Development and Characterisation of a Fibre Frequency Reference

Ya Zhang

A thesis submitted for the degree of
Doctor of Philosophy
of the Australian National University

June, 2022

Declaration

This thesis is an account of research undertaken between March 2018 and June 2022 at the Centre for Gravitational Astrophysics, Research School of Physics, The Australian National University, Canberra, Australia.

Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not been submitted in whole or part for a degree in any university.

Ya Zhang
June, 2022

Acknowledgements

your input on the scientific front and your patience and support in life. I could not have done it without you.

Abstract

In this thesis, a passive, all optical fibre frequency reference is developed and characterised. The system uses long armlength difference interferometers to measure fluctuations of laser frequency. The phase readout is handled by digital interferometry (DI), which uses spread-spectrum modulation to extract interference signals from the desired range gate with high dynamic range. The frequency stability of the fibre reference is characterised using a differential measurement between two near-identical interferometers. We achieve a relative stability of $0.1~\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ above 70 Hz Fourier frequency, which surpasses previous demonstrations of fibre optic references.

Building on prior fibre reference investigations, the new system discussed here is designed around an unbalanced Mach-Zehnder interferometer, removing the impact of first-order Rayleigh backscattering as seen in previous designs. The implementation of DI is modified to enable real-time phase reconstruction instead of at a decimated speed, reducing non-linear errors and improving readout fidelity. The new interferometers are individually enclosed in two updated dual-layer passive isolation chambers.

In system characterisation, we provide the first long-term temperature analysis of the isolation chambers. Their individual time constant is modelled and experimentally measured, at 13.2 hrs and 11.4 hrs respectively. The 14% difference between the two chambers is in alignment with the temperature independence observed in a three-reference optical measurement. We also comprehensively survey the laboratory mechanical profile, and identify the driving source for each mechanical feature in the experimental noise floor.

In addition to temperature and mechanical stability, noise limitations in other Fourier regimes are also identified and characterised. We adapt the Duan fibre thermal noise model for a long armlength interferometer, and experimentally achieve thermo-mechanical noise limited relative stability between 0.4 - 2 Hz. We also develop the first quantitative model for double Rayleigh scattering (DRS) in a fibre interferometer including effects from DI integration and suppression. The modelled contribution from DRS is shown to be in close agreement with the experimental noise floor above 70 Hz Fourier frequencies.

The achieved $0.1~\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ frequency stability represents the state-of-the-art performance for fibre references and is comparable with room temperature cavity systems. This makes our system a potential alternative for laser frequency stabilisation at short timescales, particularly in applications where the robustness of fibre systems and their intrinsic optical alignment are important considerations.

Contents

Contents

List of Figures

List of Tables

Introduction

Optical metrology is a collective term for precision measurements made on the amplitude or phase of an optical interrogation field. The interpretation and implication of these measurements are extremely diverse and spans a wide range of scientific disciplines, such as medical imagining, environmental monitoring and studies of the universe. These metrologies also utilise a variety of optical configurations. Many projects discussed hereupon fall into the category of optical interferometry, which uses the phase information of the optical beam to achieve sub-wavelength sensing precision.

The sensitivity of these systems is often dependent on the frequency stability of the optical source. Changes in the optical frequency can result in unwanted interferometric phase, obscuring the signal of interest. Besides treating the optical source, the most common frequency-noise mitigation method involves balancing the interferometer armlength. This however is not always practical for systems with extensive or variable sensing length, nor is it sufficient for projects requiring the highest performance [1, 2]. In these scenarios, pre-stabilisation of the optical source becomes necessary.

Optical frequency metrology is the field that measures, characterises and improves the stability of an optical oscillator. Here we focus on a specific type of oscillator: lasers. In the following sections, we discuss current laser manufacturing technologies and their free-running performance. We also discuss the diverse range of frequency standards and references, as well as methods to harvest their stability in the appropriate Fourier regime to further enhance laser performance.

