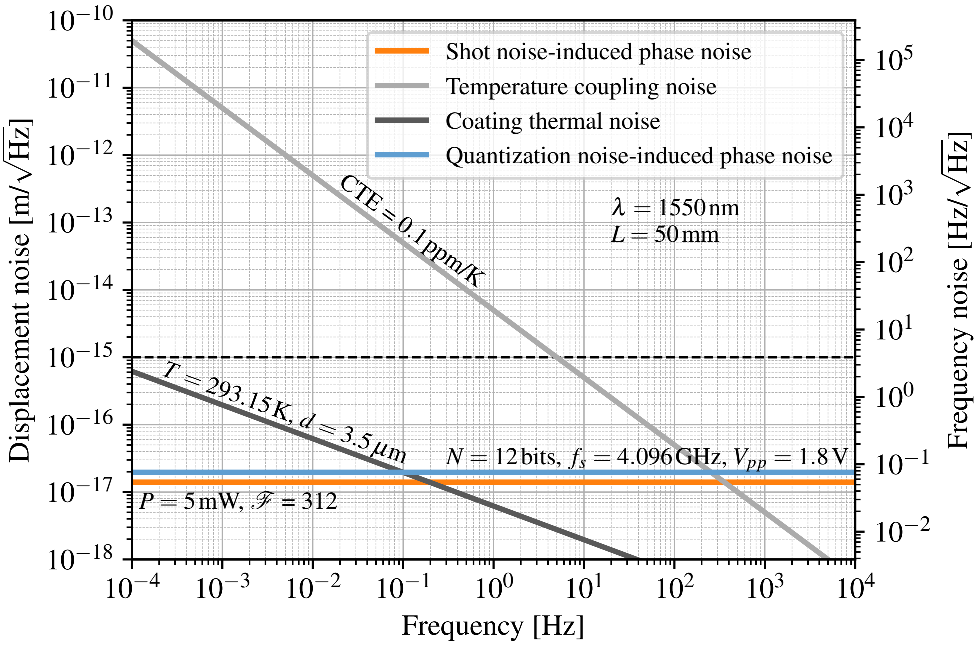

Figure 3.4: Noise budget for displacement sensing using the HS technique, considering noise sources that influence laser-locking. Contributions from all the noise sources discussed in Sec. 3.3.1 are plotted together. Key assumptions made to estimate the displacement noise contribution are also provided in the plot. The horizontal dashed line marks a displacement noise level of one femtometer.

noise will limit the displacement readout sensitivity in our frequency region of interest. To push the sensitivity below the femtometer mark, it is therefore crucial to isolate the dynamic cavity from external temperature couplings. Above a Fourier frequency of 5 Hz, the displacement readout sensitivity is already in the sub-femtometer regime.

So far, we have considered noise sources that come into play while locking the laser to the dynamic cavity. As mentioned before, the displacement sensing is also limited by the readout of the heterodyne beat note, where the displacement information is available. Below, we analyze some of these noise sources.

3.3.2 Noise sources influencing the frequency readout

ADC quantization noise

Similar to the quantization noise of the ADC digitizing the cavity reflection, the quantization noise of the ADC that digitizes the heterodyne beat note also contributes to the overall displacement sensing noise. When an analog signal of amplitude V_{in} is digitized by an ADC having a maximum input range of V_{pp}, the relative phase noise induced by ADC digitization is given by

$$\tilde{\varphi}_{\mathrm{ADC}} = \sqrt{2}\,\frac{V_{pp}}{V_{in}}\,\frac{2^{-N}}{\sqrt{6 \cdot f_s}} \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}]. \tag{3.28}$$

Equivalent frequency noise is obtained by differentiating $\tilde{\varphi}_{\mathrm{ADC}}$, which in the frequency domain is represented by multiplication by $s = 2\pi f$.

We again consider a 12-bit ADC, sampling an input signal of amplitude 1 V, with a sampling frequency of 4.096 GHz. The maximum input range of the ADC is taken as 1.8 V. However, the displacement noise contribution due to this ADC quantization noise is negligibly small. At 1 Hz, it is on the order of 10^{-24} m/$\sqrt{\mathrm{Hz}}$.

ADC sampling jitter noise

Any unwanted fluctuations in the sampling frequency or its corresponding sampling time cause undesirable phase noise, which, in turn, affects the displacement readout. The inherent variation of the ADC sampling time, the so-called sampling jitter, degrades the readout performance for signals in the megahertz regime [87].

The sampling jitter-induced phase noise $\tilde{\varphi}_{\tau_A}$ is a derivative of the sampling jitter $\tilde{\tau}_A$. For an input signal at frequency f_A, this noise can be expressed as

$$\tilde{\varphi}_{\tau_A} = 2\pi f_A \cdot \tilde{\tau}_A \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}]. \tag{3.29}$$

To obtain an initial estimation of the sampling jitter, we refer to analog front-end investigations from the LISA phasemeter[8] development, as documented in [87, Sec. 8.4]. Here, a phase noise performance of approximately 2 µrad/$\sqrt{\mathrm{Hz}}$ at 1 Hz, with a $1/\sqrt{f}$ characteristic slope, is experimentally measured for a signal frequency of 11 MHz. Taking this as a reference, we estimate the sampling jitter as

$$\tilde{\tau}_A = \frac{(2\,\mu\mathrm{rad})}{\sqrt{f}} \frac{1}{2\pi\ (11\,\mathrm{MHz})} \qquad [s/\sqrt{\mathrm{Hz}}]. \tag{3.30}$$

This value is then used to estimate the sampling jitter-induced phase noise while sampling the beat note frequency of 500 MHz. After converting into units of displacement noise, we obtain a contribution of about 2×10^{-20} m/$\sqrt{\mathrm{Hz}}$ at 1 Hz due to the ADC sampling jitter.

Shot noise-dominated additive noise

Similar to the shot noise-induced phase noise affecting laser-locking, shot noise-dominated additive noise also induces phase noise in the frequency readout. This results in an untracked phase error of the phasemeter, thus affecting the frequency readout. The phase noise induced by the additive noise (dominated by photon shot noise, which is spectrally white in phase) is given by [94]

$$\tilde{\varphi}_{\mathrm{SN}} = \sqrt{\frac{\hbar c}{2\pi\lambda\eta P_{\mathrm{sig}}}} \qquad [\mathrm{cycles}/\sqrt{\mathrm{Hz}}], \tag{3.31}$$

where η is the detection efficiency and P_{sig} is the optical power of the heterodyne beat note.

Converting this into units of displacement noise, and assuming that a 10 mW beat note is detected with ideal detection efficiency on the photodetector, we find a negligible contribution on the order of 10^{-24} m/$\sqrt{\mathrm{Hz}}$ at 1 Hz due to this noise.

[8] A phasemeter is an instrument used for the frequency readout. Chapter 5 will discuss this in detail.

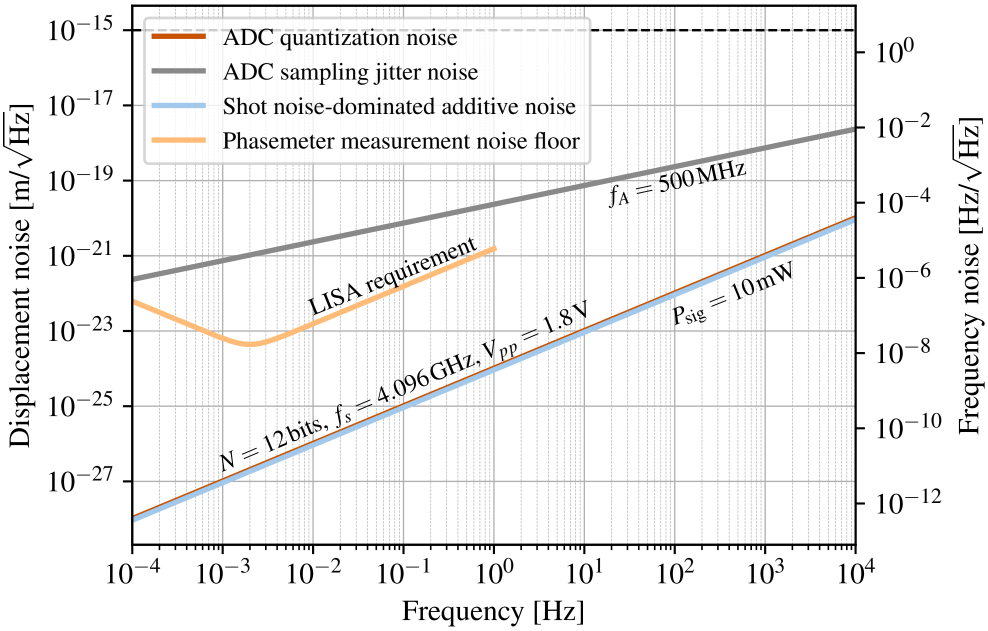

Figure 3.5: Noise budget for displacement sensing using the HS technique, considering noise sources that influence the frequency readout. Contributions from all the noise sources discussed in Sec. 3.3.2 are plotted together. Key assumptions made to estimate the displacement noise contribution are also provided in the plot. The horizontal dashed line marks a displacement noise level of one femtometer.

Phasemeter measurement noise floor

As mentioned in Sec. 2.2.2, the tracking of beat note is required when heterodyne interferometry is chosen for displacement sensing. A prime example is the LISA mission, where phasemeters are used to measure and track many heterodyne frequencies [95]. To meet the requirement of this mission, these instruments should also satisfy certain performance criteria. The required phase measurement noise floor for LISA is given by

$$\tilde{\varphi}_{\mathrm{LISA,req}}(f) = 6 \times 10^{-6} \cdot \sqrt{1 + \left(\frac{2 \times 10^{-3}}{f}\right)^4} \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}]. \tag{3.32}$$

Phasemeters meeting this requirement have already been developed [87]. Since the ADC sampling jitter noise is usually above the requirement, compensation techniques[9] are utilized to meet performance standards.

Fig. 3.5 shows all the considered noise sources that influence the displacement sensing via their coupling to the frequency readout, along with the LISA phasemeter requirement. It is evident that sub-femtometer displacement sensing is not at all limited by the chosen readout. This also emphasizes the advantage of heterodyne interferometric readout over classical two-beam homodyne interferometers, where shot noise poses a critical limitation to the readout [54].

[9] A discussion on this is available in Sec. 5.8.2.

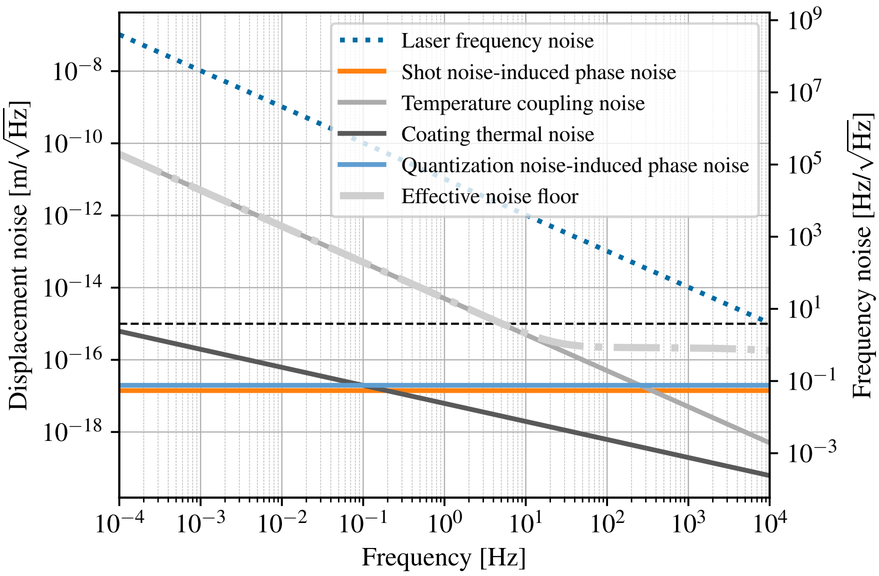

Figure 3.6: Overall noise budget for displacement sensing using the HS technique. Contributions from all the noise sources discussed in Sec. 3.3.1 and Sec. 3.3.2 are considered to calculate the achievable effective noise floor with a closed-loop control system having a unity gain frequency (UGF) of 300 kHz. As the noise sources influencing the frequency readout are negligibly small, they are not plotted. The horizontal dashed line marks a displacement noise level of one femtometer.

3.3.3 Overall noise budget

So far, we have not considered the contribution of laser frequency noise, even though the influence of laser frequency noise on displacement sensing has already been discussed in Sec. 2.3. This omission is due to the fact that the HS technique involves active control of the laser frequency. Therefore, this noise does not limit the displacement readout, assuming that the control loop gain in the measurement band is sufficiently high to suppress this noise below other noise sources.

However, widely tunable CTLs exhibit comparatively higher frequency noise than the NPRO Nd:YAG lasers, particularly in the kilohertz frequency band. Although this frequency band is above our measurement range, this noise contributes to the RMS error of the control loop, affecting its stability. This aspect is discussed further in Sec. 5.7.3. Additionally, higher noise at kilohertz frequencies makes it more challenging to keep the laser locked to the dynamic cavity in a stable manner.

For the noise budget, a free-running laser frequency noise contribution of 40 kHz/$\sqrt{\text{Hz}}$ at 1 Hz with a $1/f$ slope is considered from the CTL. In terms of displacement, this corresponds to approximately 10 pm/$\sqrt{\text{Hz}}$ at 1 Hz.

Fig. 3.6 shows the noise contributions from the different sources discussed so far (in units of both displacement and frequency). Among all the noise sources considered, laser frequency noise $\tilde{S}_{f,\text{free}}$ and temperature coupling noise exceed the femtometer mark in our frequency region of interest.

When a control loop with open-loop gain $G(s)$ is introduced to establish a lock between the laser and the dynamic cavity, it suppresses the contribution from the free-running laser frequency noise by a factor of $(1 + G(s))$. However, all other noise sources that have a direct influence on the sensor (dynamic cavity) or the heterodyne frequency readout will not be suppressed by the control loop. Their individual contributions scale with $G(s)/(1+G(s))$. We then calculate the quadratic sum of individual noise contributions to arrive at the effective frequency noise

$$\tilde{S}^2_{f,\mathrm{eff}} = \left(\frac{\tilde{S}_{f,\mathrm{free}}}{1+G(s)}\right)^2 + \sum_i \left(\tilde{S}_i \cdot \frac{G(s)}{1+G(s)}\right)^2 \qquad [\mathrm{Hz}^2/\mathrm{Hz}], \tag{3.33}$$

where i stands for the different noise sources with corresponding frequency noise contributions of $\tilde{S}_i$ (in units of $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$). In Fig. 3.6, a control loop bandwidth of 300 kHz is considered to determine the effective noise floor. Having a control loop with sufficient bandwidth and suppressing temperature coupling noise are the two main requirements to achieve sub-femtometer displacement sensitivity according to the current noise analysis.

However, it is important to note that in actual implementations, parasitic reflections can play a significant role in limiting the sensitivity of the displacement readout. Additionally, the laser actuator noise, which is not considered in this noise budget, could prove to be a non-negligible noise source.

3.4 Characterizing the frequency noise of a CTL

Before starting with the HS experiments, initial attempts were made to characterize the free-running laser frequency noise of the lasers used. The CTLs have a fiber-coupled laser output. The maximum output power of the laser is 50 mW just out of the FiberDock[10]. Due to a fiber coupling efficiency of slightly less than 80%, the target output is about 40 mW at 1570 nm. By iterating over the X- and Y-fiber alignment screws and lens adjustment screws, an output power of about 35 mW was achieved for both lasers at 1550 nm. It is possible to directly tune the laser power by changing the 'set current,' which varies the laser diode current. The lasers were primarily operated close to their maximum power during the experiments.

3.4.1 Generation and detection of beat note

As mentioned in Sec. 3.2.3, direct detection of the laser frequency is not feasible. Therefore, to characterize the frequency noise of one laser, the standard approach we followed is to interfere two laser beams and observe the beat note between them. The two lasers used are slightly frequency-offset from one another. To tune the frequency of the laser, we can either adjust the voltage of the piezoelectric elements attached to the grating or use wide-scan tuning with mechanical motors. When the beat note is within the bandwidth of the photodetector used, a sinusoidal signal oscillating at the beat frequency can be observed on an oscilloscope. On a spectrum analyzer, this appears as a prominent single peak at the beat frequency.

The most straightforward way to generate the beat note between two fiber-coupled lasers is by using a 50:50 fiber coupler/splitter. It is also possible to overlap the two beams in free space using a beam-splitter, allowing us to observe the beat note at the output ports of the beam-splitter. We then measure (or track) the beat note over a period of time and assume that both lasers contribute equally to the overall noise in the measured time series.

[10] A compact fiber coupler from Toptica Photonics AG, mounted to the laser system.

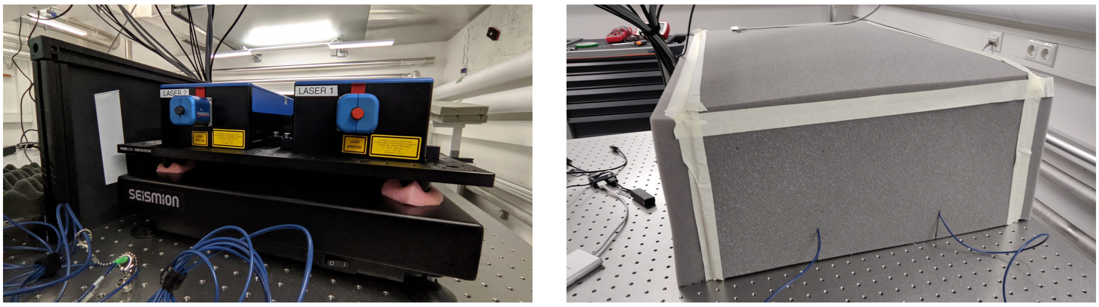

Figure 3.7: Photographs of the mechanical and acoustic isolation of the laser systems. Mechanical isolation is achieved using active vibration isolator in conjunction with a passive platform. Foam-based enclosures provide acoustic isolation.

3.4.2 Mechanical and acoustic isolation of the laser systems

The wide tunability of the CTL comes at the cost of a relatively high frequency noise and susceptibility to non-stationary excess noise. While observing the beat note between the two free-running lasers on a spectrum analyzer, we identified that the laser systems are very sensitive to external mechanical and acoustic noise coupling. To decouple the unwanted mechanical disturbances, the laser systems were placed on a separate optical breadboard that stands on four Sorbothane feet.

For acoustic isolation, an external enclosure was built around the laser systems. The enclosure is made of foam-cored black hardboard, with the interior filled with foam. Two additional foam layers were later added from the outside to improve isolation.

At a later point in time, with the intention of further enhancing mechanical isolation, a portable active vibration isolator was introduced. The ‘Seismion Reactio plus’ [96] active isolation platform was placed on the optical table, and the prior mechanical isolation system was moved on top of this platform. Fig. 3.7 shows photographs of these isolation attempts.

Before the installation of the active isolation platform, its performance was tested using two Trillium seismometers. One seismometer was placed on the platform, while the other was positioned on the ground next to it. The measurement results are shown in Fig. 3.8. In all three directions, at 20 Hz, the active isolation system suppresses the ground motion by 2 to 3 orders of magnitude. However, this comes with increased noise at lower frequencies in the two horizontal directions. The excess noise below 1 Hz can be attributed to the geophones or to tilt couplings. Since the platform is placed on top of the passively isolated optical table, and we are interested in decoupling mechanical vibrations at higher frequencies, the decision to install this platform was made.

3.4.3 Initial attempts of beat note measurement

Once the beat note is detected using a photodetector, the next task is to measure it over a period of time. The frequency time series thus obtained can be used to calculate the laser frequency noise. The most convenient way to obtain such a time series is by using frequency-tracking instruments such as phasemeters. Other options include frequency counting techniques, indirect frequency measurement using interferometric readout, and so on.

Initially, our choice was to use the phasemeter instrument on a commercial Moku:Lab platform [97]. It can track input frequencies between 1 kHz and 200 MHz, with a maximum tracking bandwidth of

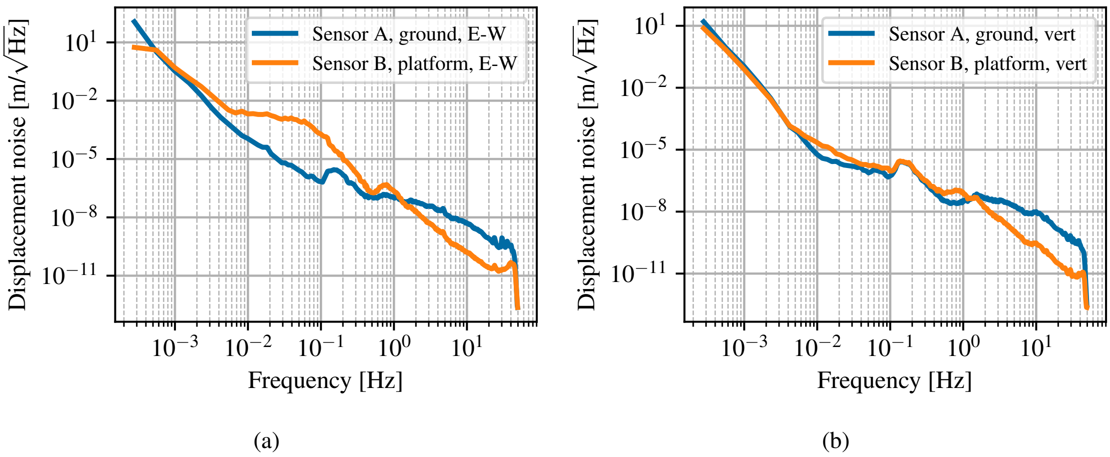

Figure 3.8: Performance test of the 'Reactio plus' active vibration isolation platform. (a) Displacement noise in the East-West horizontal direction. The other (South-North) horizontal direction exhibits very similar performance (not shown here). (b) Displacement noise in the vertical direction. We gain isolation at frequencies above 1 Hz, with a trade-off of increased noise at low frequencies in the two horizontal directions. (Measurement and analysis credits: Dr. Artem Basalaev.)

100 kHz. The tracking bandwidth defines the frequency range up to which the phasemeter can reliably measure the phase of an input signal. Unfortunately, because of this limited tracking bandwidth, the Moku:Lab phasemeter could not catch and lock onto the beat note created between two CTLs. The main reason for this is that the non-stationary frequency fluctuations of the lasers exceed the tracking bandwidth of the phasemeter. This prompted us to seek alternatives for measuring the beat note and characterizing the frequency noise of the lasers.

Frequency counter

The Siglent SDG2122X arbitrary waveform generator comes with a frequency counter feature [98]. This frequency counter can measure frequencies between 0.1 Hz and 200 MHz. However, it cannot be used directly off-the-shelf, as the measured frequency is only displayed on the screen, and there is no explicit way to record the time series data. Therefore, communication with the instrument was established using standard commands for programmable instruments (SCPI) over a local area network (LAN) interface. The Python code used to send and receive SCPI commands for data recording is provided in Appendix B.

Before using the frequency counter to measure the beat note between the lasers, its measurement noise floor was characterized by measuring a fixed tone generated from the same instrument. The sampling frequency of the frequency counter is fixed at approximately 1 Hz, and there is no option to change this. It was also found that the maximum measured variation from the fixed input tone is 1 Hz, which could result from a bit flip in the digital domain. Assuming this value as the value of the LSB, the measurement noise floor due to digitization can be obtained as

$$\tilde{S}_{f,\mathrm{meas}} = \frac{f_{\mathrm{LSB}}}{\sqrt{6 \cdot f_s}} \qquad [\mathrm{Hz}/\sqrt{\mathrm{Hz}}]. \tag{3.34}$$

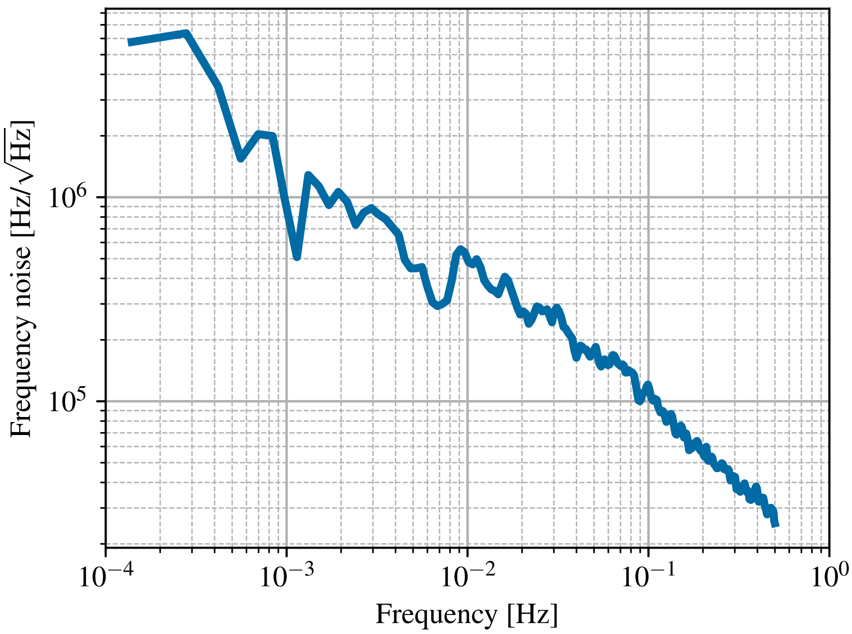

Figure 3.9: Frequency noise of a free-running CTL. The beat note between the two CTLs was measured for 2 h using the frequency counter of the SDG2122X instrument. The spectrum was calculated using the logarithmic power spectral density (LPSD) estimation technique [100] after removing the quadratic trend from the recorded time series.

Putting the values together, we obtain $\tilde{S}_{f,\mathrm{meas}} \approx 0.4\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$. The observed spectrum for the measurement of the fixed tone was in agreement with this calculation.

According to the Nyquist-Shannon sampling theorem, the sample rate must be at least twice the bandwidth of the signal to avoid aliasing [99]. Therefore, due to the fixed sampling rate of the frequency counter, we will not be able to capture frequency noise above 0.5 Hz. Additionally, the relatively high measurement noise floor will limit the applicability of the frequency counter as a long-term solution for measuring the heterodyne beat note.

Nevertheless, the initial characterization of the frequency noise of CTLs was conducted using the frequency counter readout. Fig. 3.9 shows the frequency noise spectrum for the free-running laser. The observed noise of $100\,\mathrm{kHz}/\sqrt{\mathrm{Hz}}$ at 0.1 Hz and the overall $1/f$ shape are consistent with the expected behavior of a free-running laser.

Delayed self-homodyne interferometric readout

Another measurement technique explored to characterize the frequency noise of the CTL is utilizing the delayed self-homodyne interferometer. This technique is an established method for measuring the line-width of diode lasers [101]. The basic idea is to convert frequency fluctuations into variations of light intensity within a Mach-Zehnder type interferometer. Our approach to realizing such an interferometer using single-mode polarization-maintaining fibers is sketched in Fig. 3.10. In the interferometer, the laser field interferes with a delayed replica of itself, and the interference signal is detected using a photodetector.

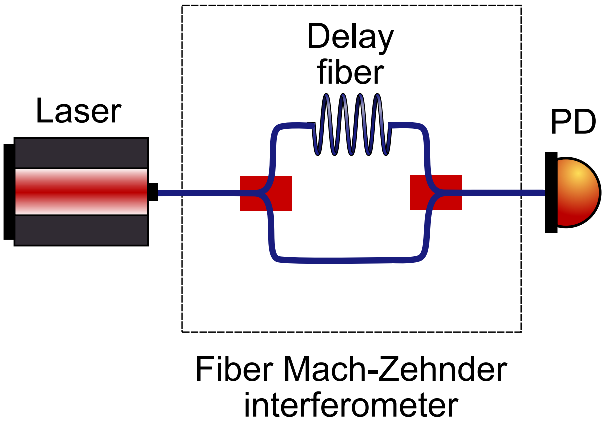

Figure 3.10: Experimental setup of the delayed self-homodyne interferometer. The output from the laser is split using a 50:50 fiber splitter. One path is delayed using a 10 m long polarization-maintaining patch cable before being combined with the other path. The output of the 50:50 fiber coupler is detected using a photodetector. (PD: photodetector.)

To characterize the free-running laser frequency noise and to probe this readout technique for laser frequency stabilization methods, the initial plan was to record the photodetector output voltage. From this time series, voltage noise spectrum is calculated. Later, the spectrum is converted into frequency noise using an appropriate scaling, which is discussed below.

However, once the fiber Mach-Zehnder interferometer was set up, it was observed that over a span of less than 30 s, the free-running drift itself exceeded a single interferometric fringe. As a result, a short data recording (approximately 1 s) at a high sampling rate (of 1 MHz) was the only feasible option.

Initially, the laser frequency is scanned by applying a triangular voltage signal to the piezoelectric elements attached to the grating. The scan amplitude is set such that it covers one complete interferometric fringe. From Eq. 2.2, we know that such a scan should result in a sinusoidal variation of the interferometer output. Thus, the recorded photodetector output voltage is fitted with a sine function, and the period of that fit provides the FSR in units of time. The gradient of the photodetector output with respect to time dV/dt is also calculated. Fig. 3.11(a) shows the measurement data along with the fit and its gradient.

The analytical expression for the FSR of the fiber Mach-Zehnder setup is given by $c/(nL)$, where L is the length of the delay fiber (which is 10 m) and n is the refractive index of the fiber. Using the approximation $n = 1.5$ for glass, we obtain an FSR of 20 MHz for the current experimental setup.

To estimate the free-running laser frequency noise, the photodetector output is recorded (using the Moku:Lab data logger instrument) when it is close to the mid-fringe value. From the recorded voltage time series, the noise spectrum $\tilde{S}_{V,\mathrm{free}}$ in units of V/$\sqrt{\mathrm{Hz}}$ is calculated (see Fig. 3.11(b)).

Now, to obtain the free-running laser frequency noise $\tilde{S}_{f,\mathrm{free}}$ in units of Hz/$\sqrt{\mathrm{Hz}}$, we utilize the FSR in two different units along with the gradient of the photodetector output at mid-fringe. That is,

$$\tilde{S}_{f,\mathrm{free}} = \tilde{S}_{V,\mathrm{free}} \cdot \left|\frac{dV}{dt}\right|^{-1}_{\mathrm{mid-fringe}} \cdot \frac{\mathrm{FSR(in\ Hz)}}{\mathrm{FSR(in\ s)}}. \tag{3.35}$$

In the actual experiments, since the short recording of the interferometric output is not exactly at

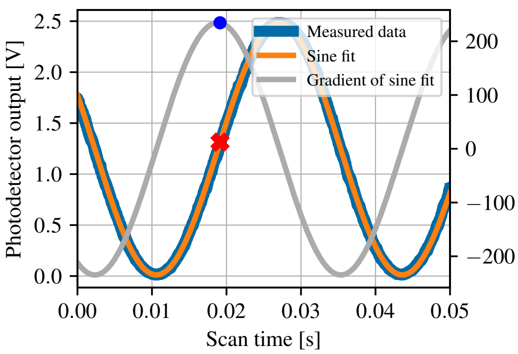

(a)

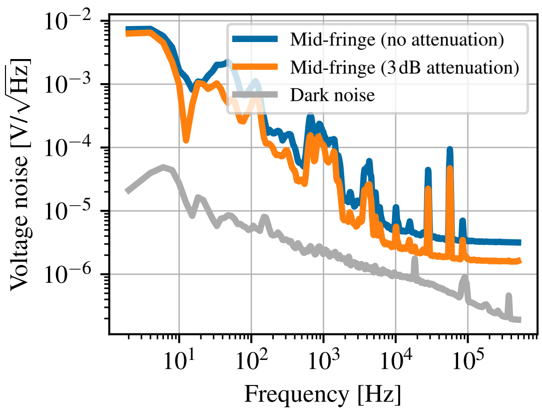

(b)

Figure 3.11: (a) Scan of the fiber Mach-Zehnder interferometer over a fringe. The measured data is fitted with a sine function, and its gradient is calculated. The red cross-mark on the sine fit corresponds to the average value of the interferometer readout, close to mid-fringe. The corresponding gradient value is indicated by a blue dot. (b) Voltage noise spectra of two separate measurements taken close to the interferometer mid-fringe. Before calculating the spectrum, each time series is detrended by subtracting its mean value. The dark noise of the measurement device is also plotted.

mid-fringe, the gradient of the fit corresponding to the mean value of the recorded time series is considered.

However, initial experimental result showed a white noise floor in the spectrum for Fourier frequencies above 10 kHz. It was confirmed that this was not due to the measurement noise floor. The dark noise[11] was not a limitation when measuring with the free-running laser. When a 3 dB optical attenuation is introduced to the interferometric output, we see in Fig. 3.11(b) that the white noise floor decreases as well. Comparing this with the square root of the mean mid-fringe values in both cases, it was concluded that the white noise is mostly due to shot noise during the measurement.

Fig. 3.12 shows the laser frequency noise spectra of the same laser obtained from two different measurements. As the measurement band is now centered around kilohertz, the piezoelectric resonances and all other mechanical and acoustic noises coupling to the laser system become apparent in the spectrum. The noise coupling was tested and verified by intentionally injecting noise into the experimental setup.

Although the delayed self-homodyne interferometer is an option for reading out the laser frequency noise, using this method for heterodyne cavity-tracking is not so trivial. Additional interferometer setups and long delay fibers are susceptible to extra noise couplings affecting the readout. The dependency on the intensity of the interferometer output increases the shot noise contribution, which appears to limit the measurement at high frequencies. Additionally, this method of delayed self-homodyne readout requires one complete scan of the fringe to get the scaling parameter. Any changes to the interferometer output necessitate new scan data. Even with that, the calculation and scaling methods are not exact and involve some approximations (e.g., optical length inside the delay fiber).

Both frequency readout methods explored here have their own limitations. Therefore, a better

[11] Dark noise is calculated using the time series values when there is no input power on the photodetector.

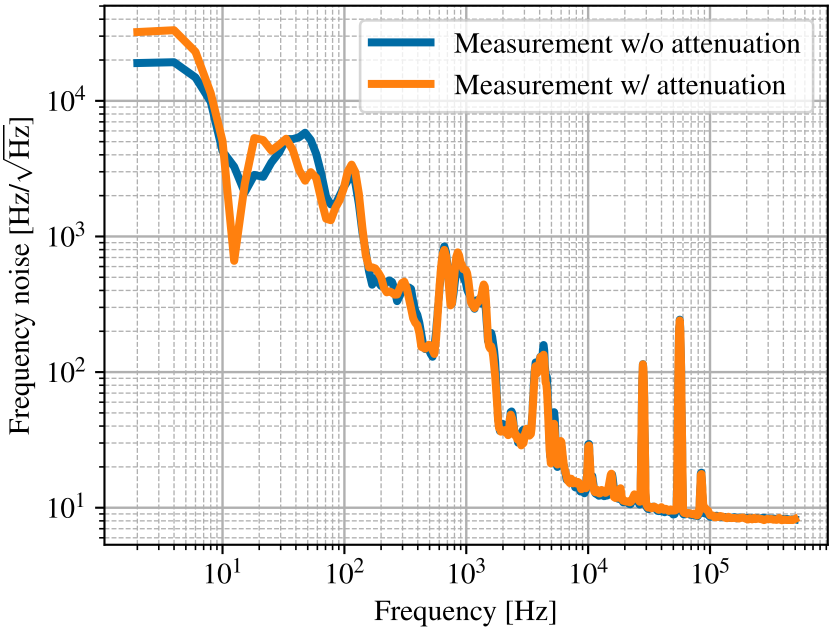

Figure 3.12: Frequency noise of a free-running CTL measured using a delayed self-homodyne interferometer. For the two spectra shown in this plot, different scaling parameters were used to convert from the voltage noise spectra. A similar measurement was conducted for the other CTL, revealing that both lasers exhibit very similar laser frequency noise spectra.

frequency readout method is anticipated for the HS scheme when used as a displacement sensor, which was developed later. With that, a proper and reliable characterization of the frequency noise of the CTL was also feasible and is presented in Sec. 5.8.1.

3.5 HS using analog electronics

After the initial characterization attempts of the laser frequency noise of CTLs, we proceeded to experimentally demonstrate the HS concept. In [85], such a demonstration was performed, where the heterodyne demodulation and laser-locking are handled in the digital domain using FPGAs. We did not yet have such a DSP system, but we were interested in testing and understanding the HS concept. Therefore, our initial approach involved heterodyne demodulation carried out using analog electronic components.

The fiber-coupled CTL outputs are interfered in a 50:50 fiber coupler/splitter (PN1550R5A2 from Thorlabs). A small portion of the interfered field is picked off using a 90:10 fiber coupler/splitter (PN1550R2A1 from Thorlabs). This pick-off is later used for demodulation and readout. Two fiber collimators were used to direct the superposed laser fields into free space.

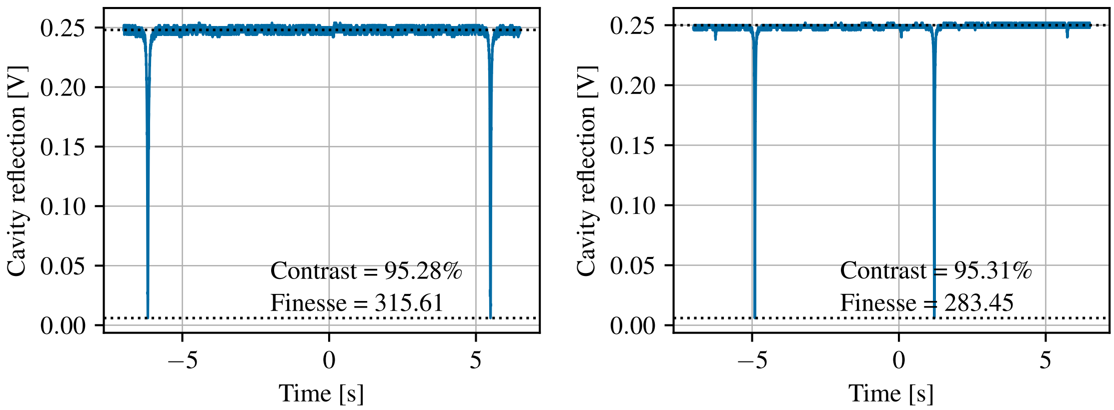

(a) 5 cm long cavity

(b) 10 cm long cavity

Figure 3.13: Table-top cavity responses after the alignment. The dotted horizontal lines represent the maximum and minimum values considered for calculating the fringe visibility. The cavity finesse value was calculated for both peaks in each case, and the average value is noted on the plots.

3.5.1 Cavity alignment and characterization

The demonstration of HS requires two optical cavities to which the two lasers are locked. For this purpose, two table-top cavities were aligned using commercial mirrors and mirror mounts. To achieve resonance of a laser field in an optical cavity, it is important to have precise spatial matching between the input laser field and the cavity modes. The laser beam profiles emerging from the two fiber collimators are measured separately and mode-matched to the fundamental TEM_{00} mode of their respective cavities. A detailed explanation of these initial preparations and mode-matching is provided in Appendix C.

Once the cavities are aligned, they are characterized by calculating the fringe visibility and finesse values. For this purpose, the cavity reflection is recorded when the laser frequency is tuned over an FSR. Fringe visibility or contrast can easily be calculated using Eq. 2.9. For estimating the cavity finesse, the Python package ‘OpenQlab’ [102] was used. The responses of both aligned cavities are shown in Fig. 3.13. Finesse values near the theoretical value of 312 confirm that there are no major losses in the cavity. With these relatively low-finesse cavities, we have achieved sufficient fringe visibility, exceeding 95%, using standard commercial mirrors without any additional polishing to demonstrate our interferometer scheme.

3.5.2 Experimental setup

The full experimental setup for realizing the HS using analog electronics is shown in Fig. 3.14. The beat note between the lasers is detected using a high-speed fiber-coupled InGaAs detector (DET08CFC/M from Thorlabs) with a bandwidth of 5 GHz. Similar free-space photodetectors with lenses (DET08CL/M from Thorlabs) are used to detect cavity reflections.

The beams exiting the fiber collimators are adjusted to be p-polarized. Because of this, upon interacting with the PBS, the entire beam passes through. The QWP then converts the linearly polarized beam into an elliptically polarized one. Afterwards, the interaction with the cavity introduces an

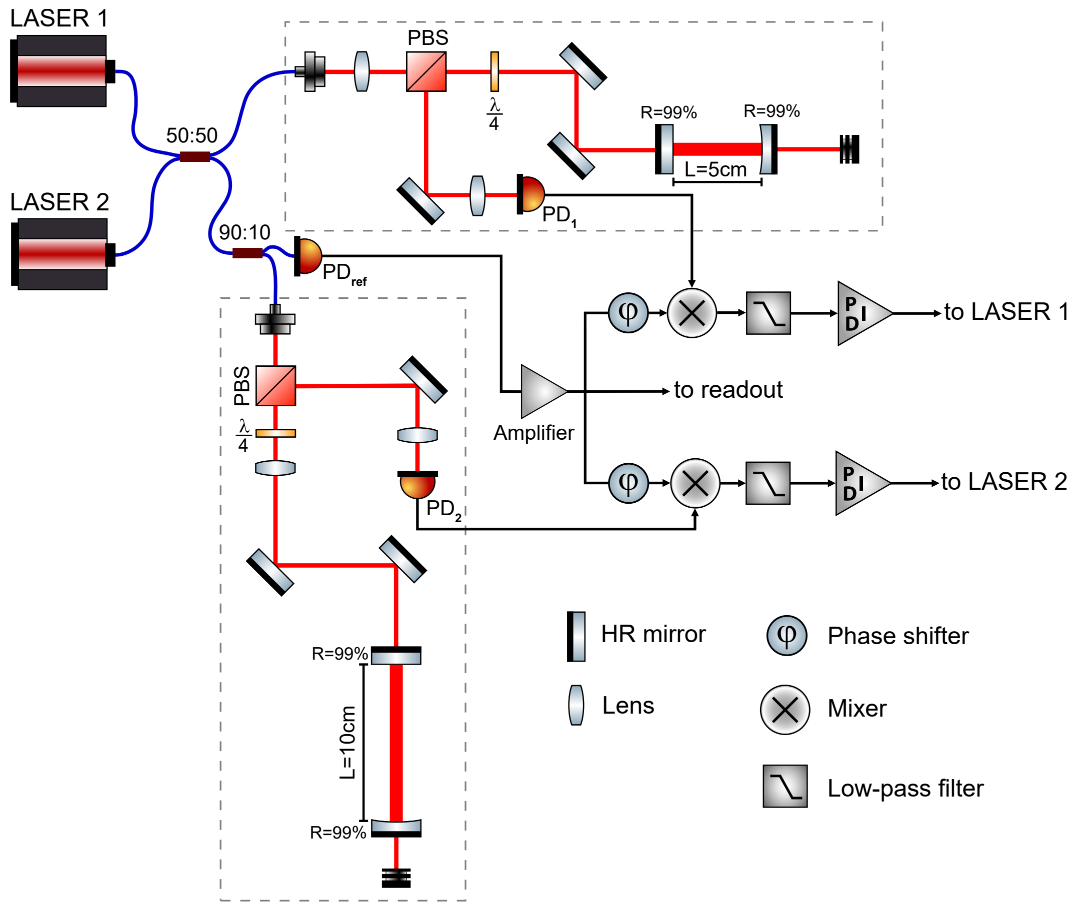

Figure 3.14: Experimental setup of the HS using analog electronics. Reflections from the cavities are demodulated with the reference beat note using analog mixer and low-pass filter components. The demodulated signal in each case is fed to a servo (proportional-integral-derivative (PID) controller) to control the laser frequency. In this way, both lasers are locked to the resonances of their respective optical cavities. The optical layout is not drawn to scale. (PBS: polarizing beam-splitter, PD: photodetector.)

additional phase shift to the laser field. The reflected beam just before the PBS is therefore s-polarized, which gets reflected by the PBS onto the photodetector.

The RF part of all the photodetector outputs is separated using bias-tees (ZFBT-4R2GW+ from Mini-Circuits), and the reflected signal from the cavity is mixed with the reference beat note. The frequency mixers used are ZX05-C24-S+ from Mini-Circuits, which have a bandwidth of 0.3–2.4 GHz. The output of the mixer is low-pass filtered (using SLP-150+ from Mini-Circuits) to remove the harmonics resulting from the analog mixing process. The obtained demodulated signal is then passed to the servo for the use in a closed-loop control.

The CTLs come with a digital controller called 'DLC pro.' This controller offers a laser lock function along with two PID control loops: one for actuation using piezoelectric voltage and the other for laser diode current. For all the lock acquisition and control attempts discussed in this thesis, we utilized these functionalities offered by the DLC pro. The generated error signals are provided as inputs to the DLC pro BNC I/O ports on the front panel.

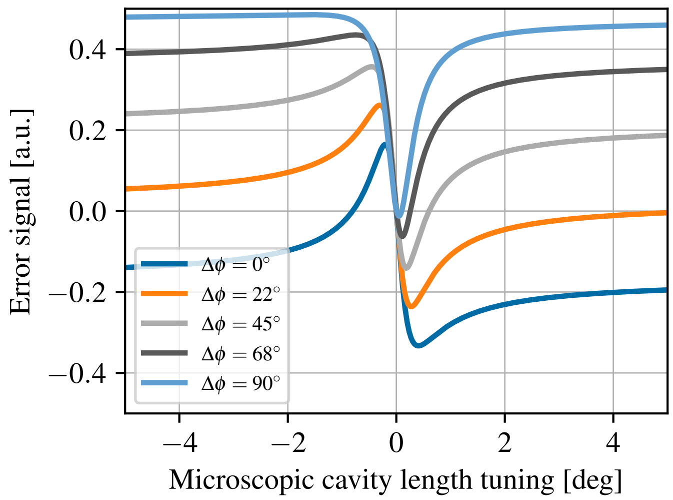

Figure 3.15: Influence of the demodulation phase on the error signal in the HS scheme. The reflected signal from the cavity is demodulated with a 500 MHz heterodyne beat note at different demodulation phases. In the Finesse3 model, the error signal is obtained using a photodetector combined with a mixer. Hence, under ideal conditions, $\Delta\phi = 0^\circ$ is the required demodulation phase for the simulation. However, for the considered simulation model, a demodulation phase of about 22° would be close to the expected value. At $\Delta\phi = 90^\circ$, the obtained error signal is almost unusable for the control loop.

However, scanning the laser frequency around the cavity resonance did not directly yield the expected error signal. The reason for this was then narrowed down to not having the proper demodulation phase.

3.5.3 Role of the demodulation phase

As explained in Sec. 3.2.1, to achieve heterodyne demodulation, a phase shift of $\pi/2$ was introduced to the reference beat note before using it as the demodulation signal. In [85], the phase of the reference beat note was initially adjusted through the placement of the reference photodetector. Additionally, to avoid changes in the demodulation phase when the beat frequency varies, the secondary laser was offset phase-locked to the primary laser such that the frequency fluctuations are common to both.

However, in the current analog experimental setup, this fixed phase shift was not properly accounted for. Hence, the heterodyne demodulation occurred at random phases, depending on the frequency of the demodulation signal. The different lengths of optical fibers and electrical cables introduce a differential group delay between the signals to be mixed, causing the demodulation phase to change.

The influence of the demodulation phase on the generation of the error signal was simulated using Finesse3 [77] for the 5 cm long cavity setup. Instead of relying on an ideal experimental setup, the simulation model takes into account the measured beam profile, the mode-matching scheme, and the associated physical distances. Fig. 3.15 shows the simulated results, where the requirement for the correct demodulation phase becomes evident. We also simulated the required demodulation phase to obtain a proper error signal for different heterodyne frequencies. In each case, the required demodulation phase $\Delta\phi$ is expressed in terms of length matching Δl using the relation

$$\Delta l = \Delta\phi \cdot \frac{c}{2\pi f_{\mathrm{het}}}. \tag{3.36}$$

Analytically, the required demodulation phase changes with the heterodyne frequency, but the length mismatch remains fixed. This also implies that a fixed phase tunability (in °/GHz) is needed to tackle this issue.

Accordingly, we introduced a variable phase shifter before the mixer (as sketched in Fig. 3.14) in both control loops. The phase shifter used (TKE-45-4-S from SHX [103]) offers a maximum tunability of 45°/GHz simply by manually adjusting the knob. Although the working principle of these phase shifters was tested individually, their influence on obtaining the proper demodulation phase was negligible.

Issues with length mismatch

Subsequently, our attempts to understand the role of the demodulation phase led us to an extensive investigation of the reflected signal from the cavity and the reference beat note. Both signals were coupled to a digital oscilloscope (RTP084 from Rohde & Schwarz) with a bandwidth of 8 GHz. When the laser frequency is off-resonance, we theoretically expect to see both signals in phase and at the same heterodyne frequency. Unfortunately, this was not the case in the analog experimental setup. A clear phase difference was observed between the reference beat note and the cavity reflection. When scanning one of the lasers off-resonance, we observed a phase difference accumulating between the two signals. It is clear from Eq. 3.8 that such an additional phase term will remain in the sine and cosine parts of the mixed signal and will not get filtered out. This results in improper heterodyne demodulation.

If this extra phase difference is because of the time delay between the two signals, then changing the length of the cables used should affect how the phase difference changes while scanning off-resonance. This was tested by introducing different cable lengths in the two paths. Depending on the length of the cable, either the phase difference increased or decreased. Therefore, it was possible to eliminate this additional phase difference due to delay by finding an optimal length combination.

For a given heterodyne frequency, the behavior of the error signal depending on the length mismatch is shown in Fig. 3.16. Unfortunately, such phase differences due to differential group delay are also frequency-dependent. Therefore, the optimal cable length combination is only valid for the current heterodyne frequency. This is a significant disadvantage when it comes to obtaining the correct demodulation phase in an analog demodulation setup.

As an example, assume a 1 m path length difference between the cavity reflection and the reference beat note, purely due to the electrical cable length. The arrival time delay between these two signals at the mixer would then be about 5 ns. If we assume a fast scan of the laser frequency covering a significant part of the FSR, say 100 MHz/s, this results in a 0.5 Hz oscillation in the error signal because of the delay. These unwanted oscillations could also be so strong, as shown in Fig. 3.16(d), that they mask the actual zero-crossing of the error signal, making it unusable for the control loop. Such effects were observed in the experiment before optimizing the length mismatch.