1.1 The Pursuit of Narrow Linewidth

The term "narrow linewidth" qualitatively equates lasers of high frequency stability and low phase noise. These sources are attractive for their compact size and high performance, requiring little to no external stabilisation subject to the specific application. Their popularity has in turn led to increased research and production of low-noise sources in recent years [3, 4, 5, 6, 7]. To better understand these lasers, we first define three terminologies used to categorise or quantify their stability.

1.1.1   Optical Stability Characterisation Methods

Linewidth

Linewidth is a convenient concept for quick assessment of source stability. It is often measured by heterodyning the source with a delayed, incoherent copy of itself [8]. The resultant spectral shape is a combination of Lorentzian and Gaussian functions, corresponding to white and $1/f$ noises respectively [9]. Through a Voigt fit, the Lorentzian component can be extracted [10]. This is referred to as the natural or intrinsic linewidth.

For sources with very high coherence, it is difficult to produce a sufficiently long delay-line without excessive technical noise. These sources are instead often characterised by their "integrated linewidth", where the full width half maximum (FWHM) of a heterodyne beat note is recorded over a specified time. As the name suggests, this linewidth is the spectral integration of frequency noise and thus a function of integration time. While it provides an intuitive classification of laser performance, as a singular point of information, it is unsuitable for quantitative or comprehensive analysis.

Frequency Spectral Density

Frequency spectral density (FSD) is the favoured approach to obtain complete insight into a laser's spectral behaviour. Its calculation involves transforming a time-domain measurement of laser frequency fluctuation into the Fourier regime. In addition, the time-domain sample period is calibrated for and referred to as the "resolution bandwidth" in the Fourier domain. The resultant spectral density can be interpreted as the optical frequency fluctuation per unit Fourier frequency, and takes on the unit of $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$. FSD is often plotted against its Fourier frequency bin in log-log scale, with the first example shown in Chapter 4.1 of this thesis (Fig. 4.2).

The majority of this thesis uses FSD for stability characterisation, however there are instances of phase spectral densities being used as intermediate steps. Both FSD and phase spectral density can be classified as amplitude spectral densities (ASDs), with the specific type determined by the time-domain measurand (i.e. optical frequency vs interferometric phase). They are related to power spectral densities (PSDs) by a square operation.

Depending on the specific convention used, there is also sometimes a $\sqrt{2}$ to $\sqrt{4\pi}$ discrepancy between different works [11]. This thesis exclusively uses the single-sided spectral density in terms of Fourier frequency f, which produces values $\sqrt{2}$ higher than its double-sided counterpart and a further $\sqrt{2\pi}$ higher than those using the angular frequency ω [12]. The computation of our experimental FSDs is carried out in MATLAB and uses the Welch method with Blackman-Harris window [13].

Allan Deviation

Allan deviation [14] is another common tool for frequency stability characterisation. Unlike the FSD, Allan deviation measures the time-domain stability by computing the sample

deviation $\sigma(\tau)$ as a function of averaging time τ. By finding the τ with the lowest deviation value, we can determine the optimal operation timescale for metrology systems and best integration period for optical oscillators.

The mathematical function for calculating Allan deviation differs by the specific version of choice. This thesis adopts the modified Allan deviation [15] for its ability to distinguish white and flicker phase noise. Another common type is the overlapping Allan deviation, which maximises data usage by forming all possible overlapping samples at each τ [16]. The overlapping variant is useful for reducing variability and enhancing statistical confidence.

1.1.2   Overview of Fibre Lasers

Modern day lasers can be broadly separated into five categories by their gain medium: gas, liquid, semiconductor, solid-state and fibre lasers [17]. In this thesis, we concentrate on metrology-grade fibre lasers, which use rare-earth-doped optical silica and other glass hosts as the active media [18]. The all optical fibre structure of these lasers offers unique advantages compared to their bulk-optic cavity counterparts, including higher optical gain and thus higher efficiency, continuously guided, superior beam quality, intrinsic alignment allowing stability in harsh environments, and small footprint [19]. The commercial fibre laser market value has rapidly risen above the billion dollar mark in the past decade, with strong, continuous growth anticipated for the upcoming years [20].

Fibre lasers are equally popular in R&D due to their unparalleled combination of form-factor, field readiness and amplitude and phase stability, already enabling unprecedented sensitivities in applications such as LiDAR, distributed sensing and spectroscopy [21, 22, 23]. To achieve narrow linewidth, these lasers typically employ a fibre Fabry-Perot (FFP) cavity with fibre Bragg gratings (FBGs) at each end. The theoretical Schawlow-Townes linewidth limit [24] is inversely proportional to the square of the cavity quality factor $Q = 2\pi\nu\tau_c$, which is further proportional to the photon cavity lifetime τ_c. Therefore increasing τ_c is key to reducing the output laser linewidth. This can be intuitively understood as increasing the intra-cavity storage of photon energy, thereby reducing the relative contribution of phase diffusion due to spontaneous emission into the lasing mode.

Fibre laser manufacturers have adopted two generic approaches to increase τ_c [25]. The first one concerns increasing the cavity length L. This is relatively easy and cost-effective to implement in an all-fibre system compared to bulk optics, however an upper length limit still applies as coupling of thermal and mechanical noises worsens. In addition, mode instabilities may occur as L is increased, due to the smaller cavity free-spectral range (FSR $= c/2nL$) and the subsequently narrower spectral spacing between resonant frequencies. A noteworthy approach in this category that circumvents these issues is the implementation of a slow-light filter [4]. This virtually extends L and τ_c, while the spectrally-narrow delay resonator allows suppression of adjacent modes, even as the cavity FSR is reduced.

The second category of approach involves increasing the reflectivity of the cavity mirrors. In the case of a FFP, this can be realised by using FBGs with strong, narrow reflections [26]. Again, the benefit of high reflectivity does not scale indefinitely as it eventually inhibits the available output power.

Utilising both approaches, today's commercial fibre lasers readily achieve a Lorentzian linewidth of several Hz to sub-Hz [27, 28]. This translates to excellent high-frequency stability, with frequency noise reduced to negligible levels by the MHz Fourier mark [28]. At low frequencies, however, fibre lasers are equally susceptible to $1/f^n$ type noises arising from thermal fluctuations and other technical challenges. Therefore projects targeting the kHz and below Fourier band may still require external laser stabilisation.

1.1.3 Where Frequency Stabilisation Comes In

Perhaps the most immediate and compelling example is the field of gravitational-wave (GW) physics. A family of large-scale, interferometric detectors are proposed or constructed to observe GW signals from astronomical events [29]. These instruments achieve extreme sensitivity through a complex suite of strategies including optical topology, mirror suspension, thermal management and seismic isolation. Currently measuring displacements less than 10^{-19} m/$\sqrt{\text{Hz}}$ in advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) [30, 31], the theme for GW detection remains further enhancement of instrument sensitivity to extend its astrophysical reach.

At the desired sensitivities, there is stringent requirement on the amount of source frequency noise injected into the interferometer between 10 Hz to 10 kHz, despite the use of a balanced Michelson configuration [32]. The source laser for aLIGO is therefore pre-stabilised to an external reference cavity and further locked to the interferometer armlength [33, 34]. For future ground-based detectors such as the Cosmic Explorer [35], a different stabilisation approach has been proposed to satisfy the project requirement [1].

On the opposite end of the Fourier spectrum are space-based GW detectors focussing on the $10^{-4} - 10^{-1}$ Hz frequency band. The Laser Interferometer Space Antenna (LISA) mission [36] is scheduled to launch in the 2030s, and will consist of three satellites positioned in an Earth-trailing heliocentric orbit [37]. When in operation, each spacecraft exchanges laser beams with the other two, and obtains a phase measurement by comparing with its local laser [38]. The vast, fluctuating inter-spacecraft separation of 2.5 million km provides a pathway for laser frequency noise to couple in. Even with the state-of-the-art lasers, their noise contribution is calculated to exceed GW signals by 7 orders of magnitude if left untreated [39]. Clearly, laser pre-stabilisation and further frequency noise mitigation techniques are vital to the success of these pioneering projects.