3.5.4 Preliminary results

Even with the limited ability to optimize the demodulation phase across all frequencies using analog electronics, a proof-of-principle demonstration of the HS was carried out. The electrical cable lengths were adjusted to optimal conditions, and later, the heterodyne beat note was consistently set to the

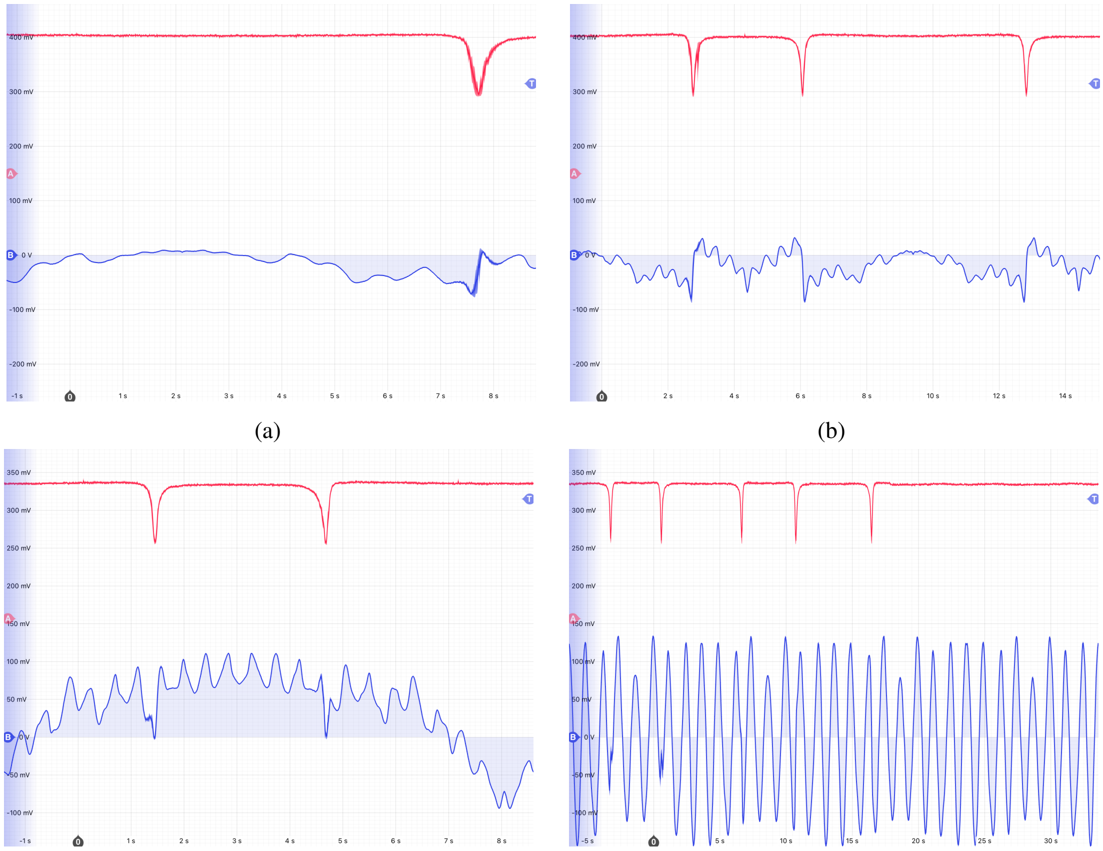

(a) (b) (c) (d)

Figure 3.16: Effect of delay between the reference beat note and cavity reflection on error signal generation. Red trace indicates the DC part of the cavity reflection, while the blue trace represents the demodulated signal.
(a) When the delay is mostly compensated using optimal cable lengths in the two paths. The matched length ensures a proper demodulation phase.
(b) When an additional short cable of length 12 cm is added to the cavity reflection path. This length mismatch results in an unwanted delay, which manifests as oscillations in the error signal. A greater length mismatch can cause oscillations that mask the actual error signal.
(c) The same cable lengths result in a different behavior when the heterodyne frequency is shifted. Here, we observe a slow oscillation superimposed on the fast oscillation. The demodulated signal has non-symmetric zero-crossings, making it unusable for the control loop.
(d) Oscillations due to delay (or length mismatch) can also be so pronounced that they nearly mask the phase shift at resonance points. Here, we observe that oscillations persist even when the laser is off-resonance and the frequency scan is turned off.

same value. To reduce the influence of external noise on the length stability of the table-top cavities, an enclosure was built around the in-air optical setup.

For the readout part, we used the phasemeter instrument on the Moku:Pro platform [97]. The upgrade from Moku:Lab to Moku:Pro comes with a maximum phasemeter tracking bandwidth of 1 MHz and a maximum input frequency of 300 MHz. While a higher tracking bandwidth offers the potential to track the beat note between the CTLs, the signal bandwidth is not sufficiently high. Due

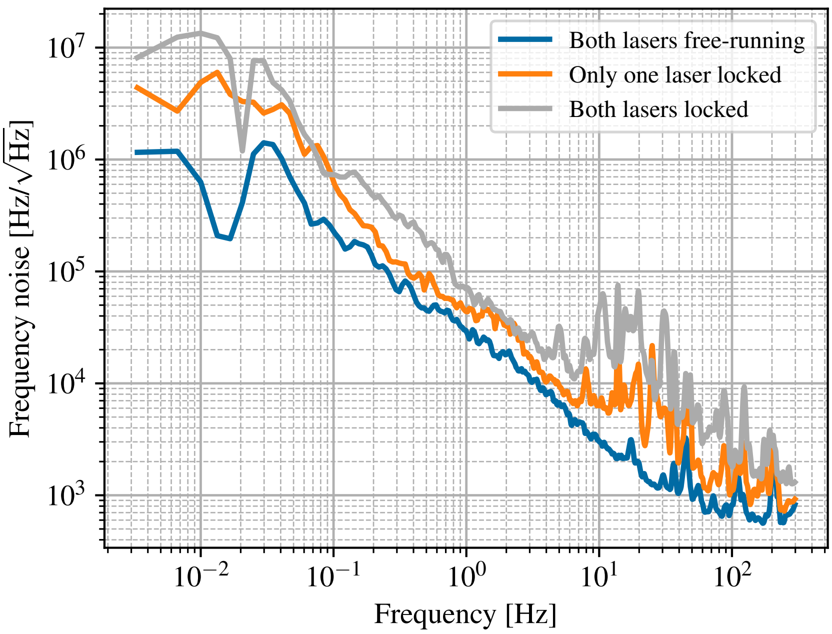

Figure 3.17: Frequency noise of the heterodyne beat note when lasers are stabilized to table-top in-air cavities using analog HS technique. The mechanical instability of the table-top cavities makes the laser frequency noisier than when it is left free. Above a Fourier frequency of 10 Hz, acoustic noise coupling is dominant. The beat note was tracked using the Moku:Pro phasemeter instrument. The spectra are calculated after removing the linear trend from all three time series.

to the bandwidths of the frequency mixers used, the beat note was maintained above 300 MHz. To enable the phasemeter to track the tone, the beat note frequency was down-converted by mixing it with a stable RF signal from the Rohde & Schwarz SMC100A signal generator.

Fig. 3.17 shows the results from one such stabilization attempt. Instead of stabilizing the laser frequency, we are increasing its noise by forcing it to follow the resonance of the table-top in-air cavity. When only one laser is locked and the other is free-running, the expectation is to see a spectrum that is $\sqrt{2}$ times lower than the noise spectrum when both lasers are free. This is because, the free-running noise of the other laser should dominate throughout. However, in our case, the mechanical instability of the table-top cavity is the dominant factor. This becomes clear when looking at the spectrum for both lasers locked to their respective cavities. The instability of both cavities shows up in the measurement, resulting in an even higher noise spectrum.

3.5.5 Limitations and the way forward

Even if laser frequency noise suppression is not achieved, experiments with the analog HS confirm the ability to read out the mechanical motion of the cavities. However, experience has shown that continuing with the analog HS scheme is not advisable when the interferometer scheme is to be used in displacement sensing applications.

The main limitation of the analog version lies in obtaining the proper demodulation phase. The additional delay, which can arise from both optical and electrical path length mismatches, causes unwanted oscillations in the error signal. These oscillations can be strong enough to mask the required zero-crossings of the error signal (see Fig. 3.16(d)). There is no simple way to set the required demodulation phase when using analog components. An analog compensation circuit would need to be implemented, which involves complex RF engineering. Even with this, the system will not be as flexible as digital solutions.

It is also suspected that when both lasers are locked to the cavities, a cross-coupling between the other laser and cavity may occur whenever there is a length mismatch in one of the heterodyne demodulation loops. This effect has not been studied or well understood. Factors such as high non-stationary noise of the lasers, mechanically unstable table-top cavities, and limited readout bandwidth have made such an investigation difficult. To overcome the challenges in the analog HS technique, a digital solution was studied instead (as shown in [85]), where the demodulation phase is easily tunable.

However, when the interferometer scheme is intended to be used as a fringe-scale displacement sensor, the readout bandwidth of the heterodyne beat note plays a major role as well. It follows from Eq. 3.16 that the maximum one-way displacement tracking range for a given optical setup can be expressed as

$$\Delta L_{\mathrm{max}} = \frac{\lambda}{2} \cdot \frac{\mathrm{BW}}{\mathrm{FSR}} = L \cdot \frac{\mathrm{BW}}{f}, \tag{3.37}$$

where λ is the wavelength of the laser with frequency f, and BW is the bandwidth of the frequency readout instrument.

From Eq. 3.37, it is evident that the operating range ΔL_{max} is proportional to the frequency readout bandwidth. Hence, the larger the bandwidth, the greater the operating range. This holds true if the secondary laser is stabilized using a different method or if the heterodyne beat is generated using another kind of frequency reference. When the HS technique is employed to lock both cavities (as done here), the maximum operating range is limited to half of the FSR.

To keep the interferometer setup compact, we would employ short cavities with lengths of a few centimeters. The commercially available readout instrument (Moku:Pro phasemeter), with a 300 MHz bandwidth, can track the 5 cm cavity displacements only up to a maximum of several tens of nanometers. Achieving a fringe-scale operating range will require a frequency readout instrument with a bandwidth in the gigahertz range. Hence, we focused on developing a high-bandwidth readout system for the heterodyne frequency based on FPGAs and subsequently integrating the HS-based digital control system with it.

Part II

High-speed digital readout and control systems

CHAPTER 4

Radio-frequency system-on-chip

In this part of the thesis, we will discuss the development of digital readout and control systems based on a radio-frequency system-on-chip (RFSoC) using field-programmable gate-array (FPGA) algorithms. Such an approach is essential for the successful demonstration of the heterodyne cavity-tracking, as discussed in Chapter 3. Before delving into the technical details of our development, a basic introduction to the platform we used and the required initial setups are provided in this chapter.

4.1 Digital signal processing (DSP)

DSP refers to the manipulation and analysis of signals in the digital domain. In our case, the signals of interest are the beat note between the two lasers and the reflected signals from the cavities. When selecting a DSP system, its speed, efficiency, and ability to handle real-time processing are important factors. After the signal of interest is converted to digital form by an analog-to-digital converter (ADC), the analysis and processing can be performed in various ways.

The conventional and well-known central processing unit (CPU) is a straightforward option for executing most processing tasks but is generally optimized for sequential processing. Thus, although CPUs excel at general-purpose computing, they are not the best choice for parallel processing tasks, which are often required in DSP applications. Today's technological advancements have enabled the use of graphics processing units (GPUs), specifically designed for parallel processing and optimized for the simultaneous handling of large data blocks. With regard to DSP, GPUs can efficiently handle complex computations, particularly those that can be parallelized.

However, in our application, we have a heterodyne frequency, which is a very fast signal typically in the gigahertz regime. The processing of this signal needs to be carried out with nanosecond delays, leaving only FPGAs as an option for the DSP.

4.1.1 Field-programmable gate-arrays (FPGAs)

An FPGA is an integrated circuit that can be programmed and reconfigured to implement various digital logic functions. They contain arrays of logic blocks that can be wired together via programmable interconnects. Unlike CPUs and GPUs, FPGAs allow for custom hardware-level designs tailored to specific tasks, optimizing performance in particular applications. Individual logic blocks can be configured to perform basic combinational logic operations like AND and XOR, and information can be stored using registers. Multiple logic blocks can then be configured in series and/or parallel to realize complex DSP algorithms [73].

FPGAs can achieve lower latency[1] in processing data compared to CPUs and even GPUs. This advantage arises because FPGAs are programmable logic (PL) and hence can be considered pure hardware devices. They also eliminate the overhead associated with operating systems and complex scheduling. Furthermore, FPGAs offer deterministic performance in terms of timing and latency, which are crucial for real-time DSP applications that we aim for.

FPGAs are programmed using code written in a hardware description language (HDL). Popular examples of HDL include VHDL and Verilog, where VHDL stands for very high speed integrated circuit (VHSIC) HDL. This code is then translated into hardware-level logic by dedicated software. All the FPGA-based developments discussed in this thesis are coded using VHDL and compiled with the Vivado Design Suite 2020.2.2 software from Xilinx Inc.[2] The chosen FPGA development platform is the Zynq™ UltraScale+™ RFSoC ZCU111 evaluation kit [104].

Integer arithmetic

The algorithms running on FPGAs are implemented using integer arithmetic, where each number is typically represented as either an unsigned or a signed N-bit integer. A signed N-bit number can represent both positive and negative integers ranging from -2^{N-1} to $2^{N-1} - 1$. Unsigned numbers can represent only positive integers between 0 and $2^N - 1$. These numbers can then be scaled in various ways; for example, scaling each integer by 2^{-N} is one of the common mappings used.

Integer addition, subtraction, and multiplication are well-supported arithmetic operations on an FPGA. However, division is not well-supported due to its inefficient consumption of logic resources [73]. Consequently, division in PL is primarily limited to inverse powers of 2, which can be efficiently achieved by bit-shifting the binary number to the right by a corresponding number of bits. Division by non-powers of 2 can be performed using a combination of multiplication and bit-shifting, but such integer arithmetic was not utilized in our implementations.

4.2 What is an RFSoC?

An RFSoC is an integrated circuit that combines RF signal processing with DSP on a single chip. Its architecture is designed to handle high-frequency analog signals and convert them directly to digital signals, and vice versa. This feature eliminates the need for separate data converter units as well as intermediate frequency stages in certain applications. RFSoCs integrate multiple system components, including ADC, digital-to-analog converter (DAC), processors, and FPGA logic fabric, all on one chip. It is the first of its kind where RF data converters and PL are integrated on a single chip, enabling direct processing of analog signals in the gigahertz regime.

Using RFSoCs instead of external data converters with traditional FPGAs presents both advantages and disadvantages. RFSoCs reduce both complexity and latency. In contrast, integrating external data converters adds complexity to the design and may require additional interface circuitry and signal conditioning. Furthermore, the additional communication overhead between the external data converter and the FPGA can introduce latency, which may be a critical factor in real-time applications.

[1] Low latency is the ability of a system or network to provide fast responses with minimal delay.

[2] Now acquired by Advanced Micro Devices, Inc.

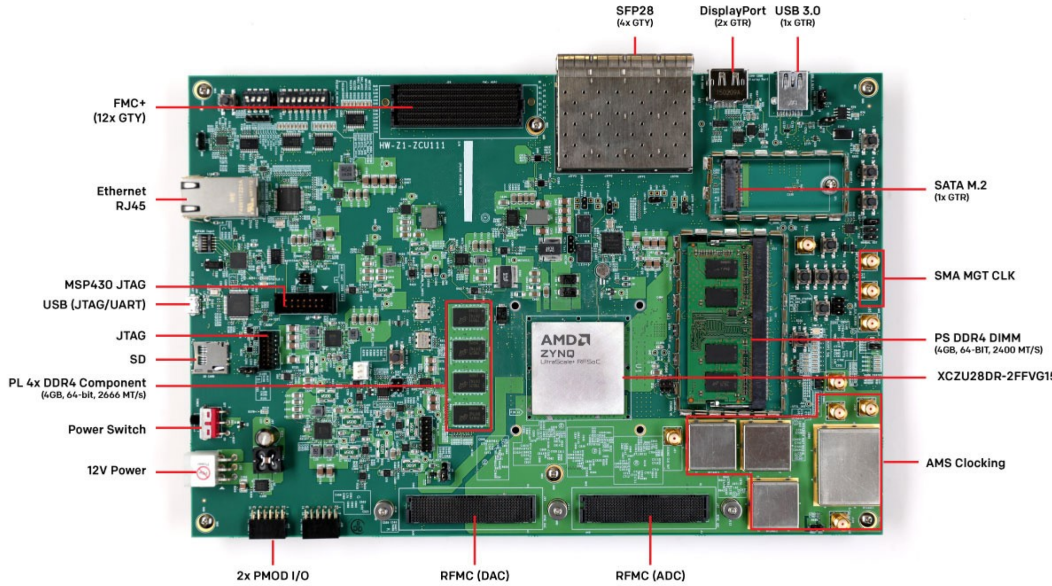

Figure 4.1: Overview of the ZCU111 evaluation board. The RF data converters are accessed via an analog interface, specifically the XM500 RFMC balun transformer add-on card. This card is attached to the RFMC expansion connectors available on the board. (Image credit: Xilinx Inc.)

However, when external data converters are chosen, a wide range of specialized converters is available, which may better suit specific requirements in terms of resolution or input voltage range. Additionally, external data converters can be easily upgraded in the future without requiring changes of FPGA itself.

4.3 Zynq UltraScale+ ZCU111

For our digital readout and control application, the Zynq UltraScale+ RFSoC ZCU111 evaluation board was chosen as the basis [104]. When tracking the beat note between the lasers, having a large bandwidth is advantageous in many regards. In our particular displacement sensing use case, such a bandwidth is necessary to push the operating range from narrowband to fringe-scale.

The ZCU111 board features an XCZU28DR RFSoC, which combines both a processing system (PS) and PL in the same device. The RFSoC supports eight 12-bit resolution 4.096 giga samples per second (GSPS) ADCs and eight 14-bit 6.554 GSPS DACs. As for the PS, it includes the Arm® flagship Cortex®-A53 64-bit quad-core processor and Cortex-R5 dual-core real-time subsystems on chip. A 4 GB 64-bit DDR4 SODIMM is attached to the PS, serving as random access memory (RAM). Similarly, the PL has direct access to a separate 4 GB 64-bit DDR4 component. A full overview of the ZCU111 evaluation board is shown in Fig. 4.1.

The eight RF-ADCs are arranged across four banks[3] (Tile 224 to Tile 227), with each bank containing

[3] In this context, the term ‘bank’ refers to a grouping of related components. Each group can operate independently or collectively, depending on the chosen design.

two ADC input channels. The eight RF-DACs are distributed across two banks (Tile 228 and Tile 229), with each bank having four DAC output channels. The ZCU111 board provides a pair of expansion connectors for the clock and signals of these data converters. An additional add-on card is typically used to interface with these converters.

4.3.1 Analog interface to RF data converters

The FMC-XM500 balun[4] transformer is a plug-in card that connects to the ZCU111 board. The card includes on-board high-frequency and low-frequency baluns, as well as SMA connectors that are then in direct contact with the RFSoC data converter banks. When the XM500 add-on card is used as analog interface to the RFSoC data converters, it comes with its own connectivity features:

- Two ADCs from Tile 224 (ADC224_T0_Ch0 and ADC224_T0_Ch1) and two DACs from Tile 229 (DAC229_T1_Ch2 and DAC229_T1_Ch3) are routed to low-frequency baluns and then to SMAs. Thus, the differential signals of the data converter are converted to single-ended signals and are limited to the 0–1 GHz range.
- Two ADCs from Tile 225 (ADC225_T1_Ch0 and ADC225_T1_Ch1) and two DACs from Tile 229 (DAC229_T1_Ch0 and DAC229_T1_Ch1) are routed to high-frequency baluns and then to SMAs. Consequently, we obtain single-ended signals within the 1–4 GHz range.
- The remaining four RF-ADCs and four RF-DACs are routed to SMAs as differential signals without any internal baluns or filters, providing access to the full frequency range.

Although these features may be beneficial for certain applications, the internal baluns of the XM500 are not useful for our purpose. Given that handling of RF electronics in the gigahertz regime is quite sensitive and challenging, we did not modify the front-end of the XM500.

4.3.2 RF data converter clocking

The ZCU111 board provides both fixed and variable clock sources for the RFSoC. We have used the fixed 100 MHz clock from the on-board Silicon Labs SI5341B clock generator to clock the PS. All these fixed and variable clock sources are limited to frequencies around 800 MHz. Hence, to enable the RF data converters to function properly, a different clocking scheme is essential. Initially, getting the clocking scheme operational proved to be challenging. Below is a brief summary of the clocking structure and how it was implemented in our work.

To generate RF data converter sample clocks, it is necessary to use the primary on-board reference phase-locked loop (PLL) (LMK04208 from Texas Instruments, Inc.) and the on-board RF PLLs (LMX2594 from Texas Instruments, Inc.). The ZCU111 features one LMK04208 PLL and three LMX2594 PLLs. The former is not directly connected to the data converter banks; instead, one of its clock outputs, named ‘CLKout4,’ serves as the reference input to all three LMX2594 PLLs.

The LMX2594 clocks can be configured either as direct RF clocks or as reference clock sources for the internal PLL contained within the RFSoC data converter tile. Each LMX2594 has two outputs: ‘RF outA’ and ‘RF outB.’ Of the three available LMX2594 RF PLLs, the first two clock the four ADC banks, while the remaining one clocks the two DAC banks. Clocking can be configured via the

[4] Originally standing for ‘balanced to unbalanced’; now derived from ‘balancing unit.’

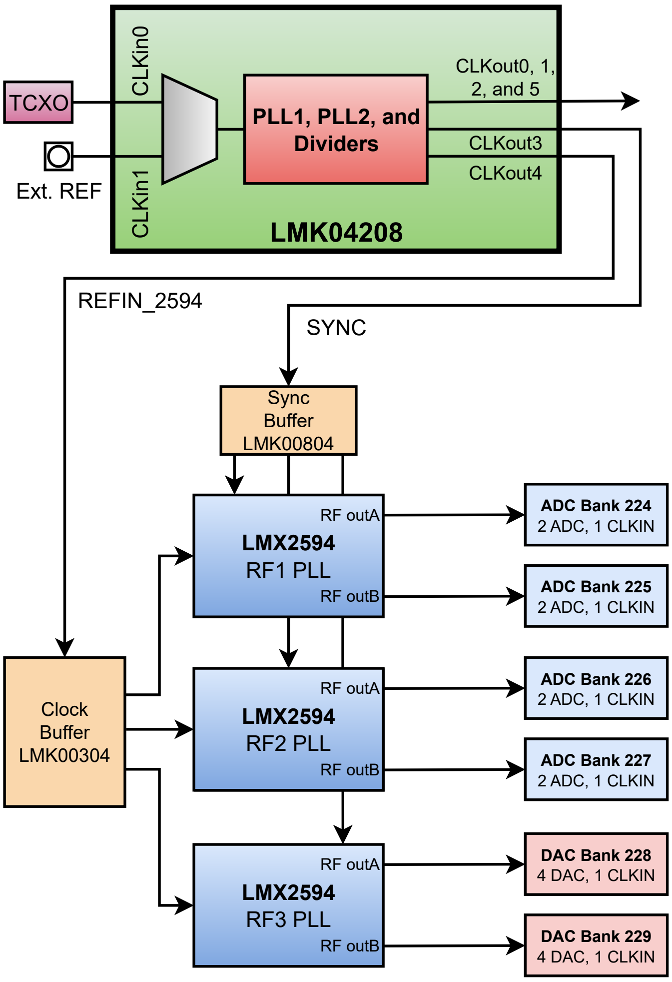

Figure 4.2: RF clocking structure of the ZCU111 for data converter banks. The reference PLL LMK04208 generates a reference input for all three LMX2594 RF PLLs. All LMX2594 units are configured to provide a reference to the PLLs contained within the data converter tiles. (Modified from [105].)

available I2C to SPI bridge. Fig. 4.2 shows the simplified clocking structure of the ZCU111 for the RF ADC and DAC banks.

In our implementation, the reference PLL LMK04208 is configured to receive a 10 MHz external reference. This PLL then produces a 245.76 MHz reference to all three LMX2594 RF-PLLs, which are configured to generate two RF outputs at 512 MHz. A detailed explanation of configuring both LMK04208 and LMX2594 is provided in Appendix D.

4.3.3 Programming the ZCU111

In the context of FPGAs, the binary representation of the configuration data used to program the FPGA hardware is called a bitstream. It contains information about the specific design, logic functions, and interconnections of the FPGA's resources. A detailed explanation of how to design the target module on an FPGA can be found in [106].

Our general programming structure can be summarized as follows: we program the PL part, the so-called firmware, with the specific DSP design and the required interfaces, while the PS is programmed with high-level software and an operating system from which the firmware is controlled. Data from the PL is also read out via the PS.

After preparing the required clock configuration files, compiling the bitstream corresponding to the project, and exporting the hardware platform using the Vivado Design Suite, we have explored two methods for programming the ZCU111. One method involves Vitis™ Software Platform, where a platform project is created and built first, followed by the creation and building of a corresponding application project. Before running the design, it is important to ensure that the LMK04208 and LMX2594 PLLs are properly configured. For this purpose, one can utilize the 'ZCU111 - Board User Interface' (SCUI) software, where hexadecimal register values corresponding to the two clock configurations can be set. After this, running the generated design on Vitis IDE should successfully program the ZCU111. Such a workflow is followed in the starter design[5] provided by the manufacturer.

The other, preferred method of programming the ZCU111 involves directly using PetaLinux[6] running on the PS part of the RFSoC. PetaLinux provides a base for running custom applications and loading custom bitstreams to the PL. All PetaLinux board support packages (BSPs) include pre-configured boot loaders, system images, and bitstreams. It also allows for custom device drivers, applications, libraries, and BSP configurations.

The bitstream of the project, along with the device-tree, board information, etc., are exported as an .xsa file, which is used by PetaLinux to generate installable packages with the bitstream and to compile the corresponding device-tree overlay. Other custom applications, including clock configuration, can be added by creating recipes that are used by BitBake to create Linux-specific RPM packages. Updating the PetaLinux running on the PS of the RFSoC using a RPM package manager will then allow us to program the ZCU111 as intended.

4.4 Initial validations

Even after gaining some basic understanding of the design flow and the specific details of the current evaluation board, it was not trivial to get a project running on the RFSoC. The complexity involved in configuring the RF data converter clocks and the PL operating with different clock domains resulted in unexpected errors at first. Consequently, some initial validation tests were conducted to verify the functionality of the RFSoC.

The first successful execution of a project on the ZCU111 was based on the aforementioned starter design from Xilinx. In this example design, a 10 MHz sine wave generated by the Direct Digital Synthesizer (DDS) Compiler IP core was provided as input to an RF-DAC, which samples

[5] `https://www.amd.com/en/products/adaptive-socs-and-fpgas/soc/zynq-ultrascale-plus-rfsoc.html#tabs-9c5228a5dd-item-2966fdd1d1-tab`; accessed on January 29, 2025.

[6] An AMD-developed customized Linux distribution created using the Yocto Project.

at 1.47456 GHz. The output of the DAC is looped back to an ADC using an external SMA cable. The ADC output is sent to an Integrated Logic Analyzer (ILA) IP core for display in the Hardware Manager of the Vivado Design Suite. In this example project, the LMX2594 was configured to directly provide the sampling clock for both data converters.

As the next step, the starter design was modified to receive an external signal from a signal generator while looping the output of the ADC internally to a DAC. The DAC output was then visualized on an oscilloscope. Successful attempts were made to include a general-purpose I/O LED that blinks based on a counter value (as per the custom script written in VHDL) or to utilize the onboard dual in-line package switch position.

Multiple data converter channels were then configured to sample at different rates, and signals between them were looped using external SMA cables and internal digital nets. All of these tests resulted in the expected behavior. These initial testing efforts helped us understand the overall clock configuration scheme and how to handle RF data converter signals successfully. After these validations, we began work on a high-speed frequency- (or phase-) tracking instrument based on the ZCU111.

CHAPTER 5

GHz Phasemeter

The interferometer scheme studied in this thesis required the implementation of two algorithms in the RFSoC: a frequency- (or phase-) tracking instrument with high bandwidth and a digital laser-locking system. The implementation of the former is described in this chapter.

In our displacement sensing application, information regarding the motion of the cavity mirror is encoded in the beat note between the two lasers. This beat note can range from several megahertz to several gigahertz in a typical experimental setup. To precisely measure and efficiently track changes in the frequency and phase of laser frequencies, an RFSoC-based ultra-fast phasemeter with gigahertz bandwidth was developed. Key results from this chapter have been published in IEEE Transactions on Instrumentation and Measurement [66].

5.1 Introduction to phase measurement

We classify instruments that measure and track the phase and frequency of an incoming electrical signal as phasemeters. There are several methods to perform such measurements. The most common approaches include zero-crossing detection [107], frequency counting [108], in-phase (I) and quadrature (Q) demodulation, and phase-locked loops (PLLs) [109, 110]. The first two approaches are limited to input signals with a high signal-to-noise ratio (SNR) and without the presence of additional tones. Here, we focus on phase-tracking using a digital local oscillator, which provides the highest possible SNR.

5.1.1 In-phase and quadrature demodulation

This technique measures the phase of the input signal by calculating the inverse tangent (or arctangent) of its orthogonal in-phase and quadrature components. Consider the input signal

$$s_{in}(t) = A\cos(\omega_{in}t + \phi_{in}), \tag{5.1}$$

where A is the amplitude, ω_{in} is the angular frequency, and ϕ_{in} is the phase. To obtain the orthogonal components of $s_{in}(t)$, we demodulate it using the orthogonal components of a local oscillator (LO).

Let these components be

$$s_{LO,1}(t) = B\sin(\omega_{in}t + \phi_{LO}) \quad \text{and} \tag{5.2}$$
$$s_{LO,2}(t) = B\cos(\omega_{in}t + \phi_{LO}), \tag{5.3}$$

with the amplitude B, the common LO phase ϕ_{LO}, but the same angular frequency ω_{in} as the input signal.

The demodulation[1] of the input signal with the orthogonal LO components results in

$$s_{in}(t) \times s_{LO,1}(t) \overset{\text{LPF}}{=} \frac{AB}{2}\sin(\Delta\phi) = s_Q(t) \quad \text{and} \tag{5.4}$$
$$s_{in}(t) \times s_{LO,2}(t) \overset{\text{LPF}}{=} \frac{AB}{2}\cos(\Delta\phi) = s_I(t), \tag{5.5}$$

where $\Delta\phi = \phi_{LO} - \phi_{in}$. The phase difference can then be recovered by calculating the arctangent of $s_Q(t)$ and $s_I(t)$.

It should be noted that the calculated phase $\Delta\phi$ has discontinuities at $\pm\pi$ because of the asymptotic limits of the inverse tangent function. Whenever $\Delta\phi$ exceeds these limits, it will wrap around by $\mp 2\pi$, jumping to the opposite boundary [73]. Therefore, phase unwrapping is required to reconstruct a continuous phase measurement of the input signal.

Although this IQ-demodulation technique is highly effective for phase measurement, it cannot be fully implemented in an FPGA. The division and phase unwrapping involved limit its performance in hard real-time applications, where fixed (or predictable) and short processing delays are required. Additionally, since the demodulation is carried out at a fixed LO frequency, this method is suitable only for input signals with a small dynamic range, as demonstrated in the case of LISA Pathfinder [111]. For signals with a large dynamic range, the demodulation and signal frequency will quickly deviate from each other, introducing either a phase error or a complete signal loss [87].

5.1.2 Phase-locked loop (PLL)

The limitations of IQ-demodulation indicate the need for low-latency feedback of phase information. This necessitates embedding the IQ-demodulation within a control loop, resulting in a complete PLL. Instead of having a fixed LO phase as in the former approach, the LO phase is continuously updated in the latter by using the phase error between the input signal and the LO. Fig. 5.1 sketches the overview of such a PLL.

For the PLL to operate, an initial frequency f_0 close to the input signal frequency is required. The current frequency of the LO is stored in the phase increment register (PIR), and by continuously accumulating the frequency values, we derive the phase value of the LO, which is stored in the phase accumulator (PA). This phase value is then used to generate two orthogonal components for demodulation purpose. The quadrature signal (Q) resulting from the demodulation is amplified in a servo controller and used to correct the LO phase.

In this manner, when the PLL is locked, the relative phase difference between the input and the LO

[1] Multiplication followed by low-pass filtering.

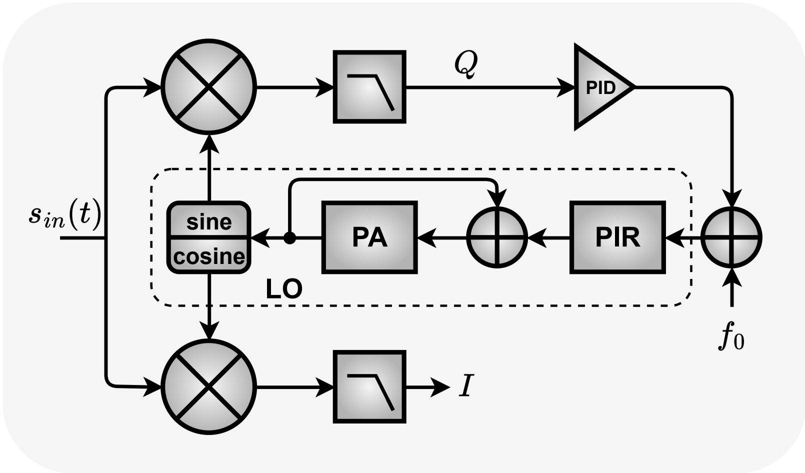

Figure 5.1: Block diagram of a PLL. The input signal is demodulated with an LO signal, and the quadrature (Q) signal is utilized as feedback in a closed-loop control system to correct the LO phase. For the PLL lock acquisition, an initial frequency estimation f_0 is required. (PIR: phase increment register, PA: phase accumulator, LO: local oscillator. Created based on [87, Fig. 2.3].)

is very small (i.e., $\Delta\phi \ll 1$). This assumption allows us to express Eq. 5.4 as[2]

$$Q \equiv s_Q(t) = \frac{AB}{2}\sin(\Delta\phi) \approx \frac{AB}{2}\Delta\phi. \tag{5.6}$$

Similarly,

$$I \equiv s_I(t) = \frac{AB}{2}\cos(\Delta\phi) \approx \frac{AB}{2}. \tag{5.7}$$

Thus, the phase information is entirely contained within Q, while the amplitude information is in I. Therefore, the advantage of a PLL over IQ-demodulation is that there are no division or arctangent calculations involved. This makes PLLs well-suited for FPGA implementations.

5.2 Overview of phasemeter developments

To measure phase and frequency with high precision, digitizing the signal and embedding PLLs in FPGAs is a well-studied approach [95, 112–115]. The core algorithm of such a digital implementation is an all-digital phase-locked loop (ADPLL). The LO of a PLL in a digital context is referred to as a numerically controlled oscillator (NCO). Digital phasemeters have been extensively studied and developed for laser interferometry, particularly in the context of the LISA mission [95, 113]. Phasemeters are also available as commercial instruments [97], which we used for the experiments discussed in Sec. 3.5.

These phasemeters are designed to track one or multiple tones in the input signal, providing information on how the phase and/or frequency of those tones change over time. Determining the absolute phase of the measured tones relative to a reference is not the primary motivation.

[2] The small-angle approximation for sine: $\sin(\theta) \approx \theta$ for $\theta \ll 1$.

General signal processing instruments like oscilloscopes or spectrum analyzers also provide frequency measurements of stable tones. However, they cannot measure rapidly changing frequencies and do not achieve precision levels in the milli- to microradian range.

5.2.1 Phasemeter features

The performance of the PLL can be characterized in different ways. One way is to examine its sensitivity, bandwidth, and throughput [73]. In this section, we will summarize some key features of the various phasemeters developed so far.

Signal bandwidth

One of the critical parameters of digital phasemeters is their signal (or detection) bandwidth, defined as the range of RF frequencies over which the PLL can track phase. Several previous studies have covered different ranges of signal bandwidth. In LISA, the RF signal is generated by beating two lasers with an effective carrier wavelength of about 1 µm, with one beam carrying the Doppler shift associated with the distance changes between two spacecraft. The beat note frequencies typically range from about 5 MHz to 25 MHz. Hence, the LISA phasemeter is designed to have a signal bandwidth above this range, with most implementations providing 40 MHz [95]. Similar signal bandwidth values are needed for other space-based missions such as Taiji [116]. For a high-speed multiplexed heterodyne interferometry, a phase measurement system with a signal bandwidth of 156.25 MHz has been demonstrated [117]. The commercial Moku:Pro phasemeter has a signal bandwidth of 300 MHz [97], making it the phasemeter with the highest reported signal bandwidth so far.

Tracking bandwidth

The tracking (or loop) bandwidth of digital phasemeters is one of their core design parameters. It is determined by the open-loop gain of the PLL, thus defining the frequency range up to which closed-loop suppression is achievable. The tracking bandwidth also determines whether the phase lock is stable and whether the loop operates linearly. When the PLL exhibits non-linear effects, it can lead to non-linear phase and amplitude measurements, cycle slips, or loss of lock [118].

The maximum achievable tracking bandwidth of previously developed phasemeters has not been discussed in detail. This is partly because a lower loop bandwidth was optimal for LISA [115], but it is also inherently limited by the clock speed of the PL.

Another related aspect that explicitly depends on the clock speed of the PL is the throughput of the phasemeter. This describes the frequency with which the LO phase of the PLL is updated. Many previous developments are limited to a throughput of less than 100 MHz. The LISA phasemeters mainly operate at 80 MHz throughput, while some other phasemeter implementations achieve only a couple of megahertz [73].

5.2.2 Motivation for further development

In our heterodyne cavity-tracking application, the signal bandwidth of the readout instrument dictates the displacement sensing operating range (see Eq. 3.37). For the considered experimental setup[3], to

[3] A laser wavelength of 1550 nm and a dynamic cavity with a physical length of 5 cm.

achieve a fringe-scale operating range, we need a phasemeter with a signal bandwidth in the gigahertz regime.

High tracking bandwidth is essential for accurately capturing highly dynamic signals with minimal additive noise. Since laser interferometry uses relatively short wavelengths, there are instances where frequency changes occur rapidly, emphasizing the need to study phasemeters with high tracking bandwidth [119]. In cases where attempts to measure the beat note between two CTLs using a low tracking bandwidth Moku:Lab phasemeter were made, it was not possible to track the tone. Therefore, with higher tracking bandwidth, it becomes feasible to use lasers with higher frequency noise and line-width, enabling the use of less stable and more cost-effective laser types for interferometry.

Digital phasemeters can measure frequencies across many orders of magnitude of dynamic range. For example, the LISA phasemeter has demonstrated a dynamic range of more than ten orders of magnitude [120]. This capability enables ultra-linear measurements across a variety of frequency ranges. By definition, the lower limit of this dynamic frequency measurement range is the phase readout noise floor, while the upper limit is the signal bandwidth of the phasemeter. The signal bandwidth of the phasemeter limits the maximum relative speeds that can be tracked in applications such as Doppler-ranging. With implementations featuring higher detection bandwidth, more applications can leverage the exceptional linearity of digital phasemeters.

Given these motivations, we have developed and tested a digital phasemeter, the GHz Phasemeter, based on RFSoC with benchmark levels of signal and tracking bandwidth. To achieve this, the phasemeter is realized by adapting and optimizing the ADPLL architecture in the PL at the hardware limits. This allowed us to develop a high signal bandwidth (> 2 GHz) and high tracking bandwidth (2 MHz) multi-channel phasemeter. We used the phasemeter technologies developed in the context of space-based laser interferometric experiments as a starting point. However, it is crucial to note that having hardware with ADCs capable of high sampling rates is not sufficient for our implementation. The benchmark parameters of the phasemeter are achieved by operating the PL at its timing limits and introducing a novel multi-demodulation and phase accumulation scheme. A detailed discussion of our implementation is covered in Sec. 5.4. Before delving into that, we will briefly review some relevant fundamentals of DSP.

5.3 Some fundamentals of DSP

5.3.1 Discrete-time transfer function model

Digital systems heavily depend on their sampling time t_s, which is generally clock-dependent. In the case of RFSoC, the ADC sampling rate is 4.096 GSPS, but the AXI4-Stream data output is clocked at 512 MHz. Thus, the DSP computations discussed in this thesis are primarily performed at 512 MHz.

The PLL can be analyzed using a combination of transfer function and time-domain models. Each component or block of the PLL is represented as either a gain or a discrete-time transfer function in the z-domain, which is the discrete time analog of the Laplace domain. The frequency-dependent behavior of the discrete-time system is described relative to the sampling frequency using the parameter

$$z = e^{i2\pi \frac{f}{f_s}}, \tag{5.8}$$

where f represents the continuous frequency variable from 0 to f_s. The inverse of the z-parameter defines a single delay at the sampling frequency. Hence, z^{-1} is also known as a delay operator.

5.3.2 Quantization

The DSP operates with discrete numbers. Each signal is represented by a finite number of distinct values with a constant, non-zero separation between them. This results in quantization noise, which is introduced into the system when an analog signal is digitized or when a digital signal is truncated [121]. The latter case is sometimes also referred to as truncation noise.

Quantization noise is related to the signal's least significant bit (LSB). Let x be a random variable representing the LSB of a signal. The standard deviation of x is given by

$$\sigma = \sqrt{\int_{-\infty}^{\infty} \rho(x)(x-\bar{x})^2 \, dx}\,, \tag{5.9}$$

where $\bar{x}$ is the statistical mean of x and $\rho(x)$ is its probability density function. Assuming that the quantization noise is uniformly distributed across $-\mathrm{LSB}/2$ and $\mathrm{LSB}/2$, and is zero elsewhere, the standard deviation becomes [73]

$$\sigma = \sqrt{\int_{-\mathrm{LSB}/2}^{\mathrm{LSB}/2} \frac{1}{\mathrm{LSB}} x^2 \, dx} \; = \; \frac{\mathrm{LSB}}{\sqrt{12}}. \tag{5.10}$$

The variance (σ^2) of random noise digitized at a sampling frequency f_s is related to the single-sided PSD by

$$\sigma^2 = \int_0^{f_s/2} \tilde{S}_{xx}(f) \, df \; = \; \frac{\tilde{S}_{xx}(f) f_s}{2}, \tag{5.11}$$

when the PSD is constant, i.e., behaving as white noise across the spectrum. Therefore, in terms of ASD, the quantization or truncation noise can be expressed as

$$\tilde{S}_{x,\mathrm{trunc}}(f) = \frac{\mathrm{LSB}}{\sqrt{6 \cdot f_s}} \qquad [1/\sqrt{\mathrm{Hz}}]. \tag{5.12}$$

However, the assumption that quantization noise is the same for all frequencies is an idealization. It is not always necessarily white; for example, in the case of a digitized sine wave. This assumption holds for signals that move through a significant range of digital values without any coherent relationship to the sampling frequency [87].

5.3.3 Dither and rounding

The process of forcing the quantization noise to appear as white noise by introducing intentional pseudo-random noise is known as dithering. In our FPGA implementation, this intentional noise floor is added to digital signals before truncation, with triangular dither being the preferred method [122]. When a signal is truncated by N bits, the dither signal is kept at a length of $N + 1$ bits.

Our implementation utilizes a triangular dither generator (from [87]) that operates by subtracting the outputs of two independent linear feedback shift registers, which generate pseudo-random noise. However, introducing dither to achieve a white quantization or truncation noise comes at the cost of a slightly higher noise floor. The additional dither, followed by truncation, results in a noise contribution of

$$\tilde{S}_{x,\text{dith+trunc}}(f) = \sqrt{3} \cdot \tilde{S}_{x,\text{trunc}}(f) \qquad [1/\sqrt{\text{Hz}}]. \tag{5.13}$$

The $\sqrt{3}$ increase in noise is tolerable, as any unwanted behavior resulting from truncation is now suppressed.

Truncation can also introduce small signal offsets due to rounding errors. Therefore, in our implementations, each truncation is followed by an offset-free rounding algorithm based on integer arithmetic. This specific rounding block was also adapted from [87] for use in our implementations.

5.4 Implementation of the GHz Phasemeter

The signal processing methods investigated and applied for the implementation of the GHz Phasemeter in FPGAs utilize integer arithmetic, which is crucial for the low-latency readout and control of laser frequencies. Our implementation uses the Zynq UltraScale+ RFSoC ZCU111 evaluation board [104] as the foundation. We explored two different approaches to implement the GHz Phasemeter. The first approach utilized the internal NCO of the RF data converter block, while the final implementation featured a custom algorithm for the entire PLL.

5.4.1 Initial implementation

The initial implementation plan was to make use of the internal NCO of the RF data converter. Following the basic PLL block diagram in Fig. 5.1, it is then necessary to update the NCO frequency of the digital complex mixers with an externally calculated PIR value in the PL. Such an implementation was successfully realized and is briefly described in Appendix E.

However, using the existing hardware infrastructure limits us in terms of internal delays, preventing us from achieving the maximum possible PLL tracking bandwidth. To overcome this internal hardware limitation and fully utilize the capabilities of the RFSoC, one must manage the NCO using a custom algorithm within the PL instead of relying on digital down converters (DDCs). Such an approach, as realized below, will allow us to maintain the overall throughput at maximum speed and utilize all eight samples from the RF-ADC without internal decimation, thus retaining the effective sample rate at maximum levels.

5.4.2 Final implementation

The RF-ADCs sampling the input signals at 4.096 GHz are configured to provide eight data samples in parallel at 512 MHz. The required AXI4-Stream clock for this purpose is generated by the Clocking Wizard IP core, which receives the 256 MHz input clock from the RF data converter IP core. The AXI4-Lite clock, obtained from the PS, is set to 100 MHz. Most of the DSP running on the PL is clocked at 512 MHz, which presents challenges and requires fine-tuning of registers within the PLL to meet the timing requirements, which are already strained at this speed due to simple routing delays within the fabric. However, maintaining this processing speed is crucial for achieving high throughput and high tracking bandwidth.

In the basic ADPLL topology, as sketched in Fig. 5.1, only one sample of the input signal is available to be demodulated against the NCO signal at any given instant. However, the input signal sampled at 4.096 GHz is now output as eight samples in parallel, which can be combined to reconstruct the actual

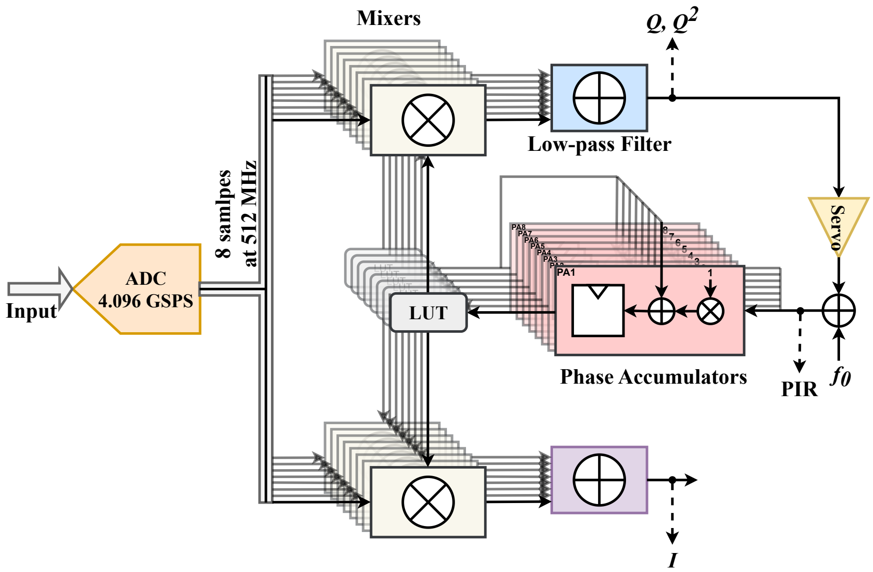

Figure 5.2: The GHz Phasemeter topology with multi-demodulation and multi-phase accumulation stages. The incoming electrical signal is digitized at 4.096 GHz, resulting in eight samples output at 512 MHz. Markers for signal readout points are shown. (PIR: phase increment register, PA: phase accumulator, LUT: look-up table.)

input tone. Therefore, handling these eight parallel samples without decimation is critical to retaining the effective sampling rate of 4.096 GHz.

Consequently, the traditional demodulation and phase accumulation stages of the ADPLL are extended into a novel multi-demodulation scheme, followed by a multi-phase accumulation stage. Without the proposed modifications to the signal processing scheme, achieving reliable operation with gigahertz regime signal bandwidth using RFSoC would not be feasible. Fig. 5.2 shows an overview of the GHz PLL, including the aforementioned processing modifications.

As each ADC channel provides eight data samples for each clock cycle, we need eight consecutive phase values in the NCO for multi-demodulation. Thus, the initial frequency estimate f_0 stored in the PIR is split and then successively incremented. Mathematically, this means multiplying each copy of the PIR by the integers 1 to 8, respectively. These values are then accumulated individually, after adding the largest value from previous products to all the resulting products to ensure effective continuity. This accumulation yields the required phase values for the eight look-up tables (LUTs). The LUTs[4] generate sine and cosine signals of the specified phase, and together with the previous multi-phase accumulators, they form the NCO of the phasemeter.

The digitized input signal samples are each IQ-demodulated at the NCO frequencies by multiplying them with the outputs of the eight LUTs. When the logic operates at 512 MHz, performing multiplication of two signals on the PL is not simple and straightforward. Allowing software to implement multiplication at this speed resulted in unreliable outputs. Therefore, each multiplication is handled manually using the DSP48 Macro IP core, which allows for direct configuration of the DSP Slices[5] within the FPGA fabric.

[4] The LUTs are realized using the DDS Compiler IP cores.

[5] These are specific multiply-accumulate hardware accelerators that speed up the execution of signal processing functions.

Higher harmonics resulting from the mixing process are filtered out using a rolling average of 16 samples. Adding all 16 samples at once created timing constraints, so this was implemented through consecutive additions. The so-generated Q value acts as the error signal for the closed-loop control and is fed into the proportional-integral (PI) controller. The output, the actuation signal, then corrects the initial starting frequency to match the input signal frequency. Consequently, when the PLL is locked, the PIR contains a copy of the input electrical signal frequency. The value of Q measures the accuracy of the tracking, while I corresponds to the amplitude of the input signal.

The available eight ADCs, with a maximum sample rate of 4.096 GSPS, allows us to implement eight readout channels, each with 2.048 GHz of signal bandwidth. However, because of the used XM500 add-on card, the first two phasemeter channels (which utilize ADC Tile 224) are limited to single-ended signals of frequencies up to 1 GHz. Similarly, the next two channels (which utilize ADC Tile 225) are limited to frequencies above 1 GHz. The remaining four channels are available for the entire signal bandwidth, with differential inputs.