A similar case can be found outside the GW arena, where laser frequency fluctuations often present an issue for armlength unbalanced interferometers. The Gravity Recovery and Climate Experiment Follow-On (GRACE-FO) mission [40] is one such instance. Launched in 2018, GRACE-FO features twin satellites deployed in a low-earth orbit, and monitors Earth's large-scale gravitational dynamics through continuous interferometeric measurement. The orbital separation results in an average round-trip armlength imbalance of 440 km [41] and a subsequent source stability requirement of 30 Hz/$\sqrt{\text{Hz}}$ in the mHz - Hz band [42]. This exceeds the best available free-running performance, necessitating some form of stabilisation for the primary light source.

Even in applications where the commercial state-of-the-art laser performance does suffice, due to the cost of these instruments [43], there are instances where developing

an in-house stabilisation system is overall more economical. Indeed, custom laser stabilisation solutions can be found in a wide array of research and industrial activities [44, 45, 46, 47, 48]. In addition to being cost-effective, these systems can be individually adjusted to provide optimal performance at critical measurement frequencies.

Fundamentally, laser frequency stabilisation is a practice of stability transfer. An external reference of higher stability is introduced, and its performance is inherited by the laser over the appropriate Fourier frequency band. Techniques for achieving this transfer can be separated into two categories: feedback and feedforward [49, 50]. The feedback approach involves generating an error signal to which the laser is stabilised, with Pound-Drever-Hall (PDH) [51, 52] and tilt locking [53, 54, 55] being two common locking techniques. Feedforward, on the other hand, utilises a different paradigm. A direct measurement of free-running laser frequency noise is made. This information is then passed onto downstream operations and removed from subsequent measurements. The feedforward design has a simpler architecture and higher bandwidth flexibility, but requires the reference system to have sufficient dynamic range and bandwidth to accurately track laser frequency fluctuations. Ultimately, the choice between these two techniques is application specific and should be determined by the type of reference and the interested Fourier regime.

1.2   Frequency Standards and References

We now discuss the types of references available in frequency metrology. Using Allan Deviation to characterise their fractional stability as a function of integration time, we can obtain the optimal operation timescale of these systems, and separate them into three broad categories. These references can be used individually to address a specific Fourier frequency band, or combined to provide both long and short term stability.

1.2.1   Molecule and Atomic Transitions

Inherently stable across time and space, molecule and atomic transitions are widely used frequency standards and clocks with outstanding precision and accuracy. These references typically achieve their best performance over the course of minutes to days [56]. Amongst them, the caesium atom's natural frequency is currently recognised as the international unit of time. The stability of this oscillation makes the second the single most accurately measured physical quantity, upon which many other quantities are defined. Since the caesium definition in 1967, various optical clocks have emerged and surpassed the performance of their microwave counterparts [57, 58, 59]. As they continue to improve in reliability, compactness and availability, the tipping point for a redefinition of the second may soon be reached [60, 61].

The main advantage of an optical clock lies in its higher intrinsic frequency, allowing the reduction of fractional instability. Promising clock candidates in this category include ionic species such as Yb+ [62] and Al+ [63], and lattice clocks using the Sr [64, 65, 66,

67, 68, 69, 70, 71] or Yb [72] atom. With these systems, fractional stability on the order of 10^{-18} has been achieved.

Outside timekeeping, hyperfine optical transitions are also used for laser frequency stabilisation and spectroscopic measurements. Examples include the molecular iodine reference for space and geodesy missions [73, 74], achieving a typical fractional stability on the order of 10^{-15} [75, 76]. At telecommunication wavelengths, HCN and Acetylene are well established standards reaching a typical stability of 10^{-12} [77].

1.2.2   Optical Frequency Combs

The turn of the century saw the birth of the optical frequency comb (OFC) [78, 79], an invention manifesting collective advances in ultra-fast laser generation [80, 81], nonlinear optics [82, 83], spectroscopy [84] and frequency metrology [85]. An ultra-fast, mode-locked laser generates an array of equidistant spectral teeth over a wide Fourier frequency span. The OFC's spectral arrangement is solely controlled by two key parameters, pulse repetition rate and carrier envelop offset frequency. Both are RF frequencies stabilised to RF standards such as the rubidium or caesium clock. The result is a stable gearwheel that links optical and microwave frequencies.

The flexibility and stability of OFCs have quickly found them incorporated in various aspects of frequency metrology. Comparison and distribution of stability across different wavelengths is made possible [86, 87], and the simplicity and efficiency of absolute frequency measurements are significantly improved [88, 89, 90]. OFC related measurements have also created new research areas such as comb spectroscopy [91, 92, 93] and comb-assisted LiDAR [94, 95, 96].

As OFCs inherent the stability of their atomic references, they also inherit the Allan Deviation profile of their references. A rubidium clock achieves a typical fractional stability of 10^{-12} upon 100 seconds of integration [97], while the state-of-the-art caesium fountain clock reaches 10^{-16} over the course of a day [98].

1.2.3   Cavities and Fibre Interferometers

The last category of references consists of cavity and interferometric systems. These operate at a shorter timescale of less than a second. Unlike the previous two categories, these references only provide a relative measurement and cannot be used for absolute frequency determination. Nonetheless, they still provide critical stability in the high Fourier frequency regime, improving laser frequency noise by at least two orders of magnitude.

The ultra-low-expansion (ULE) cavity is the standard approach for laser frequency stabilisation. This involves comparing the laser frequency with the cavity resonance and subsequently reading out and feeding back the difference. Depending on the amount of thermal and environmental isolation, ULE cavity systems achieve a typical frequency stability ranging from 0.1 Hz/$\sqrt{\text{Hz}}$ to 1 Hz/$\sqrt{\text{Hz}}$ (fractional $\sim 10^{-15}$) at $\sim$1 Hz Fourier frequency [99, 100, 101, 102]. With systematic efforts in isolation, the state-of-the-art

ULE cavity is able to reach 0.04 Hz/$\sqrt{\text{Hz}}$ frequency stability (fractional $\sim 10^{-16}$) at 1 Hz [103, 7].

One immediate downside to bulk cavity systems is their stringent requirement on laboratory conditions. Many of these systems for example are limited by thermal noise, requiring active temperature cooling to further enhance performance. This becomes an increasingly difficult task as demands for ultra-low-noise lasers increase, particularly in harsh environments such as space interferometry.

On the other end of the spectrum, fibre frequency references (FFRs) present an interesting, but comparatively under-studied alternative to these Fabry-Perot type interferometers at a similar timescale. These systems typically utilise an armlength mismatched interferometer, allowing direct readout of laser frequency fluctuations through the differential phase of the two optical beams. The broadband readout of FFRs enables greater frequency agility, while the all-fibre configuration allows lowered cost, intrinsic alignment and better suitability for field deployment.

However, FFRs do suffer from several limitations unique to guided-wave optics. These include non-linear effects such as Brillioun scattering [104], Raman scattering [105] and Kerr self-phase modulation [106] which limit the optical power levels and coherence lengths that can be used. Linear limitations include Rayleigh back scattering (RBS) and photo-thermal effects [107, 108, 109], both of which must be contended with in order to optimise the phase fidelity of the interferometer. While photo-thermal effects can be effectively reduced by restricting the sensing optical power to the minimum required for shot noise limited sensitivity, or by implementing modest laser intensity stabilisation, first order Rayleigh scattering remains an issue for interferometers operating in reflection.