Our implementation of the phasemeter provides two ADPLLs per ADC channel: one to track the main tone and the other to track a pilot tone. This configuration allows us to correct for ADC timing jitter in post-processing [95].[6] Thus, at its maximum capacity, the ZCU111 evaluation board operates as an eight-channel phasemeter with a total of 16 ADPLLs. Each PLL can track a distinct tone and is controlled individually by the PS by setting registers in the on-chip memory using a Python interface. The PL is programmed to retrieve corresponding values, such as the starting frequency and controller gains, from the assigned memory addresses.

5.5 Linear model of the ADPLL

When the PLL is locked to the input signal, we expect it to operate in a linear regime that is achieved when Q is small. With this assumption, it is possible to model the PLL reliably by understanding the transfer function of each block in the loop. A block diagram of the linearized PLL model is sketched in Fig. 5.3. The specific linear model discussed below is an adaptation of the ADPLL models presented in [87] and [73].

5.5.1 Digital multiplier as phase detector

Assuming an input signal with a peak amplitude V_{in}, digitized using an ADC with a maximum peak-to-peak voltage range of V_{pp}, the digitized signal can be expressed as

$$s_{in}[n] = \frac{V_{in}}{V_{pp}} \cos(\omega_0 n + \varepsilon_{in}[n]), \tag{5.14}$$

with $\varepsilon_{in}[n]$ as the input phase, the quantity that is sensed and actuated. It should be noted that the input angular frequency $\omega_0 = 2\pi f_0$ is only required for the initial lock acquisition, and not for the linear model.

Similarly, the signal from the NCO can be written as

$$s_{LO}[n] = \frac{1}{2} \sin(\omega_0 n + \varepsilon_{LO}[n]), \tag{5.15}$$

[6] The use of pilot tone correction (PTC) scheme is demonstrated in Sec. 5.8.2.

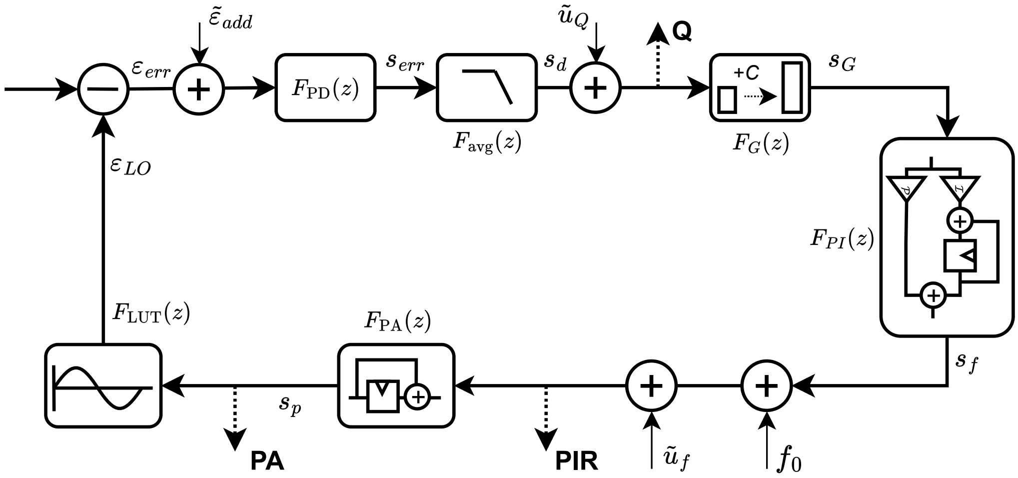

Figure 5.3: Block diagram of the linearized PLL model, considering phase (ε) as the sensed and actuated quantity. Markers for signal readout and possible noise addition points are indicated. ($\tilde{\varepsilon}_{add}$: phase noise due to additive noise, $\tilde{u}$: truncation noise-induced phase noise.)

with $\varepsilon_{LO}[n]$ as the NCO output phase. As the difference between the input and NCO phase is not directly accessible, a digital multiplier (also known as a mixer) is used to estimate the phase error between the input and the NCO signals.

The output of the mixer, which serves as the error signal of the loop, is then given by

$$\begin{aligned} s_{err}[n] &= s_{in}[n] \times s_{LO}[n] \\ &\approx \frac{V_{in}}{4V_{pp}}(\varepsilon_{LO}[n] - \varepsilon_{in}[n]) = \frac{A}{4}\varepsilon_{err}[n], \end{aligned} \tag{5.16}$$

where two linearizations are introduced to arrive at the linear model. One assumption is that the phase error is very small (i.e., $\varepsilon_{err}[n] \ll 1$), and the other is the proper suppression of the second harmonic term resulting from the mixing process.

Thus, the linear transfer function of the phase detector can be described as

$$F_{\mathrm{PD}}(z) = \frac{s_{err}(z)}{\varepsilon_{err}(z)} = \frac{A}{4}, \tag{5.17}$$

where $0 \leq A = V_{in}/V_{pp} < 0.5$ and z is the discrete-time operator defined in Eq. 5.8. From the above equation, it is clear that the overall gain of the PLL is dependent on the input signal amplitude. This could potentially lead to unstable behavior if the controller of the PLL is not configured properly. It is possible to overcome this dependency by implementing an automatic gain algorithm to scale the controller's gain proportional to the input signal amplitude. Such an algorithm was not used in our implementation.

In the case of GHz PLL, even though phase detection is performed using a multi-demodulation scheme, its linear model can be represented with a single digital multiplier operating at $f_s = 4.096$ GHz.

The V_{pp} value for the RFSoC ADCs is 1.8 V. Since the 12-bit resolution ADC is internally most significant bit (MSB)-aligned to create a 16-bit data path, we truncated 4 bits from the LSB side in our custom algorithm without dithering. Each 12-bit sample is then multiplied with the corresponding 16-bit long signal from the NCO, generating 28-bit long mixer outputs.

5.5.2 Low-pass filter

As stated in the previous section, a rate-reducing low-pass filter is used after the phase detector to suppress the second harmonic generated by the mixer. In our implementation, the multi-demodulation stage produces 8 signals at 512 MHz. One straightforward method of low-pass filtering is to simply add these 8 signals together. This is where we decimate from 4.096 GHz to 512 MHz. Adding 8 samples is equivalent to an finite impulse response (FIR) filter with 7 taps, in which all coefficients are equal to unity.

In general, the transfer function of an N-tap FIR filter can be expressed in the z-domain as

$$F_{\mathrm{FIR}}(z) = \sum_{i=0}^{N} a_i z^{-i}. \tag{5.18}$$

When all a_i are 1, we are essentially averaging the data. The filter coefficients can be chosen to follow a triangular window, which can be implemented in FPGA via successive bit-shifting. Fig. 5.4 shows the response of an FIR filter for different number of taps.

As a compromise between the required suppression below 256 MHz and the feasibility of implementation on the fabric while satisfying timing constraints, we chose an FIR filter with 15 taps for low-pass filtering the mixer output using a rolling average. Such a linear-phase (or symmetric) FIR filter introduces a delay given by

$$D = \frac{\text{order}}{2} = \frac{N-1}{2}, \tag{5.19}$$

which is an integer number of samples in our case.

For the PLL linear model, the gain contribution from the low-pass filtering stage can therefore be expressed as

$$F_{\mathrm{avg}}(z) = \frac{s_d(z)}{s_{err}(z)} = \frac{1}{N}\sum_{i=0}^{N} a_i z^{-i}, \tag{5.20}$$

with the sampling frequency associated with the z-parameter being 4.096 GHz.

In the actual implementation, to properly accommodate the addition of samples, the 28-bit long mixer outputs are resized to 31-bit before addition. The filtered signal is then dithered and truncated to 20 bits before being passed to the servo stage.

5.5.3 Controller

The in-loop controller determines the open-loop gain of the PLL. To achieve a stable loop, an overall gain reduction may be necessary, depending on the implementation. In our case, a constant gain reduction was performed before the servo to prevent any overflow in the digital accumulators. A simple method for implementing such a gain reduction is by bit-shifting or resizing. If C bits are

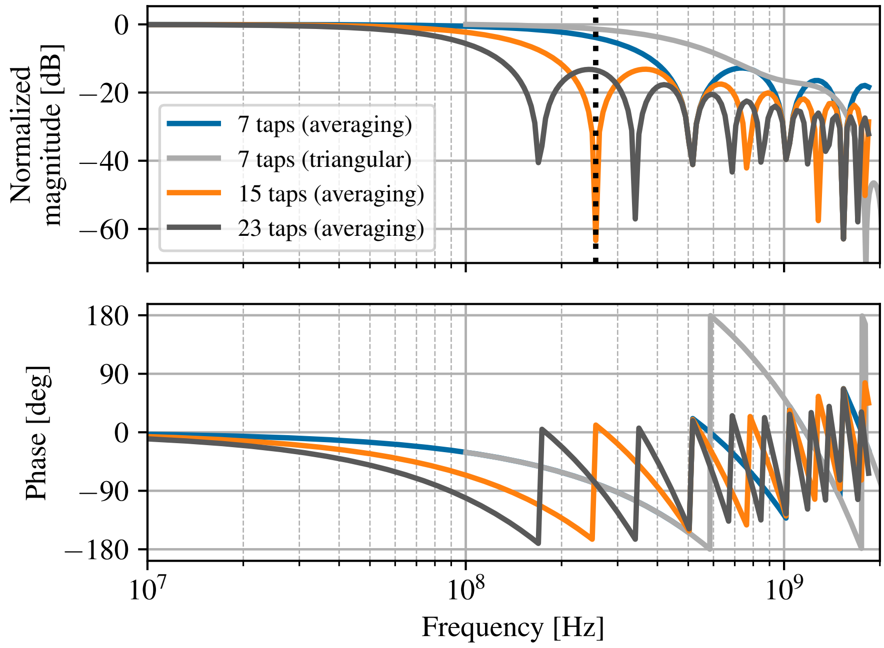

Figure 5.4: The simulated response of an FIR filter for different number of taps. The sampling frequency is set to 4.096 GHz. In the case of filter coefficients resembling a triangular window, higher suppression is achieved only at the high-frequency end. The dotted vertical line at 256 MHz represents the Nyquist frequency of the loop operating at 512 MHz. The simulation was carried out using the Python package ‘python-control’ [123].

added from the MSB side to the signal, it results in an overall gain due to in-loop bit growth given by

$$F_G(z) = \frac{s_G(z)}{s_d(z)} = 2^{-C}. \tag{5.21}$$

The actual servo of the PLL includes proportional and integral controllers. The proportional controller responds immediately and proportionally to the magnitude of the phase disturbances detected. The integral controller provides long-term feedback and eliminates steady-state error [73]. It is possible for the frequency of the input signal to change over time, and the initial frequency estimate may also be incorrect. In such cases, a constant frequency error must be compensated by the integral controller. When steady-state errors persist over a longer time, it gets time integrated in the accumulator, and the integral controller applies feedback that is proportional to this accumulated error. Thus, the longer the error persists, the higher the feedback magnitude.

The transfer function of the PI controller can therefore be written as

$$F_{PI}(z) = \frac{s_f(z)}{s_G(z)} = \mathcal{P} + \mathcal{I}\,\frac{z^{-1}}{1 - z^{-1}}, \tag{5.22}$$

with the sampling frequency associated with the z-parameter being 512 MHz. By tuning the controller’s gains $\mathcal{P}$ and $\mathcal{I}$, the desired gain, bandwidth, and loop response of the PLL can be achieved. In the GHz PLL implementation, the in-loop bit growth C was chosen to be 25, resulting in an output of the controller that is 45 bits long.

The output of the servo, which scales with the instantaneous frequency error, is then added to the initial frequency estimate f_0 and stored in the PIR.

5.5.4 Numerically controlled oscillator (NCO)

The NCO of the PLL model consists of an in-loop integrator, called the PA, and a LUT. To generate a phase value that drives the NCO, the in-loop integrator accumulates the PIR at the full sampling rate f_s over a finite range. This accumulation causes the signal to 'wrap' periodically, generating a signal that is proportional to the instantaneous phase of the NCO modulo 2π [73].

The transfer function of the PA is determined by the integrator, given by

$$F_{\mathrm{PA}}(z) = \frac{s_p(z)}{s_f(z)} = \frac{z^{-1}}{1 - z^{-1}}, \tag{5.23}$$

again with the sampling frequency associated with the z-parameter being 512 MHz.

Having the PIR as a 45-bit long number is unnecessary and complicates the accumulation processing at 512 MHz. Therefore, the PIR is dithered and truncated to a 16-bit long number before being passed to the multi-phase accumulation stage of the GHz PLL.

The generated phase value is then used as input to a sine and cosine LUT. Each address in the LUT corresponds to a particular phase of the sine or cosine wave between 0 and 2π. The transfer function of the LUT can thus be expressed as

$$F_{\mathrm{LUT}}(z) = \frac{\varepsilon_{LO}(z)}{s_p(z)} = 2\pi. \tag{5.24}$$

5.5.5 Open-loop transfer function

With the transfer function of each block known, the open-loop gain of the modeled PLL is simply the multiplication of all the individual stages. In the discussions above, we have not included the delays associated with signal processing. These delays can be directly computed from the number of registers used in the PL. For a total delay of D clock cycles, the contribution from the delay is represented as z^{-D}. When parts of the loop operate at different frequencies, the delays should be counted accordingly.

The open-loop transfer function $G(z)$ is given by

$$G(z) = \frac{\varepsilon_{LO}(z)}{\varepsilon_{err}(z)} = F_{\mathrm{LUT}}(z) \cdot F_{\mathrm{PA}}(z) \cdot F_{PI}(z) \cdot F_G(z) \cdot F_{\mathrm{avg}}(z) \cdot F_{\mathrm{PD}}(z) \cdot z^{-D}. \tag{5.25}$$

The open-loop gain allows us to determine parameters such as loop stability and noise suppression.

The highest stable loop bandwidth is limited only by the delays within the ADPLL that are caused by the registers. Nevertheless, these delays are necessary to realize various processing steps. The utilization and settings of dedicated DSP Slices in the PL allowed us to minimize the number of registers in the loop. Our DSP operating at 512 MHz requires more registers in the loop to meet timing requirements, but still uses less than twice the amount of registers needed to operate at 256 MHz. Hence, the current implementation achieves the highest possible loop gain with the chosen device.

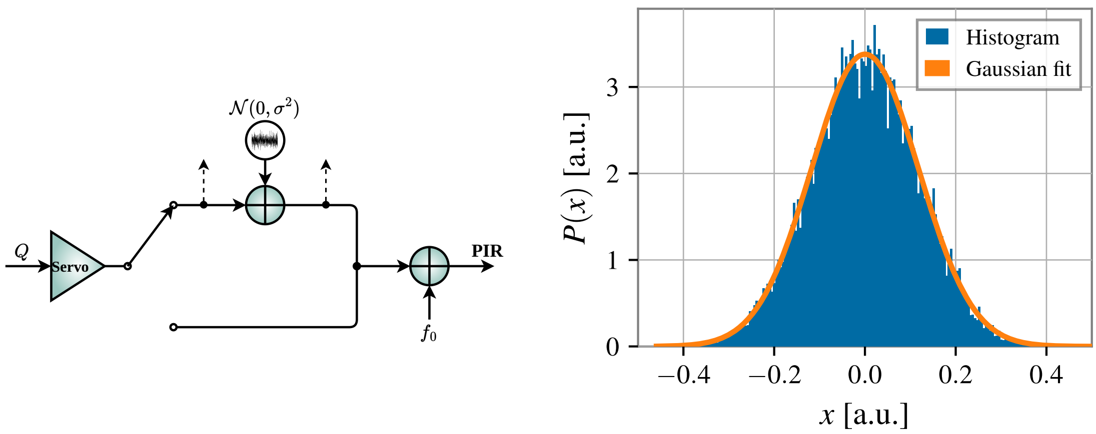

(a) (b)

Figure 5.5: (a) Implementation of a Gaussian noise generator to enable direct measurement of the PLL open-loop transfer function. Depending on the switch position, Gaussian noise is injected into the output of the servo. Measurement probes are provided to monitor the two signals before and after the addition of Gaussian noise. (b) Histogram of the generated noise along with a Gaussian fit to the recorded data. The overall probability distribution is normalized to 1.

5.5.6 Direct measurement of the loop transfer function

We have implemented and integrated a direct measurement feature for the PLL transfer function within our phasemeter. This functionality aids in measuring the achieved open-loop gain and allows for tuning the tracking bandwidth as required. Fig. 5.5(a) illustrates the implementation of this feature.

When the associated switch position is set to enable noise injection, Gaussian noise is added to the output of the servo. The switch position is controllable from the PS side. Using an ILA, it is possible to directly monitor the two signals before and after the addition of Gaussian noise at full processing speed. With these two signals, it is then straightforward to compute the transfer function of the PLL.

Gaussian noise generation is carried out in the PL using 8 triangular dither generators with different random starting values. Care is taken to ensure that the superposition of these 8 triangular shapes results in a Gaussian profile. Fig. 5.5(b) shows the histogram of the generated noise at an instant captured using an ILA IP core.

The generated noise is 11 bits long but is resized to 45 bits before being added to the servo output. The strength of the injected noise plays a role in efficiently measuring the loop gain. Depending on the loop's stability, it may be necessary to vary the injected noise amplitude. Our implementation allows for noise generation with externally adjustable amplitude. The gain of the noise can be set in the software and is correspondingly implemented in the PL by bit-shifting. A noise gain of 25 was used in our measurements.

Comparing the direct measurement with the linear model

Using the transfer functions of the individual stages of the PLL, the overall open-loop gain (as given by Eq. 5.25) is simulated using the Python package 'python-control' [123]. The number of delays within the loop at a processing speed of 512 MHz was counted, with a maximum assumed to be 27.

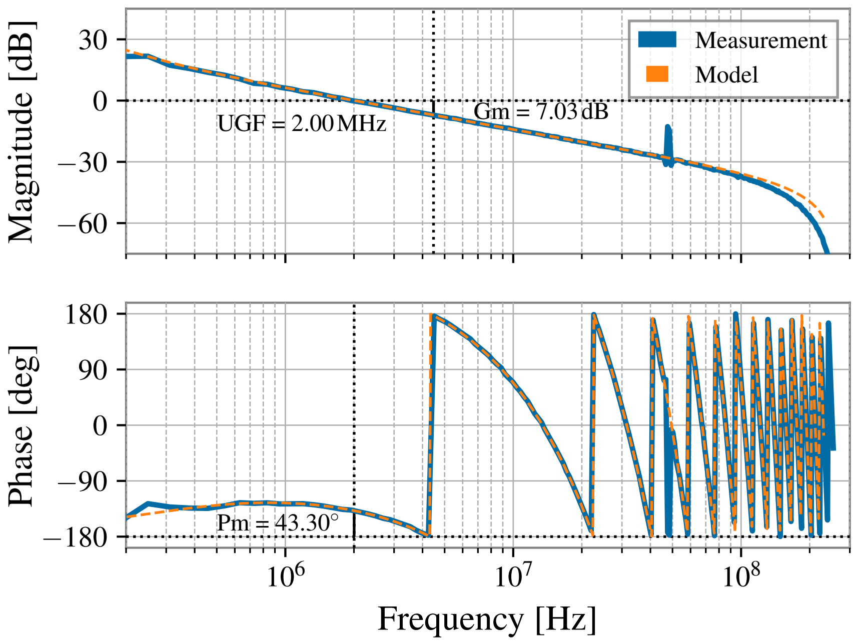

Figure 5.6: Modeled (dashed trace) and measured (solid trace) open-loop gain of the GHz Phasemeter. Both traces are in very good agreement with each other. The peak appearing at 48 MHz is the aliased second harmonic of the input signal. The deviation at the high-frequency end is not understood but is considered irrelevant for the analysis. (UGF: unity gain frequency, Gm: gain margin, Pm: phase margin.)

By tuning the servo gains, a maximum open-loop gain of 2 MHz could be achieved, leaving sufficient gain and phase margins for stable loop operation.

The servo gains obtained from the loop modeling were then applied in a real measurement, where the PLL is locked to a tone generated using a signal generator. Gaussian noise is injected into the loop, and in the Vivado Design Suite, the ILA is triggered to capture data in the block RAM. The open-loop transfer function is then calculated using the Python package 'Spicypy' [124]. Fig. 5.6 shows the measured open-loop gain compared to the modeled one.

A strong agreement between the measurement and the model confirms that the GHz Phasemeter is operating as expected, and its various components and their functions are well understood. Below a corner frequency of 300 kHz, the integral controller of the servo increases the loop gain even more rapidly, achieving a loop suppression of about 12 orders of magnitude at 1 Hz. It is possible to include an additional integrator corner, leading to a PII^2 controller, to further enhance the loop gain at low frequencies.

With the current implementation, the highest stable tracking bandwidth of the GHz Phasemeter is around 2 MHz. This loop bandwidth can easily be tuned to lower values by adapting the servo gains accordingly. The new feature of direct measurement of the loop transfer function aids in tuning the tracking bandwidth to realize optimal tracking conditions based on the input signal dynamics. The loop model can be used to adjust the PLL servo gains to achieve the desired tracking bandwidth. These controller gain values are then directly used to set the corresponding register values from the PS. This allows the tracking bandwidth of the phasemeter to be tunable on the go.

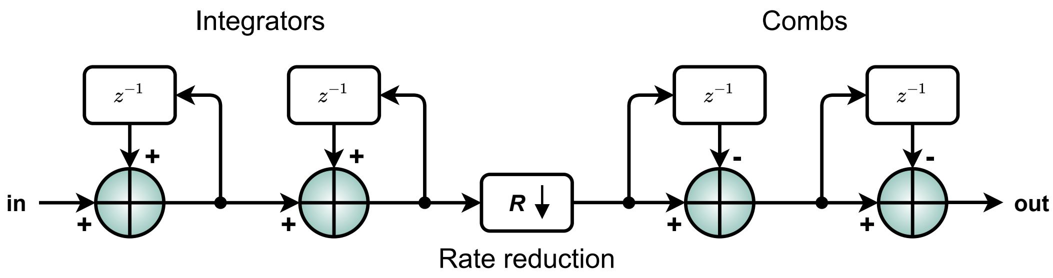

Figure 5.7: Block diagram of a second-order CIC filter. The input is integrated through two series stages using discrete-time accumulators. The output of the integrator stages is then re-sampled at f_s/R before being passed to the two consecutive differentiators, also known as combs. z^{-1} is the delay operator introduced in Sec. 5.3.1

5.6 Data acquisition

Various signals of the ADPLL are processed at 512 MHz, and each of these signals carries relevant information. Therefore, in an ideal scenario, we aim to read out and store most of these signals at the processing speed itself. However, the collection, storage, and processing of such large amounts of data over measurement times that can vary from minutes to several hours or days, depending on the requirement, is unrealistic. Consequently, only selected signals are read out from the GHz Phasemeter after being decimated to lower sampling rates, typically on the order of tens or hundreds of hertz. The signals we read out from the GHz Phasemeter are marked in Fig. 5.2. The specific implementation of decimation on the FPGA fabric is explained first, followed by a discussion of the readout signals.

5.6.1 Decimation

Decimation, or down-sampling, is a process of re-sampling within a DSP system. It is used to reduce the sampling rate of a signal. Decimation can be implemented as a single-stage process or in several steps. It is also possible to carry out decimation using different computation methods, either on the PL or the PS. In our implementation, initial decimation is performed in two steps, both taking place in the FPGA fabric. The type of filter used is a constant coefficient FIR second-order cascaded integrator-comb (CIC) filter, selected for its linear phase response and simplicity of implementation in FPGA [125]. Such filters also provide notches of suppression (or nulls) exactly at the most critical frequencies, which would otherwise alias to very low frequencies [87].

CIC filter

Fig. 5.7 shows the block diagram of a second-order CIC filter. If an implementation requires a first-order or an N^{th}-order filter, the number of integrator and differentiator stages must be adapted accordingly. Differentiators are also referred to as ‘combs’ because their frequency-domain transfer functions resemble hair combs. To keep the discussion generalized, we will consider an N^{th}-order CIC filter for the mathematical description below.

The input signal is successively integrated N times. The transfer function of N consecutive discrete-time integrators, all running at the full sampling rate f_s, is given by

$$F_{\text{int}}(z) = \left(\frac{1}{1 - z^{-1}}\right)^N . \tag{5.26}$$

The output of the integrator stage is then re-sampled with a rate reduction of R. Hence, the new sampling frequency is f_s/R. For convenience, the integer rate change between the two sampling frequencies is typically chosen such that $R = 2^X$, with X being a positive integer. Now, the transfer function of the rate-reduced discrete-time differentiators operating at f_s/R is given by

$$F_{\text{comb}}(z) = \left(1 - z^{-R}\right)^N . \tag{5.27}$$

The product of the above two transfer functions yields the overall transfer function of the CIC filter. That is,

$$F_{\text{CIC}}(z) = \left(\frac{1 - z^{-R}}{1 - z^{-1}}\right)^N . \tag{5.28}$$

The gain of an N^{th}-order CIC filter at frequencies $f \ll f_s$ is equal to R^N.

Decimation in the GHz Phasemeter

Operating the entire DSP at a speed of 512 MHz in the FPGA is not trivial. At such speeds, the custom algorithms may result in unreliable or completely incorrect behavior simply because of the routing delays within the PL. However, running the GHz PLL at this speed is very important for the overall measurement scheme.

Nevertheless, when it comes to data readout, there is no need to clock the corresponding DSP parts at these high speeds. Hence, the decimation in the GHz Phasemeter is implemented in two stages, both utilizing second-order CIC filtering, with the second stage is clocked by a slower clock. Such an approach was necessary for optimizing the processing of logical operations and meeting the timing requirements.

The first pre-filtering stage has a rate reduction of 2^4, resulting in a pre-filter output decimated to 32 MHz. Consequently, the PL can operate at a slower clock domain, such as 32 MHz. Thus, the next stage of filtering is clocked with a 32 MHz clock[7], with a rate reduction of 2^{20}. This sets the lower bound for the data readout of the GHz Phasemeter at approximately 30.5 Hz.

It is possible to achieve a higher readout sampling frequency. For this purpose, we read out the CIC filter output more frequently from the PS side, necessitating a decimation rate smaller than those hard-coded in the two filtering stages. Consequently, the readout data must be scaled with an appropriate correction factor. The correction factor is defined as

$$K = \left(\frac{2^{(X_{\text{pre}}+X)}}{D_{\text{pre}} \cdot D}\right)^N , \tag{5.29}$$

where X_{pre} and X are the hard-coded values for the rate reduction purpose in the two CIC filtering

[7] Using the Clocking Wizard IP core, this clock is derived from the 256 MHz output clock of RF data converter IP core.

stages, D_{pre} and D are the chosen decimation rates for the readout, and N is the order of the CIC filter.

In our implementation, we have $X_{\mathrm{pre}} = 4$, $X = 20$, and $N = 2$. Typically, the readout decimation rate of the pre-filtering stage remains unchanged, so $D_{\mathrm{pre}} = 2^4$ as well. The remaining parameter D can then be selected based on the desired readout frequency. For example, if a readout frequency of 500 Hz is expected, then $D = (512\,\mathrm{MHz}/2^4)/500\,\mathrm{Hz}$.

The current implementation has a maximum readout rate of 10 kHz, which is because of the current software in the onboard PS. This rate could potentially be increased by optimizing the performance of the readout code, for instance, by introducing process priorities.

The involved clock domain transfer between the two filtering stages was necessary to manage the PL with 16 ADPLLs running at a processing rate of 512 MHz. However, crossing clock domains can be tricky and may lead to unwanted errors that are difficult to detect. Even though the slower clock is derived from the fast clock, it should be separated and synthesized manually, despite being treated and analyzed as a synchronous system. We explored clock domain transfer using a two-stage synchronizer (implemented with a custom-written VHDL algorithm), as well as AXI4-Stream Clock Converter IP core, to ensure that user logic operates without errors.

The actual data acquisition is handled using relevant IP cores related to memory access. The decimated data stream is first packetized and then sent to the AXI Direct Memory Access IP core. From here, an AXI block RAM controller takes control of the data and communicates with the local block RAM of the PS.

5.6.2 Frequency readout

Despite the nomenclature used for the instrument, the frequency value stored in the PIR is used for data acquisition in general. The phase value is later reconstructed from this frequency through a separate integration during the post-processing, if required. The reason for not reading out the PA value (or the s_p signal in Fig. 5.3) is that a constant frequency offset of the input signal will result in a constant phase ramp in the ADPLL. Due to this large ramp, PA will quickly overflow, making it unsuitable for data acquisition. Decimating such a saw-tooth function is quite difficult and requires special care [87]. The non-overflowing PIR value allows for standard decimation and filtering algorithms to be implemented for the readout.

In the GHz Phasemeter implementation, the decimation rates of the filtering stages are chosen to be consistent with the PIR readout. The 16-bit PIR value is resized to 24 bits after passing through the second-order CIC pre-filtering stage with $R_{\mathrm{pre}} = 2^{X_{\mathrm{pre}}} = 2^4$. The subsequent second-order CIC filter with $R = 2^X = 2^{20}$ brings the PIR data to 64 bits long, which matches the data length used for data acquisition. Consequently, there is no truncation of the PIR in the decimation stages, minimizing additional noise contribution due to truncation.

5.6.3 IQ readout

The Q and I signals of the PLL carry relevant information. As mentioned previously, Q indicates the quality of the tracking, while I scales with the amplitude of the tone being tracked. Therefore, these two signals are included in the phasemeter readout. The IQ readout is also important for compensating for the untracked phase error (also known as the residual phase) ε_{err} of the loop. By utilizing both quadrature components, additional phase reconstruction can be performed [87]. However, the high tracking bandwidth of the GHz PLL was sufficient to achieve adequate signal suppression at low frequencies, thus making the aforementioned post-processing technique unnecessary.

In our implementation, both Q and I values, which are initially 31 bits long, are dithered and truncated to 20 bits. This choice was made to keep the phase noise due to truncation below a level of sub-nanoradians, as explained in Sec. 5.7.2. As a result, 4 bits will be truncated in the pre-filtering stage, leaving the output data stream at 64 bits. At this point, the advantage of clock domain transfer can be emphasized again. If the second-order CIC filter with a decimation rate of 2^{20} is clocked at 512 MHz, the integrator stages would need to add two 64-bit numbers at each clock cycle, leading to timing errors in the PL due to routing delays in the fabric. Conversely, running the same signal processing at 32 MHz significantly simplifies the logic.

In addition to reading out the PIR, Q, and I values, an additional novel readout of Q^2 has been implemented in the GHz Phasemeter. This implementation is discussed separately in Sec. 5.7.4.

5.7 Residual phase error

The residual phase error ε_{err} within the ADPLL determines its stability and linearity. To calculate it, one must consider the relevant spectra of each noise source and the signal of interest, as well as understand how they are suppressed by the loop transfer function. Instead of examining the spectrum of the residual phase error, denoted as $\tilde{\varepsilon}_{err}$, we compute the RMS value by integrating the spectra to obtain a simple metric for the loop. The lower the RMS of $\tilde{\varepsilon}_{err}$, the better the loop behavior and signal tracking. Hence, not only is the low-frequency value of Q important, but its RMS value also needs to stay low, which requires a high tracking bandwidth of the PLL. This is especially crucial here because of the quickly changing frequency of the signal and the high laser frequency noise. Below is a brief discussion of two major noise sources in the GHz PLL. Fig. 5.3 shows these noise injection points in the loop.

5.7.1 Input additive and phase noise

Here we account for the contributions of additive noise sources such as electronic noise from the ADC, amplitude fluctuations of the input signal, and the input phase noise. For this analysis, the input signal given in Eq. 5.14 is extended to include an additive noise term $\tilde{A}$ and phase noise $\tilde{\varepsilon}$. Thus, we have

$$s_{in}[n] = \tilde{A}[n] + A\ \cos(\omega_0 n + \varepsilon_{in}[n]). \tag{5.30}$$

The phase term ε_{in} includes both the noise term $\tilde{\varepsilon}$ and the signal ε_s, as the PLL cannot distinguish between them.

To obtain the untracked residual phase error, we compute the standard deviation of the residuals as [87]

$$\sigma_{\text{phase}} = \int_0^\infty \tilde{\varepsilon}_{in}(z) \cdot \frac{1}{1 + G(z)}\, df, \tag{5.31}$$

where $G(z)$ is the PLL open-loop gain and $\tilde{\varepsilon}_{in}$ is the ASD of the phase of the input signal.

Similarly, for the error generated by the input additive noise, we have

$$\sigma_{\text{add}} = \int_0^\infty \frac{\sqrt{2}\tilde{A}}{A} \cdot \frac{G(z)}{1 + G(z)}\, df, \tag{5.32}$$

where the factor of $\sqrt{2}$ arises from the mixing process.

From these two equations, it is evident that while both additive and phase noise are present in the input signal, the loop treats them differently. The input phase noise is suppressed by the loop gain, whereas the additive noise is converted into phase noise by the action of the phase detector. Thus, it is essential to account for the contribution of additive noise to the overall residual phase error.

As discussed in Sec. 5.3.2, an N-bit signal sampled with a sampling frequency of f_s will result in quantization phase noise of

$$\tilde{u}_N = \frac{2^{-N}}{\sqrt{6 \cdot f_s}} \qquad [1/\sqrt{\mathrm{Hz}}]. \tag{5.33}$$

The ADCs of the RFSoC have an effective number of bits (ENOB) above 10.5, which factors in the ADC noise floor as well [126]. Therefore, for the additive noise analysis, we assume that the total contribution for $\tilde{A}$ comes from $\tilde{u}_{\mathrm{ENOB}}$. Since the quantization noise is white in spectrum, Eq. 5.32 simplifies to

$$\tilde{\varepsilon}_{\mathrm{ENOB}} = \frac{\sqrt{2} \cdot \tilde{u}_{10.5}}{A} \cdot \frac{G(z)}{1 + G(z)} \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}], \tag{5.34}$$

with A as signal amplitude normalized to the maximum peak amplitude possible for these data converters. This represents the effective phase noise contribution due to the considered additive noise.

5.7.2 Truncation noise

Truncations within the ADPLL are unavoidable. They are introduced to reduce the overall bit-width, which also helps keep processing delays in check. The linear model of the ADPLL can be utilized to understand the impact of in-loop truncation noise on the residual phase error and, consequently, on phase-tracking performance. In our implementation, truncations are performed at two points in the loop, as depicted in Fig. 5.3. Contributions from these two truncations are considered separately as follows.

Truncation of the PIR

As the PIR value is of primary interest when using the PLL as a frequency-tracking instrument, it might seem advantageous to maintain the highest number of bits for the PIR. However, it turns out that the truncation noise of the PIR is directly suppressed by the open-loop gain. Consequently, the effective phase noise due to PIR truncation can be easily reduced to below the microradian regime. Additionally, having a higher number of bits for the PIR complicates the operation of the subsequent PA stage at high processing speeds.

In the case of the GHz PLL implementation, truncating the PIR to 16 bits after dithering introduces truncation noise (refer to Eq. 5.13) given by

$$\tilde{u}_f = \frac{\sqrt{3} \cdot 2^{-16}}{\sqrt{6 \cdot 512 \times 10^6}}. \tag{5.35}$$

This truncation noise manifests as phase noise, which can be expressed as

$$\tilde{\varepsilon}_{\tilde{u}_f} = \tilde{u}_f \cdot \frac{F_{\mathrm{PA}}(z) \cdot F_{\mathrm{LUT}}(z)}{1 + G(z)} \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}]. \tag{5.36}$$

Truncation of Q

Truncating the Q signal was necessary in our implementation of the GHz Phasemeter. It is therefore essential to ensure that the induced phase noise does not limit overall performance. The choice of the required number of bits for Q is made such that the phase noise due to truncation of Q is negligible in comparison to the phase noise caused by additive input noise. This is because, after the phase detector block in the PLL model, it becomes impossible to distinguish between the two aforementioned noise contributions.

From Eq. 5.34, we have

$$\tilde{\varepsilon}_{\mathrm{ENOB}} = \frac{\sqrt{2} \cdot 2^{-10.5}}{A\sqrt{6 \cdot 4.096 \times 10^9}} \cdot \frac{G(z)}{1+G(z)} \approx \frac{6.23 \times 10^{-9}}{A} \cdot \frac{G(z)}{1+G(z)} \ \mathrm{rad}/\sqrt{\mathrm{Hz}}. \tag{5.37}$$

Now, the phase noise induced by the truncation of Q to 20 bits (after dithering) is given by

$$\tilde{\varepsilon}_{\tilde{u}_Q} = \frac{4}{A} \frac{\sqrt{3} \cdot 2^{-20}}{\sqrt{6 \cdot 512 \times 10^6}} \cdot \frac{G(z)}{1+G(z)} \approx \frac{1.19 \times 10^{-10}}{A} \cdot \frac{G(z)}{1+G(z)} \ \mathrm{rad}/\sqrt{\mathrm{Hz}}. \tag{5.38}$$

Here, we use the transfer function of the phase detector, $A/4$, thereby scaling the truncation noise with the closed-loop gain.

Comparing the two contributions, it is clear that the additive noise at the input will dominate over the noise resulting from the truncation of Q. That is,

$$\sqrt{(\tilde{\varepsilon}_{\mathrm{ENOB}})^2 + (\tilde{\varepsilon}_{\tilde{u}_Q})^2} \approx \tilde{\varepsilon}_{\mathrm{ENOB}}. \tag{5.39}$$

Hence, the noise contribution due to the truncation of Q can be discarded from the further analysis.

5.7.3 Effective residual phase error

With the known noise sources and a proper understanding of the ADPLL model, it becomes possible to realize how the relevant noise spectra are shaped by the loop and how they effectively contribute to the residual phase error ε_{err} of the PLL.

As input signal, the beat note between the two CTLs was used. To obtain an input signal spectrum covering a wide range of Fourier frequencies, four different measurements were carried out in various Fourier frequency domains. The PIR values were also read out at 512 MHz using an ILA IP core with the maximum sample data depth available. The recorded PIR values were then integrated in post-processing to derive the phase values. Subsequently, to calculate the contribution from the signal, we use a spectrum of its dynamics that is suppressed by the loop error transfer function, $1/(1 + G(z))$.

Fig. 5.8 shows the spectrum of the signal with and without loop suppression. The relevant noise spectra, shaped by the loop and contributing to the effective residual phase error noise, are also plotted. It is evident that, while tracking high dynamic signals such as the beat note between the CTLs, the contribution from input phase dynamics dominates. Therefore, in such applications, the objective is to maximize the tracking bandwidth of the PLL. The achieved 2 MHz tracking bandwidth thus plays a major role.

Moreover, the input signal spectrum plotted in Fig. 5.8 is of great importance, as it captures the CTL phase noise spectrum over a Fourier frequency range from a few hertz to a quarter of a gigahertz.

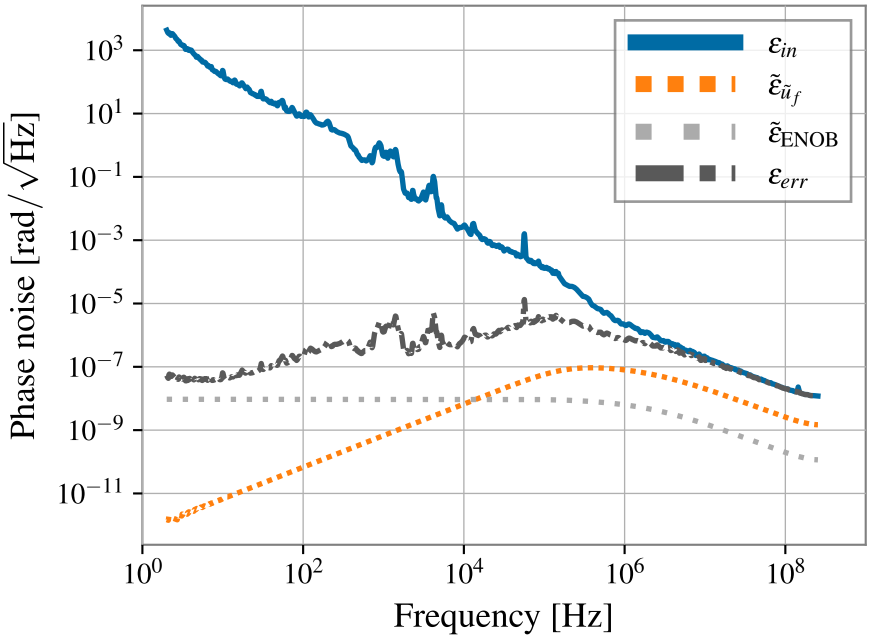

Figure 5.8: Spectrum of the residual phase error along with the relevant noise spectra shaped by the loop. The plot shows the measured phase noise spectrum of the beat note between two CTLs (solid trace) and its attenuation by a PLL with 2 MHz loop bandwidth (dash-dotted trace). The loop attenuation of the input is calculated using the linear ADPLL model. The spectrum is constructed from four measurements carried out in different Fourier frequency domains. The dotted traces represent the effective phase noise contributions from the input additive noise and the PIR truncation noise. Notations used in the plot are according to those used in Fig. 5.3.

Such a measurement was possible only because of the development of the GHz Phasemeter.

As mentioned at the beginning of this section, another quantity of interest is the RMS of the residual phase error. This value can be directly used to quantify the behavior of the PLL. To demonstrate this, we utilize a very similar input signal as that in Fig. 5.8, and its spectrum is calculated. The ADPLL model is tuned once to achieve the highest tracking bandwidth of 2 MHz and another time to a lower tracking bandwidth of 100 kHz. The resulting spectra of the residual phase error are calculated by considering the open-loop gain in these two different configurations. The RMS of the residual phase error is then determined by integrating the spectrum. We used the Python package ‘Spicypy’ [124] for this purpose.

Fig. 5.9 illustrates the resulting spectra and the associated RMS values. It is quite clear that for the same input phase dynamics, the PLL with a higher tracking bandwidth achieves better suppression, resulting in a lower RMS value. Usually, the higher the tracking bandwidth is, the better is the tracking performance. The ‘noise’ present in the input signal spectrum at a few tens of kilohertz contributes to the RMS value of the residual phase error when the loop bandwidth is not high enough. Lower RMS values correspond to a more stable and linear loop performance.

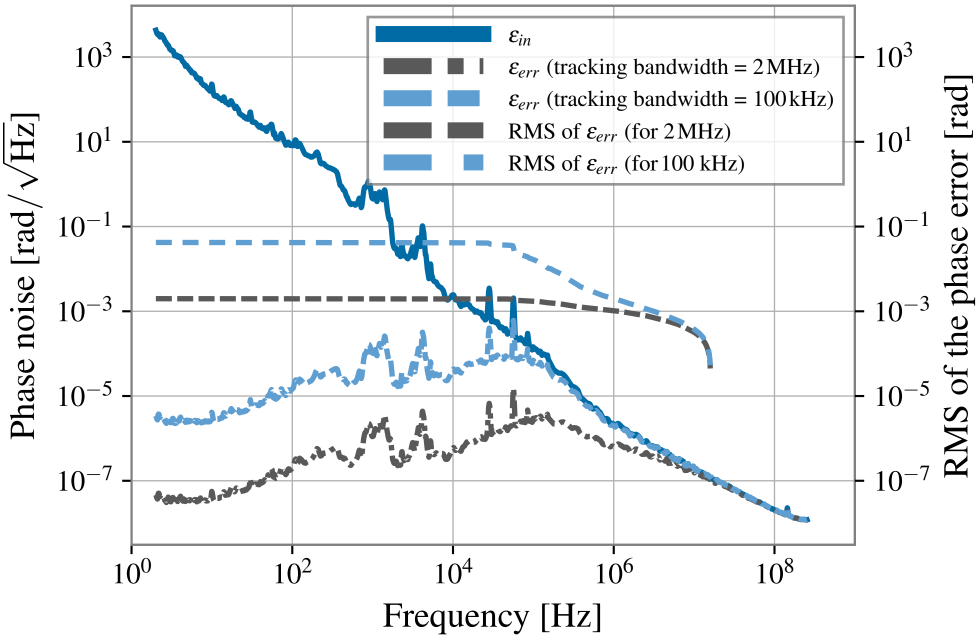

Figure 5.9: Spectra and RMS values of the residual phase error for two different tracking bandwidths of the PLL. The loop attenuation of the input (dash-dotted trace) is calculated using the loop transfer function model. The dashed traces indicate the RMS of the residual phase error, calculated starting from a Fourier frequency of 16 MHz and extending towards lower frequencies.

5.7.4 Q^2-readout

Because of the importance of the RMS value of the residual phase error in quantifying the loop performance, we implemented a real-time measurement of this quantity. This implementation is novel in the context of general ADPLL systems and is also beneficial for other applications, such as LISA phasemeters. To obtain a direct measurement of the RMS of the residual phase error, we want to calculate the quantity

$$\frac{1}{N}\sum_{i=1}^{N}\left(\varepsilon_{err,i}-\bar{\varepsilon}_{err}\right)^2, \tag{5.40}$$

where $\bar{\varepsilon}_{err}$ is the expectation value, which is equal to 0. Consequently, monitoring the Q^2 value over a specific duration in the ADPLL allows us to derive ε_{err}^2 after appropriate scaling.

Implementation of the Q^2-readout

The 20-bit long Q signal, processed at 512 MHz in the PL, is initially filtered using a second-order CIC filter with a decimation factor of 2^4. The output of this filtering stage is maintained at 24 bits, and a clock domain transfer to a slower clock frequency of 32 MHz is performed, as described in Sec. 5.6.1. The pre-filter output at the slower clock domain is multiplied by itself[8], generating a 48-bit

[8] In this case, manual handling of the multiplication process in the PL using the DSP48 Marco IP core is not required, as we operate in a slower clock domain.

Q^2 signal. This signal is then filtered with a first-order CIC filter, decimated by 2^{20}, producing 64-bit data for the readout.

The first-order CIC filter is chosen to act as the mean value calculator for Q^2. Consequently, the recorded Q^2 is scaled using the correction factor from Eq. 5.29 with $N = 1$. Finally, to obtain the estimation for the RMS of the residual phase error, the Q^2-readout must be divided by $A/4$, the transfer function of the phase detector stage.

Confirmation of the Q^2-readout

The confirmation of the Q^2-readout is accomplished by comparing the mean value of this readout against the loop model calculations under different bandwidth conditions. To do this, the procedure followed in Fig. 5.9 is repeated, but this time with the Q^2 values also recorded.

The beat note between two CTLs was tracked under four different PLL bandwidth conditions. From the recorded frequency values, a phase noise spectrum is calculated, and the residual phase error spectra for each scenario is obtained using the loop model. Since the Q^2-readout contains information up to a Fourier frequency of 16 MHz, the RMS of the residual phase error is calculated from the spectra beginning at the same Fourier frequency and going towards lower frequencies. We then compare the mean value of the Q^2 time series with the calculations from the loop model. Fig. 5.10 presents the corresponding results. It is important to note that each point on this plot is supported by an associated spectrum.

There is a strong agreement between the RMS of the modeled residual phase error and the one predicted using the Q^2-readout. The discrepancies may arise from our lasers, which exhibit higher non-stationary noise. However, it is quite clear that this novel readout can be utilized to estimate the residual phase error of the loop directly. The loop model approach to determining the residual phase error is quite tedious, as it relies on accurate loop modeling and signal dynamics, and cannot be applied in real-time to improve the tracking performance. The Q^2-readout feature enables us to deduce and minimize the residual phase error while tracking a signal by tuning the PLL servo gains in real-time. This capability is particularly valuable when a model of the signal dynamics is not available, or when the spectrum of the dynamics is non-stationary.

During measurements, the GHz Phasemeter is programmed to display the PIR, Q, I, and Q^2 values on an external personal computer (PC) terminal, from which the PS of the RFSoC is controlled. This setup allows direct improvement of tracking performance by adjusting the servo gains to minimize the Q^2 value. The maximum and minimum amplitudes of the input signal are also displayed to avoid exceeding the input range of the ADC channel under consideration.

5.8 Performance of the GHz Phasemeter

The functionalities of the phasemeter were initially tested using behavioral simulations. Test benches were written in VHDL for individual blocks of the PLL, particularly focusing on the multi-demodulation and phase accumulation stages. After these functional tests, further measurements and characterizations were conducted.

Performance measurement of the GHz Phasemeter is important for understanding the phase readout noise floor of the instrument. This, in turn, provides the corresponding displacement readout noise, as discussed in Sec. 3.3.2, when the GHz Phasemeter is used to track the beat note in our interferometric

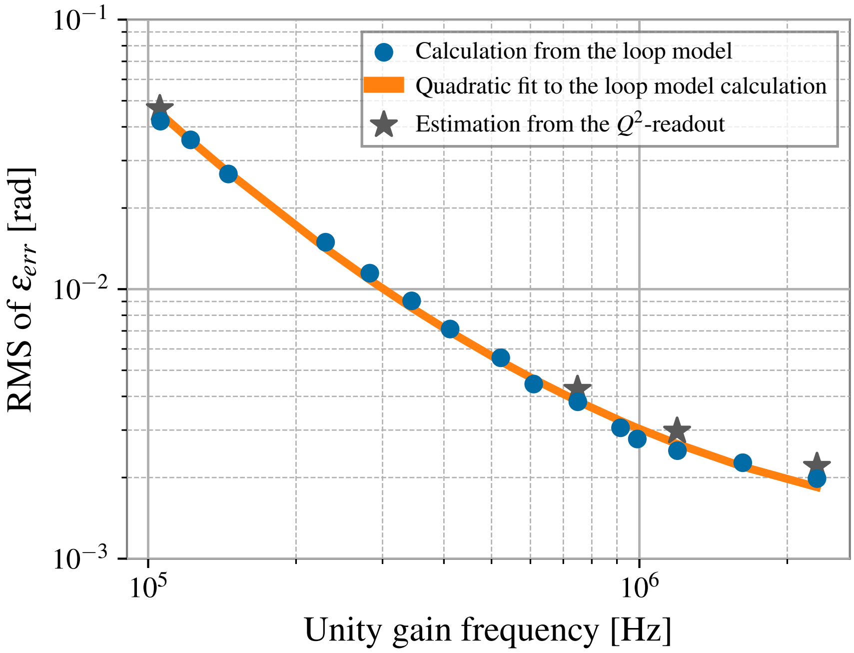

Figure 5.10: The RMS of the residual phase error for different UGFs (or tracking bandwidths) of the PLL. For a set of servo gain combinations, the loop suppression is calculated for the laser phase noise spectrum shown in Fig. 5.9. The RMS of this residual phase error is calculated similarly and is marked as a dot in the plot. A quadratic fit in logarithmic scale to all the calculated points provides insight into how the RMS of the residual phase error varies with the UGF of the PLL. For four of the loop bandwidth conditions, the RMS of the residual phase error is estimated using the novel Q^2-readout.

scheme. In addition to understanding the phasemeter measurement noise floor, we also demonstrate the advantage of its high tracking bandwidth and characterize the instrument for its tracking speed and capture range.