1.3 Fibre Frequency References

1.3.1 A Brief Review

A number of experiments have investigated the performance of fibre-based frequency references. While most of them utilised a Michelson optical configuration [110, 111, 112, 113, 114], a few opted for the Mach-Zehnder interferometer to avoid first-order RBS limitations [115, 116]. The length of the fibre delay-line in these demonstrations ranged from 100 m to 10 km, displaying a variety of optical gains and cut-off frequencies.

In most of the experiments discussed here, the feedback approach is used for laser stabilisation. Experimental characterisation of the FFR performance requires a test laser to be locked to the interferometer, and its stability compared with a second laser of equal or higher stability, achieved using either a second FFR [112, 111] or a separate high-finesse Fabry-Perot cavity [110, 113]. In this feedback architecture, at least two locking loops are required for performance characterisation. This increases the overall system complexity and potentially introduces bandwidth penalty and control noise. In contrast, the work presented in this thesis adopts the feedforward approach. Our FFR is entirely passive and serves one simple function of tracking laser frequency fluctuation. Its noise floor is measured by subtracting the readout from a second identical FFR.

Regardless of operation principle, all FFR work discussed here utilise a high readout frequency. Therefore frequency spectral density, instead of Allan Deviation, is a more appropriate and comprehensive sensitivity characterisation. Amongst these systems, a noise floor of 1 to 10 Hz/$\sqrt{\text{Hz}}$ at Fourier frequencies above 100 Hz is commonly achieved [117, 115, 118, 119, 120, 116, 114]. Surpassing the 1 Hz/$\sqrt{\text{Hz}}$ sensitivity limit in a similar Fourier frequency range has been a more significant undertaking and only demonstrated experimentally in a few instances [113, 112, 111]. The landmark sensitivity in this field however belongs to the work carried out by Kefelian et al, where a laser was stabilised to a 1 km unbalanced Michelson interferometer, achieving a final stability better than 0.1 Hz/$\sqrt{\text{Hz}}$ between 1-4 kHz [110]. This work demonstrated the feasibility of FFRs as an alternative to Fabry-Perot cavities, and to date still represents the state-of-the-art FFR performance.

1.3.2 Research Objectives

The work we present in this thesis aims to further investigate fibre-based frequency references as an alternative to optical cavities. The primary proposed application remains future space missions, where the robustness and simplicity of an all-fibre system are highly desirable. As a proof-of-concept experiment, we use the popular telecommunication wavelength of 1550 nm, allowing the instrumentation and scientific discoveries discussed here to be transferable to the broader fields of laser frequency stabilisation and fibre sensing.

To fully harvest the advantages of an all-fibre system, we apply Digital Interferometry (DI) to maximally simplify the optical hardware. In doing so, the entire system can be made passive while the interferometric readout remains robust and high dynamic range. We will detail the operating principles of DI in the next chapter.

The ultimate goal for this experiment is to extend the state-of-the-art, 0.1 Hz/$\sqrt{\text{Hz}}$ stability beyond the Fourier frequency range achieved by Kefelian et al. Specifically, we develop a new optical apparatus including a redesign and rebuild of isolation chambers and table optics. We also upgrade and optimise the DI signal processing algorithm to match the new hardware parameters. Lastly, we identify and characterise technical limitations such as thermal, mechanical and scattering noises, and discuss their implications and mitigation solutions.

1.4 Publications

Below is a list of peer-reviewed journal publications, conference proceedings and patents resulting from the research undertaken during my PhD program.

1.4.1 Journal Publications

- Ya Zhang, Chathura P. Bandutunga, Terry G. McRae, Malcolm B. Gray, and Jong H. Chow, "Double Rayleigh scattering in a digitally enhanced, all-fiber optical frequency reference," *Opt. Express* 29, 26319-26331 (2021).

- Ya Zhang, Chathura P. Bandutunga, Malcolm B. Gray, and Jong H. Chow, "Multi-target CW interferometric acoustic measurements on a single optical beam," *Opt. Express* 27, 18477-18483 (2019).

- Chathura P. Bandutunga, Terry G. McRae, Ya Zhang, Malcolm B. Gray, and Jong H. Chow, "Infrasonic performance of a passively stabilized, all-fiber, optical frequency reference," *Opt. Express* 28, 9280-9287 (2020).