Before starting with these analyses, we will revisit the characterization of the frequency noise of the CTL, as we now have the necessary instrument at our disposal.

5.8.1 Characterizing the frequency noise of a CTL

A reliable characterization of the frequency noise of the CTL was not possible before because of the unavailability of a proper frequency readout method (refer to Sec. 3.4). Using the developed GHz Phasemeter, it is now possible to track the beat note between two CTLs with ease, thanks to its large tracking bandwidth and throughput. Fig. 5.11 displays the free-running laser frequency noise of a CTL, starting at a Fourier frequency of one millihertz and extending up to a quarter of a gigahertz. The internal piezoelectric resonances around a kilohertz are visible in the spectrum. The measured noise of $10\,\mathrm{kHz}/\sqrt{\mathrm{Hz}}$ at 1 Hz aligns with the expected frequency noise value for NPRO lasers. However, in the case of the CTL, this measurement should be considered as being taken during a quiet stretch. The observed white noise above 1 MHz is attributed to the quantization noise of the ADC.

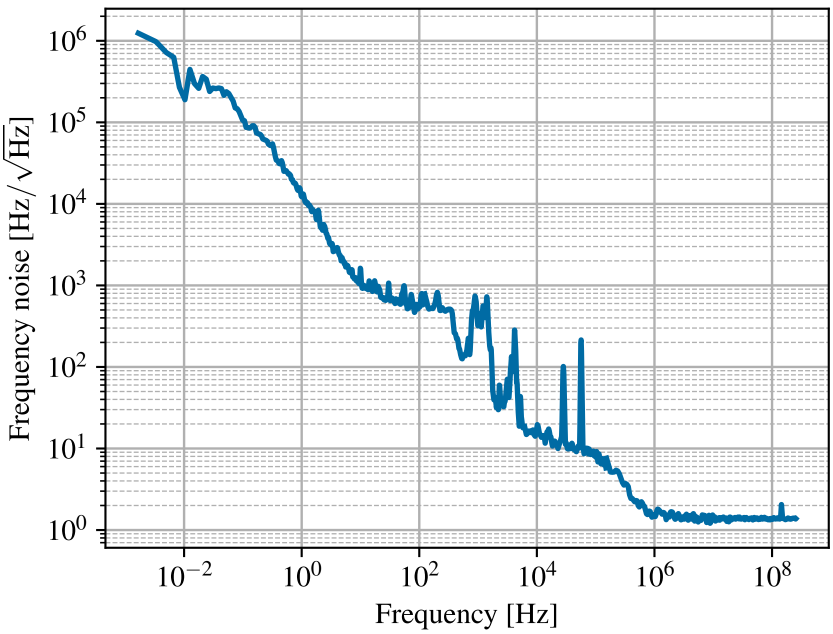

Figure 5.11: Frequency noise of a free-running CTL over a wide range of Fourier frequencies. The beat note between the two CTLs was measured using the GHz Phasemeter across four different Fourier domains. The plotted spectrum is constructed using these four measurements.

5.8.2 Measurement noise floor

Methodology

To probe the phase measurement noise floor, the conventional split-measurement technique is employed. In a split-measurement test setup, a signal is split into two (or more) measurement chains. Each segment is then coupled to a phasemeter channel, resulting in the two measured phases differing only due to effects in the measurement chains. Any noise that is common to the two phasemeter channels is subtracted, making the dynamics or stability of the input signal itself irrelevant. A schematic sketch of the measurement setup is shown in Fig. 5.12.

When such a measurement is conducted, we record time series data $\varphi(t)$ from the two channels and focus on understanding the self-noise $\tilde{n}$. We quantify the self-noise by calculating its PSD, $\tilde{S}_{nn}(f)$. Without loss of generality, it can be assumed that $H_A(s) = H_B(s)$, as the two phasemeter channels operating on the same board are identical.

Thus, the PSD of the recorded time series difference will contain information solely about the self-noise of each channel. That is,

$$\tilde{S}^2_{\Delta\varphi}(f) = \tilde{S}_{nn,A}(f) + \tilde{S}_{nn,B}(f) \qquad [\text{rad}^2/\text{Hz}]. \tag{5.41}$$

Assuming that the self-noise in each phasemeter channel is the same, we can arrive at the self-noise or

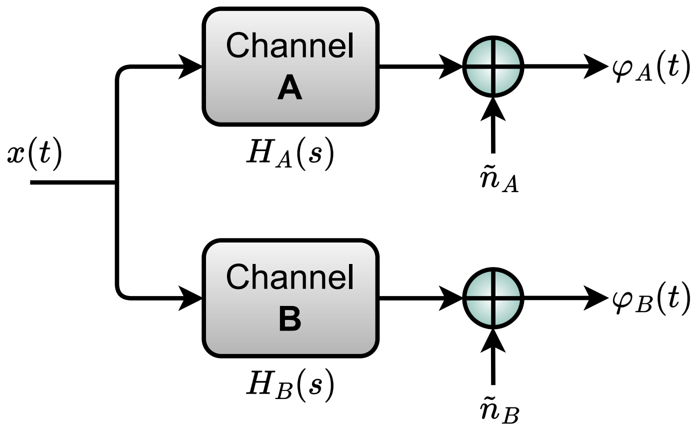

Figure 5.12: Schematic sketch of the setup used for characterizing the phase measurement noise floor. The input signal is split and fed into two phasemeter channels: A and B. $\tilde{n}_i$ represents the self-noise of Channel i, which is our quantity of interest. $H_i(s)$ is the transfer function of Channel i in the Laplace domain.

the measurement noise floor of a phasemeter channel as

$$\tilde{S}_n(f) = \frac{1}{\sqrt{2}}\,\tilde{S}_{\Delta\varphi}(f) \qquad [\mathrm{rad}/\sqrt{\mathrm{Hz}}]. \tag{5.42}$$

This means that by calculating the spectrum of the difference, we can characterize the measurement noise floor.

The split-measurement method described above makes one critical assumption: the self-noise floor of both channels is identical. While this may be true, we cannot confirm it at this stage. Also, the common noise is not visible in the spectrum of the difference, as it is subtracted out. To characterize the measurement noise floor of an individual channel, a coherent subtraction method with transfer function estimation can be employed [127, 128]. For a system sketched in Fig. 5.12, it follows that

$$\tilde{S}_{nn,A}(f) = \tilde{S}_{\varphi_A\varphi_A}(f) - \tilde{S}_{\varphi_B\varphi_A}(f)\frac{H_A(s)}{H_B(s)} \qquad [\mathrm{rad}^2/\mathrm{Hz}]. \tag{5.43}$$

Here, $\tilde{S}_{\varphi_A\varphi_A}(f)$ is the PSD of $\varphi_A(t)$, and $\tilde{S}_{\varphi_B\varphi_A}(f)$ is the cross spectral density (CSD) between the two output time series.

In this method, the only assumption required is that the two noises are uncorrelated, meaning $\tilde{S}_{n_An_B}(f) = \tilde{S}_{n_Bn_A}(f) = 0$. Any critical assumptions about the transfer functions of the two channels being identical or their self-noises being the same are not necessary here.

ADC sampling jitter is another prominent limitation in digital phasemeter implementations. This leads to frequency-dependent phase noise, as described by Eq. 3.29. Since this sampling jitter couples into all tones present in the ADC input, it is possible to correct for it. A reference tone, the so-called pilot tone, is added to the input, and its measured phase variations are used to correct the phase of the main tone of interest. This technique, referred to as pilot tone correction (PTC), is utilized in LISA phasemeter developments [115].

Consider a pilot tone f_P present in the input signal, with a PLL assigned to track its phase φ_P. The tracked phase φ_A of the main tone f_A can then be corrected as

$$\varphi_{A,\mathrm{corr}} = \varphi_A - \varphi_P \frac{f_A}{f_P}. \tag{5.44}$$

Such a PTC scheme is investigated in our measurement noise floor characterization to account for the effects of ADC sampling jitter.

Experimental characterization

As mentioned previously, the GHz Phasemeter implementation includes two ADPLLs per ADC channel. For the phase readout performance testing, two tones (referred to as main and pilot tones) of different frequencies were generated using external signal generators. These tones are first combined and then split using RF power splitters/combiners (ZX10R-14-S+ from Mini-Circuits) before being fed into two phasemeter channels, labeled A and B. Each channel uses one ADPLL to track the main and the other to track the pilot tone.

From the recorded frequency values (readout of the PIR), the measured phase for each channel is calculated. The phase measurement noise floor is then determined by subtracting the phase values from the two channels and calculating the spectral density. This method provides a measure of the incoherent sum of the noise present in both channels. Later, we also used the coherent subtraction method with transfer function estimation (as described in Eq. 5.43) to get the measurement noise floor of each phasemeter channel independently.

Fig. 5.13(a) compares the measurement noise floor results from both methods (the split-measurement technique, also known as zero-measurement, and the coherent subtraction method with transfer function estimation) for tracking a tone in the LISA frequency band. Both methods estimate similar noise levels, indicating no significant phase noise differences between the channels. One possible candidate for the current limitations is the phase noise originating from the analog front-end, often driven by temperature and humidity fluctuations [129]. To reach the performance requirements, it would require active temperature stabilization of the analog front-end and passive heat exchanging from ADCs and FPGAs [130]. Such an extensive performance measurement was not carried out for the GHz Phasemeter, as the current noise floor is low enough for our interferometry application.

The measurement noise floor results, with and without the PTC, are plotted in Fig. 5.13(b). A 100 MHz pilot tone was used to correct the phase fluctuations of the 2 GHz main tone. The results show that we are currently limited by ADC timing jitter at frequencies below 2 Hz for higher signal frequencies. For frequencies up to 2 GHz, we achieve a noise floor below 1 mrad/$\sqrt{\mathrm{Hz}}$ at nearly all readout frequencies, even without optimizing the analog front-end. This measurement validates the phasemeter's capability to directly track signals in the gigahertz regime with high precision.

In the LISA band measurements shown in Fig. 5.13(a), we observed no performance improvement when using an additional 38.2 MHz pilot tone and the corresponding correction. Instead, the phase noise maintains its $1/\sqrt{f}$ characteristic, indicating that an unknown Flicker-type noise is the current limit for a LISA-type phasemeter implementation in these RFSoCs.

5.8.3 Role of high tracking bandwidth

The advantages of having a high tracking bandwidth have been discussed in Sec. 5.7.3 and quantitatively illustrated in Fig. 5.9. To further demonstrate the benefits of high loop bandwidth, we simultaneously

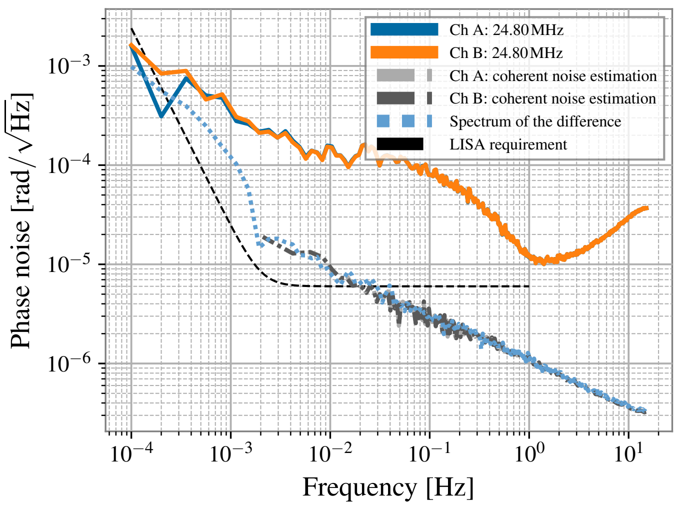

(a)

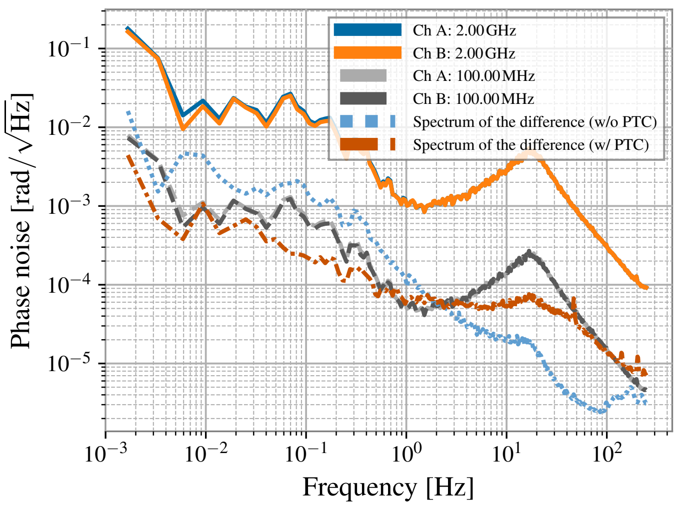

(b)

Figure 5.13: Measurement noise floor of two of the GHz Phasemeter channels. (a) Self-noise levels calculated using the coherent subtraction method with transfer function estimation (dash-dotted trace). The input signal contained a 24.8 MHz main tone. Both channels exhibit very similar measurement noise floors. In the split-measurement technique, two time series values were subtracted and the spectrum of the difference was evaluated (dotted trace). Above a Fourier frequency of about 20 mHz, the phase noise floor is below the LISA requirement (dashed trace). (b) Two PLLs were set to track a 2 GHz main tone and a 100 MHz pilot tone. During post-processing, each channel's main tone phase readout values were corrected using the corresponding pilot phase. Above a Fourier frequency of about 0.1 Hz, the phase noise floor is in the sub-milliradian regime. (PTC: pilot tone correction.)

tracked the beat note between two CTLs under two different tracking bandwidth conditions. As mentioned previously, these lasers have a relatively high phase noise spectrum and are sensitive to non-stationary excess noise, making effective phase- (or frequency-) tracking challenging with

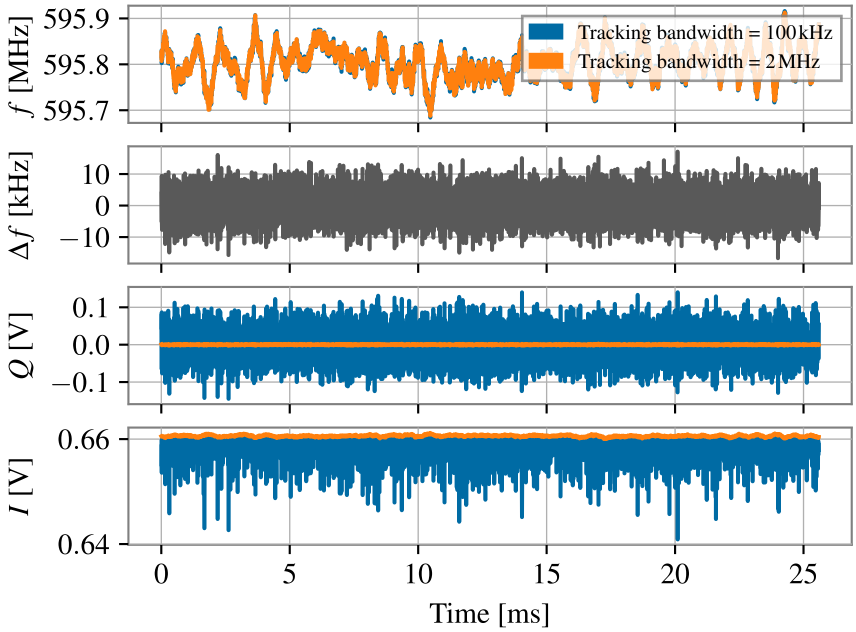

Figure 5.14: Experimental demonstration of the role of PLL's high tracking bandwidth when tracking 'noisy' or highly dynamic signals. The beat note between the two CTLs was tracked simultaneously using the GHz Phasemeter, one PLL with a tracking bandwidth of 100 kHz and the other with 2 MHz. The phasemeter readouts: PIR, Q, and I, are plotted as time series. The second subplot displays the difference in the tracked frequency as measured by the two ADPLLs.

commercially available phasemeter solutions or a LISA-type phasemeter.

We tracked the laser beat note using two ADPLLs operating on a single ADC channel, each configured with different servo gains to create the two tracking bandwidth conditions depicted in Fig. 5.9. Fig. 5.14 shows the measured signal frequency f, Q, and signal amplitude I. The measured difference between the two frequency estimates from high and low loop gains is also plotted.

Examining their behavior highlights the role of high tracking bandwidth. Tracking with low loop bandwidth introduces significant frequency errors, leading to an expected increase in the RMS for Q. For the amplitude I, an underestimation is observed, exhibiting one-sided noise behavior resulting from non-linear effects in the I demodulation and decimation [87]. Furthermore, limited tracking bandwidth can lead to the ADPLL temporarily or permanently losing track of the input signal.

5.8.4 Tracking highly dynamic signals

Tracking speed

A PLL can track signals only up to certain dynamics. The maximum tracking speed of a PLL defines the highest rate of change of input frequency that the PLL can reliably track. To characterize this value for the GHz Phasemeter, it was initially locked to a fixed frequency from an external signal generator. By introducing a frequency modulation at various modulation frequencies, we assessed the behavior of the phasemeter. With a modulation frequency of 30 kHz and a frequency deviation of

4 MHz, the GHz Phasemeter could track the frequency changes reliably, demonstrating its ability to follow signals that change at a speed of 240 GHz/s. It is worth noting that this is not the maximum tracking speed of our ADPLL; rather, it was constrained by the available signal generators and the data readout frequency of the phasemeter.

Acquisition range

The acquisition range, also known as pull-in range or capture range, of a PLL is critical for locking the phasemeter to highly dynamic signals. It defines how far the initial frequency estimate f_0 can deviate from the input signal frequency for a successful lock of the PLL. Using a straightforward testing method where f_0 is maintained at a fixed value while the input tone is varied, we determined that our PLL has an acquisition range of 4.1 MHz. Since the acquisition range scales with the tracking bandwidth, it is beneficial that our phasemeter provides both high tracking and detection bandwidths. The lower limit of its signal bandwidth is around 10 kHz. Below this input frequency, the tracking behavior and precision began to degrade significantly.

It is possible to implement a fast Fourier transform (FFT) calculation of the ADC data (in either the PL or PS) to identify the desired locking frequency. This approach allows for an automated lock acquisition procedure similar to that developed for space-based laser interferometry [131] to be applied to the GHz Phasemeter as well.

5.9 Conclusion

We have developed and demonstrated an 8-channel GHz Phasemeter based on an RFSoC system. Although this development of a frequency-readout system began as part of a displacement sensing scheme, it has evolved into a standalone instrument. The GHz Phasemeter features a signal bandwidth of 2.048 GHz, a tracking bandwidth of 2 MHz, a capture range of 4.1 MHz, and a tracking speed of more than 240 GHz/s.

Achieving a measurement noise floor in the sub-milliradian regime when comparing two signals, even at gigahertz frequencies, represents a commendable performance anticipated by many applications. The readout noise floor is comparable to that of the LISA phasemeter at measurement frequencies around 1 Hz for input signal frequencies around 25 MHz, reaching a phase readout precision of microradian. For signal frequencies of 2 GHz, this precision reduces to approximately 100 µrad (see Fig. 5.13).

Given the information about signal bandwidth and precision, we can determine the dynamic range. For the LISA phasemeter, this results in a maximum dynamic range on the order of about $(25 \times 10^6)/(1 \times 10^{-6}) = 2.5 \times 10^{13}$. With the GHz Phasemeter, we maintain a similar order of dynamic range but shift it to a different frequency range. We find a dynamic range of $(2.048 \times 10^9)/(1 \times 10^{-4}) = 2.048 \times 10^{13}$ for the readout, assuming the worst-case precision demonstrated.

With this phasemeter, tracking the beat note between two CTLs becomes significantly easier compared to other phasemeter solutions. This phasemeter can be utilized to realize laser interferometric readout schemes, such as heterodyne cavity-tracking, which require a high range and dynamics of frequency changes.

CHAPTER 6

Digital Heterodyne Controller

The choice of heterodyne stabilization (HS) technique for the laser frequency control was made to achieve very high displacement sensitivity with a fringe-scale dynamic range. In this scheme, the information of interest lies in the heterodyne frequency, which is also used for demodulation. Therefore, having just a high-bandwidth digital readout system is insufficient because we have seen in Sec. 3.5 that performing the heterodyne demodulation using analog electronics is hard and prone to unwanted effects. Consequently, we have integrated the laser-locking control loop together with the GHz Phasemeter into the PL of the same RFSoC using an FPGA algorithm. This implementation is referred to as the Digital Heterodyne Controller (DHC).

Such a digital implementation using FPGAs has been demonstrated by Eichholz et al. [85]. In this work, we created an implementation with a significantly larger bandwidth capable of handling varying heterodyne frequencies over a wide range. To achieve heterodyne cavity-tracking with a high dynamic range, such an implementation is necessary.

6.1 Implementation

In an all-digital laser-locking scheme using the HS technique, the error signal, as described in Eq. 3.9, is generated similarly to the analog implementation. One can implement a direct digital demodulation, in which the cavity reflection and reference beat signals are digitized and multiplied in a digital mixer. The demodulated signal can then be low-pass filtered, passed through a servo stage, and input to a DAC. This analog signal, the output of DAC, acts as the actuation signal for the laser frequency.

However, just as in the analog HS case, a correct demodulation phase is critical for the locking scheme to function properly. Direct digital demodulation poses challenges, as the correct demodulation phase must be achieved either through physical delays or by using digital registers before the mixer. This approach to setting the correct demodulation phase will only work for a given beat frequency. Consequently, it brings us back to the issues encountered with the analog implementation of HS, making it difficult to maintain compatibility with arbitrary beat frequency. For this reason, we decided against using direct digital demodulation for our implementation.

In our approach, the reference beat note is tracked by a GHz Phasemeter channel. The PIR value of the GHz PLL is then used to generate the demodulation signal, leading to an indirect demodulation technique. This allows us to tune the phase of the demodulation signal as needed, because there will

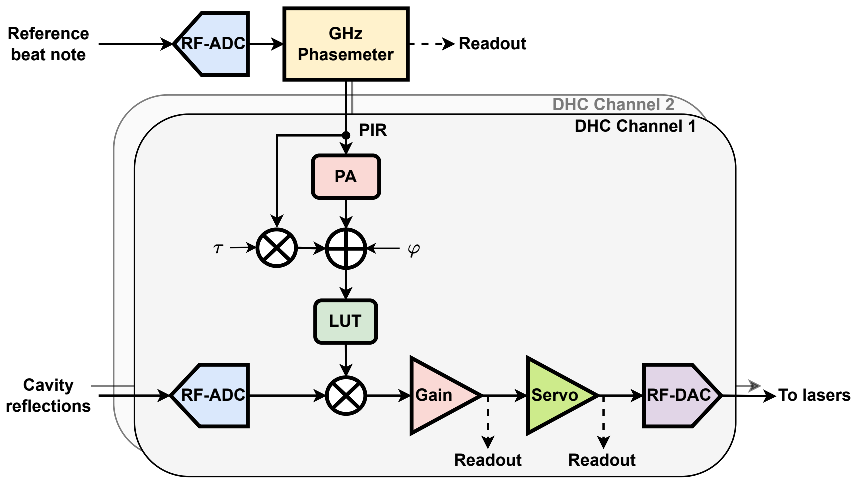

Figure 6.1: Block diagram of the DHC. The reference beat note and reflections from the cavities serve as inputs to the RF-ADCs. The actuation signals for the laser frequencies emerge from the RF-DACs. Although a simple mixer is shown for heterodyne demodulation, it represents the multi-demodulation stage followed by a rolling average (as seen in Sec. 5.4.2 for the implementation of the GHz Phasemeter). The same applies to the PA stage. (PIR: phase increment register, PA: phase accumulator, LUT: look-up table, τ: frequency-dependent demodulation phase correction term, φ: frequency-independent demodulation phase correction term.)

be a secondary NCO instance in addition to the one in the PLL. Due to the higher gain and bandwidth of the phasemeter channel compared to the laser-locking loop, the second NCO acts as a real-time copy of the reference beat but with a correct demodulation phase [132].

The block diagram of our DHC implementation is sketched in Fig. 6.1. The beat note between the lasers is coupled to an ADC on the ZCU111 evaluation board. This tone is then tracked by the GHz Phasemeter, which is our readout system. The current frequency value of the PLL, stored in the PIR, acts as the reference for the heterodyne demodulation. The cavity reflection is demodulated against a secondary NCO instance derived from the PIR of the GHz PLL, corrected for the proper demodulation phase. The subsequent low-pass filtering and feedback control implementation in the FPGA is a straightforward process, handled in a manner very similar to that in the ADPLL (refer to Sec. 5.4.2). The number generated by the servo is converted to a voltage by a DAC, which is then passed to the laser frequency modulation port.

6.1.1 Demodulation phase correction

While we were analyzing the HS scheme, to arrive at Eq. 3.3, we neglected the z-dependency and focused solely on the time-dependent term. The generalized form of Eq. 3.3, representing the reference beat, can be expressed as

$$Y(t) = A_1^2 + A_2^2 + 2A_1A_2\cos(\Delta\omega t + \Delta\phi(z)), \tag{6.1}$$

where the propagation phases are included into the phase difference

$$\Delta\phi(z) = -(k_1 - k_2)z + \phi_1 - \phi_2. \tag{6.2}$$

For correct demodulation, we require $\Delta\phi(z) = -\pi/2$ [85].

However, when considering contributions from both signals in the heterodyne demodulation, the correct demodulation phase can be expressed as

$$\Delta\phi_Y(z) - \Delta\phi_X(z) = (k_2 - k_1)\Delta z = -\pi/2, \tag{6.3}$$

indicating that Δz must be non-zero. This intentional distance offset between the sampling points of the cavity reflection and the reference beat is almost always present in an experimental setup. However, since

$$k_2 - k_1 = \frac{2\pi}{c}(f_2 - f_1), \tag{6.4}$$

the demodulation phase is affected by any changes in the beat frequency, $f_2 - f_1$.

This dependency can be entirely eliminated in a digital implementation. The phase of the demodulation signal can be freely adjusted by adding phase offsets before the 'look-up' process for sinusoidal signal generation. Our implementation of the DHC in the FPGA algorithm includes two methods to correct for the heterodyne demodulation phase, as explained below.

Frequency-independent demodulation phase correction

This is a constant phase value (denoted as φ in Fig. 6.1) that is added to all eight outputs of the multi-phase accumulation stage. This straightforward implementation helps maintain the correct demodulation phase for the heterodyne demodulation. In an ideal scenario, when indirect digital demodulation with a secondary NCO is chosen, one can aim to achieve $\Delta z = 0$, making the demodulation phase entirely independent of the beat frequency. In this case, φ determines the correct demodulation phase.

Frequency-dependent demodulation phase correction

In the actual experimental setups, achieving perfect path length matching to make $\Delta z = 0$ is very difficult. Any residual Δz causes the demodulation phase to change with variations in the beat frequency. This can be accounted for by introducing a frequency-dependent correction term. The path length mismatch Δz leads to a time delay τ. Therefore, for a change of δf in the beat frequency, the demodulation phase must be adjusted by $\delta\phi = \tau \times \delta f$.

Such a phase correction implementation will take an external value τ and multiply it by the current beat frequency (as given by the PIR). The product is then used as the correction term for the PA outputs.

6.1.2 Implementation on the ZCU111

Different clocks associated with the PL and the clocking of RF data converter IP core are configured as described in Sec. 5.4.2. The RF-ADCs are set to a sample rate of 4.096 GSPS, while the RF-DACs operate at a sampling rate of 0.512 GSPS. The chosen sampling rate is sufficient for the DACs, as the logic runs at a maximum speed of 512 MHz.

To demonstrate HS, we need two laser-locking loops to operate simultaneously. To facilitate this, two identical channels of the DHC are implemented on the RFSoC. Each channel includes one ADC to digitize the cavity reflection and one DAC to output the laser actuation signal. The reference beat is coupled to a separate ADC, which is tracked by a GHz PLL. The phasemeter acts as the readout system, and its PIR value is made available to the DHC channels at the maximum processing speed of 512 MHz.

As the RF-ADC digitizing the cavity reflection produces eight data samples for each 512 MHz clock cycle, we need eight consecutive phase values in the NCO instance of the DHC as well. Only then the multi-demodulation scheme can be utilized. Hence, the PIR value is split and incremented successively to generate an effectively continuous phase value. This multi-phase accumulation stage is identical to the process in the GHz PLL. In theory, one could directly take the multi-phase accumulator values from the GHz PLL as well.

Demodulation phase correction

Instead of using the generated phase values as inputs to the LUTs, we need to perform the demodulation phase correction. The correction terms τ and φ can be set externally from the PS side by writing into dedicated registers. The PL is programmed to look at the corresponding registers and use those values for logical operations. Although a 32-bit register is available for τ, its lower 27 bits are multiplied by the 16-bit PIR value using DSP48 Macro IP core[1]. Since this multiplication occurs in the 512 MHz clock domain, such an implementation was necessary to meet timing requirements. From the 43-bit output of this IP core, we select a 16-bit value for the frequency-dependent demodulation phase correction[2]. Similarly, a 16-bit register is available for φ, and this value is directly added to the multi-phase accumulator outputs alongside the previous frequency-dependent correction value. The corrected 16-bit long phase values are then passed to the LUTs, which generate 16-bit sinusoids for demodulation.

Heterodyne demodulation

The heterodyne demodulation scheme is very similar to the IQ-demodulation in the GHz PLL. The top 12 bits of each ADC output are multiplied by the corresponding 16-bit LUT output. The multi-mixer output, after being filtered by a rolling average of 16 samples, is 31 bits long. It is then dithered and truncated to 20 bits, which represents the error signal for the laser-locking control loop. We have implemented an error signal boosting stage where the error signal can be amplified by multiplying it with a desired external value. The boosted error signal then passes through a PI controller, again implemented similarly to the servo in the GHz PLL, to generate a 20-bit actuation signal. The controller gains can be set externally using dedicated 16-bit registers.

The configuration of the RF-DAC requires two data samples per AXI4-Stream clock cycle. Thus, the top 16 bits of the actuation signal are concatenated with themselves in the current implementation and passed to the DACs after a clock domain transfer from 512 MHz to 256 MHz. There is also a provision for directly adding an externally tunable offset to this actuation signal, if necessary.

Along with reading out the GHz Phasemeter parameters, the boosted error signal and the actuation signal are also read out. The decimation of these signals is handled as discussed in Sec. 5.6.1.

[1] The maximum input signal length it can handle is 27 bits.

[2] We chose to start from bit 16 to bit 31, allowing the PIR value to be adjusted as needed.

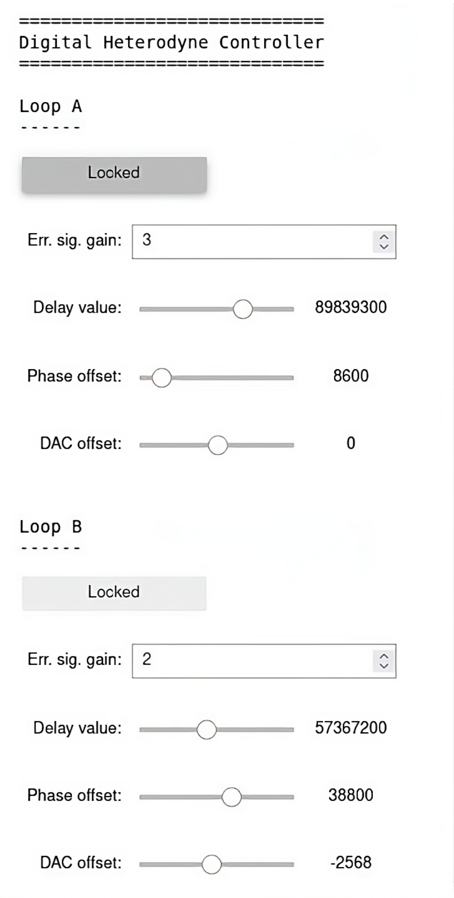

Figure 6.2: Screenshot of the simple graphical user interface (GUI) for the DHC. It is based on the well-known 'Jupyter Widgets' Python package. Sliders are provided for tuning the DHC parameters. The demodulated signal can be digitally amplified by entering an integer-valued gain in the corresponding field. The values displayed on the GUI are the integer values of the corresponding registers, and scaling to the physical units is not performed. The functionalities of the GUI can similarly be extended to include the tuning of servo parameters.

6.1.3 A simple graphical user interface

From the implementation of the DHC, it is evident that the smooth functioning of the laser-locking control loops require several parameters to be set externally. Since the demodulation phase correction terms need to be tuned for each experimental trial, a fast and user-friendly method of writing the register values is appreciated. For this purpose, a simple graphical user interface (GUI) based on the Python package 'Jupyter Widgets' [133] was made available. Fig. 6.2 shows a screenshot of this interface.

It is therefore easy to tune the DHC parameters in real-time while conducting laser-locking experiments. Such straightforward and fast tuning of the demodulation phase correction is essential for our experimental trials. One can adjust the delay value (τ) and the phase offset (φ) until a proper demodulated signal is achieved. Once the lasers are locked, the implemented toggle button for that loop can be activated. Subsequently, tuning the slider will not update the corresponding registers on the RFSoC. This feature prevents accidental or unintentional adjustments to the sliders when the lasers

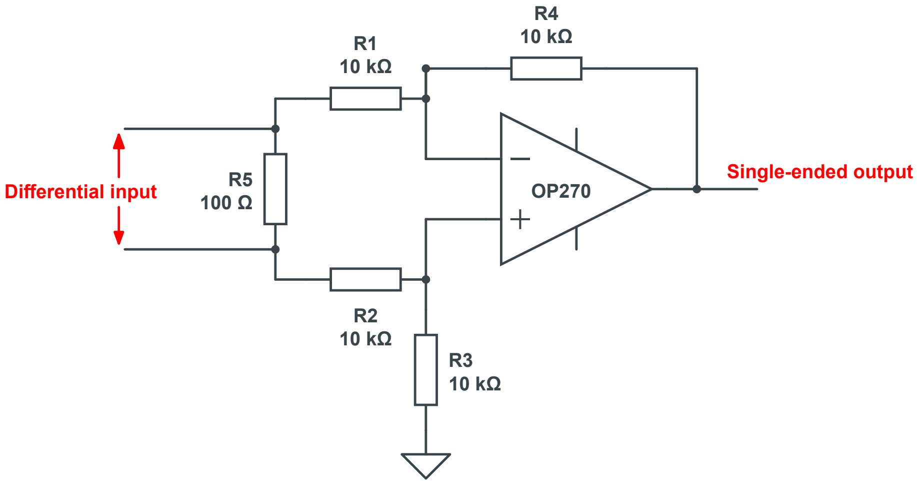

Figure 6.3: Circuit diagram of the differential amplifier acting as a balun. The gain of the differential amplifier is kept at unity, as the amplification can be introduced digitally in the FPGA. If analog gain is desired, the feedback resistor R4 in this circuit must be modified. The parallel resistor R5 at the input is chosen to be 100 Ω to match the impedance with the on-chip calibrated 100 Ω termination [135].

are locked or measurements are being taken. For the GUI to be accessible, the Jupyter server must be hosted on the PS of the RFSoC, and we need to connect to it from an external PC.

6.2 Analog front-end electronics for the XM500

All the ADCs chosen for the implementation of the DHC are from Tile 226 and Tile 227, which have differential input ports on the XM500 plug-in card (refer to Sec. 4.3.1). However, the input signals to the board come from the corresponding photodetectors, which are single-ended in nature. Therefore, these single-ended signals must be converted into differential signals before being fed to the XM500 analog interface. For this purpose, we have selected the TRF1208 evaluation module from Texas Instruments [134]. The TRF1208 is a high-performance RF amplifier suitable for AC-coupled applications requiring single-ended to differential conversion when driving an ADC.

Similarly, the laser actuation signals are also expected to be single-ended. However, the chosen DACs from Tile 228 have differential output ports on the XM500. A circuit called balun can be utilized to convert this differential (balanced) signal into a single-ended (unbalanced) signal. Since the output of the DAC is expected to be a DC signal, a general low-noise operational amplifier will suffice. Such a balun was prepared in-house, and Fig. 6.3 shows the associated circuit diagram.

6.3 Initial validations

Before using the FPGA-based DHC in heterodyne laser-locking experiments, the functionality of the DHC was validated with analog signals from external signal generators. By introducing a frequency difference between the two signals, proper demodulation was confirmed first. Subsequently, both signals used for heterodyne demodulation were kept the same, and the functionality of the frequency-independent demodulation phase correction was tested. A full scan of this external phase offset was conducted by changing the corresponding register value every 0.1 s. It was ensured that the error signal covered a full cycle during this scan, demonstrating that the frequency-independent external phase offset tuning works as expected.

To investigate the frequency-dependent demodulation phase correction, an intentional path length mismatch was introduced by using a longer electrical cable in front of one of the ADCs. The external signal generator was then frequency modulated, and attempts were made to keep the error signal close to zero. It was found out that when the input frequency to the GHz Phasemeter was modulated with a frequency deviation of several megahertz, the tracked signal amplitude (given by the I value of the PLL) also exhibited a similar modulation. Small amplitude variations are acceptable under such circumstances, but the effect we observed was prominent. The origin of this behavior is likely related to the RF-ADC itself, or, more specifically, to the analog front-end of the ADC.

While this behavior was observed in the GHz Phasemeter channel, it is expected to persist in the ADC that digitizes the cavity reflection. From the standpoint of the GHz Phasemeter, the signal of interest is the frequency, so the amplitude variation can be neglected. However, for the heterodyne demodulation in the laser-locking loop, such additional effects are not acceptable, as they may influence performance. This led us to properly measure the gain of the RF-ADC transfer function and rectify the issues.

6.4 Gain of the RF-ADC transfer function

The most straightforward way to measure the gain of an RF-ADC in the ZCU111 is to use the GHz Phasemeter. We lock the PLL to the input tone and then introduce a frequency sweep to the input. The tracked I value indicates how the gain of the transfer function varies with frequency. However, the external signal generator used implements the RF sweep as multiple frequency steps, each with a specified dwell time[3]. Such an input signal consisting of steps is not ideal for tracking with an ADPLL. Therefore, a reliable gain measurement of the RF-ADC was not possible with this approach.

The new approach involved running multiple phasemeter measurements, each for a short duration at a different input frequency. By using the mean value of the tracked I value from each measurement, we obtained the gain of the ADC transfer function. The signal generator frequency was tuned from 10 MHz to 2 GHz with a total of 200 logarithmic steps. The output of the signal generator was coupled to the ADC (ADC226_T2_Ch0) via the TRF1208EVM as an RF-balun. Each frequency value was tracked by the PLL for 5 s. Fig. 6.4 shows the measured gain of the ADC transfer function using this approach.

The steel-blue trace (labeled as ‘Before modifying TRF1208EVM’) of Fig. 6.4 shows the initial measurement when the purchased analog front-end (TRF1208EVM) was used as it is. The large ripples in the gigahertz regime were unexpected and were found to be consistent even when the measurement was repeated with additional points at those frequencies.

When different circuit elements are not impedance-matched, it can lead to reflections of the RF signal, resulting in this type of behavior. This phenomenon is analogous to parasitic reflections causing low-frequency oscillations in an interferometer setup. The RF-balun has a 0 Ω output resistor at its differential outputs, whereas the RF-ADC features a dedicated high-speed, high-performance, differential input buffer with an on-chip calibrated 100 Ω termination [135], leading to impedance mismatch. Such impedance mismatches should be avoided, especially when handling RF signals.

[3] Dwell time is the period during which the output frequency remains constant before jumping to the next frequency.

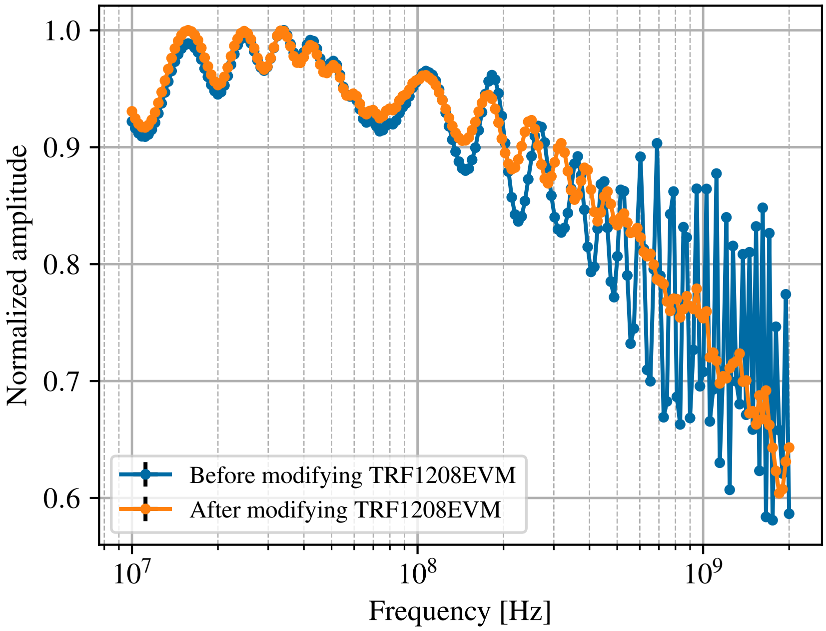

Figure 6.4: Gain measurement of the RF-ADC transfer function using GHz Phasemeter readout. Multiple phasemeter measurements were conducted for a short duration at different input frequencies. The tracked amplitude was used to generate this plot. The calculated errors, based on the standard deviation of the measured data, were very small and thus are not visible in the plot. It can be observed that even after modifying the analog front-end, small ripples are still present.

We swapped the output resistors of TRF1208EVM from $0\,\Omega$ to $49.9\,\Omega$,[4] and repeated the measurement. The dark-orange trace (labeled as 'After modifying TRF1208EVM') in Fig. 6.4 now shows significantly fewer ripples at the high-frequency end, indicating improved impedance-matching. The remaining ripples on top of the general drop-off at higher frequencies suggest that RF back-reflections due to impedance-mismatch are still present to some extent. This residual mismatch will limit our ability to match the demodulation phase and to maintain constant controller gains for large frequency changes. In theory, this transfer function measurement can be utilized to correct the DHC servo gains appropriately.

With both the digital readout and control systems in place, the next step is to demonstrate the HS for displacement sensing. One advantage of the present interferometric scheme is that it integrates both readout and control architectures. Accordingly, our digital implementation of both the readout and control systems on the same evaluation board makes the overall experimental setup compact. Fig. 6.5 shows a photograph of the evaluation board along with its analog front-end electronics.

We know that HS-based displacement sensing relies on the frequency stability of the other laser used to generate the beat note. Therefore, it is important to explore laser frequency references that can help maintain the stability of this laser's frequency. We will discuss such frequency references in the following part of the thesis and then proceed to the experimental validations of displacement sensing.

[4] It was the closest value to $50\,\Omega$ among commercially available compatible resistors for the evaluation module.

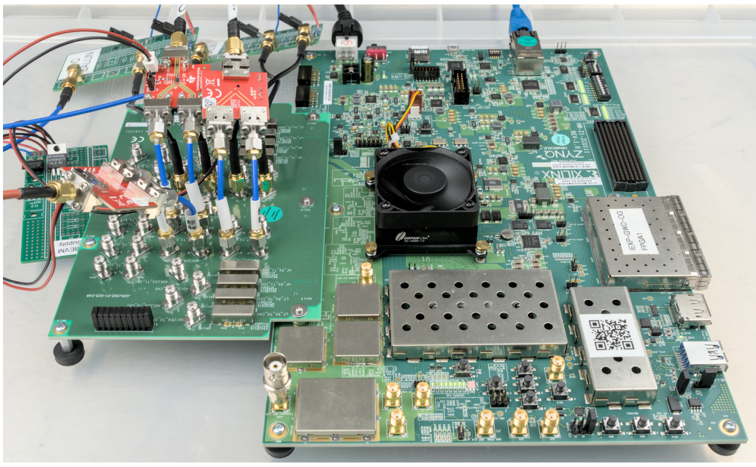

Figure 6.5: Photograph of the ZCU111 evaluation board with all the associated analog front-end electronics. (Photo credit: Dr. Christian Darsow-Fromm.)

Part III

Laser frequency references

CHAPTER 7

Probing athermal etalon as frequency reference

In this chapter, we study the possibility of using a special chalcogenide glass as a potential laser frequency reference. Before that, it is important to recall the requirements for frequency references and the conventional solutions in this field.

The displacement sensing scheme presented in this thesis involves high-precision relative displacement measurement. As discussed in Sec. 3.2.3, the precision and accuracy of the dynamic cavity displacement readout are limited by the stability of the reference to which the secondary laser is locked. This leads us to investigate possible frequency references and explore novel alternatives.

7.1 Conventional solutions for frequency references

Optical frequency references based on Doppler-free spectroscopy of molecular iodine are one of the well-proven technologies [136]. The other widely investigated and employed method in laboratories around the world is the use of stable optical cavities. For instance, it is currently assumed that the LISA lasers must have a pre-stabilized frequency stability of $30\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ over LISA's most sensitive frequency band. This is planned to be accomplished by employing high-finesse optical cavities based on Ultra-Low Expansion (ULE) glass spacers [137, 138].

Studies have been carried out on the geometry of such cavities and the limiting noise sources [139]. Visible lasers with sub-hertz line-widths have been reported for use as local oscillators in optical frequency standards [140]. For rigid cavities, thermal fluctuations become an inevitable noise source, and at room temperature, there is a limiting thermal-noise wall around $0.1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 1 Hz [139]. This value coincides with the frequency stability of a rigid cavity used for the Virgo gravitational-wave detector [141].

One of the main aspects to keep in mind when building rigid optical cavities is their sensitivity to ambient temperature changes. If changes in temperature influence the cavity length, the stability we expect from such a reference is compromised. The system's ability to maintain performance while undergoing temperature changes is an important required property. Therefore, the construction of rigid cavities is usually based on glass ceramics with very low coefficient of thermal expansion (CTE). Zerodur from Schott [142], ULE glass from Corning [89], and Clearceram-Z from Ohara [143] are some available options. To further improve the stability of these cavities, it is common to place them inside multiple layers of thermal shielding [136]. To isolate them from external influences, such cavities are also placed in separate vacuum chambers. Such optical reference cavity systems are commercially available as well, where the cavity is mounted in a sealed vacuum housing, engineered for exceptional temperature stability [144].

The disadvantage of such a setup for a frequency reference is that it is neither compact nor easily portable. Thus, when the HS is chosen for compact displacement sensing applications, using such frequency references is not a convenient solution. There have also been attempts to realize miniaturized, compact references. For example, an ultra-high Q-factor micro-rod-based reference in the ambient environment achieved a frequency noise of about $10\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 10 Hz, essentially limited by the predicted fundamental thermal noise of the system [145].

In this chapter, instead of the standard glass ceramics with near-zero CTE as listed above, special chalcogenide glass samples are probed as compact, in-air frequency references. Their 'athermal' nature is intriguing for characterizing them as potential frequency reference candidates.

7.2 Athermal glass

At first, we will try to understand what makes these chalcogenide glass samples so special and promising as frequency references. Our understanding of these special chalcogenide glass samples is based on the work by Davis et al. [146]. In this section, the results from that work that are interesting and important for our purpose are presented.

Materials with relevant optical figure-of-merit (FOM) that are temperature independent are classified as athermal. Certain glasses in the Arsenic-Sulfur (As-S) and Arsenic-Selenium (As-Se) binary systems, as well as in the Arsenic-Sulfur-Selenium (As-S-Se) ternary combination, satisfy the optical and physical properties necessary for an athermal solid etalon. It was found that when the temperature was changed in the 25–35 °C range, the interference fringes at a wavelength of about 1.5 µm exhibited very little movement in the resonant frequency of these etalon samples. This makes these infrared-transmitting, inorganic chalcogenide glass compositions perfectly suitable as frequency references.

For a solid etalon of length L, the fractional change in its optical path length with respect to temperature T is given by

$$\frac{1}{nL}\frac{d(nL)}{dT} = \alpha + \frac{1}{n}\frac{dn}{dT}, \tag{7.1}$$

where n is the refractive index of the etalon material and α is the linear CTE. The term dn/dT refers to the absolute change in the refractive index with respect to temperature. If such an etalon is used as a frequency reference, we anticipate that this fractional change in the optical path length will be close to zero over a significant temperature range. This fractional change is what is defined as the FOM of the etalon[1].

The athermal glass etalons have higher linear CTE compared to other glass ceramics such as ULE. The linear CTE of athermal glasses lies in the range of 35–40 ppm/°C. At the same time, their unique 'thermo-optic' property gives dn/dT values of −80 to −100 ppm/°C. Additionally, the refractive index of athermal samples is within the range 2.2 to 2.6. By putting all these values together, we can see from Eq. 7.1 that the FOM of athermal etalons falls very close to zero. For a general understanding, it

[1] There are other FOM values for different optical configurations. We focus on the one corresponding to the optical lengths within an etalon, which is relevant for laser frequency stabilization.

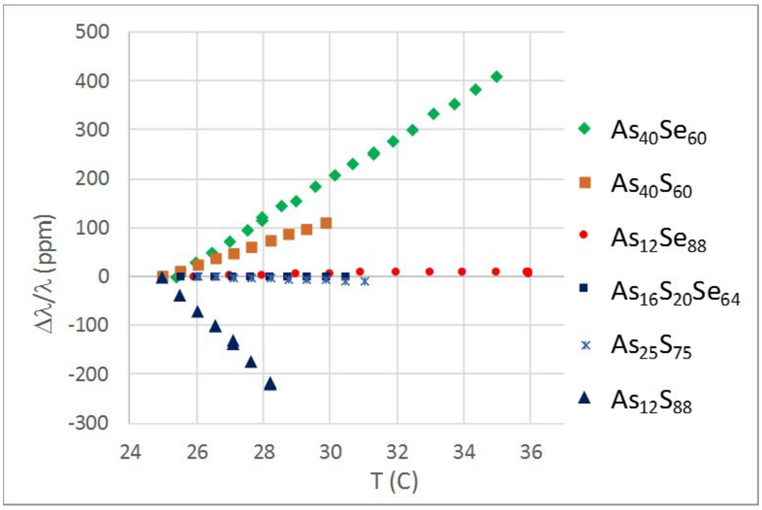

Figure 7.1: FOM characterization of different athermal etalon samples. Based on the chemical composition of the glass, the fractional change in transmission wavelength varies as the temperature changes. The ternary As-S-Se composition exhibits a flat response in the considered temperature range. (Taken from [146].)

can be stated that when the temperature of the etalon changes, the effect due to CTE is compensated by the thermo-optic coefficient, thus keeping the optical path length inside the etalon constant.

7.2.1 Characterizing the FOM

Many chalcogenide samples with different combinations of As, S, and Se were studied by Davis et al. To characterize the FOM, they used the interference fringe measurement method. The wavelength corresponding to a transmission spectrum of the etalon varies with temperature as

$$\frac{1}{\lambda_0}\frac{d\lambda}{dT} = \frac{1}{nL}\frac{d(nL)}{dT} = \text{FOM}, \tag{7.2}$$

where λ_0 is the initial wavelength corresponding to a transmission peak, which was kept at around 1.5 µm. Thus, monitoring the change in peak transmission wavelength for any resonance directly enables the characterization of the FOM. Detailed discussions on the experimental setup and results are available in [146].

Fig. 7.1 shows how the FOM changes as a function of temperature for a variety of chalcogenide glasses. The slope of a trace on this graph provides the FOM of the corresponding chalcogenide sample. The binary As-S or As-Se systems result in either positive or negative FOM, while certain glasses in the ternary As-S-Se system exhibit near-zero FOM. For example, the ternary system with 16 mol% As, 20 mol% S, and 64 mol% Se, represented as $As_{16}S_{20}Se_{64}$, has a FOM of −0.1 ppm/K. This was the first report of inorganic glasses having near-zero FOM. The ratio of As, S, and Se atoms in the substrate defines the FOM of the etalon, and fine-tuning the chemical composition is an option.

When such athermal etalons are used as frequency references, the achievable laser frequency stability is given by

$$\frac{\tilde{S}_f}{f_0} = \frac{\tilde{S}_L}{L} = \text{FOM}_{\text{etalon}} \times \tilde{S}_T. \tag{7.3}$$

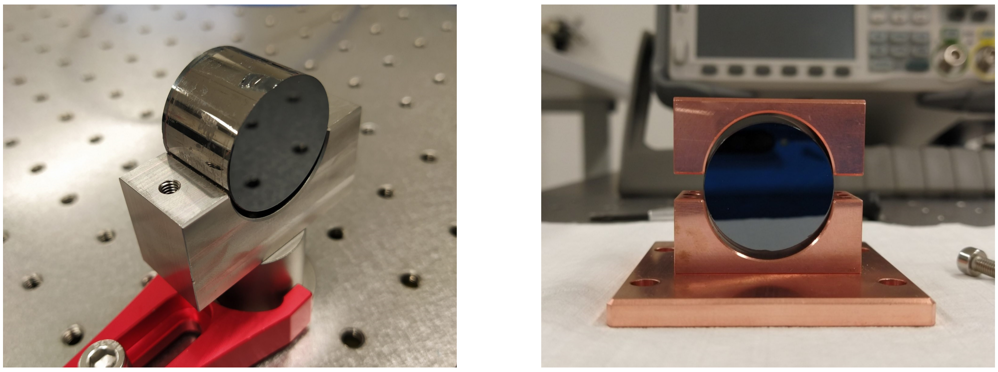

Figure 7.2: Photographs of athermal etalons at the time of purchase.

Therefore, to decouple the influence of temperature noise $\tilde{S}_T$, we need $\mathrm{FOM}_{\mathrm{etalon}}$ to be as small as possible—in an ideal case, zero. If we refer back to Fig. 3.6, we realize that such a temperature coupling noise was limiting our displacement readout sensitivity (refer to Sec. 3.3.3).

The athermal etalon can compensate for the effects of various noises on frequency stability by maintaining a constant optical path length. This capability may eliminate the need for vacuum shielding. With a refractive index of around 2.5, a 2 cm long etalon will result in an optical path length of about 5 cm, making athermal glass a perfect candidate for compact and portable frequency references.

7.2.2 Chosen samples for the initial testing

After recognizing the potential of athermal etalons as candidates for laser frequency stabilization, three etalons were purchased from Schott NA, Inc. We selected two of them to be of the same chemical composition (IRA-14) to enable the full characterization of the etalons under the HS scheme. The third etalon (IRA-15) had a slightly different composition compared to the former. IRA-14 has the composition $As_{21}S_{39}Se_{40}$, whereas IRA-15 is the previously mentioned $As_{16}S_{20}Se_{64}$.

All the etalons were 20 mm in length and a little more than an inch in diameter. Both facets of the etalons were polished to a flatness of $\lambda/10$ at 633 nm, with a surface quality of scratch–dig 10–5 over the clear aperture. Fig. 7.2 shows the photographs of the two etalons at the time of purchase.

7.3 Athermal etalon responses

At the interface of two media with different refractive indices, n_1 and n_2, the reflectivity is given by

$$R = \left| \frac{n_1 - n_2}{n_1 + n_2} \right|^2 . \tag{7.4}$$

The refractive indices of both IRA-14 and IRA-15 are close to 2.5. Because of this refractive index difference compared to air, an optical beam propagating from air to the athermal etalon will experience a reflectivity of about 18%.

Consequently, the samples we have for the initial testing are treated as impedance-matched etalons, with input and end facet reflectivities of about 18%. This results in an etalon with very low finesse,

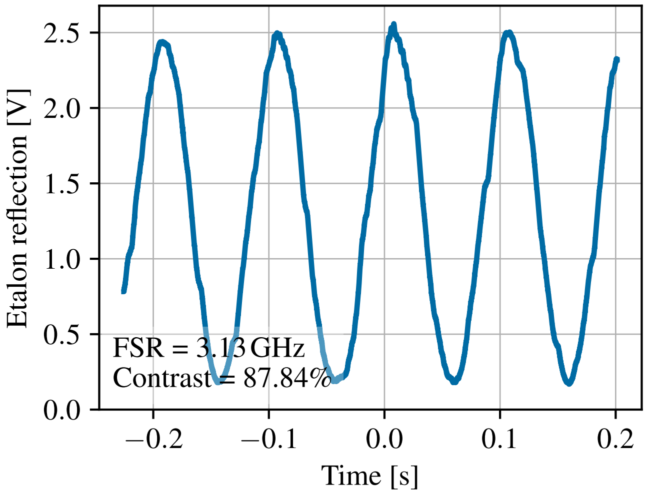

Figure 7.3: Athermal etalon response for the scanning of the input laser wavelength. The CTL wavelength was scanned at a speed of 0.25 nm/s using the wide-scan feature. The achieved contrast is not the maximum possible; as this was an initial test of the sample, alignment was stopped once a satisfactory contrast was achieved.

whose reflection appears almost sinusoidal instead of displaying sharp resonance peaks when the input laser frequency (or the microscopic cavity length) is scanned.

Fig. 7.3 shows the response of one of the etalons. The reflection from the etalon was separated from the input beam using a fiber circulator and detected using an adjustable gain photodetector. The recorded data was then fitted with a sine function. The period of the fit yields the FSR of the etalon in units of nanometers when scaled using the scan speed. This value can then be converted to frequency units and was found to be approximately 3.13 GHz. The theoretical FSR is 3 GHz if the refractive index is taken as 2.5.[2]

Such an etalon setup was simulated using Finesse3 [77], and it was found that the observed etalon response agrees with the simulation results. The achievable contrast and the FSR value falling within the expected range were motivating factors for using the etalon in laser frequency stabilization trials.

7.4 Laser frequency stabilization trials using athermal etalons

When testing athermal etalons for stabilizing laser frequency noise, RFSoC-based readout and control systems had not yet been developed. The available readout options were a frequency counter and the Moku:Lab phasemeter. Subsequent attempts were made to investigate whether such a low-finesse etalon can be used as a frequency reference.

7.4.1 HS using athermal etalons

The beat note generation and cavity reflection separation were carried out using a fiber splitter/coupler and a fiber circulator, respectively. The heterodyne demodulation was managed using analog electronic

[2] In reality, the samples have slightly different refractive indices, and the value is not exactly equal to 2.5.

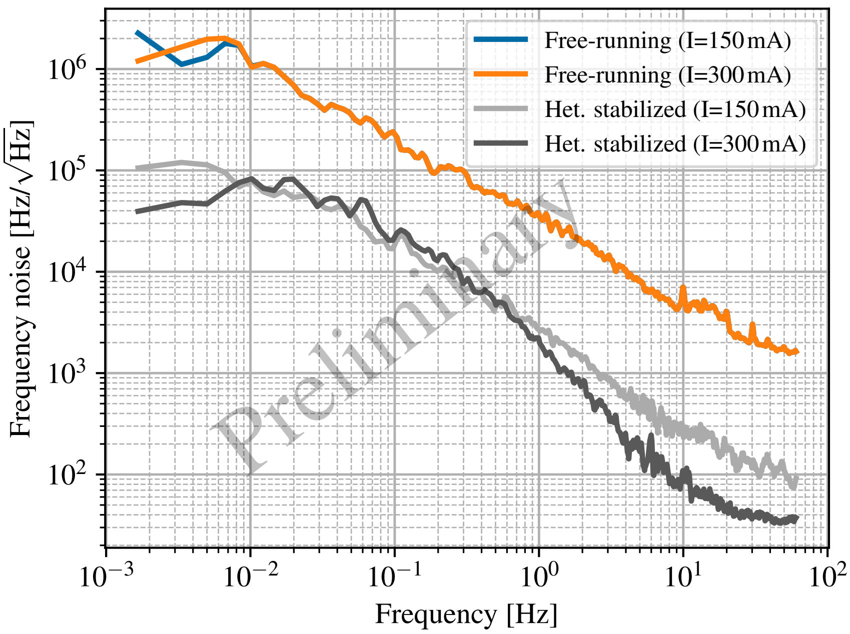

Figure 7.4: Preliminary results from the analog HS trials using athermal etalons. To verify the reproducibility of the results from the chosen readout and locking schemes, the experiment was conducted at two different set current values of the CTLs. The spectra from two different measurements largely agree with each other. All spectra were calculated after appropriately detrending the time series, which aids in isolating the short-term fluctuations in the data.

components, as described in Sec. 3.5. Previously, tracking the beat note with the Moku:Lab phasemeter had not been successful. However, when the lasers were locked to the athermal etalons, we were able to track the beat note using the Moku:Lab phasemeter. To verify the tracking, a simultaneous measurement with both the frequency counter and the phasemeter was conducted.

Fig. 7.4 shows the results of laser frequency stabilization attempts using athermal etalons. These results are considered preliminary because a rigorous investigation of the laser-locking loop was not conducted. The shoulder-like feature (the so-called 'scattered-light shoulder') in the spectrum when the lasers were stabilized may be due to parasitic reflections in the optical fiber components used. When the signal strengths are small and the finesse value is minimal, such parasitic effects couple more strongly compared to a higher finesse cavity. Nevertheless, these preliminary results appear promising for using athermal etalons as frequency references.

A further attempt of transition to free-space optical setup, thereby eliminating the fiber circulators, was carried out. However, the achieved stabilizations were not as effective as before. The generation of the error signal itself proved to be quite challenging, which became the starting point of our investigation into the analog demodulation process and related aspects. At that time, the role of the demodulation phase and the importance of length matching (refer to Sec. 3.5.3) were not properly understood.

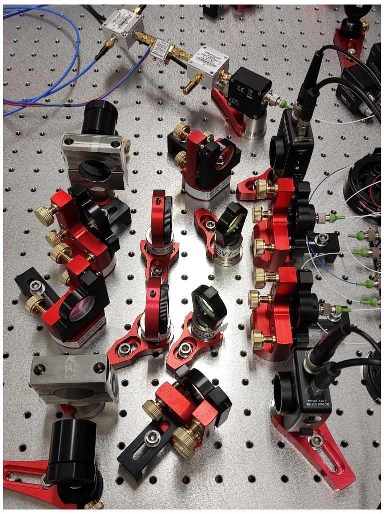

(a)

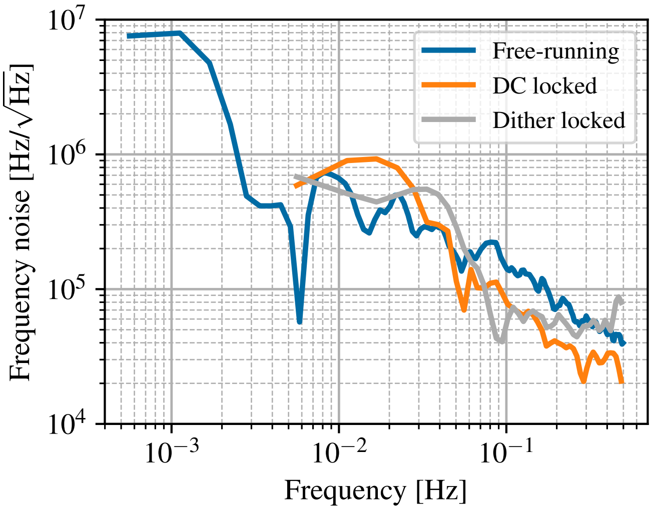

(b)

Figure 7.5: (a) Photograph of the experimental setup used for DC and dither locking of the CTLs to two athermal etalons. The reflected beam is detected in free-space propagation, and the two athermal etalons were held in custom-made aluminum holders. (b) Frequency noise of a laser with and without stabilization. The heterodyne beat note measured using the frequency counter is detrended, the spectrum is calculated, and then divided by $\sqrt{2}$ to obtain the performance of a single laser.

7.4.2 Side-of-fringe and dither locking attempts

To characterize the athermal glass samples as frequency references, we continued laser frequency stabilization attempts using other resonator locking methods discussed in Sec. 2.5.1.

Each laser was initially side-of-fringe locked using the reflected signal from the corresponding athermal etalon. A photograph of the experimental setup is shown in Fig. 7.5(a). The beat note was generated using a fiber splitter/coupler. By employing a combination of a PBS and a QWP, the reflected beam from the etalon was separated from the input laser beam. In such setups, one high-reflective mirror in front of the etalon and another in front of the photodetector helps achieve a decent contrast and prevents clipping. Since the active area of the photodetector used was large enough, no focusing lens was necessary. The heterodyne frequency between the lasers was measured using the frequency counter.

Fig. 7.5(b) shows the stabilization achieved when lasers were locked to athermal etalons. In the case of side-of-fringe locking, a suppression factor of about 2 to 3 times in the laser frequency noise is observed for frequencies above 50 mHz. Unfortunately, the maximum sampling frequency of the readout is preventing a proper characterization of the athermal etalons.

To determine if another locking scheme results in a very different behavior, we attempted to stabilize the lasers to corresponding etalons using dither locking technique. In this method, the reflected signal from the etalon was used as a Lock-In signal for top-of-fringe locking. The DLC pro lock module internally modulates the laser at a given frequency, and the Lock-In-out signal is generated by demodulating the Lock-In signal at the same frequency. Although locking the lasers was successful, the achieved noise suppression was not significantly better than (or much different from) before.

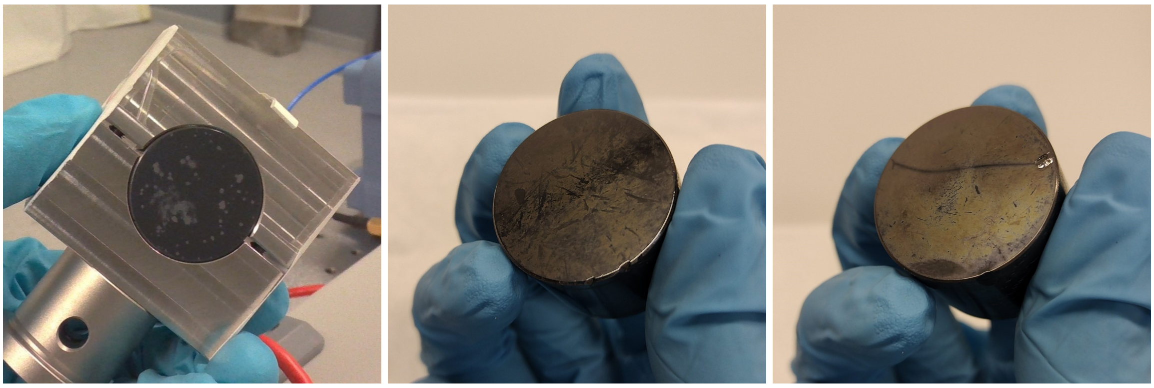

Figure 7.6: Photographs of athermal etalons after the facets degraded due to hydrophilic reactions. The first photo was taken when we first observed the stains on the etalon facet, while the other two photos were taken 8 months later.

We also attempted side-of-fringe locking of one of the lasers to an athermal etalon and using the delayed self-homodyne readout method. Between 10 Hz and 100 Hz, the frequency noise of the laser was suppressed by about one order of magnitude. However, at higher frequencies, there was no improvement in stabilization. As pointed out in Sec. 3.4.3, this method is not a reliable readout because of the high non-stationary noise of the laser, and long-term measurements are not possible. Therefore, what happened at lower frequencies remained unknown.

Another approach we took involved DC/dither locking and etalon readout. In this method, the laser beam is split into two and coupled to two different etalons. The reflection from one etalon is used for locking, while the reflection from the other is used for DC readout. Unfortunately, we later discovered that such a readout using very low finesse etalons was limited by the dark noise of the instrument used to record the photodetector output.

7.5 Degradation of etalon facets

About a month or two after we started experimenting with the athermal etalons, we began to observe a layer of haze on the facets. Subsequently, stains began to form on the facets, as shown in the first photograph of Fig. 7.6. After discussing this issue with the manufacturers, it was interpreted that this is a chemical change resulting from exposure to humidity. The manufacturer's internal solution was to perform a surface polish on the material. The recommended storage conditions were in a vacuum desiccator or a dry nitrogen storage box.

After that, the etalons were stored in a vacuum desiccator for about one and a half months. The hope was that the water molecules might diffuse out and that the etalons could be reusable in the experiments. However, the surface quality of the etalons continued to degrade, and unwanted effects such as absorption of the input laser power and birefringence were observed.

7.5.1 Birefringence effects in IRA-15

While continuing experiments using stained athermal etalons, we observed that even when only p-polarized light was coupled to the etalon, the reflected light appeared to contain both p- and

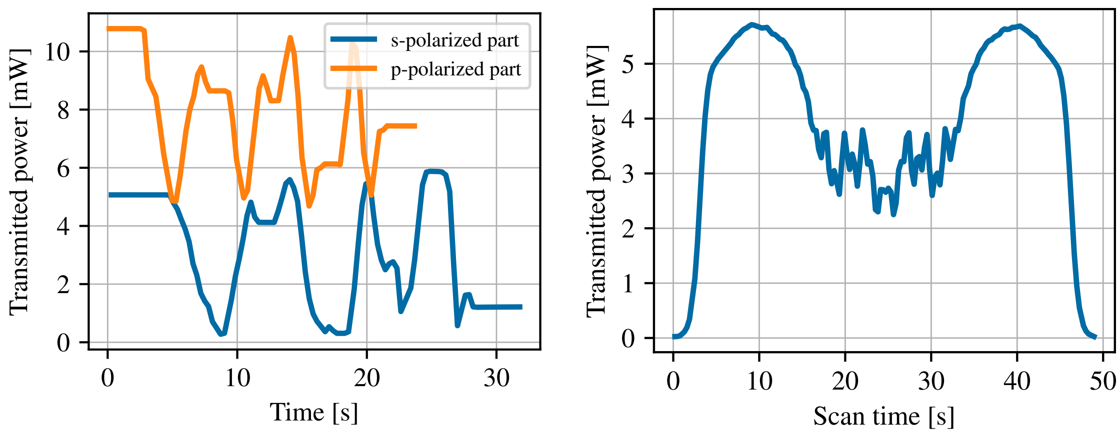

(a)

(b)

Figure 7.7: Investigation of the birefringence properties of the stained athermal etalon. (a) The s- and p-polarized components of the transmitted power from IRA-15 for a p-polarized input. The etalon was manually rotated around the laser propagation axis during the measurement, with the plot showing two independent measurements. (b) The s-polarized component of the transmitted power from IRA-15 when the p-polarized laser input undergoes a wide-scan of its wavelength. The scan starts at 1510 nm and reaches 1630 nm after about 24 s, and then returns to 1510 nm. The mirrored response confirms the reproducibility of the observed behavior.

s-polarized components. The athermal material should be isotropic; hence, such a behavior is not expected. It was confirmed that this was not purely stress-related birefringence.

For a proper investigation, the input laser was passed through a PBS, ensuring the laser impinging on the etalon is p-polarized. The transmitted light from the etalon was then passed through another PBS, and the two output ports of the PBS were measured using a power meter. The laser wavelength was maintained at 1550 nm.

During the measurement, the IRA-15 etalon was manually rotated around the laser propagation axis. Since only p-polarized light was incident on the etalon, we expect that nothing will be reflected from the PBS behind the etalon. Regardless of the etalon's rotation, the transmitted power from that PBS should remain constant. Nevertheless, what we observed with IRA-15 is shown in Fig. 7.7(a). Only a specific orientation of the etalon removes the s-polarized component from its transmittance, behaving like an anisotropic material.

The influence of laser wavelength on the birefringence properties was also tested. The IRA-15 etalon was kept in a fixed orientation while the laser wavelength was scanned from 1510 nm to 1630 nm and back to 1510 nm. The s-polarized component of the transmitted light from IRA-15 is plotted in Fig. 7.7(b). Once again, IRA-15 exhibits unexpected behavior. Interestingly, such a birefringence effect was not observed in either of the IRA-14 samples. The exact reason for this is discrepancy is still unknown but is expected to originate from differences in chemical composition.

7.6 Preparations for the athermal glass characterization

Our long-term plan is to properly characterize athermal glass samples in a HS setup. Both lasers will be locked to two different athermal samples, and then the FOM will be characterized by tuning the temperature of one sample while keeping the other at a fixed temperature. Similar characterization is planned by tuning the wavelength of the laser source, which is one of the reasons for using widely tunable lasers like CTLs in our experiments.

7.6.1 Temperature control box

To efficiently control and modulate the temperature of the etalons, two temperature control boxes were designed. The computer-aided design (CAD) model is shown in Fig. 7.8. To reduce stress on the athermal sample, it is held from the top and bottom by two separate copper pieces. To ensure good thermal conductivity between these two pieces and, in turn, a homogeneous temperature distribution throughout the athermal glass, four fine copper rods and two copper plates connect the top and bottom copper pieces. Another copper baseplate then attaches this copper assembly to the Peltier plate of the commercial Direct-to-Air thermoelectric cooler (TEC) unit (DA-045-24-02 from Laird Thermal Systems [147]).

To further improve the thermal isolation, the athermal sample along with its copper assembly is kept inside a box made of polyoxymethylene (POM). The box is designed to directly clamp onto an optical table or an optical breadboard. The front and back panels of the box have holes at a height of 5 cm to couple the laser beam onto the athermal glass and to provide access to the transmitted beam.

7.6.2 Temperature sensing and control

The commercial TC-XX-PR-59 temperature controller [148] was used to control the TEC assembly. This controller can accommodate three temperature sensors. A thermistor, TC-NTC-1, was designated as the main sensor for closed-loop feedback control, while the other sensors (TC-NTC-2) were used as witness or out-of-loop sensors. An RS-232 interface was utilized for communication and real-time control. Additionally, there is also a GUI called ‘LT Interface’ from the manufacturer that can be used to control and log data. For our convenience, and to enable control via Python scripts, the RS-232 interface was externally converted to USB using an appropriate adapter. The Python code snippet used to communicate with the controller and record temperature sensor data is provided in Appendix F. For characterization and testing, the temperature sensors were kept in contact with the Peltier plate of the TEC unit.

Heating and cooling gains of the regulator

When we aimed to use PR-59 with the TEC assembly to control the temperature at a given set point, it was important to establish appropriate heating and cooling gains for the regulator. Since these values for the DA-045-24-02 unit were not directly available, we experimentally measured the required gain values as follows.

Set point was kept at 1 K above the thermal equilibrium state, and only the proportional gain was set in the controller’s GUI. Data were recorded, and the slope of the initial rise was calculated from the time series. The same procedure was repeated by setting the set point below the thermal equilibrium state, and the slope of the initial drop was calculated. The heating and cooling gains were adjusted so that the ratio of these two slopes was close to 1. With a proportional gain of 32, our experimental setup required a cooling gain of 2.6, compared to a heating gain of 1.

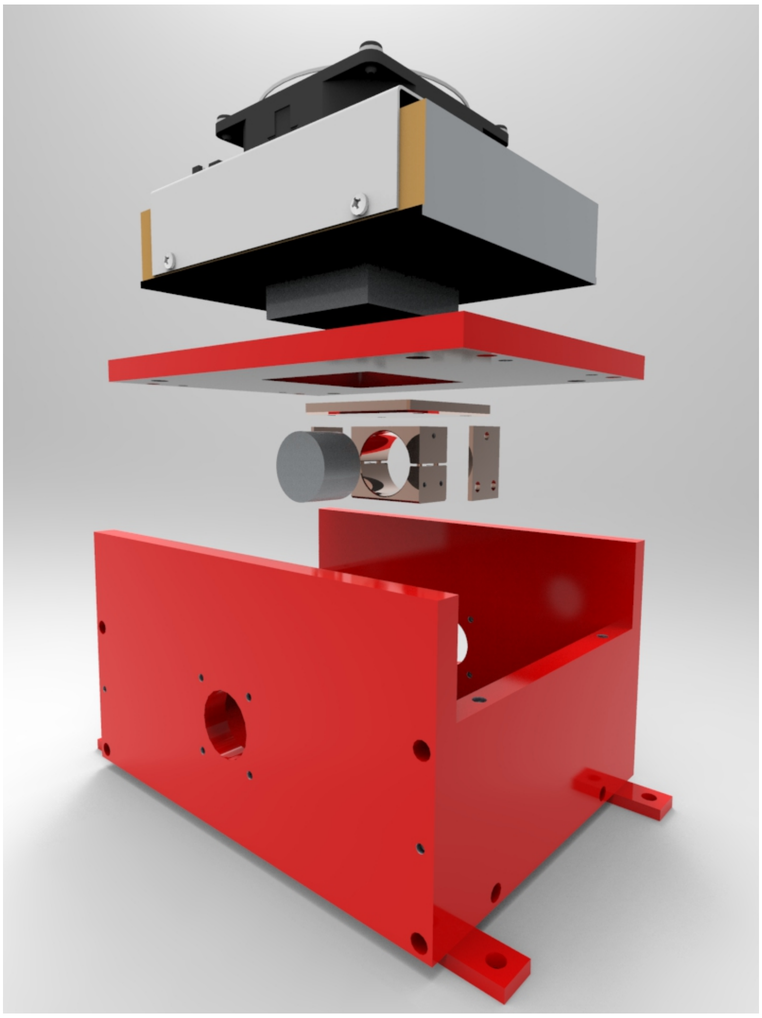

Figure 7.8: CAD model of the temperature control box for the athermal glass. The athermal glass sample is held in a copper holder that is attached to the TEC assembly. This assembly is then placed inside a box made of POM. (Created using Solid Edge software from Siemens Industry Software, Inc.)

Characterization of temperature noise

Our measurement trials revealed that the TC-NTC-2 temperature sensors have a resolution of about 0.1 K. Consequently, while using TC-NTC-1 as the in-loop sensor and TC-NTC-2 as the out-of-loop sensor, a proper characterization of temperature noise was not possible. Therefore, we utilized two PR-59 units, designating one of them as the in-loop sensor for temperature stabilization. The temperature set point was maintained at 25 °C. For control and readout, the LT Interface GUI was used[3]. The time series data and spectra are shown in Fig. 7.9. The self-noise of TC-NTC-1 was found to be about 1 mK/$\sqrt{\text{Hz}}$ in the 0.1–1 Hz range. A Wheatstone bridge-based µK-precision temperature sensor was later developed as part of a bachelor's project [149], and switching to these temperature sensors in the future is an option.

[3] PID parameters used on LT Interface: kP=33, kI=0.45, and kD=500.

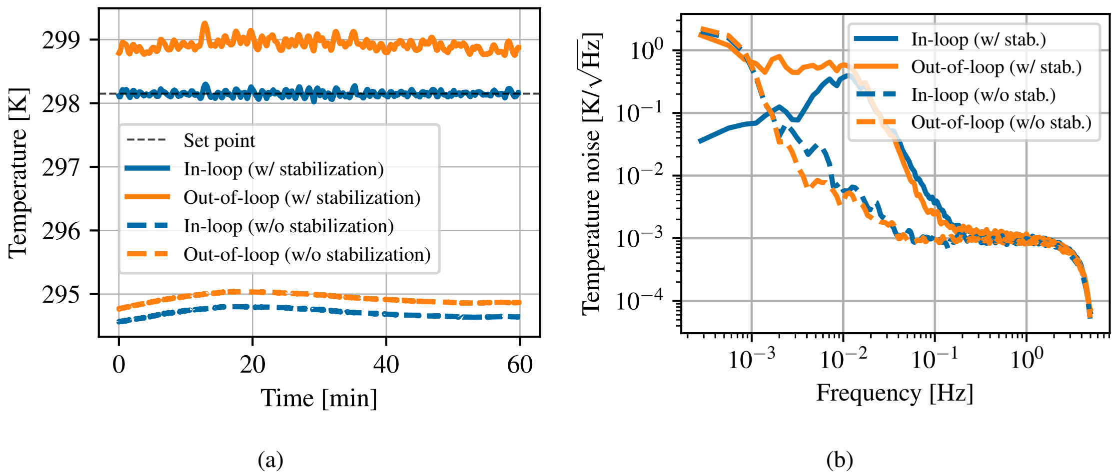

Figure 7.9: Characterization of temperature noise using TC-NTC-1 sensors. Two independent measurements were taken: one with the controller active and one without, to compare the performance. (a) Time series data. There is a slight offset between the two sensors, which is expected to be due to calibration differences. It is evident that the in-loop sensor follows the set point when stabilization is active. (b) Spectra of the corresponding measurements, calculated after detrending the average value from each time series. Now, as expected, the in-loop and out-of-loop sensors agree well with each other. The controller adds significant noise below 0.2 Hz, indicating the potential for optimizing the control loop parameters.

Transfer function measurement

For a single-input (x) single-output (y) system, the closed-loop transfer function $H(f)$ can be estimated as

$$H(f) = \frac{\tilde{S}_{yx}(f)}{\tilde{S}_{xx}(f)}, \tag{7.5}$$

where $\tilde{S}_{yx}(f)$ is the CSD of x and y, and $\tilde{S}_{xx}(f)$ is the PSD of x. This estimate assumes that the noise is uncorrelated with the system input. The open-loop transfer function $G(f)$ is then calculated as $H(f)/(1-H(f))$.

To measure the transfer function of the temperature controller system, Gaussian white noise was added on top of the set point of 20 °C. The PID gains used were the same as before. The set points were changed roughly at intervals of 30 s. The logged sensor data is plotted in Fig. 7.10(a). From the time series data, it is then straightforward to obtain the Bode plot shown in Fig. 7.10(b). We used Python package ‘SciPy’ [150] to obtain the results.

Temperature modulation

For the actual characterization of athermal glass samples, it is important to sinusoidally modulate the temperature of the sample while a laser is locked to it. This would allow us to compare the recorded beat note time series with that of the temperature modulation. Consequently, an automated modulation and readout method has been developed via serial command interface (SCI) using Python programming. The code snippet is provided in Appendix F.

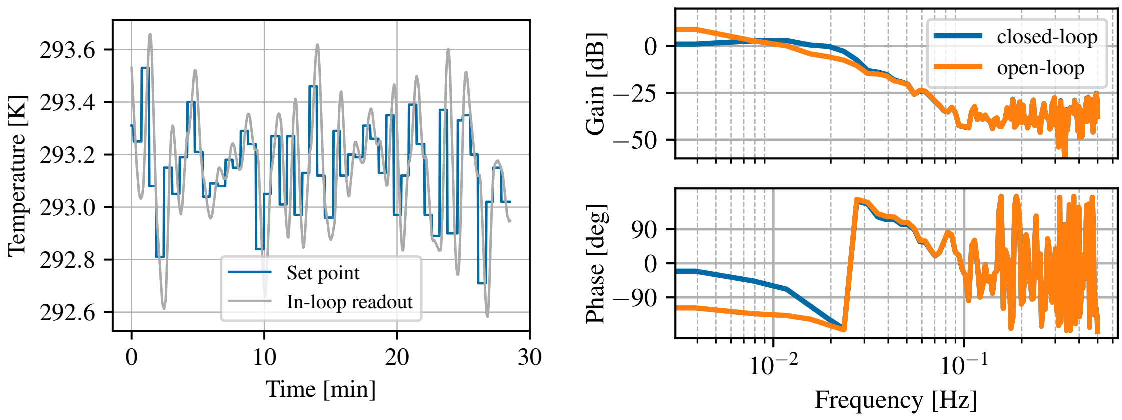

(a) (b)

Figure 7.10: Transfer function measurement of the temperature controller system. (a) The recorded time series involving the actual temperature measured by the in-loop sensor and the specified set point values. It should be noted that temperature control loops are generally very slow; hence, a 30 s interval may not be sufficient to fully settle. (b) The Bode plot containing both the open-loop and closed-loop transfer functions. The UGF is found to be approximately 10 mHz.

The performance of the temperature modulation scheme was studied for different step lengths[4]. With slow modulation, the actual temperature followed the set point values quite well. However, once the modulation frequency exceeds the 10–15 mHz range, we observed reduced amplitude and a larger delay between the actual values and set points, even though the modulation was as expected. Fig. 7.11 shows the controller behavior in two of the modulation tests conducted.

It is reassuring that control of the TEC assembly is achievable as desired. There is uniformity in the achieved temperature modulation cycles, making the system reliable when used with athermal glass samples.

7.7 Current status

The initial plan was to test the athermal etalons and, after gaining a thorough understanding of the material, convert the etalons into a stable cavity by introducing a curved surface to one of the facets. To increase the cavity finesse to a desirable value, the plan was to coat the hemispherical monolithic samples with high-reflective coating layers at 1550 nm.

However, it is now crucial to address the issue of surface degradation first. While chalcogenide glasses are suitable for molding, the machining characteristics of these brittle materials are largely unknown [151]. Choosing the molding process was not feasible because of the high cost for just three samples, as molding is typically carried out for batches of samples on an industrial scale.

Moreover, because of the presence of Arsenic in these samples, the dust produced during the machining process is carcinogenic. As a result, many industrial companies worldwide were unwilling to machine the etalons into a hemispherical cavity.

[4] The time difference between two set point values.

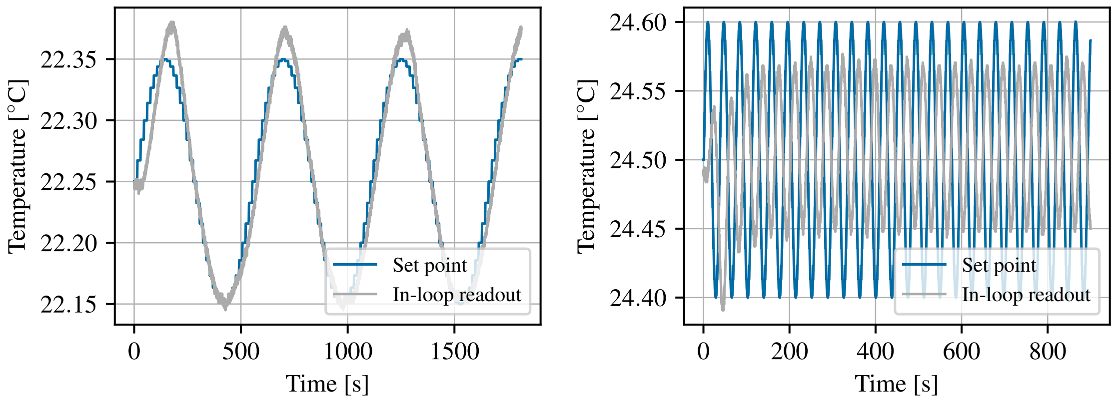

(a) Step length of 15 s

(b) Step length of 1 s

Figure 7.11: Testing temperature modulation of the TEC unit using PR-59 temperature controller. When the step length is large enough (or equivalently, with a lower modulation frequency, as shown in (a)), the controller can ensure that the modulation is roughly as expected. However, for very small step lengths or higher modulation frequencies (as shown in (b)), the actual in-loop value no longer aligns with the expected modulation. It is possible to further optimize the heating/cooling gains and PID values to achieve better performance. Both the set point and in-loop sensor values were read out from the PR-59 controller using a Python script.

After more than a year of contacting various companies and machinists, we found the possibility of ultra-precision diamond machining of chalcogenide glasses, as reported in [151]. This led us to engage with the Instrumentation Development Group of the William States Lee College of Engineering at the University of North Carolina at Charlotte (UNCC). A formal project agreement was established between the two institutions, and the samples were handed over to them in June 2023.

As of the present date[5], the etalons are undergoing single-point diamond turning to create a flat facet at one end and a 10 cm curved facet on the other. After the machining process, the hemispherical athermal samples will receive a protective coating on all surfaces to prevent from future surface degradation. To later convert them into high-finesse cavities, end surfaces will be coated with alternating tantala and silica layers, resulting in about 99% reflectivity for 1550 nm. Such a coating needs to be applied at low temperatures due to the low glass transition temperature of the chalcogenide samples[6]. Since we do not have information on the absorption of the athermal samples, the idea is to choose an over-coupled cavity configuration. Coating both end facets with the same reflectivity will turn the sample into an under-coupled cavity in the presence of absorption, making it unsuitable for use as a frequency reference, as pointed out in Sec. 2.4.

An ideal geometry for these samples would involve reducing the sample diameter from its current value of about an inch to something like 10 mm. This reduction will help decrease the thermal capacity of the monolithic sample and facilitate faster temperature control.

The prolonged search for machinists, coupled with a breakdown of the instrument at UNCC, has hindered the continuation of frequency stabilization experiments with athermal samples and their characterization at different temperatures and wavelengths. However, the initial results appear promising, and we hope this will lead to exciting outcomes in the future.

[5] January 29, 2025.

[6] Schott NA, Inc.: The glass transition temperatures are 121 °C and 112 °C for IRA-14 and IRA-15, respectively.

CHAPTER 8

Optical cavities using Ultra-Low Expansion glass spacers

As the anticipated frequency reference made of athermal glass could not be employed in the HS experimental setup, an intermediate solution had to be realized. The standard frequency references are made from materials with very low CTE, as discussed in Sec. 7.1. Following this, we decided to construct two stable cavities by gluing mirrors onto the sides of ULE glass spacers.

This optical cavity construction was carried out as a bachelor's project by Mika Brennecke. Only an overview of the process and the results are presented in this chapter. For a detailed discussion, one can refer to [149].

8.1 Initial preparations

Based on our experiences with the alignment of cavities in a flat-curved configuration with finesse values of about 312 using off-the-shelf commercial mirrors, we decided to adopt a very similar geometry. The major difference is that instead of two mirror mounts holding the cavity mirrors, a spacer made of ULE glass is used. This ensures relatively high mechanical stability, and due to its very low CTE, it better reduces temperature noise coupling to length stability.

The parallelism between the cavity mirrors plays a major role in realizing a stable cavity and achieving high fringe visibility. Tilted mirrors will allow higher-order modes to be resonant within the cavity, which is not expected [80]. The effect of tilt between cavity mirrors was simulated using Finesse3 [77], and the results are plotted in Fig. 8.1. It was found that a parallelism within 100 µrad is roughly what we should aim for.

Accordingly, the ULE rods were machined to have an outer diameter of 40 mm and a 10 mm hole drilled through their length to allow for laser beam propagation. To prevent air traps from forming in this hole during use in vacuum environments, a 2 mm venting hole was also drilled from the side. One of the spacers was machined to a length of 5 cm, while the other measured 10 cm.

To facilitate the use of these cavities inside a vacuum chamber, vacuum-compatible fiber collimators (60FC-4-A4.5S-45-Ti from Schäfter+Kirchhoff GmbH) were used to couple the laser beams. Therefore, the mode-matching scheme must be adapted according to the beam profile of these collimators. Following a similar procedure to that in Appendix C, we measured the beam profile and arrived at a suitable mode-transfer plan.

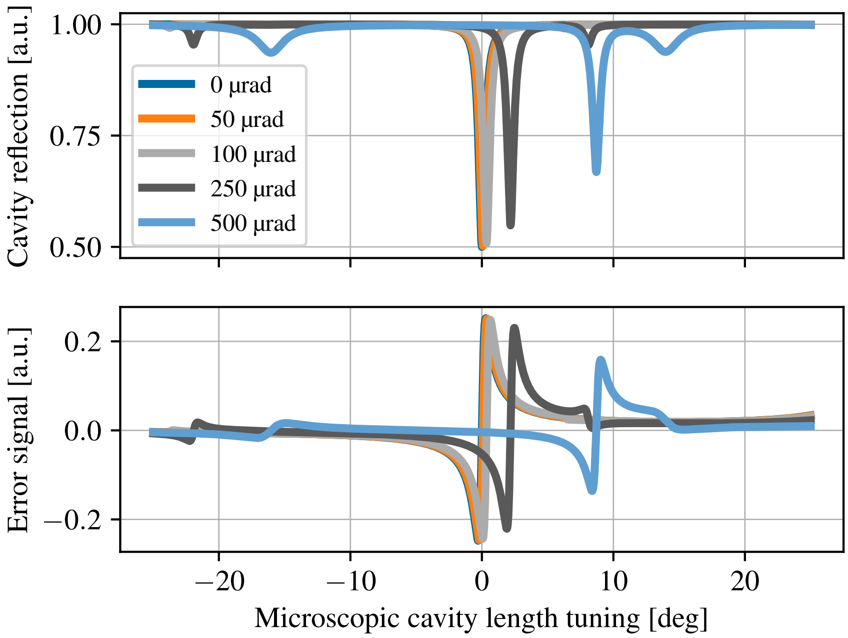

Figure 8.1: Simulation of the effect of cavity mirror parallelism on the cavity response. Standard 5 cm long cavity in a HS configuration is used in the simulation setup. By introducing a tilt to the end mirror of the cavity, the behavior of the cavity reflection and the demodulated signal are observed. From this, we can determine the tolerance for the parallelism of the ULE spacer facets. Above a tilt value of 100 µrad, the cavity response and the generated error signal begin to deviate significantly from the ideal response.

A plano-convex lens with a focal length of 250 mm, placed 14.6 cm in front of the fiber collimator, resulted in greater than 99.6% achievable coupling efficiency in the case of the 5 cm long ULE cavity. Similarly, using a plano-convex lens with a focal length of 300 mm at 14.3 cm in front of the collimator provided over 99.8% contrast for the 10 cm long ULE cavity. In both cases, the input mirror of the cavity was located 30 cm from the collimator. It was also ensured that a ±1 cm difference in the positioning of the mode-matching lens does not reduce the coupling efficiency below 99%.

8.2 Construction setup

Now the task is to retain the achievable fringe visibility even after the mirrors are glued to the spacer. The best way to achieve this is to ensure that the expected contrast is obtained before finalizing the gluing procedure. Therefore, the choice of ultra-violet (UV)-curing optical adhesives was made. These adhesives gel in seconds when exposed to UV light and cure fully in minutes, resulting in a tough and resilient bond. This allows for the position of the mirror to be adjusted after applying the glue between the spacer facet and the mirror until a good contrast is reached. Only then is UV light shined, fixing the mirror's position.

Keeping the cavity in the usual horizontal position is convenient for beam steering and mode matching. However, it is not the preferred method for adjusting the cavity mirrors after the adhesive is applied. Therefore, a cavity construction setup was designed in which the spacer is held in a vertical

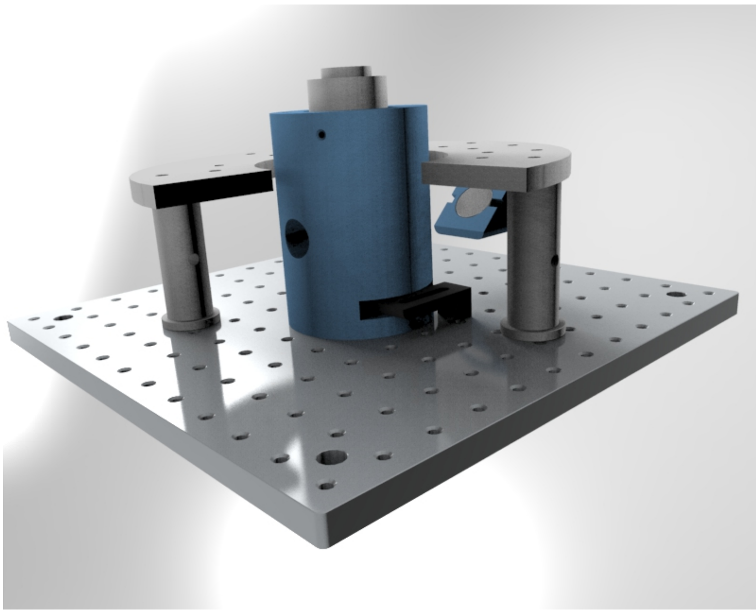

Figure 8.2: CAD model of the cavity construction setup. It features a cavity holder with the cavity spacer positioned in the middle and a detachable mirror holder from behind. The circular ring held above the optical breadboard can accommodate additional linear translation stages with ‘pointy fingers,’ allowing for fine-tuning of the mirror position. (Created using Solid Edge software from Siemens Industry Software, Inc.)

position. This arrangement makes it easier to glue the mirror from the top and adjust its position. The gravitational force acts evenly on the mirror, ensuring that the adhesive is evenly distributed and eliminating the need to press the mirror against the spacer facet. The challenge, then, is to ensure that the laser beam is also coupled to the cavity during the gluing process while retaining mode-matching.

Fig. 8.2 shows the CAD model of the cavity construction setup. The optical breadboard hosts the fiber collimator, beam-steering mirrors, and mode-matching lens; however, these components are not shown in the CAD model. The cavity holder is clamped down to the optical breadboard. As we want to scan the laser frequency during the gluing process to enable us to achieve high contrast, the laser beam will also need to travel in a vertical direction, shooting upward from the breadboard. Therefore, the cavity holder includes a separate detachable mirror holder to reflect the laser beam onto the cavity. The interior of the cavity holder is designed such that spacer is held from the sides and on the outer edge of the facet.

The laser beam that enters through the front hole of the cavity holder hits the high-reflective mirror on the detachable mirror holder at an angle of 45°. It then couples to the cavity mirror at normal incidence. This ultimately allows for real-time monitoring of the cavity response during the gluing process.

The cavity construction setup also features a ring-shaped platform held above the optical breadboard. This platform accommodates linear translation stages that can move in the x-, y-, and z- directions, respectively. Attached to each of the translation stages is a so-called ‘pointy finger’[1]. Such a setup aids in aligning the cavity mirrors. Fine-tuning of the mirror position can be achieved by rotating the micrometer screws of the linear translation stages, which, in turn, push the mirror via the attached

[1] A metallic needle-like structure with a ball-point end.

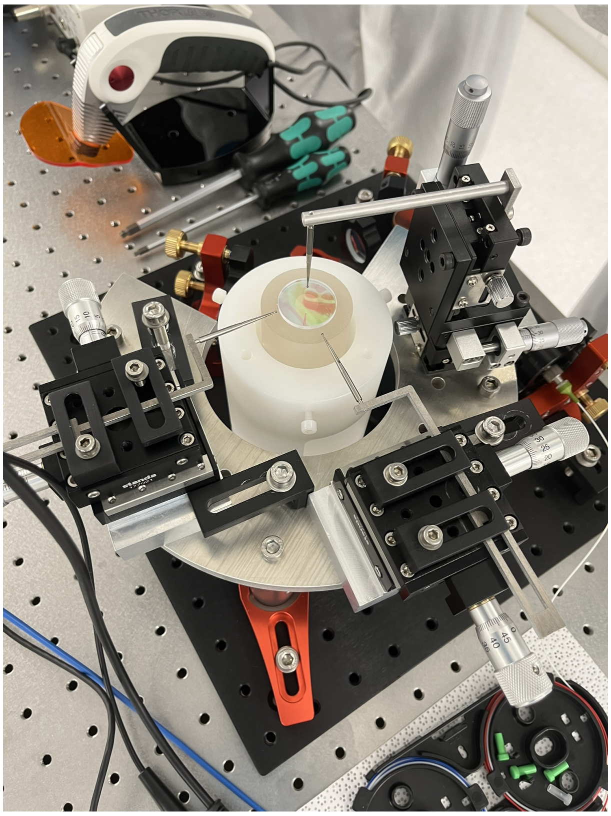

Figure 8.3: Photograph of the cavity construction setup. The fiber collimator is located in the bottom right corner of the optical breadboard. Two beam-steering mirrors and one mode-matching lens mounted on the optical breadboard can also be identified. The additional ring-shaped platform hosts three linear translation stages with pointy fingers for fine adjustments of the cavity mirror. In the top right corner of the photograph, you can see the UV-curing LED source.

pointy fingers. A photograph of the fully assembled cavity construction setup is shown in Fig. 8.3.

8.3 Construction procedure

For the gluing of the flat mirror, the cavity construction setup is not actually required. The alignment of the first mirror is not crucial because the second mirror can always be aligned parallel to the first one. Hence, the machined ULE spacer was kept upright on top of a flat surface, such as an optical table. The UV-curing adhesive was applied to the surface of the spacer facet by placing small droplets using a pipette. Care was taken to ensure that an even amount of glue is distributed and is not too close to the inner edge of the spacer facet. This precaution helps prevent the glue from diffusing to the center of the optic during the later steps.

On top of these adhesive droplets, the flat mirror was placed centrally onto the spacer facet and allowed to settle properly for a minute. Once the position and flatness of the mirror were satisfactory, the UV-curing LED system was turned on, exposing the adhesive to UV light for about 20 s.