- C. P. Bandutunga, Y. Zhang, T. G. McRae, M. B. Gray and J. H. Chow, "Coherent Rayleigh Backscatter Phase Noise in Digitally Enhanced Fiber Interferometers," in *Journal of Lightwave Technology*, vol. 39, no. 8, pp. 2625-2630, 15 April15, 2021, doi: 10.1109/JLT.2021.3049567.

- Justin Wong, Chathura P. Bandutunga, Ya Zhang, Malcolm B. Gray, and Jong H. Chow, "Digitally enhanced molecular dispersion spectroscopy," *Opt. Lett.* 45, 6290-6293 (2020).

- Anneshwa Dey, Ya Zhang, Justin Wong, Paul G. Sibley, Chathura P. Bandutunga, Malcolm B. Gray, and Jong H. Chow, "Algebraic cancellation of inter-channel cross talk in multiplexed heterodyne interferometry," *Opt. Lett.* 46, 5830-5833 (2021).

1.4.2   Conference Proceedings

- Y. Zhang, C. P. Bandutunga, T. G. McRae, M. B. Gray, and J. H. Chow, "0.1 Hz / Hz Frequency Stability in a Passive, Optical Fiber Frequency Reference," in *Frontiers in Optics + Laser Science 2021*, paper JTh5A.68.

- Y. Zhang, Y. Zhao, C. P. Bandutunga, M. B. Gray and J. H. Chow, "Stereoscopic audio-band vibrometry with source triangulation and interferometric sensitivity," *2020 Conference on Lasers and Electro-Optics Pacific Rim (CLEO-PR)*, 2020, pp. 1-2, doi: 10.1364/CLEOPR.2020.P3_27.

- Y. Zhang, C. P. Bandutunga, Terry G. McRae, M. B. Gray, and J. H. Chow, "A Passive, Fiber Optic Frequency Reference with 0.1 Hz/$\sqrt{\text{Hz}}$ Stability," *14th Edoardo Amaldi Conference on Gravitational Waves* (2021).

- J. Wong, C. P. Bandutunga, Y. Zhang, M. B. Gray and J. H. Chow, "Digitally Enhanced Homodyne Dispersion Spectrometer," *2020 Conference on Lasers and Electro-Optics Pacific Rim (CLEO-PR)*, 2020, pp. 1-2, doi: 10.1364/CLEOPR.2020.C6D_4.

- J. Wong, C. P. Bandutunga, Y. Zhang, M. B. Gray, and J. H. Chow, "Molecular Gas Sensing in the Near Infrared using Digitally Enhanced Dispersion Spectroscopy," in *OSA Optical Sensors and Sensing Congress 2021*, paper ETh1A.4.

1.4.3   Patents

- Systems and methods for optical interferometric sensing, A. Dey, C.P. Bandutunga, Y. Zhang, M. B. Gray, J. H. Chow, J. C. T. Wong, P.G. Sibley, Australian Provisional Patent Application No. 2021902822. Date Filed: 31 Aug 2021.

Digitally Enhanced Homodyne Interferometry

In this chapter, we describe the readout technique used for this experiment, Digitally Enhanced Homodyne Interferometry (DEHoI) [121]. DEHoI represents the homodyne variant of Digital Interferometry (DI) [122], and was originated from its heterodyne counterpart, Digitally Enhanced Heterodyne Interferometry (DEHeI). To describe its working principles, we start with a comparison of the two types of interferometry in Section 2.1. This is followed by Section 2.2 which details DEHoI code operation and signal recovery.

From these two sections, we will notice the higher complexity of DEHoI compared to DEHeI due to the use of a multi-level code. This makes the heterodyne system the preferred testbed for redesigning and optimising DI operation for various purposes. In the past two years, several pieces of research have occurred on this front [123, 124, 125]. In Section 2.3, we highlight notable progress and discuss possibilities of transferring these improvements into DEHoI.