After finishing the gluing of the flat mirror, the spacer was placed in the cavity holder of the

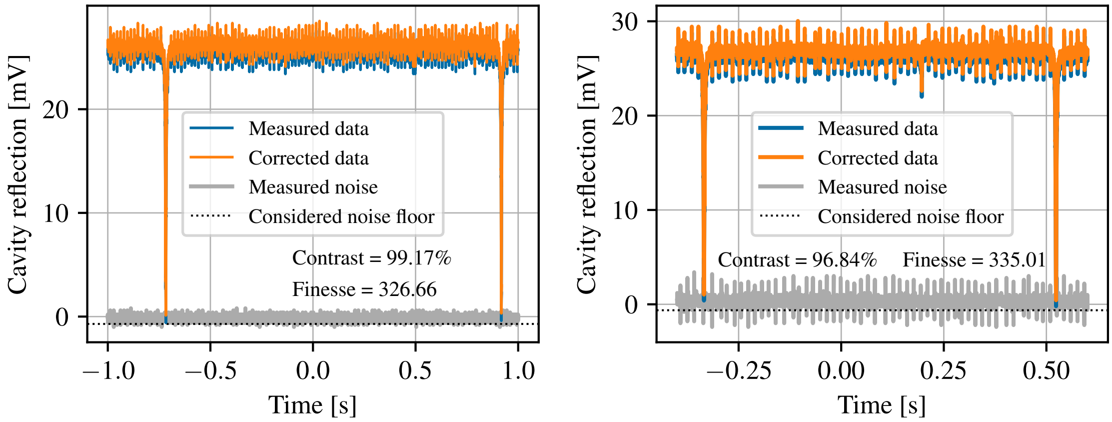

(a) 5 cm long ULE cavity

(b) 10 cm long ULE cavity

Figure 8.4: Cavity responses of the two stable optical cavities constructed using the ULE spacer. Since the measurement data recorded from the oscilloscope traces are comparatively noisy and the resonance peaks are immersed in the dark noise of the instrument, an exact calculation of the fringe visibility is not possible. The values written on the plot present the worst-case scenario.

construction setup with the flat mirror facing down. The spacer was secured by tightening the three screws located on the side of the cavity holder. The adhesive droplets were then applied again, and the curved mirror (with a 0.3 m radius of curvature) was positioned centrally. Minor adjustments to the beam-steering mirrors were made until the reflection from the cavity is coupled back to the fiber collimator. Using a 50:50 fiber splitter/coupler behind the collimator, we visualized the reflected signal on an oscilloscope.

The relatively low finesse value of the cavity made it somewhat easier to obtain the cavity resonances by manually adjusting the position of the curved mirror. There was no real need for fine adjustments using the manual translation stages. The photodetector output showed the fundamental Gaussian mode along with some higher-order modes resonating within the cavity immediately after the proper placement of the curved mirror. Later, by improving the beam steering, it was possible to achieve fringe visibility very close to 100%. Once such an alignment was confirmed, UV curing was carried out.

8.4 Evaluation of the cavities

Two stable ULE cavities were built using the construction procedure. In both cases, a very high fringe visibility was achieved, which assures the reproducibility of the construction technique followed. However, a proper evaluation of these cavities was carried out later when the cavities were held in a horizontal position in the HS experimental setup.

The evaluation of both cavities was conducted similarly to the discussions presented in Sec. 3.5.1. Fig. 8.4 shows the reflected signal from the ULE cavities when the laser frequency is scanned over an FSR. It is clear that the measured resonant points are buried in the measurement noise floor, and hence an exact value for the contrast cannot be determined. However, it is evident that a contrast close to the theoretical maximum has already been achieved.

The approach taken here to estimate the contrast in the worst-case scenario involved fitting the measured dark noise of the instrument with a Gaussian profile. From the mean value (μ) of the Gaussian fit, five times the standard deviation (5σ) was subtracted. The resultant value was treated as the noise floor, and the measured cavity reflection was corrected by subtracting this value. Referring to Fig. 8.4, we can see that the considered noise floor is nearly at the lower end of the dark noise; thus, the measured contrast value will serve as the lower bound for the actual value.

For calculating the contrast, given by Eq. 2.9, the median of the corrected data was taken as the off-resonance reflected power, while the minimum of the corrected data was used as the reflected power at resonance. The finesse of the cavities was calculated for both peaks in each case, and the average value is displayed on the plot. With the estimated finesse values falling within the expected range and the measured worst-case contrast values exceeding 95%, we can confidently assert that very good reference cavities, close to the ideal scenario, have been constructed. We expect that the actual fringe visibility values of these cavities will be close to 99%. For contrast values above 99%, the mirrors would need to be super-polished to reduce scattering losses. We have used standard commercial off-the-shelf mirrors. Achieving such high coupling efficiency is also challenging when cavities are constructed by gluing mirrors. For example, in [152], a maximum achievable contrast of 70% is documented.

The evaluation results validate our construction approach for creating well-aligned, stable reference cavities with moderate finesse values. We anticipate that this construction setup will also work for high-finesse cavity constructions. In such cases, the linear translation stages with micrometer screws will be quite helpful for alignment. These in-house constructed frequency references were subsequently used in the experiments, and their length stabilities were characterized.

Part IV

Experimental validation of displacement sensing

CHAPTER 9

Displacement sensitivity via heterodyne cavity-locking

In this chapter, we demonstrate and characterize the heterodyne cavity-locking using the fully functional high-bandwidth frequency-tracking instrument and the digital laser-locking system implemented on the RFSoC. With this, we can probe the sensitivity of our displacement sensing scheme. During previous attempts using table-top cavities (refer to Sec. 3.5), both mechanical instability in the cavities and an incorrect demodulation phase limited performance. The corresponding results are discussed in Sec. 3.5.4.

As a proof-of-principle demonstration of the concept and to test the digital readout and control systems, our first approach is to stabilize the laser frequencies to the same table-top cavities while utilizing the FPGA infrastructure for readout and control. Subsequently, improvements to the overall experimental setup were implemented.

9.1 HS using table-top in-air cavities

The optical part of the experimental setup remains the same as shown in Fig. 3.14, but the entire heterodyne demodulation is now handled using the DHC, significantly simplifying the electronics involved in the experimental setup. Fig. 9.1 illustrates the current experimental setup.

The two PID control loops of the DLC pro were again used to generate the laser frequency actuation signals. In the meantime, the table-top experimental setup was relocated, necessitating a realignment of the cavities to achieve a fringe visibility of about 95% in both cases. Using a beam profiler, we ensured that the fundamental Gaussian mode (TEM_{00}) was resonating within the cavities.

The finesse values for the cavities were also re-measured to confirm the absence of major losses. The acoustic enclosure for the laser heads was improved, and the table-top optical setup was also provided with its own enclosure. Optical powers coupling to the cavities and reflected onto the photodetectors were in the 10 mW regime.

The experimental procedure for achieving laser-locking via DHC can be outlined as follows:

- Initially, both laser frequencies are brought close to one of the resonances of their respective optical cavities. It is important to ensure that the heterodyne frequency is within the bandwidth of the GHz Phasemeter. Laser frequency tuning is accomplished by changing the offset voltage

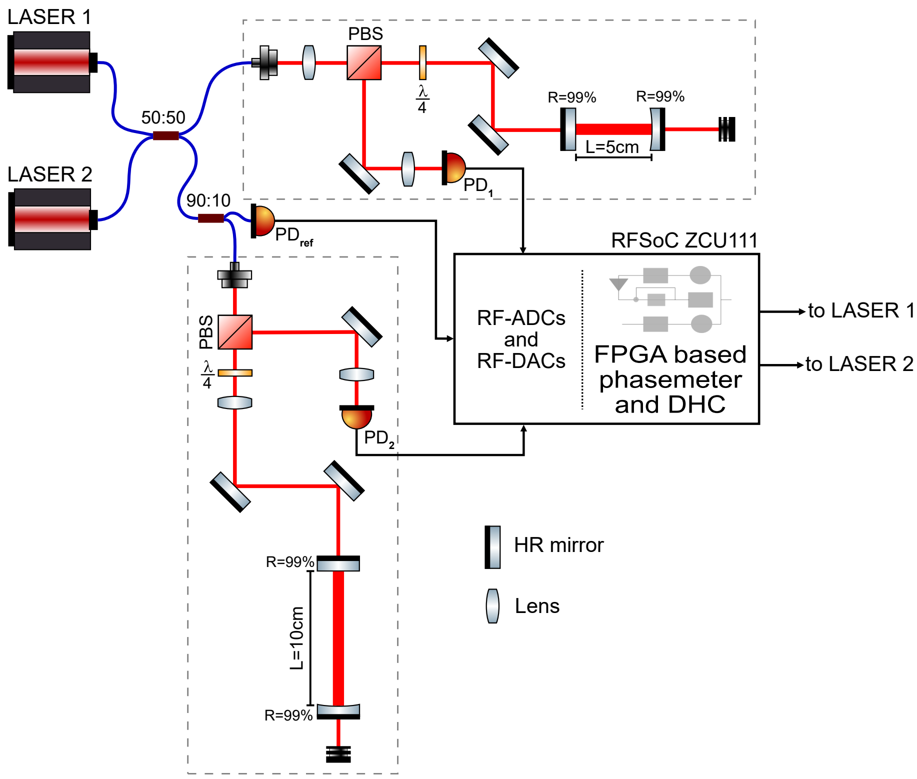

Figure 9.1: Experimental setup of the heterodyne cavity-locking using table-top in-air cavities and digital demodulation. The photodetector outputs are coupled to the ADCs of the ZCU111, and the error signals for closed-loop control are obtained from the DACs. The optical layout is not drawn to scale. (PBS: polarizing beam-splitter, PD: photodetector, DHC: Digital Heterodyne Controller.)

applied to the piezoelectric element in contact with the ECDL grating. By observing the DC component of the cavity reflection on an oscilloscope, the resonance point can easily be identified.

- Next, only one of the laser frequencies is scanned around the associated cavity resonance by activating the piezoelectric voltage scan with an amplitude of about 1 V. The demodulated signal generated in the DHC is passed to the DAC[1] and is visualized on an oscilloscope.

- Typically, this signal will have a superposition of unwanted oscillations that arise due to path length mismatch. The objective here is to correct the demodulation phase to produce an error signal close to the ideal one (see Fig. 3.2). The tuning of the demodulation phase begins with adjusting the delay value, followed by fine-tuning with the frequency-independent phase offset.

- Once an error signal with expected zero-crossing is obtained, the signal is passed to the laser controller. Using its lock module, the laser frequency is then locked to the resonance of the associated cavity.

[1] This can be achieved by setting the proportional gain of the servo to 1, while keeping the integral gain at 0.

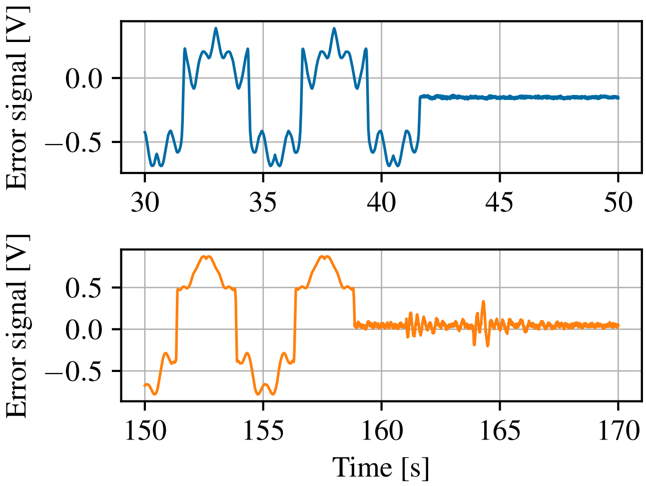

Figure 9.2: Achieved error signals using DHC for demonstrating the HS scheme with table-top in-air cavities. The overall shape of the error signal deviates from its expected form, as unwanted artifacts appear at the beginning and end of the triangular piezoelectric voltage scan. The exact reason for this behavior remains unclear. However, the linear portion of the error signal across the cavity resonance can be used for locking. Hence, a rigorous investigation of this artifact was not a primary concern.

- A similar procedure is then repeated for the second laser. It is sometimes observed that continuously scanning the frequency of the second laser can disrupt the lock of the first loop. Such cross-coupling should not occur in an ideal setup. In these cases, one can disengage the lock of the first laser and lock both lasers together once the demodulation phases for both the loops are finalized.

However, achieving an error signal in its fully expected shape was not feasible due to a more complex phase-over-frequency dependency. Some example error signals used in our HS experiment are plotted in Fig. 9.2. The shape of the error signal resembles what we achieved through analog demodulation by matching cable lengths. Nonetheless, with the DHC implementation, tuning the demodulation phase became much easier. The linear slope in the error signal across the cavity resonance allowed this signal to be effectively used for closed-loop control.

Fig. 9.3 shows the results from one such HS attempt. There was no noticeable laser frequency stabilization achieved with the table-top in-air cavities. Additionally, the frequency noise did not increase significantly. To investigate the acoustic noise coupling, an intentional audio tone at 100 Hz was played outside the enclosure of the optical setup, allowing the injected noise to couple to the length of the cavity, which is visible along with its harmonics in the beat note between the lasers.

These preliminary results confirm that the heterodyne cavity-locking scheme is operational. When the lasers are locked to the table-top cavities and measurements are recorded during quiet periods in the laboratory, the influence of the laboratory air conditioning system on the cavity length is evident in the spectrum. Similarly, during intentional noise injection, the length noise of the cavity is observed as expected. However, stable cavities are essential to fully demonstrate the laser frequency stabilization. When conducting the experiment in-air, cavities made of athermal glass would be a promising candidate. If optical cavities are constructed using commercial mirrors held in mirror mounts, a vacuum environment is necessary to maintain the cavity length stability and effectively

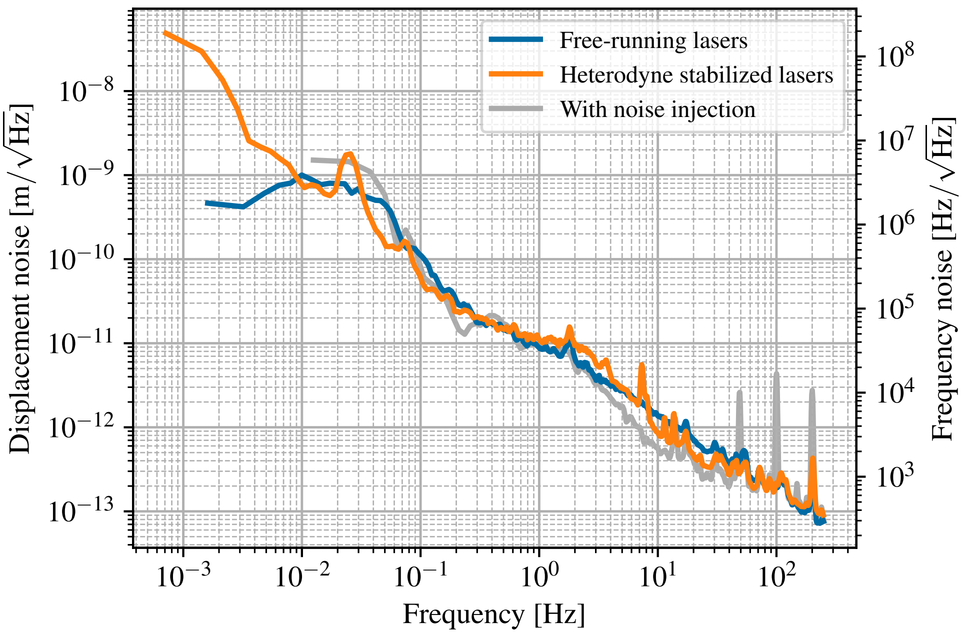

Figure 9.3: Heterodyne stabilization results for lasers locked to table-top in-air cavities using DHC. The 10 cm long cavity serves as the reference, and the displacement noise of the 5 cm long cavity is plotted (using Eq. 3.16). The frequency time series is detrended by subtracting the result of a linear least-squares fit to eliminate the long-term trends prior to calculating the spectrum. The peak appearing at about 25 mHz was consistent in all measurements and was attributed to the air conditioning system in the laboratory. For noise injection, a sinusoidal audio tone at 100 Hz was played outside the enclosure of the experiment.

demonstrate frequency stabilization. A vacuum environment would significantly decouple temperature and acoustic effects.

9.2 HS using mirror mount-based cavities in vacuum

9.2.1 Experimental facility

The experimental facility consists of a customized vacuum chamber[2] made of aluminum. For pumping the air out of the chamber, we use a dry compressing vacuum pump (SCROLLVAC 7 plus from Leybold GmbH), which is placed in a separate room outside the laboratory, thereby reducing acoustic and mechanical noise couplings to the experiment. The vacuum chamber is designed to accommodate a turbo pump connected to one of the flanges, which can help achieve higher vacuum levels; however, we have not used the turbo pump for our experiments. An active sensor (PENNINGVAC PTR 90 RN from Leybold GmbH) is installed for measuring the achieved vacuum level and is connected to a live display. Based on our experiences with such a vacuum system, a vacuum level of about 10^{-2} to 10^{-3} mbar can be reached. However, during most of the work presented in this thesis, the vacuum level was about 1.8×10^{-2} mbar, which, while not very high, was sufficient for our initial experiments.

[2] Designed by Marcel Beck and manufactured by Hositrad Deutschland.

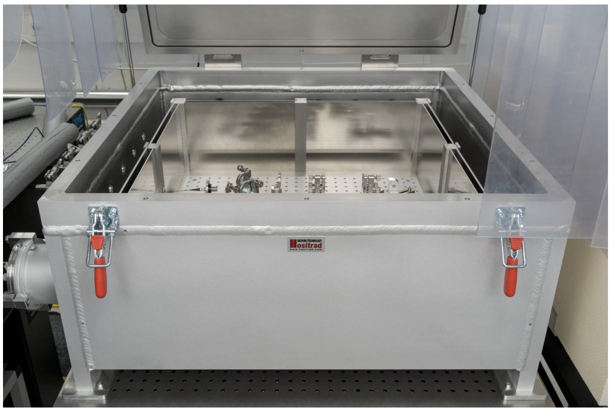

Figure 9.4: Photograph of the experimental facility where the HS scheme is demonstrated. The facility consists of a vacuum chamber that hosts a thermal shield inside. The experimental setup is assembled on the optical breadboard, which can then be placed inside this thermal enclosure. (Photo credit: Dr. Christian Darsow-Fromm.)

To improve thermal isolation, inside the vacuum chamber is an aluminum enclosure consisting of a baseplate, two-layered aluminum walls, and a lid. This enclosure is expected to act as a thermal shield for the experimental setup, which can be assembled on an additional aluminum optical breadboard. To further reduce temperature coupling from the outside environment to the experimental setup, spacers made of polyether ether ketone (PEEK) are used between the chamber and the thermal shield, as well as the thermal shield and the breadboard. PEEK material was chosen because of its high temperature resistance or low thermal conductivity [153].

This experimental facility is located below a laminar flow box to ensure stable temperature surroundings. The overall temperature and humidity of the laboratory are also controlled. A photograph of the vacuum chamber, including the thermal shield and an experimental setup inside, is shown in Fig. 9.4.

9.2.2 Experimental setup

Although the basic experimental setup remains the same, there are changes to the optical layout because of the use of the vacuum facility. The approach involves using optical fiber-based components to generate the beat note and obtain the cavity reflections onto the photodetectors. All fibers and fiber optics use polarization-maintaining single-mode fibers with angled physical contact (APC) connectors. These fiber optics and photodetectors are placed outside the vacuum chamber. The vacuum chamber hosts fiber collimators, which bring the laser beam to free space, beam-steering mirrors, mode-matching lenses, and the static cavities constructed using commercial mirror mounts. Fig. 9.5 illustrates this experimental setup.

For the static cavities to which the lasers are locked, we used high-reflective mirrors mounted in commercial kinematic mounts (POLARIS-K1E from Thorlabs) to form impedance-matched cavities with specifications similar to those used previously. A new mode-matching scheme was necessary because the fiber collimators used in this setup are vacuum compatible and differ from those used previously. The cavity alignment and mode-matching procedures remain unchanged. We ensured that sufficient spacing is maintained between the fundamental cavity mode to which we lock and the

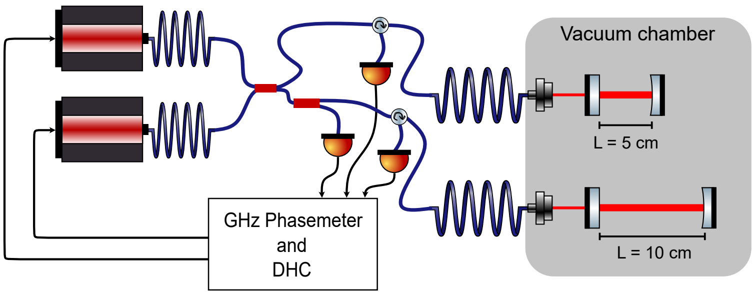

Figure 9.5: Sketch of the experimental setup for the demonstration of the HS scheme using mirror mount-based cavities in vacuum. This sketch emphasizes the compactness and simplicity of the overall interferometric scheme used for displacement sensing. The fiber-optic components and photodetectors are outside the vacuum chamber. The beam-steering mirrors and mode-matching lenses used inside the vacuum chamber are not shown. (DHC: Digital Heterodyne Controller.)

higher-order modes. Thus, even if higher-order modes are excited due to minor misalignment, they do not significantly influence the locking [80].

Fiber feed-throughs are used to bring the beat note between the lasers into the vacuum chamber. Fiber-optic circulators (CIR1550PM-APC from Thorlabs), which are non-reciprocating, unidirectional, three-port devices, behind the feed-throughs help distinguish between the laser input and the reflected light from the cavities.

9.2.3 Results

Fig. 9.6 shows the experimental results from lasers stabilized to mirror mount-based cavities in vacuum using the heterodyne locking technique. As mentioned in the previous section, the experimental setup included fiber circulators to distinguish the cavity reflection from the input laser beam. From many measurements with the fiber circulators in the setup, we observed that they introduce parasitic beams that interfere with the main interference, resulting in parasitic etalons. These produce significant frequency instability, limiting length stability at low frequencies [154]. The shoulder-like feature around the Fourier frequency of 0.1 Hz (refer to the dark-orange trace labeled as 'HS (w/ fiber circulators)' in Fig. 9.6) is primarily due to such effects.

Additionally, the peak at around 25 mHz, attributed to the air conditioning system of the laboratory, is also visible in this stabilization measurement but appears about 45 times less prominent compared to Fig. 9.3. This reduction is evident because of the vacuum environment around the cavities. At these low frequencies, temperature and humidity variations significantly contribute to the noise spectrum as well.

The aforementioned effects of parasitic beams apply to other fiber-optic components, but they were found to be particularly pronounced with our fiber circulators. Consequently, we replaced the fiber circulators with fiber couplers/splitters, which resulted in reduced reflected power detected by the photodetectors. Along with this change to the experimental setup, all three photodetectors and the fiber-optic components were isolated to further reduce the temperature and humidity influences. This isolation was achieved by placing them inside a plastic box lined with foam and covered in aluminum foil. The box is then placed on top of the vacuum chamber, beneath the laminar flow box.

These changes led to an immediate improvement in sensitivity at low frequencies (refer to the dim-

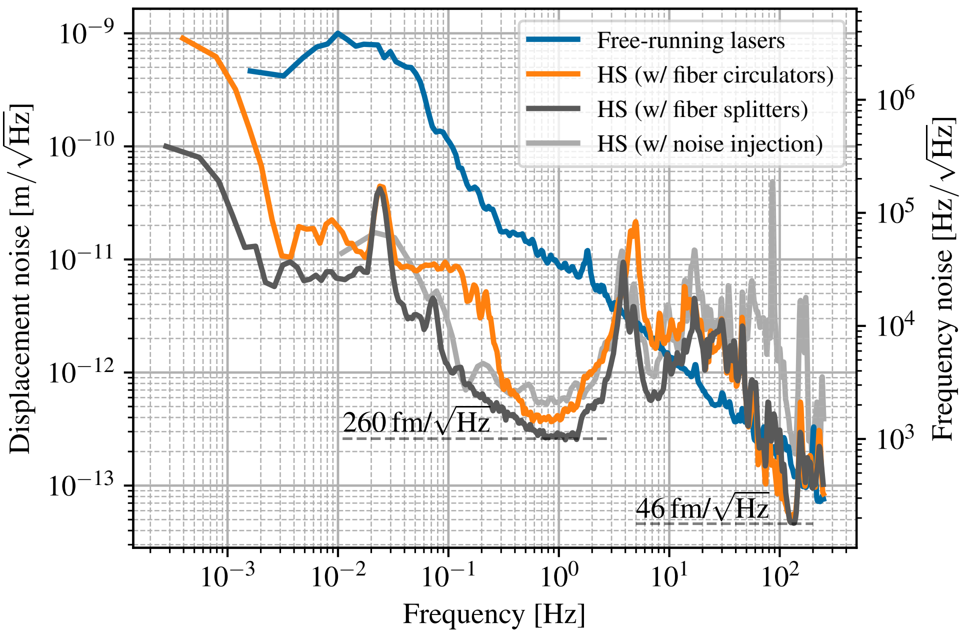

Figure 9.6: The HS results from lasers locked to mirror mount-based cavities in vacuum using DHC. The 10 cm long cavity serves as the reference, while the displacement noise of the 5 cm long cavity is plotted. The frequency readout is detrended to remove the linear trend before calculating the spectrum. The current setup achieved a maximum sensitivity of 46 fm/$\sqrt{\text{Hz}}$ at around 130 Hz. For noise injection, an intentional sinusoidal audio tone at 86.1 Hz was played in the proximity of the vacuum chamber.

gray trace labeled as 'HS (w/ fiber splitters)' in Fig. 9.6), significantly reducing the 'scatter-shoulder' in the spectrum. Between 0.1 Hz and 3 Hz, the displacement sensitivity reached the sub-picometer level. We also observed a peak sensitivity of 46 fm/$\sqrt{\text{Hz}}$ at around 130 Hz in the current setup.

At Fourier frequencies above 1.5 Hz, the length instability of the kinematic mirror mount-based cavities continues to pose a major issue. The mechanical and acoustic disturbances influencing the cavity length are limiting our ability to read out the actual achievable sensitivity at these frequencies. Instead of measuring the length noise when the cavities are static, we likely detect actual cavity length fluctuations at these frequencies due to unwanted external influence.

The acoustic and mechanical noise couplings were confirmed by noise injection. This time, a sinusoidal audio tone at 86.1 Hz was played using an external speaker placed below the vacuum chamber on the same optical table. As a result, we observed significant noise coupling at and around that frequency (refer to the dark-gray trace labeled as 'HS (w/ noise injection)' in Fig. 9.6). While this highlights the current major limitation of unstable cavities, it also cross-verifies the displacements sensed in the tens and hundreds of femtometers. This is promising, considering that tens of femtometer-level displacement sensing is possible with a 'basic' experimental setup consisting of cavities made from kinematic mirror mounts. This indicates that this displacement readout technique could achieve very high precision, potentially reaching sub-femtometer levels at 1–10 Hz.

The implementation of the DHC (as discussed in Chapter 6) and the achieved displacement sensitivity with mirror mount-based cavities in vacuum have been published in Optics Express [67].

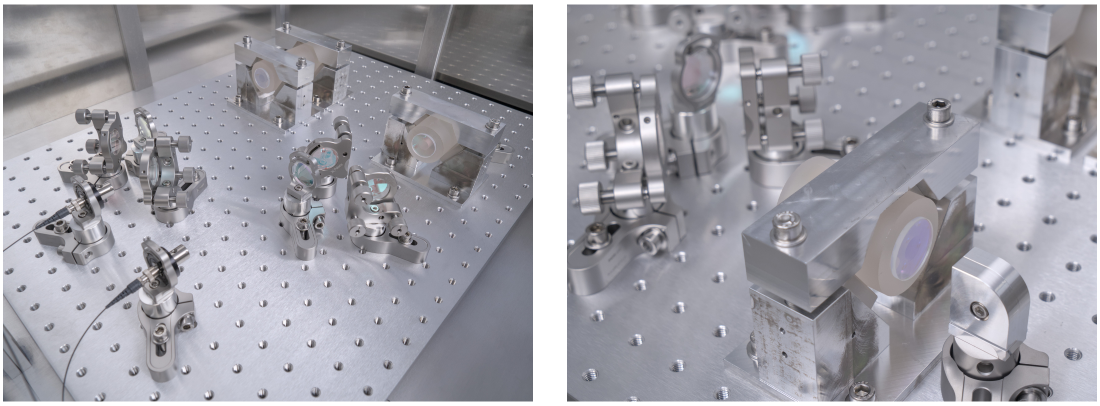

(a) (b)

Figure 9.7: Photographs of the HS experimental setup with ULE glass-based cavities inside the vacuum chamber. The cavities are held in a custom aluminum holder. To minimize thermal conduction to the cavity spacer, PEEK materials are used at the contact points. (Photo credits: Dr. Christian Darsow-Fromm.)

9.3 HS using ULE glass-based cavities

As the HS attempts with kinematic mirror mount-based cavities were limited by cavity length stability, the true displacement measurement noise floor of this interferometric scheme could not be explored experimentally. Our next approach involved using ULE glass-based optical cavities, the construction and characterization of which are discussed in Chapter 8. The presence of a ULE glass spacer between the cavity mirrors is expected to directly enhance cavity length stability, with the very low CTE of the material providing an additional advantage by reducing susceptibility to temperature fluctuations.

The experimental setup largely remains the same as illustrated in Fig. 9.5, but the mirror mount-based cavities have been replaced with ULE glass-based cavities. Accordingly, the focal lengths and positions of the mode-matching lenses have been adapted to ensure optimal mode matching and high fringe visibility. Photographs of the experimental setup are shown in Fig. 9.7. The cavity alignment results of this setup are discussed in Sec. 8.4. Taking advantage of fiber couplers/splitters over fiber circulators, we employed the former. Fiber-optic components and photodetectors are isolated from ambient influences by placing them inside a box as explained before.

9.3.1 Experimental results

Fig. 9.8 shows the results from our HS attempts with ULE glass-based cavities. Comparing the current results with the dark-gray trace labeled as ‘HS w/ mirror mount cavities’—the best result from Fig. 9.6—the improvement from the ULE glass-based cavities in the setup is evident. Most of the acoustic and mechanical noise couplings in the audio frequency band are suppressed by roughly two orders of magnitude overall.

At the low-frequency end, the reduction in temperature noise coupling is likely to improve sensitivity. However, the scatter-shoulder due to parasitic beams is now dominating at those frequencies. The parasitic beams resulting from the fiber components are expected to be non-stationary, and therefore they contribute significantly to instability. This is clear from two different stabilization measurements with ULE glass-based cavities. The precision with which we can read out the cavity length change of

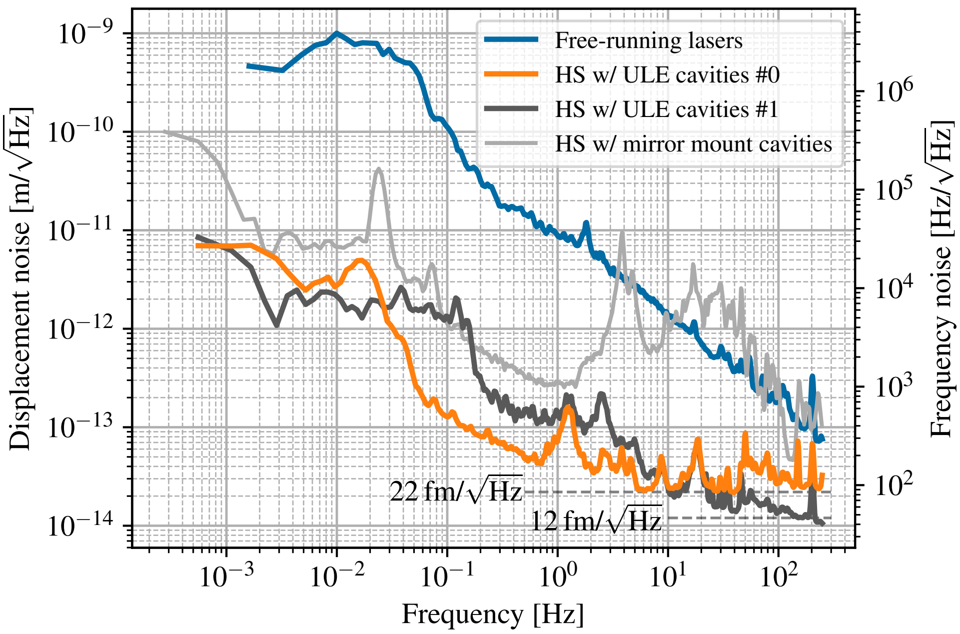

Figure 9.8: The HS results for lasers locked to ULE glass-based cavities in vacuum using DHC. The 10 cm long cavity serves as the reference, while the displacement noise of the 5 cm long cavity is plotted. The frequency readout is detrended by subtracting the result of a linear least-squares fit before calculating the spectrum. Spectra from two measurements with different open-loop gains for the laser-locking control loop are shown. At Fourier frequencies above 5 Hz, the overall displacement sensitivity is in the range of 20 fm/$\sqrt{\text{Hz}}$.

the 5 cm long cavity, compared to the stability of the 10 cm long cavity, is in the 20 fm/$\sqrt{\text{Hz}}$ regime for Fourier frequencies above 5 Hz.

9.3.2 Optimizing the laser-locking control loops

In addition to enhancing the experimental setup for better isolation from external noise, we have also focused on optimizing the laser-locking control loops. A general procedure used to optimize the PI gains of both fast and slow controllers within the DLC pro is to minimize the standard deviation of the error signal. However, this approach complicates the quantification of the control loop behavior. Therefore, we added the capability for open-loop transfer function measurement to the laser-locking control loop. Fig. 9.9 illustrates the required modifications to the existing experimental setup.

The transfer function measurement option is integrated to facilitate the real-time optimization of our control loops. The error signal generated from the DHC is split into two parts. One part is sent to the Frequency Response Analyzer, which generates and puts out a frequency sweep with specified parameters. This sweep is added to the other part of the error signal. The resulting signal is then split again, with one part going to the laser controller (DLC pro, in our case) and the other returning to the Frequency Response Analyzer. This allows us to obtain the open-loop transfer function by comparing the two inputs in the Frequency Response Analyzer.

Once the laser is locked to its respective cavity, we initiate the transfer function measurement and tune the servo gains of the slow controller that actuates the piezoelectric voltage. We achieved a UGF

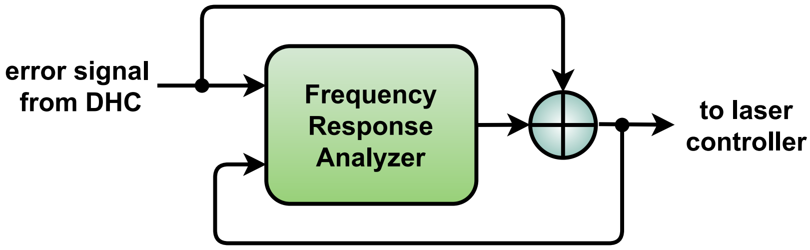

Figure 9.9: Sketch of the open-loop transfer function measurement setup for the laser-locking control loops. We used a Frequency Response Analyzer instrument on the Moku:Pro platform [155], with the subsequent addition processed directly within the platform. By comparing the two inputs of the Frequency Response Analyzer, we can derive the loop transfer function. The setup for only one laser-locking loop is shown here; an identical setup was also implemented for the second laser-locking loop.

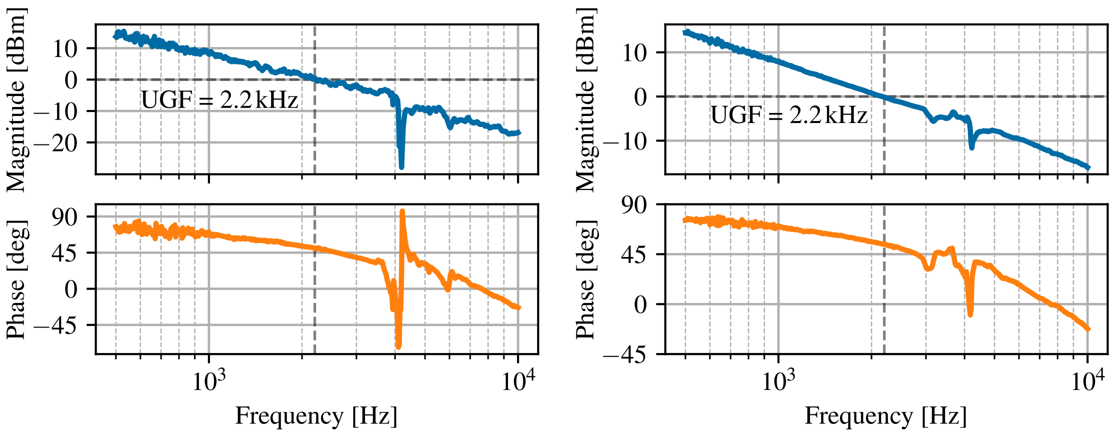

(a) 5 cm long cavity loop

(b) 10 cm long cavity loop

Figure 9.10: Bode plots corresponding to the open-loop transfer functions of two laser-locking control loops. The data were collected using the Frequency Response Analyzer instrument of the Moku:Pro platform. Both loops were optimized to achieve a UGF of about 2.2 kHz.

of about 200–300 Hz without exciting the piezoelectric resonances. A further increase in control loop bandwidth can be attained by utilizing the fast controller that actuates directly on the laser diode current. By adjusting the PI gains of this fast controller, we have the flexibility to tune the control loop bandwidth.

In the experimental results shown in Fig. 9.8, the UGF of the control loops was initially not measured. We suspect that the control loop bandwidth was relatively low, leading to a noise floor limited to 22 fm/$\sqrt{\text{Hz}}$ (refer to dark-orange trace labeled as ‘HS w/ ULE cavities #0’ in Fig. 9.8). Following this, we integrated the transfer function measurement option and optimized the servo gains to achieve a UGF of about 2.2 kHz for both laser-locking control loops. The Bode plots of the corresponding measurements are shown in Fig. 9.10.

The resulting stabilization spectrum after optimizing the control loops for this bandwidth is represented by the dim-gray trace labeled as ‘HS w/ ULE cavities #1’ in Fig. 9.8. We observe an improvement in the sensitivity above 10 Hz; however, at mid-frequency ranges, the noise increased. This increase can primarily be attributed to non-stationary effects caused by parasitic beams. The increased bandwidth also implies higher gain values at those frequencies, which, in turn, enhances the responsiveness of the laser frequency to cavity length fluctuations. Thus, the lasers more closely follow cavity length variations, resulting in increased displacement noise. However, further increasing the UGF to 6 kHz did not alter the performance, confirming that we are not limited by the gain of the control loop.

Our stabilization attempts provided a slight indication that achieving the correct demodulation phase, and thus a proper error signal for closed-loop control, may impact the stabilization performance. However, this aspect has neither been fully understood nor thoroughly explored during the course of this thesis. The ripples in the ADC front-end (see Fig. 6.4), residual amplitude modulation fluctuations of the laser source, and thermal noise (Brownian motion) in the cavities are additional possible limitations at this time. Sec. 10.3 discusses more about the current experimental limitations.

9.4 Noise discussion

In principle, the proposed technique for displacement sensing can achieve sub-femtometer level performance at frequencies above 1 Hz, as discussed in Sec. 3.3. Such an extremely low noise floor is possible because contributions from shot noise, electronic noise, and ADC noise in the laser-locking loop are strongly suppressed by the optical cavity, while contributions from ADC noise and electronic noise become negligible for the heterodyne frequency readout. The parameter of interest is the signal frequency rather than the signal power.

We can separate the noise discussion into two parts: the noise associated with locking the lasers to the cavities and the noise associated with the frequency readout of the laser beat note. A detailed discussion of several noise sources that appear in laser-locking loops is provided in [156] for the digital implementation of PDH locking. This analysis demonstrates a noise floor below $1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$. We expect a very similar noise coupling in the HS laser-locking loops implemented digitally.

Now, turning to the frequency readout noise, a 1 Hz readout precision of the beat note corresponds to a displacement readout noise contribution of $< 10^{-15}$ m because of the phasemeter in our experimental setup. However, the developed GHz Phasemeter is capable of resolving signal frequencies with millihertz precision even for gigahertz signal frequencies, as shown in Fig. 5.13. Accordingly, a millihertz precision will induce a displacement readout noise of $< 10^{-18}$ m. In all practical implementations, other noise sources will be orders of magnitude above this contribution, limiting the displacement sensitivity to higher levels. However, the jitter of the phasemeter clock and differential phase noise between the reference beat and the cavity reflection, which does not result from laser-cavity interaction, are potential noise sources that can affect displacement sensing [85]. The use of impedance-matched cavities and more stable reference clocks can help reduce these influences.

The ADC noise from the DHC channels and the shot noise from the photodetectors that receive reflected light from the cavities still contribute to the overall noise budget. With the chosen cavity parameters and input laser power levels of about 10 mW, the displacement noise induced by shot noise falls within the $10^{-17}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ range, as discussed in Sec. 3.3.1. This fundamental noise floor can be lowered by increasing laser power and utilizing higher finesse cavities. Employing cavities with higher finesse values also benefits the HS technique (see Eq. 3.12) and helps suppress the influence of unwanted parasitic reflections. Other potential noise sources are not rigorously analyzed in this work, but numerous descriptions in the context of PDH locking are available in the literature [139, 157, 158].

9.4.1 Influence of external phase noise

Unlike the PDH stabilization scheme, the HS technique relies on reference phase detection. This makes the system susceptible to external phase noise, such as phase drifts in the low frequency regime. This is especially problematic when the experimental setup is realized using fiber optics.

The phase noise that occurs when splitting the laser into different paths and then directing it to each optical cavity, along with the reference beat detection, will show up as measurement noise. This is true even if the lasers are locked to their respective cavities using, e.g., PDH for a similar type of heterodyne detection. If such an effect becomes critical, thermal stabilization of the fiber setup or transitioning to a stable free-beam setup maybe required.

For a quantitative understanding, if the phase noise due to the fiber paths limits our performance, then to drive the noise levels in the sub-picometer regime, the frequency noise introduced by the fiber paths would need to be on the order of $10^3\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ (see Fig. 9.8). At 1 Hz, this corresponds to a phase noise of $10^3 \times 2\pi\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$ or, in terms of displacement[3], about $1.55\,\mathrm{mm}/\sqrt{\mathrm{Hz}}$. Since frequency is the differential of phase, the coupling is stronger for higher frequencies. However, the spectrum of the phase noise is arbitrary.

The second aspect is the influence of phase drifts on the heterodyne demodulation phase. This is the phase between the beat note measured on the reference photodetector and the signal detected by the photodetector in the reflection port of the cavity. Since these are RF signals in the gigahertz range, their wavelengths are large (>10 cm), and small optical phase or length changes in the fibers have a minimal effect on their relative phase. They both experience almost the same exact phase behavior in the fibers, keeping their relative phase virtually constant.

Nevertheless, the demonstrated readout noise floor of about $260\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$ at 1 Hz with kinematic mirror mount-based cavities and below $20\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$ for frequencies above 5 Hz with ULE glass-based cavities are very promising results. By employing stable reference cavities and enhancing thermal and mechanical isolation, we can further improve the displacement readout sensitivity and work towards demonstrating sub-femtometer noise levels.

For the demonstration of heterodyne cavity-locking, we used static cavities, meaning that the physical lengths of both cavities were maintained at a constant value. This approach allowed us to probe the displacement readout sensitivity of our scheme. Following this successful demonstration, we now proceed to the heterodyne cavity-tracking experiment to investigate the operating range of this displacement sensing technique.

[3] Assuming a laser wavelength of 1550 nm.

CHAPTER 10

Operating range via heterodyne cavity-tracking

For any sensor, it is advantageous to have a high dynamic range. To achieve a high dynamic range, both the sensitivity and operating range of the sensor should be as high as possible. Hence, after verifying the heterodyne cavity-locking technique and understanding the current displacement readout noise floor, the operating range must now be probed using the GHz Phasemeter and DHC. For this purpose, the plan is to introduce a dynamic cavity into the experimental setup instead of using two static cavities.

10.1 Experimental setup

The incorporation of a dynamic cavity into the experimental setup required the installation of the end mirror of the 5 cm long cavity using a kinematic mirror mount equipped with three piezoelectric adjusters (POLARIS-K1S3P from Thorlabs). Using the mechanical adjustment screws of the mount, the cavity was initially aligned to achieve a fringe visibility of about 97%. A photograph of this optical setup is shown in Fig. 10.1.

The piezoelectric elements are controlled using a 3-channel, open-loop piezo controller (MDT693B from Thorlabs). These elements can accept an input voltage range of 0–150 V. A constant voltage of about 75 V is supplied to all three adjusters using the internal master scan of the controller. To modulate the cavity length, a sinusoidal modulation signal is generated using an external signal generator and fed into the external master scan of the piezo controller. The amplitude of the modulation signal determines the displacement range of the piezoelectric elements. By simultaneously applying the same voltage to all three piezoelectric adjusters, we rely on the linear motion of the cavity end mirror along the interferometric axis. However, as we will discuss in the following section, this approach might have resulted in significant tilts of the cavity end mirror during modulation in the current experimental setup.

As the cavity is chosen to have a flat-curved configuration with a relatively large radius of curvature, the movement of the end mirror by a fraction of a micron will not significantly affect the mode-matching criterion. The beam waist is still expected to be positioned at the cavity input mirror with nearly the same beam waist size. This ensures good fringe visibility of the dynamic cavity throughout the length modulation. However, care should be taken when choosing the cavity geometry to ensure that higher-order modes are not degenerate with the fundamental mode. This is to guarantee that

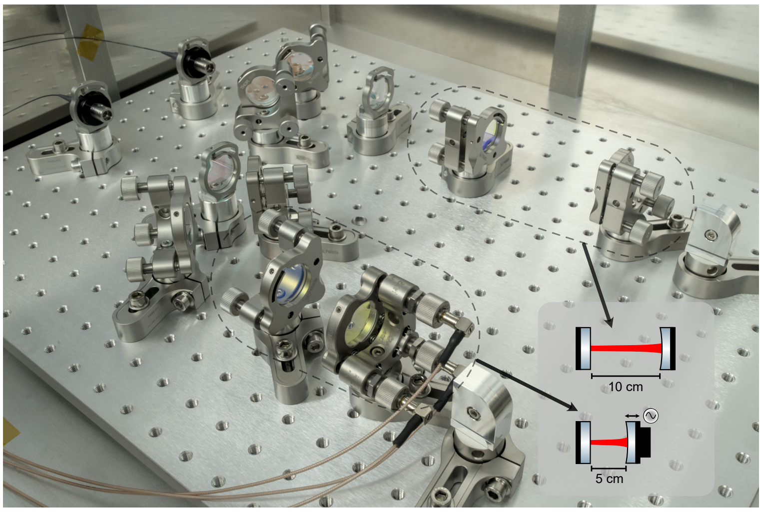

Figure 10.1: Experimental setup for the heterodyne cavity-tracking. This version of the experiment utilizes kinematic mirror mount-based optical cavities. The construction of the ULE glass-based cavities happened later in the timeline. The aligned experimental setup is placed inside a thermal shield, which goes into the vacuum chamber. (Photo credit: Dr. Christian Darsow-Fromm.)

the fundamental mode will not be affected by the higher-order mode under minor misalignment. In contrast to the current configuration, in the case of confocal cavities, even small displacements can move higher-order modes through the fundamental mode, disturbing the generated error signal.

10.2 Results

Once the lasers are locked to these kinematic mirror mount-based cavities, a sinusoidal signal is applied to all piezoelectric adjusters. This, in turn, modulates the cavity length, and the corresponding laser is expected to follow this motion. We started with an external modulation signal amplitude of 50 mV and gradually increased it, ensuring that the laser still followed the cavity resonance. By tracking the beat note, we can then use Eq. 3.16 to obtain the one-way displacement of the dynamic cavity end mirror with respect to the static cavity length reference.

When examining the experimental results of heterodyne cavity-tracking, the recorded frequency time series contains relevant information that is intuitive to interpret. Time series data from one such measurement is shown in Fig. 10.2. The deviation from the sinusoidal shape is likely due to the non-uniform expansion and compression of the piezoelectric elements relative to one another, resulting in additional tilt effects that create an effective slow modulation of the cavity length, which is visible here as a decreasing modulation depth. We also expect contributions from the hysteresis of piezoelectric actuators. Additionally, the aforementioned non-flat phase and amplitude response of the ADC front-end transfer function can also influence this.

When the cavity mirror moves by an order of one-tenth of the wavelength used, the beat note shifts by a fraction of a gigahertz. The high-bandwidth readout and control system is capable of keeping the lasers locked, thereby functioning as a displacement sensor with a fringe-scale operating range. Here, we demonstrated a maximum operating range of about 0.15 µm. The laser could not maintain lock

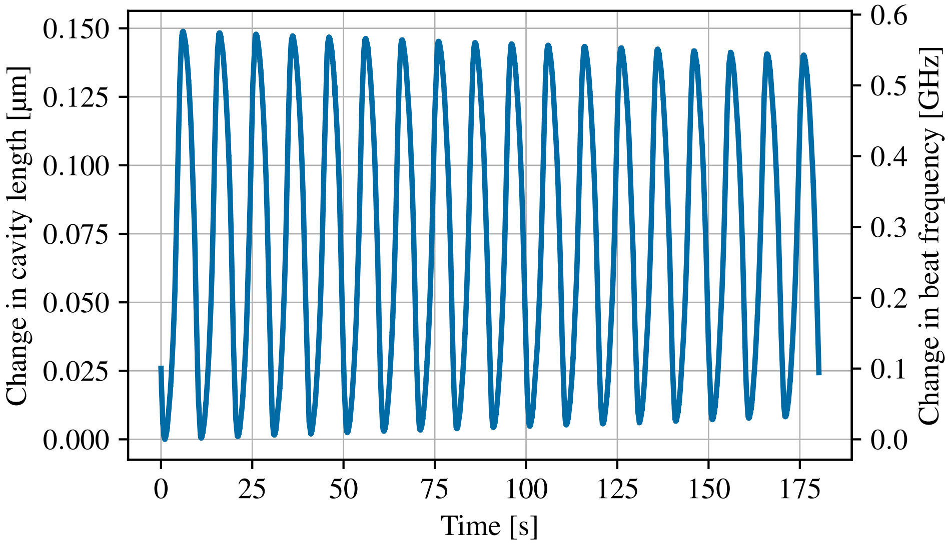

Figure 10.2: Experimental result from the heterodyne cavity-tracking demonstration. The beat note between the lasers was tracked while the cavity length was modulated using piezoelectric elements. The corresponding physical length change of the cavity was calculated and is provided for reference. This particular trial used an external sinusoidal scan signal with a frequency of 0.1 Hz and an amplitude of 0.5 V.

with the dynamic cavity when the modulation depth was increased further.

10.2.1 Effect of mirror tilt on displacement measurement

The interferometric scheme investigated in this thesis is suitable only for one-dimensional displacement sensing. However, the tilt of the end mirror of the dynamic cavity can mimic a change in cavity length, introducing errors in displacement measurement. The interferometric phase should only read longitudinal distance variations; therefore, any coupling of angular or lateral motion is considered unwanted cross-talk. In this particular case, the cross-talk is referred to as tilt-to-length (TTL) coupling [159]. The TTL coupling is a considerable noise source, especially in LISA [160].

Such a TTL coupling factor was studied for the current experimental setup. By introducing different tilt values to the end mirror of the 5 cm long dynamic cavity, the phase shift relative to the cavity resonance (or to the zero-crossing of the demodulated signal) was simulated using Finesse3 [77]. This phase shift due to tilt was then converted into cavity length error. Fig. 10.3 illustrates the results of this simulation. Above a misalignment of approximately 30 µrad, the TTL coupling factor exceeds 10 pm/µrad. Therefore, to reduce the error due to end-mirror tilts, the construction of the dynamic cavity must be performed with care. For experimental demonstration, one would need both a highly stable and tunable optical high-finesse cavity, as implemented in [161], for example.

10.3 Current limitations

The current limitations of the demonstrated heterodyne cavity-locking and tracking performance are primarily the mechanical instability of our cavities, which likely dominates the measurements at the

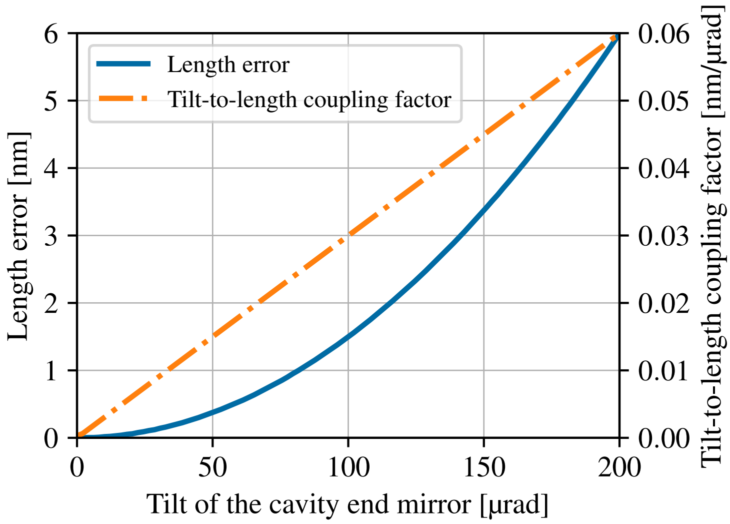

Figure 10.3: Effect of tilt of the dynamic cavity end-mirror on the displacement readout. For different tilt values of the end mirror, the corresponding length error is obtained from the simulation. The length error as a function of the introduced tilt follows a quadratic trend (solid trace). By calculating the gradient, we determine the TTL coupling factor (dash-dotted trace).

high-frequency end. The noise injection attempts clearly indicate that acoustic disturbances influence the stability of the cavity length. Using stable cavities and/or more stable frequency references is necessary to demonstrate a lower achievable noise floor, and future studies must verify that no additional, unexpected readout noise couplings arise in our technique when compared to, e.g., PDH.

Another major limitation of the currently demonstrated results arise from the fiber-optic-based experimental setup. While such an approach helps in realizing compact experimental arrangements, it incurs increased noise, especially at frequencies below 1 Hz. Length and phase noise of the fibers will also be present in the differential frequency measurements. The effects of unwanted parasitic beams, which may cause parasitic etalons, can be reduced by employing higher finesse cavities, which increase the SNR of the nominal signals relative to the parasitic signals. If the scheme utilizes an external ultra-stable reference laser provided via fiber, then corresponding fiber length stabilization can be employed to improve performance [162]. We assume no significant phase dynamics exist in our fiber setup between the two laser frequencies; however, proper follow-up and validation are required.

The demonstrated operating range of the displacement sensor is also limited by the optimization of the error signal in terms of time delay compensation. In our experiments, we were able to achieve an error signal suitable for closed-loop control. However, a complete removal of oscillations in the error signal was not possible. This suggests we may currently be limited by the step sizes of the demodulation phase correction parameters implemented in the PL. A higher bit resolution could be implemented in the future, allowing for finer tuning of these parameters. Additionally, the uncorrected effects in the analog signal digitization chain also play a role. The general drop in the ADC transfer function magnitude at higher frequencies, along with variations or ripples in its phase response (see Fig. 6.4), are not accounted for or corrected in our current implementation. This limitation currently affects our ability to match the demodulation phase and maintain constant control gains for large frequency changes, impacting wideband tuning and, consequently, the operating range.

It is also important to ensure an adequate laser-locking bandwidth for high dynamic range displacement readout (refer to Sec. 3.3.3). Increasing the control bandwidth is feasible by introducing additional actuators, such as electro-optic modulators, as demonstrated in LIGO laser stabilization [70]. The maximum control bandwidth supported by our laser systems is approximately 30 kHz [72]. From the RFSoC perspective, achieving UGFs in the hundreds of kilohertz range is possible, assuming that the control loop bandwidth is limited by the delays of the RF data converters.

With the signal bandwidth of the GHz Phasemeter being about 2 GHz, the maximum displacement tracking range can be calculated using Eq. 3.37 as $\lambda/3$ for a 5 cm long dynamic cavity. In general, the displacement sensing operating range for this type of cavity-tracking technique is limited only by the bandwidth of the frequency readout. However, in our experimental demonstration using the HS technique, the maximum displacement range is limited to half of the FSR, assuming that the secondary laser resonance is precisely positioned between the two resonance peaks of the primary to begin with. This is because, when both lasers are at resonance in the same cavity, no side-band gets reflected, resulting in a loss of lock. From the frequency readout perspective, the PLL will fall out of lock as the beat note approaches a DC value.

In spite of all the aforementioned unaddressed limitations, we have demonstrated a displacement readout sensitivity of about 20 fm/$\sqrt{\mathrm{Hz}}$ for Fourier frequencies above 5 Hz and an operating range of about 0.15 µm. This demonstrates a dynamic range between the operating range and sensitivity of more than six orders of magnitude.

CHAPTER 11

Summary and future scope

The displacement sensors of the seismic isolation systems of ground-based gravitational-wave detectors play a vital role in keeping the detectors sensitive enough to detect gravitational waves. Improving the sensitivity of these sensors at frequencies below 10 Hz is crucial for future ground-based detectors. For planned missions such as LGWA, the primary requirement is to have an inertial sensor with sub-femtometer precision in the 0.1–1 Hz range. However, such sensors need to operate over a significant range to cope with actual motions in the corresponding experiments. Finally, they must be compact to integrate with compact inertial sensors [61] and seismic isolation systems.

Keeping these requirements in mind, we have introduced a compact interferometric scheme suitable for high-precision displacement readouts, aiming for sub-femtometer noise levels at frequencies around 1 Hz and a fringe-scale operating range. We have demonstrated an overall displacement readout sensitivity of about 20 $\mathrm{fm}/\sqrt{\mathrm{Hz}}$ for frequencies above 5 Hz and have achieved below 15 $\mathrm{fm}/\sqrt{\mathrm{Hz}}$ for frequencies above 100 Hz. By probing a cavity with length actuation, we have shown about six orders of magnitude of dynamic range for displacement sensing, reaching a maximum motion of 0.15 µm.

A compact laser interferometer-based displacement sensing system with a readout noise floor in the femtometer regime at 1–100 Hz Fourier frequencies is rarely achieved, making the contributions of this work very significant in the field of precision metrology. The precision we achieved here is comparable to the exceptionally low residual displacement measurement noise of 32.1 $\mathrm{fm}/\sqrt{\mathrm{Hz}}$ attained down to much lower frequencies by the LISA Pathfinder mission [41].

This higher precision and moderate operating range are relevant for compact opto-mechanical sensors that operate in quiet environments or within in-loop seismic isolation systems [23, 62]. For instance, inertial noise of about 0.1 $\mathrm{pm}/\sqrt{\mathrm{Hz}}$ at 50 Hz will move the mechanical oscillator reported in [61] by roughly 100 nm at its resonance (assuming a Q value of about 1×10^6 [62]). Such a moderate displacement range is trackable by our heterodyne cavity-tracking scheme with very high precision. A coating thermal noise-limited displacement measurement with this oscillator, which is conceptually possible with our scheme, as described in Sec. 3.3, would be sufficient to realize a suspension thermal noise-limited inertial sensor that is more sensitive than an L-4C geophone at all frequencies between 0.1 Hz and 100 Hz. Our solution can bridge the gap between narrowband and multi-fringe displacement sensing and achieve high precision in such applications.

The use of optical cavities with a few centimeters in length and a fiber-coupled setup, along with readout and control integrated into an RFSoC evaluation board, makes our system compact and easily portable. Our scheme, when implemented as a readout of compact inertial sensors currently being studied and developed [61, 163, 164], can improve their performance. To use this as a readout for other inertial sensors, one would merely need to construct a cavity with one mirror attached to the proof mass of the sensor. While our scheme requires one laser per optical cavity being read out, we assume that the high bandwidth of our GHz Phasemeter and laser-locking loops implemented with DHC enables the use of relatively noisy and therefore comparably inexpensive lasers.

Improvements in handling the analog front-end effects and future RFSoCs with even higher bandwidth indicate that a full interferometric-fringe operating range can be achievable in the near future for cavities as short as 5 cm. By employing stable reference cavities and providing better thermal and mechanical isolation, one can further improve displacement readout sensitivity and work towards demonstrating sub-femtometer noise levels. The influence of parasitic beams can be reduced by using a free-beam setup or increasing the cavity finesse.

As pointed out in Sec. 5.9, what started as part of the displacement sensing scheme has led to the development of the GHz Phasemeter, which has grown into a standalone instrument offering high-speed, high-bandwidth phase- or frequency- tracking. Its signal bandwidth of 2.048 GHz, tracking bandwidth of 2 MHz, capture range of 4.1 MHz, and tracking speed of more than 240 GHz/s set a new benchmark for these parameters in phasemeter developments. The tracking and detection bandwidths of the digital phasemeter used to monitor the frequency will determine the maximum tolerable range and speed of mirror motion in the cavity of our displacement sensing scheme. To the best of our knowledge, the developed GHz Phasemeter is the highest bandwidth, highly stable, and fastest phase-tracking instrument reported so far. The two additional features of our phasemeter—the direct measurement of tracking bandwidth and direct estimation of residual phase error—facilitate the realization of optimal tracking conditions based on input signal dynamics. These features are novel in the field of FPGA-based phasemeter developments.

In the future, one can investigate the calibration of the phase response and transfer function of this device to enable absolute phase measurements. Our approach can be extended with techniques such as ADC interleaving to push the phasemeter detection bandwidth to even higher levels. Consequently, it will be possible to utilize ultra-high dynamic range frequency measurements of the digital phasemeters in even more applications.

The proposed displacement sensing scheme is a relative measurement. As one cannot distinguish between the length fluctuations of the two cavities, the accuracy and precision of the dynamic cavity readout depend upon the stability of the frequency reference. Hence, we have also attempted in exploring a potential frequency reference candidate, namely athermal glass. With its unique thermo-optic properties, which might eliminate the need for vacuum shielding to some extent, it has the potential to serve as a very good frequency reference when transformed into a monolithic cavity. All the necessary preparations for characterizing these athermal cavities and probing them as frequency references were carried out. Unfortunately, because of the degradation of the etalon facets (discussed in Sec. 7.5) and the extremely delayed timeline in getting the samples machined as requested, we could not continue our experiments on the athermal glass samples. Nevertheless, a cavity construction method by gluing the mirrors to the facets of a spacer material has been realized in the meantime, and the stability of the cavities has been demonstrated.

Putting together all these individual components, namely the heterodyne stabilization (HS) interferometric scheme, the RFSoC-based readout and control system, and the laser frequency reference, we have successfully demonstrated femtometer-precision, fringe-scale displacement sensing that can be utilized in future ground-based gravitational-wave detectors, other planned gravitational-wave detection missions, or in any other relevant fields.

Appendices

APPENDIX **A**

Analyzing a time series

To evaluate the noise performance of various systems, we record data over time, resulting in a time series. We then analyze this time series in the frequency domain. One of the ways is to look at the power spectral density (PSD) (also known as the power spectrum) or its square root, the amplitude spectral density (ASD). According to the Wiener-Khinchin theorem, the PSD of a stationary random process is given by the Fourier transform of the auto-correlation of that process.

Let $x(t)$ be the recorded time series of the process we are interested in, expressed in units of meters. The auto-correlation of $x(t)$ is defined as the correlation of it with a delayed copy of itself as a function of delay, which can be expressed as

$$x \star x(\tau) \equiv \int_{-\infty}^{\infty} x(t)x(t+\tau)\, dt \qquad [\mathrm{m}^2/\mathrm{Hz}]. \tag{A.1}$$

In practice, we replace the integrals from $-\infty$ to ∞ with a limit, resulting in

$$x \star x(\tau) = \lim_{T\to\infty} \frac{1}{2T} \int_{-T}^{T} x(t)x(t+\tau)\, dt \qquad [\mathrm{m}^2]. \tag{A.2}$$

For a deterministic time series, Fourier transform is the correct way to translate into the frequency domain. However, when $x(t)$ is a random time series, it is beneficial to calculate the PSD of the time series. The PSD of $x(t)$ is defined as

$$\tilde{S}_{xx}(f) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x \star x(\tau) e^{-i2\pi f\tau}\, d\tau \qquad [\mathrm{m}^2/\mathrm{Hz}]. \tag{A.3}$$

$\tilde{S}_{xx}(f)$ measures the amount of time variation in the time series occurring at frequency f. In other words, it represents the power distribution of the time series as a function of frequency, measuring the extent to which a noisy time series contains components resembling sinusoids of frequency f.

Moreover, a single-sided power spectrum $\tilde{S}^2_{xx}(f)$ is defined to consider only positive frequencies, i.e., $\tilde{S}^2_{xx}(f) \equiv 2\tilde{S}_{xx}(f)$ for $f \geq 0$, and 0 otherwise. Throughout this thesis, when the term PSD is used, it refers to this single-sided version.

We go a step further and work with the ASD, defined as

$$\tilde{S}_x(f) \equiv \sqrt{\tilde{S}^2_{xx}(f)} \qquad [\mathrm{m}/\sqrt{\mathrm{Hz}}]. \tag{A.4}$$

The advantage of dealing with the ASD is that the unit of meters matches what we actually record as time series.

In this thesis, the ASD is calculated using the LPSD estimation technique, as presented in [100]. The logarithmic frequency axis yields results that are more informative than standard spectrograms. The Python package 'LPSD' [165] is useful for directly obtaining the spectrum from the recorded time series.

For a detailed discussion on the topics briefly introduced in this appendix, one can refer to [166], for example.

APPENDIX B

Frequency counter of SDG2122X: data recording

The SDG2122X arbitrary waveform generator features a frequency counter capable of measuring frequencies from 0.1 Hz to 200 MHz. Initially, this frequency counter was employed to measure the beat note between two CTLs. As the measured frequency is displayed only on the instrument's screen, an explicit method of data recording was needed. The Socket API allows control of the SDG2122X via LAN without the need to install additional libraries, reducing the complexity of programming. Additionally, Python has a low-level networking module that provides direct access to the socket interface.

The Python code snippet below was used to establish such a communication and record the measured frequency from the frequency counter.

```
import socket
import sys
import time
from datetime import datetime

def SocketConnect():
  try:
    #create an AF_INET, STREAM socket (TCP)
    s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
  except socket.error:
    print ("Failed to create socket.")
    sys.exit();
  try:
    #Connect to remote server
    s.connect((remote_ip , port))
  except socket.error:
    print ("Failed to connect to ip " + remote_ip)
  return s

def SocketQuery(Sock, cmd):
  try :
    #Send cmd string
    Sock.sendall(cmd)
    time.sleep(1)
```

```
    except socket.error:
      #Send failed
      print ("Send failed")
      sys.exit()
    reply = Sock.recv(4096)
    return reply

def SocketClose(Sock):
    #close the socket
    Sock.close()
    time.sleep(.300)

def FreqCounter(s, ref_freq, t, exp_name=None, file_name="FreqCounter.txt"):
    """
    Frequency counter that counts the frequency of the input signal for a given
    time ‘t’ (in min) with respect to the given reference frequency ‘ref_freq’

    The measured data along with time stamps are written to a .txt file
    """
    refq = f"FCNT REFQ,{ref_freq}\n"
    refq = bytes(refq, "utf-8") # converting f-string into BytesStr
    s.sendall(refq) # Set reference frequency to given value

    s.sendall(b"FCNT STATE,ON\n") # Turn on the frequency counter

    time.sleep(3)
    start_time = datetime.now().strftime("%Y-%m-%d %H:%M:%S")

    with open(file_name,"w") as fh:
      fh.write("% This is a comment at the beginning of the file \n")
      fh.write(f"% Starting time of the measurement: {start_time} \n")
      fh.write(f"% Measurement Name: {exp_name} \n")
      t_end = time.time() + t * 60
      while time.time() <= t_end:
        fh.write(f"TIME, {time.time()}, "
            + SocketQuery(s, b"FCNT?\n").decode("utf-8"))

    s.sendall(b"FCNT STATE,OFF\n")


if __name__ == "__main__":

    remote_ip = "10.11.13.214" # instrument's IP address
    port = 5025 # SDG2xxx has 5025 as the instrument service port number

    s = SocketConnect()

    ref_freq = 10000000 # Set reference frequency
    t = 10 # Measurement time in min
    exp_name = "free-running CTL laser frequency noise"
    file_name = "20240218_CTL_free.txt"
    FreqCounter(s, ref_freq, t, exp_name, file_name)

    SocketClose(s)
```

APPENDIX **C**

Optical cavity alignment procedures

This appendix provides the experimental approaches followed to align optical cavities. Here, we consider two cavities, each constructed using two high-reflective off-the-shelf commercial mirrors held in two commercial mirror mounts. The overall procedure discussed remains mostly the same even when employing spacer-based stable cavities.

C.1 Input laser beam profiles

To achieve resonance of a laser filed in an optical cavity, it is important to have precise spatial matching between the input laser field and the cavity mode. Therefore, the laser beam profiles emerging from the two fiber collimators (TC06APC-1550 from Thorlabs) were first measured using a scanning slit optical beam profiler (BP209-IR/M from Thorlabs).

At different positions in front of the fiber collimator, the beam profile was measured, and the beam diameter at a clipping level of 13.5% (which corresponds to $1/e^2$) was noted down. The profile measurements were averaged over 10 frames. Throughout the measured distances, the beam profiles exhibited a decent Gaussian profile, with more than 90% of Gaussian intensity in both X- and Y- directions, according to the fit from the instrument. The recorded data were then fitted with a Gaussian profile using the Python package 'OpenQlab' [102], and the waist size, along with the waist position, was obtained.

Fig. C.1 shows the measurement data in both X- and Y- directions along with the Gaussian fits. One of the fiber collimators has a waist size of 557.0 ± 3.4 µm in the X-direction at a distance of 1.713 ± 0.508 cm in front of the collimator housing, and 584.1 ± 5.5 µm in the Y-direction at 4.191 ± 0.956 cm. The corresponding values for the other collimator are 544.5 ± 4.6 µm at 13.556 ± 0.824 cm and 566.9 ± 6.2 µm at 16.150 ± 1.268 cm, respectively.

C.2 Fundamental mode profiles of the cavities

The chosen cavities for the HS demonstration are impedance-matched cavities with a flat-curved configuration. Both cavities were assembled using commercial high-reflective mirrors with a reflectivity of 99 ± 0.5%. This results in a cavity finesse of about 312, which is relatively small

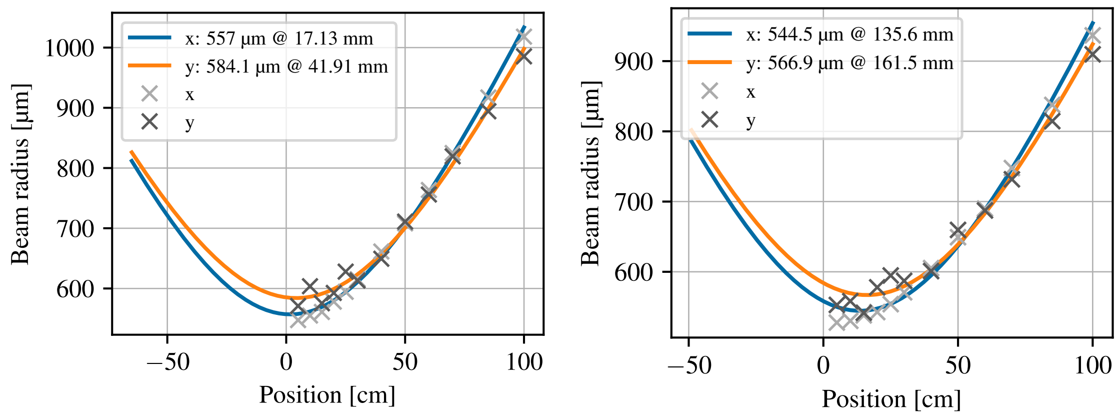

(a) First fiber collimator

(b) Second fiber collimator

Figure C.1: Beam profile measurements of the two laser fields emerging from the fiber collimators used in the table-top HS experiment. The beam radius was measured at different positions, and the data were fitted with a Gaussian profile. The waist radius and waist position in both X- and Y- directions were obtained. Although the collimators have the same specifications, a significant difference in the waist position was observed between them. This was confirmed through multiple measurements and was suspected to originate from differences in the triplet lens assembly.

compared to the nominal values used in other high-precision experiments[1]. For our purposes, cavities with such a finesse value were sufficient, as analyzed in Sec. 3.3. The input mirror of the cavity was a flat mirror, while the end mirror had a radius of curvature of 300 mm.

When considering the optical cavity parameters, it is important to ensure that the combination of reflectivities and cavity lengths results in a stable resonator. Only a stable cavity will facilitate the periodic refocusing of the intracavity beam. For a cavity of length L, an input mirror reflectivity R_1, and an end mirror reflectivity R_2, the stability criterion is given by

$$0 \leq \left(1 - \frac{L}{R_1}\right)\left(1 - \frac{L}{R_2}\right) = g_1 \cdot g_2 = g \leq 1. \tag{C.1}$$

Based on the chosen mirror reflectivities, a 5 cm long cavity will have a stability g-factor of 0.83, while a 10 cm cavity will have a g-factor of 0.67.

In the flat-curved configuration, the fundamental resonance mode (TEM_{00}) of the cavity has its waist positioned at the input mirror. For the 5 cm long cavity, the required beam radius at the input mirror was found to be 234.865 µm. In the case of the 10 cm long cavity, this parameter was 264.149 µm. These calculations were carried out using a linear cavity simulator called ‘jsCav’[2].

[1] It is quite common to employ cavities with finesse values in the tens of thousands [85, 167].

[2] S. Steinlechner, `https://sestei.github.io/jscav/`; accessed on January 29, 2025.

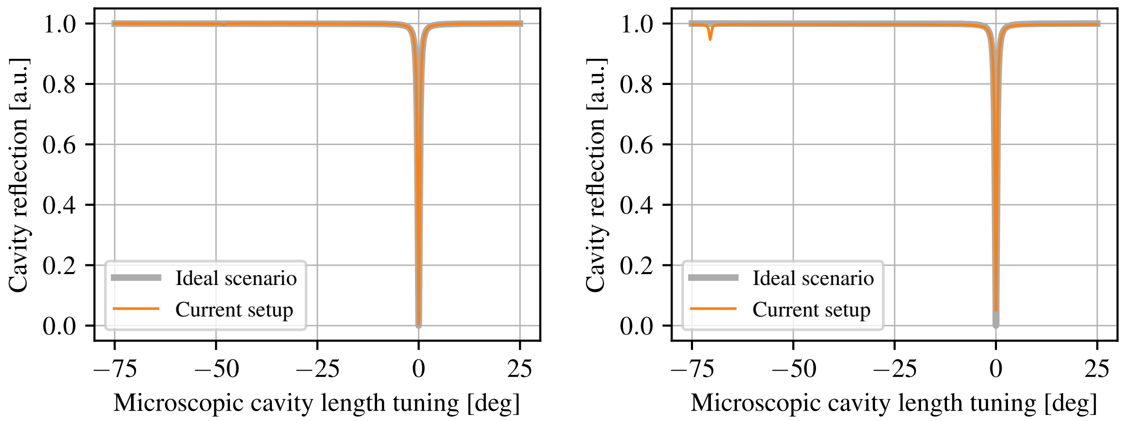

(a) 5 cm long cavity

(b) 10 cm long cavity

Figure C.2: Simulated responses of the table-top cavities under perfect mode-matched conditions and in accordance with the current mode-transfer plan. A single laser beam with the measured beam profile of the fiber collimator is used for the simulation. In the case of the 5 cm long cavity, near-ideal mode overlap can be achieved. For the 10 cm long cavity, the mode-transfer plan will prevent achieving very high fringe visibility.

C.3 Mode matching between input beam and cavity

Now, after determining the mode profile of the input laser beam and the TEM_{00} mode of the optical cavity, the task is to convert the former so that it matches the latter. This process is called mode-matching, and such mode transfer can be achieved using a lens or a combination of lenses positioned correctly between the fiber collimator and the cavity. In our case, the mode-matching parameters were obtained with the assistance of a simulator called 'Just another mode matching tool' (JamMT)[3].

Placing a plano-convex lens with a focal length of 300 mm at a distance of 2.4 cm from the first fiber collimator will convert the input beam profile into a beam with a waist size of 240.012 µm at 26.8 cm from the collimator. This mode transfer can be utilized for the 5 cm long cavity to achieve a mode overlap of more than 99%. Similarly, for the 10 cm long cavity, a plano-convex lens with the same focal length is positioned 30 cm away from the other collimator to obtain a waist size of 264.26 µm at 26.7 cm after the lens.

Once the mode-matching requirements for both cavity setups were investigated, such an optical setup was simulated using Finesse3 [77]. The simulated results are plotted in Fig. C.2 against the ideal scenario in the case of perfect mode-matching. As mentioned above, the planned mode-matching is very close to the ideal condition (i.e., 100% mode overlap) for the 5 cm long cavity. However, in the case of the 10 cm long cavity, the planned mode-transfer allows a higher-order mode to resonate inside the cavity as well. This is the result of imperfect mode-matching. Nevertheless, the fundamental mode is prominent enough to consider this option and proceed to cavity alignment.

It is worth noting that the mode-matching in the case of the 10 cm long cavity can be improved, e.g., by using two lenses or a lens with a different focal length. In the actual experiment, it was a compromise to use the already existing lens. The mode-transfer plans chosen here are not the only possible solutions; they are just one among many possible ways to achieve mode-matching. Different

[3] `https://www.sr.bham.ac.uk/dokuwiki/doku.php?id=geosim:jammt`; accessed on January 29, 2025.

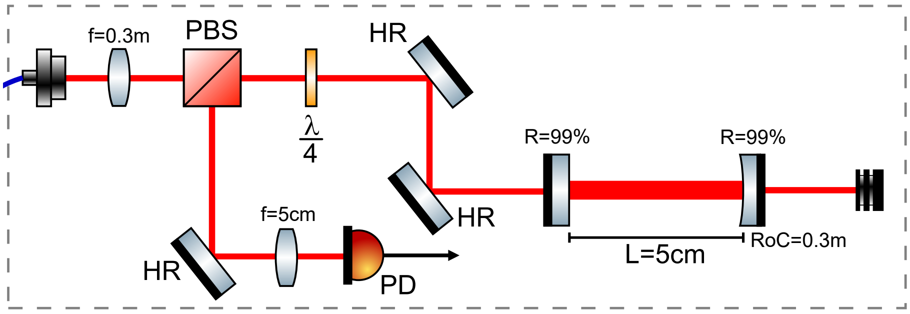

Figure C.3: Optical setup for building a table-top cavity. In this setup, free-space propagation using polarization-based optics is chosen to separate the cavity reflection from the input beam. The placement of the optics is not drawn to scale. (PBS: polarizing beam-splitter, HR: high-reflective mirror, RoC: radius of curvature, PD: photodetector.)

solutions will result in different achievable fringe visibility.

C.4 Alignment steps

With all the requirements for setting up the optical cavities satisfied, the next step is to align the cavities constructed using commercial mirrors in the mirror mounts. The goal is to achieve the desired fringe visibility for the current setups (see Fig. C.2). To properly steer the outgoing beam from the fiber collimator into the optical cavity, two high-reflective mirrors with reflectivity > 99% are introduced between the collimator and the cavity. A high-reflective mirror and a focusing lens are placed in front of the photodetector to efficiently couple the cavity reflection onto the photodetector and to avoid any possible clipping of the beam. The optical setup is sketched in Fig. C.3.

Aligning such cavities built with commercial mirror mounts is quite a tedious task. The procedure followed during our alignment runs can be summarized as follows:

- The laser output is nearly linearly polarized; hence, the fiber collimator is rotated so that the PBS treats the input as a p-polarized beam.
- The polarization-based optics (QWP and PBS) are removed from the setup temporarily.
- Proper beam-walking must be ensured over a sufficient distance using the beam-steering mirrors.
- Only then is the input mirror of the cavity introduced, ensuring that the reflection directly couples back to the fiber collimator. For this purpose, it is beneficial to have the ability to separate the input beam from the cavity reflection behind the collimator, e.g., by employing a fiber circulator. Input mirror alignment should be carried out to maximize this reflection.
- The end mirror of the cavity is placed and adjusted until the required fringe visibility is achieved.
- If the cavity reflection is separated from the cavity input beam using polarization-based optics, those must be re-included in the setup.
- Minor adjustments of the beam-steering mirrors and cavity mirrors might be required to further improve the cavity alignment.

APPENDIX D

Configuration of the RF PLLs of ZCU111

To clock the RF data converter tiles of the RFSoC, the ZCU111 evaluation board features two types of PLLs: LMK04208 and LMX2594. The overall clocking structure is shown in Fig. 4.2. In this appendix, information on how to configure the aforementioned PLLs is provided.

D.1 Configuring the reference PLL LMK04208

In our implementation, the LMK04208 is set to operate in dual PLL mode with an internal voltage-controlled oscillator (VCO). The LMK04208 is configured to receive a 10 MHz external reference via SMA input (J109 on the ZCU111). This clock input is used for PLL1, resulting in an external voltage-controlled crystal oscillator (VCXO) frequency of 122.88 MHz. When this signal is fed to PLL2, the internal VCO frequency can be set to 2949.12 MHz. This value serves as the reference for clock distribution within the LMK04208. With a clock divider value of 12, it is then possible to set ‘CLKout4’ to 245.76 MHz, operating in low-voltage differential signaling (LVDS) format.

Such a configuration of the LMK04208 was carried out using ‘TICS Pro’ software from Texas Instruments. The setup values used to achieve the expected VCO frequency are shown in Fig. D.1. The hexadecimal register values corresponding to this configuration were then obtained from TICS Pro and used to configure the LMK04208 reference PLL during the boot sequence of the ZCU111.

D.2 Configuring the RF-PLL LMX2594

A very similar configuration procedure is carried out for the LMX2594 RF-PLL as well. There are two possible methods of configuration. One involves configuring the RF outputs directly as sample clocks for the RFSoC data converters. The other method is to generate a reference clock for the internal PLLs of data converter tile. We have tested both options and found them to be functional as intended. Our final implementation uses the LMX2594 as a reference for the internal PLLs of the RFSoC data converters.

The input reference to the LMX2594 is the above generated 245.76 MHz output from the LMK04208. This reference oscillator is then used to generate an internal VCO frequency of 8192 MHz in the LMX2594. A channel divider with a value of 16 then sets both the RF outputs to 512 MHz. Fig. D.2 shows this configuration, achieved using TICS Pro.

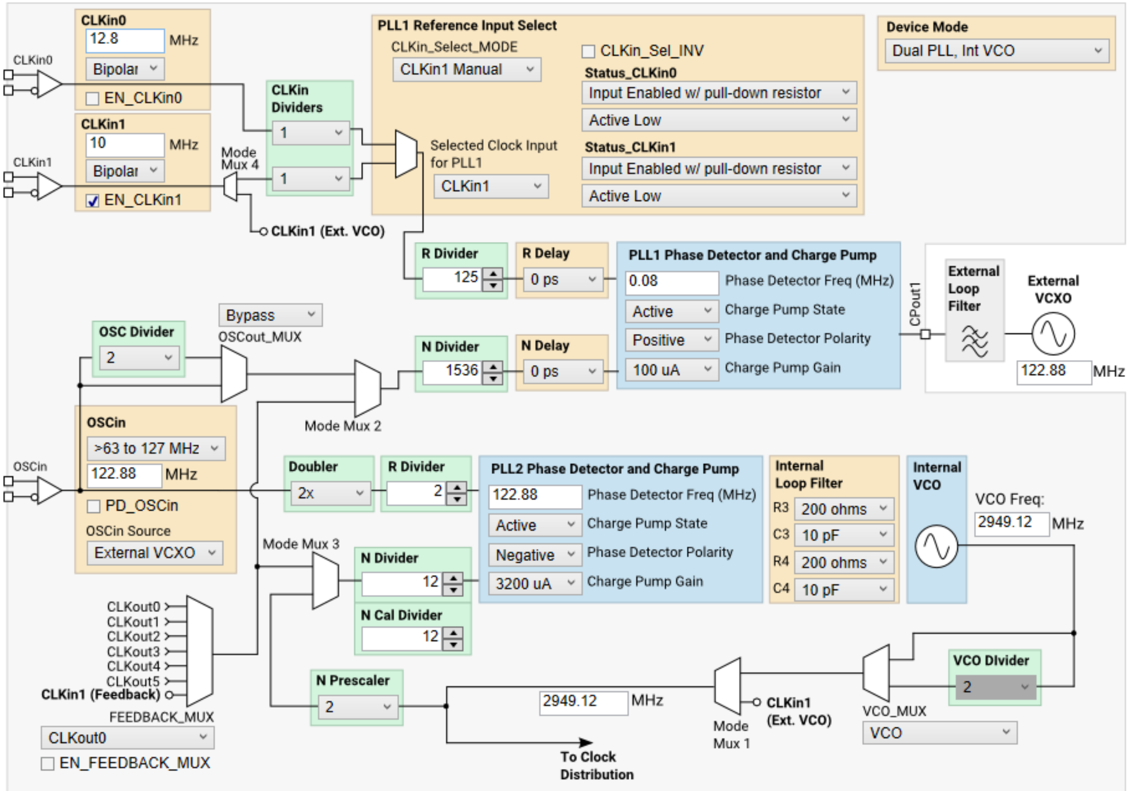

Figure D.1: Configuration of the LMK04208 clock inputs and PLLs. With the shown setup, a 10 MHz reference input results in a VCO frequency of 2949.12 MHz. This frequency is then used to generate a 245.76 MHz output at CLKout4 pin (not shown here).

If different clock reference inputs are expected for different data converter tiles, the configuration values must be adapted accordingly. It is important to note that not all PLL reference clock values are supported by the RF data converter tiles. Depending on the chosen sampling rate, the accepted PLL references are displayed in the Vivado Design Suite while configuring the RF data converter IP core. The outputs of the LMX2594 must be configured accordingly. Once again, the hexadecimal register values corresponding to the configuration of the LMX2594 were obtained from TICS Pro and used during the boot sequence of the ZCU111.

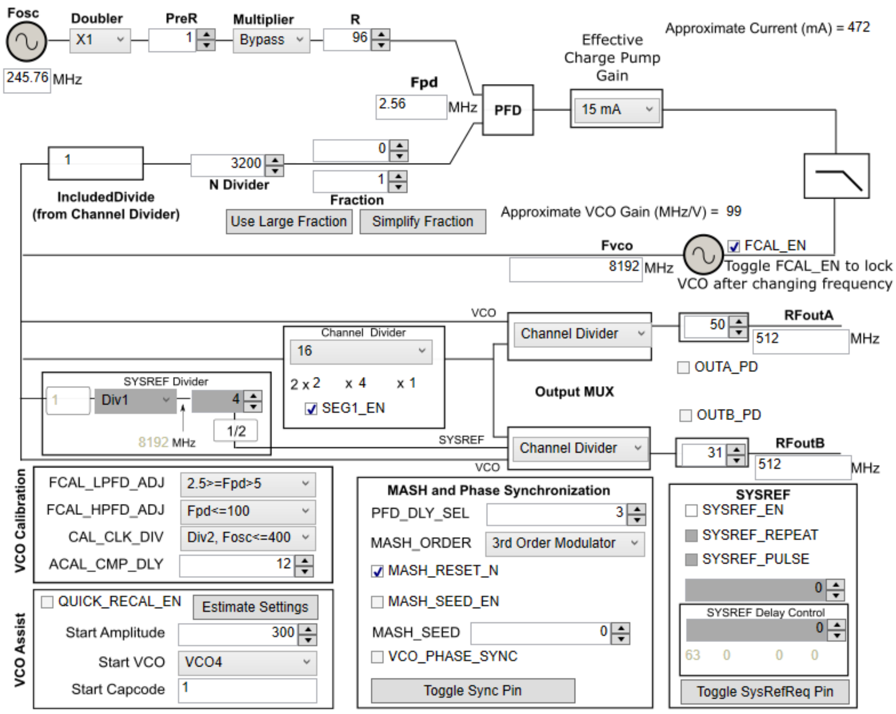

Figure D.2: Configuration of the LMX2594 RF-PLL. With the shown setup, a 245.76 MHz reference input results in a VCO frequency of 8192 MHz. This frequency is then used to generate two RF outputs at 512 MHz.

APPENDIX E

Initial implementation of the GHz Phasemeter

Before developing the final implementation of the GHz Phasemeter, an initial working implementation was explored using the internal NCO of the RF data converter. This appendix provides a brief explanation of that implementation.

The RF-ADCs of the RFSoC come with digital down converters (DDCs), which include programmable decimation rates, an NCO, and a complex mixer [135]. The two ADC channels on a single tile can be configured individually for real input signals or, as a pair, for I/Q input signals. Although the RF-ADCs have 12-bit resolution, the data stream is MSB-aligned to a 16-bit DSP data path at the output. Each digital complex mixer uses a 48-bit NCO per channel.

The RF-ADC (from Tile 226) is configured to receive a real input signal, which is then converted into I/Q signals by the internal digital complex mixer. The data stream is set to have an 8× decimation, providing one 16-bit sample per AXI4-Stream cycle. The ADC operates at its maximum sample rate (4.096 GSPS) and is clocked by the internal PLL, which receives a 512 MHz reference from the LMX2594 RF-PLL. The RF data converter tile offers a user-configurable output clock from the core to the user logic, set to a maximum of 256 MHz. The required 512 MHz clock for the AXI4-Stream interface of the RF-ADC data is generated from this output using the LogiCORE™IP Clocking Wizard. The same 512 MHz clock is used to operate the PL.

The 16-bit long Q signal coming from the RF data converter IP core is resized to a 32-bit signal. It is then passed through a PI controller, producing a 48-bit output. The integrator functionality is realized by accumulating the servo input and introducing gain through appropriate bit-shifting. The servo output is added to the initial 48-bit PIR estimation to update the NCO frequency of the RF data converter tile.

The initial PIR and PI gains can be set externally. 32-bit registers are reserved for them in RAM, accessible by both the PS and PL. Data acquisition from the PLL is handled in a similar fashion. However, the actual DSP involves filtering the data, synchronizing between different clock domains, etc.

While the internal NCO simplifies the user logic, real-time updating of the NCO value on one converter register typically takes about 45 clock cycles of the AXI4-Lite clock[1]. This hardware constraint directly limits the speed at which we can run the PLL.

[1] This clock is common to all tiles and should be driven by the system CPU used for the AXI4-Stream interconnect connected to the AXI4-Lite interface [135].

Despite this limitation, we successfully tested the acquisition and tracking of the input signal frequency using tones from external signal generators. However, since this is not the final version of the GHz Phasemeter that we aim to achieve, further performance measurements were not conducted.

APPENDIX F

PR-59 temperature controller: communication and data recording

The figure-of-merit (FOM) of athermal glass samples was planned to be characterized by tuning the temperature of one sample while keeping the other at a fixed temperature. To enable efficient temperature control, the choice was made to use TEC assemblies with PR-59 temperature controllers. Although a GUI from the manufacturer is available to control and log data, an alternative method must be implemented to achieve temperature modulation.

The PR-59 temperature controller can be controlled using serial command interface (SCI). Thus, using available Python packages such as 'Serial,' it is possible to establish communication with the PR-59.

F.1 Temperature sensor readout

In our setup, the primary sensor at the slot Sensor 1 was a TC-NTC-1, whereas Sensor 2 and Sensor 3 were of type TC-NTC-2. The Python code snippet below was used to simultaneously record temperature data from the three sensors connected to the PR-59.

```
import serial
import time
from datetime import datetime

def Read_Temp_Parallel(t, exp_name = None, file_name = "TempData.txt"):
  """
  Measures the temperature using three sensors in parallel for a given
  time interval 't' (in min) and the data along with time stamps are
  written to a .txt file
  """
  ser = serial.Serial()
  ser.baudrate = 115200
  ser.port = "COM3" # Serial port to which PR-59 is connected
  ser.open()

  start_time = datetime.now().strftime("%Y-%m-%d %H:%M:%S")
  with open(file_name,"w") as fh:
    fh.write("% This is a comment at the beginning of the file \n")
```

```
    fh.write(f"% Starting time of the measurement: {start_time} \n")
    fh.write(f"% Measurement Name: {exp_name} \n")
    fh.write("Time[s] \t Temp1[*C] \t Temp2[*C] \t Temp3[*C] \n")

    t_end = time.time() + t * 60

    while time.time() <= t_end:
      # Setting NTC Steinhart mode coefficients for TC-NTC-2 @Temp2
      ser.write(b"$R62=1.037284e-003\r\n")
      ser.write(b"$R63=2.331723e-004\r\n")
      ser.write(b"$R64=8.389566e-008\r\n")

      # Setting NTC Steinhart mode coefficients for TC-NTC-2 @Temp3
      ser.write(b"$R65=1.037284e-003\r\n")
      ser.write(b"$R66=2.331723e-004\r\n")
      ser.write(b"$R67=8.389566e-008\r\n")

      ser.write(b"$R100?\r\n")
      ser.write(b"$R101?\r\n")
      ser.write(b"$R102?\r\n")
      out = b""
      time.sleep(1)
      while ser.inWaiting() > 0:
        out += ser.readline(1)
      result = out.decode().rstrip().split()[13::3]
      fh.write(f"{time.time()}\t {result[0]}\t {result[1]}\t {result[2]}\n")
  ser.close()

if __name__ == "__main__":
  exp_name = "Parallel measurement for 3 min"
  file_name = "RoomTempParallel.txt"
  t = 3 # measurement time in min
  Read_Temp_Parallel(t, exp_name, file_name)
```

F.2 Temperature modulation and readout

The code snippet below handles the temperature modulation and sensor value readout. Initially, a set point of 22.5 °C is set, the controller is turned on, and data logging begins. A custom function samples a single period of a sine wave with an amplitude of 0.1 (corresponding to 100 mK) and a frequency of 50 mHz, resulting in 37 sampled points. These points are used as continuous set points to achieve the desired modulation. The step length—the time difference between two set point values—can be altered by adjusting the ‘sleep’ value in line 51. At the end of the given measurement time, data logging is stopped, and the controller is turned off.

```
1 import serial
2 import time
3 import numpy as np
4 from datetime import datetime
5
6 def sinfunc(c=20, f=0.005):
7   A = 0.1 # modulation amplitude 100mK
8   t = np.linspace(0,1/f,37)
```

```
    return A * np.sin(2 * np.pi * f * t) + c

ser = serial.Serial("COM4", 115200) # Serial port to which PR-59 is connected
ser.open()

temp_point = 22.5

set_points = sinfunc(c=temp_point)[:-1]

ser.write(b"$R13=6.0e+00\r\n") # Regulator mode (6=PID)
ser.write(b"$W\r\n")
time.sleep(3)
ser.write(b"$RR\r\n")

set_point_byte = bytes("$R0=" + str(temp_point) + "\r\n", "utf-8")
ser.write(set_point_byte)
out = b""
time.sleep(1)
while ser.inWaiting() > 0:
  out += ser.readline(1)
print(out.decode())
print("-#-#-#-#-")

with open("file_name.csv", "w") as fh:
  datetime_now = datetime.now().strftime("%Y-%m-%d %H:%M:%S")
  fh.write(f"% Starting time of the measurement: {datetime_now} \n")
  fh.write(f"% Measurement Name: 1s time steps \n")
  fh.write(f"View ERR Mode Tc Ta1 Ta2 Tr Ta Tp Ti Td TLP_A TLP_B \n")
  ser.write(b"$A3\r\n") # Start data-logging

  exp_start_time = time.time()
  meas_time = 15 * 60 # seconds
  exp_end_time = exp_start_time + meas_time
  i = 0
  while time.time() <= exp_end_time:
    set_point = set_points[i % len(set_points)]

    set_point_byte = bytes("$R0=" + str(set_point) + "\r\n", "utf-8")
    ser.write(set_point_byte)

    start_time = time.time()
    while time.time() < start_time + 1:
      time.sleep(1)
      out = b""
      while ser.inWaiting() > 0:
        out += ser.readline()
      result = out.decode().strip().split("\r\n")

      if i==0:
        data = result[2:]
      else:
        data = result

      for j in range(len(data)):
        fh.write(f"{data[j][1:-1]}\n")
```

```

    i += 1

ser.write(b"$A\r\n") # Stop data-logging

ser.write(b"$R13=0.0e+00\r\n") # Turn-off the controller
ser.write(b"$W\r\n")
out = b""
time.sleep(3)
while ser.inWaiting() > 0:
  out += ser.readline(1)
print(out.decode())
print("-#-#-#-#-")

ser.close()
```

The above code snippet does not include the setup for the PID gains or NTC Steinhart mode coefficients for TC-NTC-2. If configuring these parameters is anticipated, the following code snippet should be included.

```
ser.write(b"$R1=33.0e+00\r\n") # kP
ser.write(b"$R2=0.45e+00\r\n") # kI
ser.write(b"$R3=500.0e+00\r\n") # kD
ser.write(b"$R4=2.00e+00\r\n") # kTr (lowpass Tr) 0.05 to turn-off
ser.write(b"$R5=1.00e+00\r\n") # kTe (lowpass Te) 0.05 to turn-off
ser.write(b"$R7=0.0e+00\r\n") # dead-band
ser.write(b"$R10=2.6e+00\r\n") # cooling gain
ser.write(b"$R11=1.0e+00\r\n") # heating gain

# Setting NTC Steinhart mode coefficients for TC-NTC-2 @Temp2
ser.write(b"$R62=1.037284e-003\r\n")
ser.write(b"$R63=2.331723e-004\r\n")
ser.write(b"$R64=8.389566e-008\r\n")

ser.write(b"$RW\r\n")
time.sleep(3)
ser.write(b"$RR\r\n")

out = b""
time.sleep(1)
while ser.inWaiting() > 0:
  out += ser.readline(1)
print(out.decode())
print("-#-#-#-#-")
```

Bibliography

[1] B. P. Abbott et al., *Observation of Gravitational Waves from a Binary Black Hole Merger*, Physical Review Letters **116** (2016), ISSN: 1079-7114, URL: `http://dx.doi.org/10.1103/PhysRevLett.116.061102`.

[2] A. Einstein, *Die Grundlage der allgemeinen Relativitätstheorie*, Annalen der Physik **354** (1916) 769, ISSN: 1521-3889, URL: `http://dx.doi.org/10.1002/andp.19163540702`.

[3] B. P. Abbott et al., *GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs*, Physical Review X **9** (2019), ISSN: 2160-3308, URL: `http://dx.doi.org/10.1103/PhysRevX.9.031040`.

[4] B. P. Abbott et al., *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*, Physical Review Letters **119** (2017), ISSN: 1079-7114, URL: `http://dx.doi.org/10.1103/PhysRevLett.119.161101`.

[5] R. Abbott et al., *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run*, Physical Review X **13** (2023), ISSN: 2160-3308, URL: `http://dx.doi.org/10.1103/PhysRevX.13.041039`.

[6] G. Agazie et al., *The NANOGrav 15 yr Data Set: Constraints on Supermassive Black Hole Binaries from the Gravitational-wave Background*, The Astrophysical Journal Letters **952** (2023) L37, ISSN: 2041-8213, URL: `http://dx.doi.org/10.3847/2041-8213/ace18b`.

[7] A. Berlin et al., *Detecting high-frequency gravitational waves with microwave cavities*, Physical Review D **105** (2022), ISSN: 2470-0029, URL: `http://dx.doi.org/10.1103/PhysRevD.105.116011`.

[8] J. Aasi et al., *Advanced LIGO*, Classical and Quantum Gravity **32** (2015) 074001, ISSN: 1361-6382, URL: `http://dx.doi.org/10.1088/0264-9381/32/7/074001`.

[9] F. Acernese et al., *Advanced Virgo: a second-generation interferometric gravitational wave detector*, Classical and Quantum Gravity **32** (2014) 024001, ISSN: 1361-6382, URL: `http://dx.doi.org/10.1088/0264-9381/32/2/024001`.

[10] H. Lück et al., *The upgrade of GEO 600*, Journal of Physics: Conference Series **228** (2010) 012012, ISSN: 1742-6596, URL: http://dx.doi.org/10.1088/1742-6596/228/1/012012.

[11] T. Akutsu et al., *Overview of KAGRA: Detector design and construction history*, Progress of Theoretical and Experimental Physics **2021** (2020), ISSN: 2050-3911, URL: http://dx.doi.org/10.1093/ptep/ptaa125.

[12] B. P. Abbott et al., *Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA*, Living Reviews in Relativity **23** (2020), ISSN: 1433-8351, URL: http://dx.doi.org/10.1007/s41114-020-00026-9.

[13] R. X. Adhikari et al., *A cryogenic silicon interferometer for gravitational-wave detection*, Classical and Quantum Gravity **37** (2020) 165003, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/1361-6382/ab9143.

[14] M. Barsuglia et al., *Virgo nEXT Preliminary R&D Plan*, (2023), URL: https://indico.ego-gw.it/event/636/contributions/5562/attachments/2953/5207/V_nEXT_RnD-prelim-plan.pdf.

[15] M. Punturo et al., *The Einstein Telescope: a third-generation gravitational wave observatory*, Classical and Quantum Gravity **27** (2010) 194002, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/27/19/194002.

[16] D. Reitze et al., *Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO*, (2019), URL: https://arxiv.org/abs/1907.04833.

[17] E. D. Hall et al., *Gravitational-wave physics with Cosmic Explorer: Limits to low-frequency sensitivity*, Physical Review D **103** (2021), ISSN: 2470-0029, URL: http://dx.doi.org/10.1103/PhysRevD.103.122004.

[18] C. J. Moore, R. H. Cole, and C. P. L. Berry, *Gravitational-wave sensitivity curves*, Classical and Quantum Gravity **32** (2014) 015014, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/32/1/015014.

[19] K. Danzmann and A. R diger, *LISA technology concept, status, prospects*, Classical and Quantum Gravity **20** (2003) S1, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/20/10/301.

[20] J. I. Thorpe et al., *The Laser Interferometer Space Antenna: Unveiling the Millihertz Gravitational Wave Sky*, Bulletin of the AAS **51** (2019), URL: https://baas.aas.org/pub/2020n7i077.

[21] M. Armano et al., *Sub-Femto-g Free Fall for Space-Based Gravitational Wave Observatories: LISA Pathfinder Results*, Physical Review Letters **116** (2016), ISSN: 1079-7114, URL: http://dx.doi.org/10.1103/PhysRevLett.116.231101.

[22] J. Harms et al., *Lunar Gravitational-wave Antenna*, The Astrophysical Journal **910** (2021) 1, ISSN: 1538-4357, URL: http://dx.doi.org/10.3847/1538-4357/abe5a7.

[23] J. V. van Heijningen et al., *The payload of the Lunar Gravitational-wave Antenna*, Journal of Applied Physics **133** (2023), ISSN: 1089-7550, URL: `http://dx.doi.org/10.1063/5.0144687`.

[24] P. Ajith et al., *The Lunar Gravitational-wave Antenna: mission studies and science case*, Journal of Cosmology and Astroparticle Physics **2025** (2025) 108, ISSN: 1475-7516, URL: `http://dx.doi.org/10.1088/1475-7516/2025/01/108`.

[25] W.-R. Hu and Y.-L. Wu, *The Taiji Program in Space for gravitational wave physics and the nature of gravity*, National Science Review **4** (2017) 685, ISSN: 2053-714X, URL: `http://dx.doi.org/10.1093/nsr/nwx116`.

[26] J. Luo et al., *TianQin: a space-borne gravitational wave detector*, Classical and Quantum Gravity **33** (2016) 035010, ISSN: 1361-6382, URL: `http://dx.doi.org/10.1088/0264-9381/33/3/035010`.

[27] K. Jani and A. Loeb, *Gravitational-wave Lunar Observatory for Cosmology*, Journal of Cosmology and Astroparticle Physics **2021** (2021) 044, ISSN: 1475-7516, URL: `http://dx.doi.org/10.1088/1475-7516/2021/06/044`.

[28] F. Matichard et al., *Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance*, Classical and Quantum Gravity **32** (2015) 185003, ISSN: 1361-6382, URL: `http://dx.doi.org/10.1088/0264-9381/32/18/185003`.

[29] F. Matichard et al., *Advanced LIGO two-stage twelve-axis vibration isolation and positioning platform. Part 1: Design and production overview*, Precision Engineering **40** (2015) 273, ISSN: 0141-6359, URL: `http://dx.doi.org/10.1016/j.precisioneng.2014.09.010`.

[30] F. Matichard et al., *Advanced LIGO two-stage twelve-axis vibration isolation and positioning platform. Part 2: Experimental investigation and tests results*, Precision Engineering **40** (2015) 287, ISSN: 0141-6359, URL: `http://dx.doi.org/10.1016/j.precisioneng.2014.11.010`.

[31] J. J. Carter, *High precision, compact inertial sensors for use in gravitational wave detectors*, PhD thesis, 2024, URL: `https://www.repo.uni-hannover.de/handle/123456789/16013`.

[32] S. Biscans et al., *Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors*, Classical and Quantum Gravity **35** (2018) 055004, ISSN: 1361-6382, URL: `http://dx.doi.org/10.1088/1361-6382/aaa4aa`.

[33] B. Sathyaprakash et al., *Scientific objectives of Einstein Telescope*, Classical and Quantum Gravity **29** (2012) 124013, ISSN: 1361-6382, URL: `http://dx.doi.org/10.1088/0264-9381/29/12/124013`.

[34] H. Tariq et al., *The linear variable differential transformer (LVDT) position sensor for gravitational wave interferometer low-frequency controls*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **489** (2002) 570, ISSN: 0168-9002, URL: `http://dx.doi.org/10.1016/S0168-9002(02)00802-1`.

[35] V. Josselin, P. Touboul, and R. Kielbasa, *Capacitive detection scheme for space accelerometers applications*, Sensors and Actuators A: Physical **78** (1999) 92, ISSN: 0924-4247, URL: http://dx.doi.org/10.1016/S0924-4247(99)00227-7.

[36] S.-C. Wang, B.-R. Xie, and S.-M. Huang, *Design and Analysis of Small Size Eddy Current Displacement Sensor*, Sensors **22** (2022) 7444, ISSN: 1424-8220, URL: http://dx.doi.org/10.3390/s22197444.

[37] Y. Ye et al., *A Review on Applications of Capacitive Displacement Sensing for Capacitive Proximity Sensor*, IEEE Access **8** (2020) 45325, ISSN: 2169-3536, URL: http://dx.doi.org/10.1109/ACCESS.2020.2977716.

[38] S. Boonsang and R. Dewhurst, *A sensitive electromagnetic acoustic transducer for picometer-scale ultrasonic displacement measurements*, Sensors and Actuators A: Physical **127** (2006) 345, ISSN: 0924-4247, URL: http://dx.doi.org/10.1016/j.sna.2005.12.029.

[39] B. El Mansouri et al., *High-resolution MEMS inertial sensor combining large-displacement buckling behaviour with integrated capacitive readout*, Microsystems & Nanoengineering **5** (2019), ISSN: 2055-7434, URL: http://dx.doi.org/10.1038/s41378-019-0105-y.

[40] S. A. Zotov, B. R. Simon, A. A. Trusov, and A. M. Shkel, *High Quality Factor Resonant MEMS Accelerometer With Continuous Thermal Compensation*, IEEE Sensors Journal **15** (2015) 5045, ISSN: 1558-1748, URL: http://dx.doi.org/10.1109/JSEN.2015.2432021.

[41] M. Armano et al., *Sensor Noise in LISA Pathfinder: In-Flight Performance of the Optical Test Mass Readout*, Physical Review Letters **126** (2021), ISSN: 1079-7114, URL: http://dx.doi.org/10.1103/PhysRevLett.126.131103.

[42] J. Watchi, S. Cooper, B. Ding, C. M. Mow-Lowry, and C. Collette, *Contributed Review: A review of compact interferometers*, Review of Scientific Instruments **89** (2018), ISSN: 1089-7623, URL: http://dx.doi.org/10.1063/1.5052042.

[43] L. Carbone et al., *Sensors and actuators for the Advanced LIGO mirror suspensions*, Classical and Quantum Gravity **29** (2012) 115005, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/29/11/115005.

[44] S. J. Cooper et al., *Sensors and actuators for the advanced LIGO A+ upgrade*, Review of Scientific Instruments **94** (2023), ISSN: 1089-7623, URL: http://dx.doi.org/10.1063/5.0117605.

[45] T. Akutsu et al., *Compact integrated optical sensors and electromagnetic actuators for vibration isolation systems in the gravitational-wave detector KAGRA*, Review of Scientific Instruments **91** (2020), ISSN: 1089-7623, URL: http://dx.doi.org/10.1063/5.0022242.

[46] S. J. Cooper et al.,
A compact, large-range interferometer for precision measurement and inertial sensing,
Classical and Quantum Gravity **35** (2018) 095007, ISSN: 1361-6382,
URL: http://dx.doi.org/10.1088/1361-6382/aab2e9.

[47] G. M. B. Bouricius and S. F. Clifford, *An Optical Interferometer Using Polarization Coding to Obtain Quadrature Phase Components*, Review of Scientific Instruments **41** (1970) 1800, ISSN: 1089-7623, URL: http://dx.doi.org/10.1063/1.1684414.

[48] J. van Dongen et al., *Reducing control noise in gravitational wave detectors with interferometric local damping of suspended optics*,
Review of Scientific Instruments **94** (2023), ISSN: 1089-7623,
URL: http://dx.doi.org/10.1063/5.0144865.

[49] S. J. Cooper et al., *Interferometric sensing of a commercial geophone*,
Classical and Quantum Gravity **39** (2022) 075023, ISSN: 1361-6382,
URL: http://dx.doi.org/10.1088/1361-6382/ac595c.

[50] O. Gerberding, *Deep frequency modulation interferometry*, Optics Express **23** (2015) 14753, ISSN: 1094-4087, URL: http://dx.doi.org/10.1364/OE.23.014753.

[51] J. Smetana et al., *Compact Michelson Interferometers with Subpicometer Sensitivity*,
Physical Review Applied **18** (2022), ISSN: 2331-7019,
URL: http://dx.doi.org/10.1103/PhysRevApplied.18.034040.

[52] O. Gerberding and K.-S. Isleif, *Ghost Beam Suppression in Deep Frequency Modulation Interferometry for Compact On-Axis Optical Heads*, Sensors **21** (2021) 1708,
ISSN: 1424-8220, URL: http://dx.doi.org/10.3390/s21051708.

[53] T. Eckhardt and O. Gerberding, *An analytic, efficient and optimal readout algorithm for compact interferometers based on deep frequency modulation*, Scientific Reports **14** (2024), ISSN: 2045-2322, URL: http://dx.doi.org/10.1038/s41598-024-70392-9.

[54] T. Eckhardt and O. Gerberding,
Noise Limitations in Multi-Fringe Readout of Laser Interferometers and Resonators,
Metrology **2** (2022) 98, ISSN: 2673-8244,
URL: http://dx.doi.org/10.3390/metrology2010007.

[55] L. S. Collaboration, *LSC Instrument Science White Paper (2024 edition)*, tech. rep., 2024,
URL: https://dcc.ligo.org/T2300411-v1.

[56] L. Prokhorov et al.,
Design and sensitivity of a 6-axis seismometer for gravitational wave observatories,
Physical Review D **109** (2024), ISSN: 2470-0029,
URL: http://dx.doi.org/10.1103/PhysRevD.109.042007.

[57] D. J. McManus et al.,
Mechanical characterisation of the TorPeDO: a low frequency gravitational force sensor,
Classical and Quantum Gravity **34** (2017) 135002, ISSN: 1361-6382,
URL: http://dx.doi.org/10.1088/1361-6382/aa7103.

[58] S. S. Y. Chua et al.,
The torsion pendulum dual oscillator for low-frequency Newtonian noise detection,
Applied Physics Letters **122** (2023), ISSN: 1077-3118,
URL: http://dx.doi.org/10.1063/5.0145092.

[59] M. Beker et al., *The Rasnik 3-point optical alignment system*,
Journal of Instrumentation **14** (2019) P08010, ISSN: 1748-0221,
URL: http://dx.doi.org/10.1088/1748-0221/14/08/P08010.

[60] H. van der Graaf et al., *The ultimate performance of the Rasnik 3-point alignment system*,
Nuclear Instruments and Methods in Physics Research Section A: Accelerators,
Spectrometers, Detectors and Associated Equipment **1050** (2023) 168160, ISSN: 0168-9002,
URL: http://dx.doi.org/10.1016/j.nima.2023.168160.

[61] J. J. Carter, P. Birckigt, O. Gerberding, and S. M. Koehlenbeck,
Compact inertial sensors for measuring external disturbances of physics experiments,
Scientific Reports **14** (2024), ISSN: 2045-2322,
URL: http://dx.doi.org/10.1038/s41598-024-68623-0.

[62] J. J. Carter, P. Birckigt, O. Gerberding, and S. M. Koehlenbeck,
Designing gram-scale resonators for precision inertial sensors,
Physical Review Applied **22** (2024), ISSN: 2331-7019,
URL: http://dx.doi.org/10.1103/PhysRevApplied.22.014045.

[63] M.-K. Zhou et al.,
Note: A three-dimension active vibration isolator for precision atom gravimeters,
Review of Scientific Instruments **86** (2015), ISSN: 1089-7623,
URL: http://dx.doi.org/10.1063/1.4919292.

[64] D. Inaudi et al.,
Development of a displacement sensor for the CERN-LHC superconducting cryodipoles,
Measurement Science and Technology **12** (2001) 887, ISSN: 1361-6501,
URL: http://dx.doi.org/10.1088/0957-0233/12/7/324.

[65] C. C. Speake and S. M. Aston, *An interferometric sensor for satellite drag-free control*,
Classical and Quantum Gravity **22** (2005) S269, ISSN: 1361-6382,
URL: http://dx.doi.org/10.1088/0264-9381/22/10/019.

[66] S. Chalathadka Subrahmanya, C. Darsow-Fromm, and O. Gerberding,
On the Development of an RFSoC-Based Ultra-Fast Phasemeter With GHz Bandwidth,
IEEE Transactions on Instrumentation and Measurement **74** (2025), ISSN: 1557-9662,
URL: http://dx.doi.org/10.1109/TIM.2024.3509586.

[67] S. Chalathadka Subrahmanya, C. Darsow-Fromm, and O. Gerberding, *Integrating high-precision and fringe-scale displacement sensing using heterodyne cavity-tracking*,
Optics Express **33** (2025), ISSN: 1094-4087,
URL: http://dx.doi.org/10.1364/OE.540189.

[68] C. Bond, D. Brown, A. Freise, and K. A. Strain,
Interferometer techniques for gravitational-wave detection,
Living Reviews in Relativity **19** (2016), ISSN: 1433-8351,
URL: http://dx.doi.org/10.1007/s41114-016-0002-8.

[69] W. Koechner, *Solid-State Laser Engineering*, Springer New York, 2006, URL: http://dx.doi.org/10.1007/0-387-29338-8.

[70] P. Kwee et al., *Stabilized high-power laser system for the gravitational wave detector advanced LIGO*, Optics Express **20** (2012) 10617, ISSN: 1094-4087, URL: http://dx.doi.org/10.1364/OE.20.010617.

[71] K. Numata et al., "Laser system development for the LISA (Laser Interferometer Space Antenna) mission," *Solid State Lasers XXVIII: Technology and Devices*, ed. by W. A. Clarkson and R. K. Shori, SPIE, 2019 51, URL: http://dx.doi.org/10.1117/12.2508181.

[72] https://www.toptica.com/products/tunable-diode-lasers/ecdl-dfb-lasers/ctl, [Accessed 29-Jan-2025].

[73] L. E. Roberts, *Internally Sensed Optical Phased Arrays*, PhD thesis, 2016, URL: https://openresearch-repository.anu.edu.au/handle/1885/110523.

[74] P. Fritschel, T. J. Kane, and A. Jeffries, *Frequency fluctuations of a diode-pumped Nd:YAG ring laser*, Optics Letters **14** (1989) 993, ISSN: 1539-4794, URL: http://dx.doi.org/10.1364/OL.14.000993.

[75] M. Suter and P. Dietiker, *Calculation of the finesse of an ideal Fabry–Perot resonator*, Applied Optics **53** (2014) 7004, ISSN: 2155-3165, URL: http://dx.doi.org/10.1364/AO.53.007004.

[76] A. Freise et al., *Frequency-domain interferometer simulation with higher-order spatial modes*, Classical and Quantum Gravity **21** (2004) S1067, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/21/5/102.

[77] D. Brown, S. Rowlinson, S. Leavey, P. Jones, and A. Freise, *Finesse 3*, URL: https://gitlab.com/ifosim/finesse/finesse3.

[78] F. Guzmán Cervantes, L. Kumanchik, J. Pratt, and J. M. Taylor, *High sensitivity optomechanical reference accelerometer over 10 kHz*, Applied Physics Letters **104** (2014), ISSN: 1077-3118, URL: http://dx.doi.org/10.1063/1.4881936.

[79] J. Sanjuan et al., *Simultaneous laser frequency stabilization to an optical cavity and an iodine frequency reference*, Optics Letters **46** (2021) 360, ISSN: 1539-4794, URL: http://dx.doi.org/10.1364/OL.413419.

[80] J. A. Boyd and T. Lahaye, *A basic introduction to ultrastable optical cavities for laser stabilization*, American Journal of Physics **92** (2024) 50, ISSN: 1943-2909, URL: http://dx.doi.org/10.1119/5.0161369.

[81] E. D. Black, *An introduction to Pound–Drever–Hall laser frequency stabilization*, American Journal of Physics **69** (2001) 79, ISSN: 1943-2909, URL: http://dx.doi.org/10.1119/1.1286663.

[82] T. Požar, P. Gregorčič, and J. Možina, *A precise and wide-dynamic-range displacement-measuring homodyne quadrature laser interferometer*, Applied Physics B **105** (2011) 575, ISSN: 1432-0649, URL: http://dx.doi.org/10.1007/s00340-011-4512-5.

[83] Y. Zhang and F. Guzman, *Quasi-monolithic heterodyne laser interferometer for inertial sensing*, Optics Letters **47** (2022) 5120, ISSN: 1539-4794, URL: http://dx.doi.org/10.1364/OL.473476.

[84] J. Lawall and E. Kessler, *Michelson interferometry with 10 pm accuracy*, Review of Scientific Instruments **71** (2000) 2669, ISSN: 1089-7623, URL: http://dx.doi.org/10.1063/1.1150715.

[85] J. Eichholz, D. B. Tanner, and G. Mueller, *Heterodyne laser frequency stabilization for long baseline optical interferometry in space-based gravitational wave detectors*, Physical Review D **92** (2015), ISSN: 1550-2368, URL: http://dx.doi.org/10.1103/PhysRevD.92.022004.

[86] R. W. P. Drever et al., *Laser phase and frequency stabilization using an optical resonator*, Applied Physics B Photophysics and Laser Chemistry **31** (1983) 97, ISSN: 1432-0649, URL: http://dx.doi.org/10.1007/BF00702605.

[87] O. Gerberding, *Phase readout for satellite interferometry*, PhD thesis, 2014, URL: https://www.repo.uni-hannover.de/handle/123456789/8308.

[88] D. Roma-Dollase et al., *Resistive-Based Micro-Kelvin Temperature Resolution for Ultra-Stable Space Experiments*, Sensors **23** (2022) 145, ISSN: 1424-8220, URL: http://dx.doi.org/10.3390/s23010145.

[89] https://www.corning.com/worldwide/en/products/advanced-optics/product-materials/semiconductor-laser-optic-components/ultra-low-expansion-glass.html, [Accessed 29-Jan-2025].

[90] G. M. Harry et al., *Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings*, Classical and Quantum Gravity **19** (2002) 897, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/19/5/305.

[91] https://media.schott.com/api/public/content/41e799d0bf874807a0bb8e702fbb75b5?v=54856406, [Accessed 29-Jan-2025].

[92] K. Craig et al., *Mirror Coating Solution for the Cryogenic Einstein Telescope*, Physical Review Letters **122** (2019), ISSN: 1079-7114, URL: http://dx.doi.org/10.1103/PhysRevLett.122.231102.

[93] M. Principe et al., *Material loss angles from direct measurements of broadband thermal noise*, Physical Review D **91** (2015), ISSN: 1550-2368, URL: http://dx.doi.org/10.1103/PhysRevD.91.022005.

[94] C. S. Sambridge et al., *Subfemtowatt Laser Phase Tracking*, Physical Review Letters **131** (2023), ISSN: 1079-7114, URL: http://dx.doi.org/10.1103/PhysRevLett.131.193804.

[95] D. Shaddock, B. Ware, P. G. Halverson, R. E. Spero, and B. Klipstein, "Overview of the LISA Phasemeter," *AIP Conference Proceedings*, AIP, 2006, URL: http://dx.doi.org/10.1063/1.2405113.

[96] https://seismion.com/schwingungstechnik/, [Accessed 29-Jan-2025].

[97] https://liquidinstruments.com/products/integrated-instruments/phasemeter/, [Accessed 29-Jan-2025].

[98] https://siglentna.com/product/sdg2122x/, [Accessed 29-Jan-2025].

[99] *Sampling, Approximation, and Signal Analysis: Harmonic Analysis in the Spirit of J. Rowland Higgins*, Springer International Publishing, 2023, ISBN: 9783031411304, URL: http://dx.doi.org/10.1007/978-3-031-41130-4.

[100] M. Tröbs and G. Heinzel, *Improved spectrum estimation from digitized time series on a logarithmic frequency axis*, Measurement **39** (2006) 120, ISSN: 0263-2241, URL: http://dx.doi.org/10.1016/j.measurement.2005.10.010.

[101] H. Ludvigsen, M. Tossavainen, and M. Kaivola, *Laser linewidth measurements using self-homodyne detection with short delay*, Optics Communications **155** (1998) 180, ISSN: 0030-4018, URL: http://dx.doi.org/10.1016/S0030-4018(98)00355-1.

[102] https://gitlab.com/las-nq/openqlab, [Accessed 29-Jan-2025].

[103] http://www.shx-sh.com/files/en/TKE-45-4-S.pdf, [Accessed 29-Jan-2025].

[104] https://www.xilinx.com/products/boards-and-kits/zcu111.html, [Accessed 29-Jan-2025].

[105] *ZCU111 Evaluation Board User Guide (UG1271)*, tech. rep., [Accessed 29-Jan-2025]: Advanced Micro Devices, Inc., URL: https://docs.amd.com/r/en-US/ug1271-zcu111-eval-bd.

[106] T. Izumi and Y. Mitsuyama, "Design Flow and Design Tools," *Principles and Structures of FPGAs*, Springer Singapore, 2018, ISBN: 9789811308246, URL: http://dx.doi.org/10.1007/978-981-13-0824-6_4.

[107] S. E. Pollack and R. T. Stebbins, *Demonstration of the zero-crossing phasemeter with a LISA test-bed interferometer*, Classical and Quantum Gravity **23** (2006) 4189, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/23/12/014.

[108] S. Johansson, "New frequency counting principle improves resolution," *Proceedings of the 20th European Frequency and Time Forum*, 2006, URL: https://ieeexplore.ieee.org/document/6230960.

[109] G.-C. Hsieh and J. Hung, *Phase-locked loop techniques. A survey*, IEEE Transactions on Industrial Electronics **43** (1996) 609, ISSN: 0278-0046, URL: http://dx.doi.org/10.1109/41.544547.

[110] F. M. Gardner, *Phaselock Techniques*, Wiley, 2005, ISBN: 9780471732693, URL: http://dx.doi.org/10.1002/0471732699.

[111] H. Audley et al., *The LISA Pathfinder interferometry—hardware and system testing*, Classical and Quantum Gravity **28** (2011) 094003, ISSN: 1361-6382, URL: http://dx.doi.org/10.1088/0264-9381/28/9/094003.

[112] W. Lindsey and C. M. Chie, *A survey of digital phase-locked loops*, Proceedings of the IEEE **69** (1981) 410, ISSN: 0018-9219, URL: http://dx.doi.org/10.1109/PROC.1981.11986.

[113] I. Bykov, J. J. E. Delgado, A. F. G. Marín, G. Heinzel, and K. Danzmann, *LISA phasemeter development: Advanced prototyping*, Journal of Physics: Conference Series **154** (2009) 012017, ISSN: 1742-6596, URL: http://dx.doi.org/10.1088/1742-6596/154/1/012017.

[114] M. Kumm, H. Klingbeil, and P. Zipf, *An FPGA-Based Linear All-Digital Phase-Locked Loop*, IEEE Transactions on Circuits and Systems I: Regular Papers **57** (2010) 2487, ISSN: 1558-0806, URL: http://dx.doi.org/10.1109/TCSI.2010.2046237.

[115] O. Gerberding et al., *Phasemeter core for intersatellite laser heterodyne interferometry: modelling, simulations and experiments*, Classical and Quantum Gravity **30** (2013) 235029, URL: https://doi.org/10.1088/0264-9381/30/23/235029.

[116] H.-S. Liu, T. Yu, and Z.-R. Luo, *A low-noise analog frontend design for the Taiji phasemeter prototype*, Review of Scientific Instruments **92** (2021), ISSN: 1089-7623, URL: http://dx.doi.org/10.1063/5.0042249.

[117] K.-S. Isleif et al., *Highspeed multiplexed heterodyne interferometry*, Optics Express **22** (2014) 24689, ISSN: 1094-4087, URL: http://dx.doi.org/10.1364/OE.22.024689.

[118] G. Ascheid and H. Meyr, *Cycle Slips in Phase-Locked Loops: A Tutorial Survey*, IEEE Transactions on Communications **30** (1982) 2228, ISSN: 0096-2244, URL: http://dx.doi.org/10.1109/TCOM.1982.1095423.

[119] M. T. L. Hsu et al., *Subpicometer length measurement using heterodyne laser interferometry and all-digital rf phase meters*, Optics Letters **35** (2010) 4202, ISSN: 1539-4794, URL: http://dx.doi.org/10.1364/OL.35.004202.

[120] T. S. Schwarze et al., *Picometer-Stable Hexagonal Optical Bench to Verify LISA Phase Extraction Linearity and Precision*, Physical Review Letters **122** (2019), ISSN: 1079-7114, URL: http://dx.doi.org/10.1103/PhysRevLett.122.081104.

[121] A. Abramovici and J. Chapsky, *Feedback Control Systems*, Springer US, 2000, ISBN: 9781461543459, URL: http://dx.doi.org/10.1007/978-1-4615-4345-9.

[122] B. Widrow and I. Kollár, *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*, Cambridge University Press, 2008, ISBN: 9780511754661, URL: http://dx.doi.org/10.1017/CBO9780511754661.

[123] S. Fuller et al., "The Python Control Systems Library (python-control)," *2021 60th IEEE Conference on Decision and Control (CDC)*, IEEE, 2021, URL: http://dx.doi.org/10.1109/CDC45484.2021.9683368.

[124] A. Basalaev et al., *Spicypy*, 2023,
URL: `https://zenodo.org/doi/10.5281/zenodo.10033638`.

[125] E. Hogenauer, *An economical class of digital filters for decimation and interpolation*, IEEE Transactions on Acoustics, Speech, and Signal Processing **29** (1981) 155, ISSN: 0096-3518, URL: `http://dx.doi.org/10.1109/TASSP.1981.1163535`.

[126] C. Liu, M. E. Jones, and A. C. Taylor, *Characterizing the performance of high-speed data converters for RFSoC-based radio astronomy receivers*, Monthly Notices of the Royal Astronomical Society **501** (2020) 5096, ISSN: 1365-2966, URL: `http://dx.doi.org/10.1093/mnras/staa3895`.

[127] L. G. Holcomb, *A direct method for calculating instrument noise levels in side-by-side seismometer evaluations*, 1989, URL: `http://dx.doi.org/10.3133/ofr89214`.

[128] R. Kirchhoff et al., *Huddle test measurement of a near Johnson noise limited geophone*, Review of Scientific Instruments **88** (2017), ISSN: 1089-7623, URL: `http://dx.doi.org/10.1063/1.5000592`.

[129] H.-S. Liu, Y.-H. Dong, Y.-Q. Li, Z.-R. Luo, and G. Jin, *The evaluation of phasemeter prototype performance for the space gravitational waves detection*, Review of Scientific Instruments **85** (2014), ISSN: 1089-7623, URL: `http://dx.doi.org/10.1063/1.4865121`.

[130] O. Gerberding et al., *Readout for intersatellite laser interferometry: Measuring low frequency phase fluctuations of high-frequency signals with microradian precision*, Review of Scientific Instruments **86** (2015), ISSN: 1089-7623, URL: `http://dx.doi.org/10.1063/1.4927071`.

[131] B. Bachman et al., *Flight phasemeter on the Laser Ranging Interferometer on the GRACE Follow-On mission*, Journal of Physics: Conference Series **840** (2017) 012011, ISSN: 1742-6596, URL: `http://dx.doi.org/10.1088/1742-6596/840/1/012011`.

[132] J. M. Eichholz, *Digital Heterodyne Laser Frequency Stabilization for Space-Based Gravitational Wave Detectors and Measuring Coating Brownian Noise at Cryogenic Temperatures*, PhD thesis, 2015, URL: `https://ufdc.ufl.edu/UFE0049565/00001`.

[133] `https://github.com/jupyter-widgets/ipywidgets`, [Accessed 29-Jan-2025].

[134] `https://www.ti.com/tool/TRF1208EVM`, [Accessed 29-Jan-2025].

[135] *Zynq UltraScale+ RFSoC RF Data Converter v2.4 Gen 1/2/3 Product Guide (PG269)*, tech. rep., Advanced Micro Devices, Inc., URL: `https://docs.amd.com/v/u/2.4-English/pg269-rf-data-converter`.

[136] T. Schuldt et al., *High-Performance Optical Frequency References for Space*, Journal of Physics: Conference Series **723** (2016) 012047, ISSN: 1742-6596, URL: `http://dx.doi.org/10.1088/1742-6596/723/1/012047`.

[137] G. Mueller, P. McNamara, I. Thorpe, and J. Camp, *Laser Frequency Stabilization for LISA*, Technical Memorandum NASA/TM-2005-212794, NASA Goddard Space Flight Center, 2005, URL: `https://ntrs.nasa.gov/citations/20060012084`.

[138] T. Schuldt et al., *Absolute laser frequency stabilization for LISA*, International Journal of Modern Physics D **28** (2019) 1845002, ISSN: 1793-6594, URL: `http://dx.doi.org/10.1142/S0218271818450025`.

[139] K. Numata, A. Kemery, and J. Camp, *Thermal-Noise Limit in the Frequency Stabilization of Lasers with Rigid Cavities*, Physical Review Letters **93** (2004), ISSN: 1079-7114, URL: `http://dx.doi.org/10.1103/PhysRevLett.93.250602`.

[140] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist, *Visible Lasers with Subhertz Linewidths*, Physical Review Letters **82** (1999) 3799, ISSN: 1079-7114, URL: `http://dx.doi.org/10.1103/PhysRevLett.82.3799`.

[141] F. Bondu, P. Fritschel, C. N. Man, and A. Brillet, *Ultrahigh-spectral-purity laser for the VIRGO experiment*, Optics Letters **21** (1996) 582, ISSN: 1539-4794, URL: `http://dx.doi.org/10.1364/OL.21.000582`.

[142] `https://www.schott.com/en-us/products/zerodur-p1000269`, [Accessed 29-Jan-2025].

[143] `https://oharacorp.com/low-expansion-glass/`, [Accessed 29-Jan-2025].

[144] `https://www.menlosystems.com/products/ultrastable-lasers/orc-cylindric-orc-cubic/`, [Accessed 29-Jan-2025].

[145] W. Zhang, F. Baynes, S. A. Diddams, and S. B. Papp, *Microrod Optical Frequency Reference in the Ambient Environment*, Physical Review Applied **12** (2019), ISSN: 2331-7019, URL: `http://dx.doi.org/10.1103/PhysRevApplied.12.024010`.

[146] M. J. Davis and M. Kocher, "Athermal glass for infrared optics," *Current Developments in Lens Design and Optical Engineering XIX*, ed. by R. B. Johnson, V. N. Mahajan, and S. Thibault, SPIE, 2018 12, URL: `http://dx.doi.org/10.1117/12.2320708`.

[147] `https://lairdthermal.com/products/thermoelectric-cooler-assemblies/peltier-plate-assemblies/DA-045-24-02`, [Accessed 29-Jan-2025].

[148] `https://lairdthermal.com/products/product-temperature-controllers/tc-xx-pr-59-temperature-controller`, [Accessed 29-Jan-2025].

[149] M. T. Brennecke, *Construction of ULE-cavities and high precision temperature sensors for laser frequency stabilisation*, Bachelor's Thesis: Institute of Experimental Physics, University of Hamburg, 2014.

[150] P. Virtanen et al., *SciPy 1.0: fundamental algorithms for scientific computing in Python*, Nature Methods **17** (2020) 261, ISSN: 1548-7105, URL: `http://dx.doi.org/10.1038/s41592-019-0686-2`.

[151] J. Owen, M. Davies, D. Schmidt, and E. Urruti,
On the ultra-precision diamond machining of chalcogenide glass, CIRP Annals **64** (2015) 113,
ISSN: 0007-8506, URL: `http://dx.doi.org/10.1016/j.cirp.2015.04.065`.

[152] M. Tröbs, *Laser development and stabilization for the spaceborne interferometric gravitational wave detector LISA*, PhD thesis, 2005,
URL: `https://www.repo.uni-hannover.de/handle/123456789/6675`.

[153] P. Patel et al., *Mechanism of thermal decomposition of poly(ether ether ketone) (PEEK) from a review of decomposition studies*, Polymer Degradation and Stability **95** (2010) 709,
ISSN: 0141-3910,
URL: `http://dx.doi.org/10.1016/j.polymdegradstab.2010.01.024`.

[154] H. Shen, L. Li, J. Bi, J. Wang, and L. Chen, *Systematic and quantitative analysis of residual amplitude modulation in Pound-Drever-Hall frequency stabilization*,
Physical Review A **92** (2015), ISSN: 1094-1622,
URL: `http://dx.doi.org/10.1103/PhysRevA.92.063809`.

[155] `https://liquidinstruments.com/products/integrated-instruments/frequency-response-analyzer/`, [Accessed 29-Jan-2025].

[156] T. T.-Y. Lam et al., *Digital Laser Frequency Stabilization Using an Optical Cavity*,
IEEE Journal of Quantum Electronics **46** (2010) 1178, ISSN: 1558-1713,
URL: `http://dx.doi.org/10.1109/JQE.2010.2044867`.

[157] S. Amairi et al., *Reducing the effect of thermal noise in optical cavities*,
Applied Physics B **113** (2013) 233, ISSN: 1432-0649,
URL: `http://dx.doi.org/10.1007/s00340-013-5464-8`.

[158] D. Kryuchkov et al., *Compensation of residual amplitude modulation fluctuations in an optoelectronic system for laser radiation frequency stabilisation*,
Quantum Electronics **50** (2020) 590, ISSN: 1468-4799,
URL: `http://dx.doi.org/10.1070/QEL17367`.

[159] M.-S. Hartig, S. Schuster, and G. Wanner, *Geometric tilt-to-length coupling in precision interferometry: mechanisms and analytical descriptions*, Journal of Optics **24** (2022) 065601,
ISSN: 2040-8986, URL: `http://dx.doi.org/10.1088/2040-8986/ac675e`.

[160] M. Tröbs et al., *Reducing tilt-to-length coupling for the LISA test mass interferometer*,
Classical and Quantum Gravity **35** (2018) 105001, ISSN: 1361-6382,
URL: `http://dx.doi.org/10.1088/1361-6382/aab86c`.

[161] K. Möhle, E. V. Kovalchuk, K. Döringshoff, M. Nagel, and A. Peters,
Highly stable piezoelectrically tunable optical cavities, Applied Physics B **111** (2013) 223,
ISSN: 1432-0649, URL: `http://dx.doi.org/10.1007/s00340-012-5322-0`.

[162] M. Musha et al., *Robust and precise length stabilization of a 25-km long optical fiber using an optical interferometric method with a digital phase-frequency discriminator*,
Applied Physics B **82** (2006) 555, ISSN: 1432-0649,
URL: `http://dx.doi.org/10.1007/s00340-005-2125-6`.

[163] J. Carter et al., "A High Q, Quasi-Monolithic Optomechanical Inertial Sensor," *2020 IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, IEEE, 2020, URL: http://dx.doi.org/10.1109/INERTIAL48129.2020.9090085.

[164] A. Hines et al., *Compact optomechanical accelerometers for use in gravitational wave detectors*, Applied Physics Letters **122** (2023), ISSN: 1077-3118, URL: http://dx.doi.org/10.1063/5.0142108.

[165] https://gitlab.com/uhh-gwd/lpsd, [Accessed 29-Jan-2025].

[166] P. R. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors*, 2nd ed., Singapore: World Scientific Publishing Co. Pte. Ltd., 2017, ISBN: 978-981-314-307-4, URL: http://dx.doi.org/10.1142/10116.

[167] J. H. Põld and A. D. Spector, *Demonstration of a length control system for ALPS II with a high finesse 9.2 m cavity*, EPJ Techniques and Instrumentation **7** (2020), ISSN: 2195-7045, URL: http://dx.doi.org/10.1140/epjti/s40485-020-0054-8.

List of figures

Acronyms

ADC analog-to-digital converter xi, 29, 39–42, 61–65, 67, 73, 75–77, 79, 87, 88, 92, 93, 95, 96, 98, 99, 102–104, 106–109, 136, 145, 148, 150, 154, 169, 190
ADPLL all-digital phase-locked loop xi, 71, 73, 75–77, 79, 81, 83, 84, 86–91, 96, 98, 99, 102, 107
aLIGO Advanced LIGO 1, 4, 5, 8, 10–12, 38
AOM acousto-optic modulator 20
APC angled physical contact 139
API application programming interface 159
ASD amplitude spectral density 36, 74, 87, 157, 158
AXI Advanced eXtensible Interface 73, 75, 86, 104, 169

BBH binary black hole 1, 3, 7
BBSS big beam-splitter suspension 11
BNS binary neutron star 1, 7
BOSEM Birmingham Optical Sensor and Electro-Magnetic actuator 10, 11
BSP board support package 66

CAD computer-aided design 122, 123, 129, 190, 191
CE Cosmic Explorer 5
CIC cascaded integrator-comb 84–87, 91, 92, 190
COBRI Compact Balanced Readout Interferometer 11
CPU central processing unit 61, 62, 169
CSD cross spectral density 95, 124
CTE coefficient of thermal expansion 37, 113–115, 127, 142
CTL Continuously Tunable Laser x, xi, 17, 18, 29, 43–47, 49, 50, 52, 55, 73, 89, 90, 92–94, 97–99, 117–119, 122, 159, 189, 190

DAC digital-to-analog converter 62–67, 101–104, 106, 136
DDC digital down converter 75, 169

DDR	Double Data Rate 63
DDS	Direct Digital Synthesizer 66, 76
DFM	deep-frequency modulation 11, 12, 19
DHC	Digital Heterodyne Controller v, vii, 13, 101–106, 108, 135–138, 141, 143, 145, 147, 154, 190, 191
DI	Digital Interferometry 12
DRFPMI	dual-recycled Fabry-Pérot Michelson interferometer 3, 4, 19
DSP	digital signal processing 13, 35, 39, 50, 61, 62, 66, 73–75, 81, 84, 85, 169
ECDL	external cavity diode laser 17, 18, 136, 189
EMRI	extreme mass ratio inspiral 5
ENOB	effective number of bits 88
ET	Einstein Telescope 5
FFT	fast Fourier transform 99
FIR	finite impulse response 79, 80, 84, 190
FOM	figure-of-merit 114, 115, 121, 171, 190
FPGA	field-programmable gate-array v, vii, 11, 50, 57, 61–63, 66, 70, 71, 74–76, 79, 84, 85, 96, 101–103, 106, 135, 154
FSR	free spectral range 22, 25, 32, 33, 35, 37, 48, 51, 54, 57, 117, 131, 151
FWHM	full-width at half-maximum 22
GLOC	Gravitational-wave Lunar Observatory for Cosmology 7
GPU	graphics processing unit 61, 62
GSPS	giga samples per second 63, 73, 77, 103, 169
GUI	graphical user interface 105, 106, 122, 123, 171, 190
HDL	hardware description language 62
HoQI	Homodyne Quadrature Interferometer 11
HS	heterodyne stabilization x, xii, 29, 30, 33, 34, 36, 39, 40, 42–44, 50–57, 101–103, 108, 114, 116–118, 121, 127, 128, 131, 135, 137–143, 145, 146, 151, 154, 161, 162, 189–191
I2C	Inter-Integrated Circuit 65
IDE	Integrated Development Environment 66
ILA	Integrated Logic Analyzer 67, 82, 83, 89
IP	intellectual property 66, 67, 75, 76, 82, 85, 86, 89, 91, 103, 104, 166, 169
ISI	intra-vacuum seismic isolator 10, 11

KAGRA Kamioka Gravitational Wave Detector 3, 10

LAN local area network 46, 159
LAO LIGO Aundha Observatory 4
LED light emitting diode 10, 67, 130
LGWA Lunar Gravitational-Wave Antenna 6, 7, 12, 13, 153
LHO LIGO Hanford Observatory 3, 8, 189
LIGO Laser Interferometer Gravitational-Wave Observatory 1, 4, 11, 23, 151
LISA Laser Interferometer Space Antenna 5, 6, 9, 10, 21, 30, 41, 42, 70–73, 91, 95–99, 113, 149, 153
LLO LIGO Livingston Observatory 3
LO local oscillator 69–72
LPSD logarithmic power spectral density 47, 158
LSB least significant bit 39, 46, 74, 79
LUT look-up table 76, 81, 104
LVDS low-voltage differential signaling 165
LVDT linear variable differential transformer 9

MEMS micro-electro-mechanical system 9, 10
MSB most significant bit 79, 80, 169

NCO numerically controlled oscillator 71, 75–79, 81, 102–104, 169
Nd:YAG neodymium-doped yttrium aluminum garnet 17, 21, 43
NPRO non-planar ring oscillator 17, 43, 93

PA phase accumulator 70, 81, 86, 88, 102, 103
PBS polarizing beam-splitter 28, 51, 52, 119, 121, 164
PC personal computer 92, 106
PDH Pound-Drever-Hall 28, 30, 145, 146, 150, 189
PEEK polyether ether ketone 139, 142
PI proportional-integral 77, 80, 104, 143, 144, 169
PID proportional-integral-derivative 52, 123, 124, 126, 135, 174
PIR phase increment register 70, 75–77, 81, 86–90, 92, 96, 98, 101–104, 169
PL programmable logic 62, 63, 66, 72, 73, 75–77, 81, 82, 84–87, 91, 99, 101, 103, 104, 150, 169
PLL phase-locked loop xiii, 64–66, 69–73, 75–83, 85–93, 95, 97–99, 101, 102, 104, 107, 151, 165–167, 169, 190, 191
POM polyoxymethylene 122, 123
PS processing system 63, 64, 66, 75, 77, 82–86, 92, 99, 104, 106, 169

PSD	power spectral density 37, 38, 74, 94, 95, 124, 157
PTC	pilot tone correction 77, 95, 96
QWP	quarter-wave plate 28, 51, 119, 164
RAM	random access memory 63, 83, 86, 169
ReDFMI	Resonantly enhanced Deep-Frequency Modulation Interferometry 12
RF	radio-frequency x, xi, xiii, 30, 52, 56, 57, 62–67, 72, 75, 85, 96, 102–104, 106–109, 146, 151, 165–167, 169, 190, 191
RFSoC	radio-frequency system-on-chip x, 13, 61–64, 66, 69, 73, 75, 76, 79, 88, 92, 96, 99, 101, 103, 105, 106, 117, 135, 151, 153, 154, 165, 169
RIN	relative intensity noise 26
RMS	root-mean-square 8, 37, 43, 87, 90–93, 98, 190
RS	Recommended Standard 122
SCI	serial command interface 124, 171
SCPI	standard commands for programmable instruments 46
SMA	Sub-Miniature version A 64, 67, 165
SNR	signal-to-noise ratio 3, 36, 37, 39, 69, 150
SODIMM	Small Outline Dual In-line Memory Module 63
SPI	Serial Peripheral Interface 65
TEC	thermoelectric cooler 122, 123, 125, 126, 171, 191
TEM	transverse electromagnetic 51, 135, 162, 163
TTL	tilt-to-length 149, 150
UGF	unity gain frequency 43, 93, 125, 143–145, 151, 190
ULE	Ultra-Low Expansion xii, 13, 37, 113, 114, 127, 128, 130, 131, 142, 143, 146, 148, 191
USB	universal serial bus 122
UV	ultra-violet 128, 130
VCO	voltage-controlled oscillator 165–167
VCXO	voltage-controlled crystal oscillator 165
VHSIC	very high speed integrated circuit 62

Acknowledgements

I would like to express my heartfelt gratitude to Prof. Dr. Oliver Gerberding. Words can't adequately convey my appreciation for this wonderful opportunity, as well as for all the guidance, motivation, and inspiration you have provided along the way. You welcomed me despite my limited knowledge and skills, and you cultivated an environment that allowed me to grow to this stage. This means a lot to me. Being your first-generation Ph.D. student and contributing to the establishment of the laboratory from the beginning has been an amazing experience. Thank you for everything, Oliver!

Reflecting on my journey, pursuing research abroad wasn't an easy decision for me. People often say, "You need to sacrifice in order to achieve something; no pain, no gain." *Amma* (mother), in our case, you are the one who sacrificed your comfort to allow me to pursue this path. I couldn't be there for you during the times you needed me. Without you and *Appa* (father), I'm nothing. Thank you for your unwavering strength and support.

I would also like to take this opportunity to extend my acknowledgements to my sister, Dr. Meghashree C. S., my brother-in-law, Kirana B., and my beloved nephew, Skanda K. Bhat. Technological advancements have helped bridge the physical distance between us to some extent, and I have felt the warmth of closeness during each of our conversations.

Returning to my stay in Hamburg, I feel incredibly fortunate to have such a friendly, supportive working group. Although I arrived without any expectations, I never imagined that each of you would accept me in the way you did. I'll cherish the discussions, shared learnings, casual conversations in the lunchroom, movie evenings, and board game sessions for a long time.

This write-up would be incomplete without explicitly mentioning some of my colleagues. Tobias Eckhardt, thank you for being there to address my silly doubts and rants. I'm glad this five-year Ph.D. journey unfolded alongside yours. Marcel Beck, cheers to our 'space lab'! Having you as my lab mate and our collaborative attempts to understand theoretical aspects and solve experimental challenges will always be memorable experiences. Leander Weickhardt, despite our minimal overlap in projects, we have formed a deep bond. I don't know how to express my thanks to you. A special thanks to these three, along with Daniel Voigt and Sara Porras Bedmar, for their efforts in bringing me out of the 'all-time work mode.' You are all wonderful people!

My heartfelt appreciation also goes to the two post-docs, Dr. Artem Basalaev and Dr. Christian Darsow-Fromm. Interacting with both of you was such a nice experience, and I learned a lot simply from my observations and our conversations. Chris, as you pointed out, often I just needed two ears to hear my thoughts. I chose you for this role and may have disturbed you from time to time—my apologies!

I'm also thankful to Artem, Dr. Jonathan Carter, and Chris for reviewing some chapters of this thesis draft and providing valuable feedback. Thanks to Leander and Jan-Niklas Feldhusen for helping with the German abstract for this thesis.

Dr. Vanamali Shastry and Dr. Ishwara Chandra C. H., your guidance before and during my academic journey has been instrumental. I'm grateful for your insights.

Among my friends, I would like to mention Abinaya Omkumar and Fg. Offr. Shubhangi Narvekar. Irrespective of the time zone difference, you've maintained our connection. Your understandings and care have meant a lot to me.

Gabriele Hoffmann and Thomas Heese, you treated me like family and made my life in Hamburg relaxed and stress-free. I appreciate the opportunity you've created for me to experience sailing and learn the basics of swimming—activities I might not have tried otherwise! Thank you for all your help and support.

Anna Hoffmann, you and your friends are my only non-colleague contacts in Hamburg! Thank you for including me in all the meet-ups, vacation trips, and cooking sessions.

I also want to extend my gratitude to the DESY International Office, as well as the administrative, IT, and workshop staff of the university, the Quantum Universe office, and all those unknown individuals behind the scenes who made my work and life in Hamburg easier.

Lastly, thank you, the reader, whether you're reading just this page, parts of the thesis, or the entire document. I'm glad to have had the opportunity to write and make this thesis accessible to you!

Curriculum vitae

Personal details

Name	Shreevathsa Chalathadka Subrahmanya
Date of birth	February 18, 1997
Place of birth	Puttur, Karnataka, India
E-mail	shreevathsa.subrahmanya@uni-hamburg.de

Education

2012	Secondary School Leaving Certificate, Karnataka State Board (Distinction with 96.96%)
2012 - 2014	Pre-university education, Karnataka State Board (Distinction with 96.83%)
2014 - 2017	B. Sc. in Physics, Mathematics, and Electronics, Mangalore University, India (First rank with 96.02%)
2017 - 2019	M. Sc. in Physics, Mangalore University, India (Distinction with CGPA 8.44)

Professional experiences

June 2018 - July 2018	Summer Research Fellow at RRCAT, Indore, India
Sept. 2019 - Nov. 2019	Visiting Student at IIA, Bengaluru, India
Jan. 2020 - current	Research Assistant at the University of Hamburg, Germany

Languages

Kannada	Mother tongue
English	Fluent
Hindi	Basic
German	Basic

Publications

Peer-reviewed journal articles

S. Chalathadka Subrahmanya, C. Darsow-Fromm, and O. Gerberding
"On the Development of an RFSoC-Based Ultra-Fast Phasemeter With GHz Bandwidth"
IEEE Transactions on Instrumentation and Measurement, Vol. 74 (2025)
`https://doi.org/10.1109/TIM.2024.3509586`

S. Chalathadka Subrahmanya, C. Darsow-Fromm, and O. Gerberding
"Integrating high-precision and fringe-scale displacement sensing using heterodyne cavity-tracking"
Optics Express, Vol. 33, Issue 03 (2025)
`https://doi.org/10.1364/OE.540189`

J. V. van Heijningen, H. J. M. ter Brake, O. Gerberding, **S. Chalathadka Subrahmanya** et al.
"The payload of the Lunar Gravitational-wave Antenna"
Journal of Applied Physics, Vol. 133, Issue 24 (2023)
`https://doi.org/10.1063/5.0144687`

P. Ajith et al.
"The Lunar Gravitational-wave Antenna: mission studies and science case"
Journal of Cosmology and Astroparticle Physics, Vol. 2025, Issue 01 (2025)
`https://doi.org/10.1088/1475-7516/2025/01/108`

25+ other publications as a co-author of the LIGO Scientific Collaboration.

Submitted manuscript

M. Beck, **S. Chalathadka Subrahmanya**, and O. Gerberding
"Adjustable picometer-stable interferometers for testing space-based gravitational wave detectors"
`https://doi.org/10.48550/arXiv.2502.01212`
(Submitted to Classical and Quantum Gravity)

Co-supervision

Projects co-supervised:

- Nurmi Sawlanski (2024)
  Bachelor thesis
  Unequal arm-length interferometer for laser frequency pre-stabilization

- Mika Tjorben Brennecke (2024)
  Bachelor thesis
  Construction of ULE-cavities and high precision temperature sensors for laser frequency stabilisation

Declaration on oath

I hereby declare and affirm that this doctoral dissertation is my own work and that I have not used any aids and sources other than those indicated.

If electronic resources based on generative artificial intelligence (gAI) were used in the course of writing this dissertation, I confirm that my own work was the main and value-adding contribution and that complete documentation of all resources used is available in accordance with good scientific practice. I am responsible for any erroneous or distorted content, incorrect references, violations of data protection and copyright law or plagiarism that may have been generated by the gAI.

19.04.2025	
Date	Signature of doctoral candidate