2.1 Heterodyne and Homodyne Interferometry

The difference between heterodyne and homodyne interferometry lies in the existence of a frequency offset between the two interfering beams. Figure 2.1 (a) shows a typical heterodyne Mach-Zehnder interferometer. Here, the top beam is frequency up-shifted by an acousto-optic modulator (AOM), while the bottom beam remains at its original optical frequency. We can write out the electric fields of the beams as:

$$E_{\text{sig}}(t) = A_{\text{sig}}\, e^{-i(\omega+\omega_h)t+i\phi_{\text{sig}}(t)}$$
$$E_{\text{ref}}(t) = A_{\text{ref}}\, e^{-i\omega t+i\phi_{\text{ref}}(t)}$$

(2.1)

where A_{sig} and A_{ref} are the respective amplitude of the two fields, ω is the angular optical frequency and ω_h is the offset, or heterodyne, frequency. $\phi_{\text{sig}}(t)$ and $\phi_{\text{ref}}(t)$ represent the

phase term of each arm. Upon re-combination, the interference pattern on the photo-detector can be written as:

$$
\begin{aligned}
\mathrm{PD} &= (E_{\mathrm{sig}} + E_{\mathrm{ref}})\,(E_{\mathrm{sig}} + E_{\mathrm{ref}})^* \\
&= A_{\mathrm{sig}}{}^2 + A_{\mathrm{ref}}{}^2 + 2\cos\left[\omega_h t - (\phi_{\mathrm{sig}}(t) - \phi_{\mathrm{ref}}(t))\right]
\end{aligned}
\tag{2.2}
$$

We can see that the interferometric phase $\phi_{\mathrm{sig}}(t) - \phi_{\mathrm{ref}}(t)$ is carried on the heterodyne frequency ω_h. In the Fourier frequency domain, the photo-detector spectrum is centred around the heterodyne frequency. Extraction of the interferometric signal therefore requires removal of this heterodyne frequency. This can be done by mixing with a sine and cosine function of the same angular frequency ω_h, and subsequently low-pass filtering to remove the second harmonic. This process is commonly referred to as IQ demodulation, as it extracts the in-phase (I) and quadrature (Q) components of the desired phase signal. Following this, an arctangent and unwrapping operation is used to recover the interferometric phase readout.
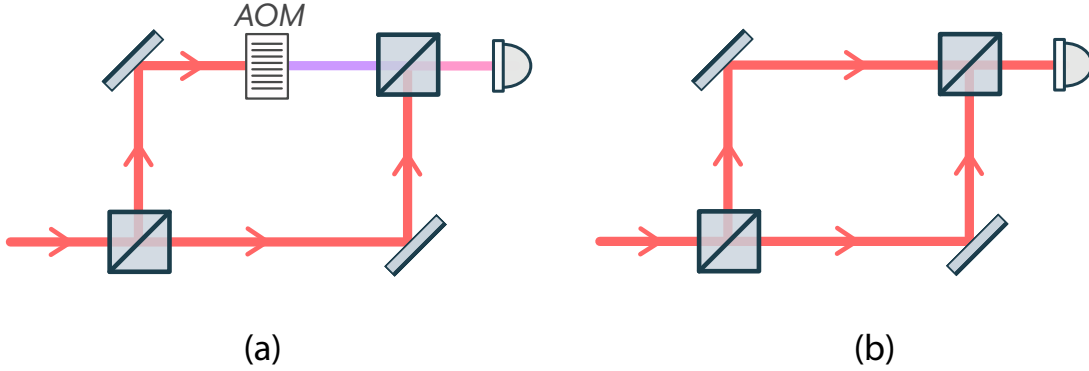
Figure 2.1: The key difference between heterodyne and homodyne interferometry lies in whether the two beams interfere at the same optical frequency. In the case of heterodyne interferometry (a), one arm is typically frequency shifted using an AOM, creating a RF beat note at the photo-detector. With homodyne interferometry (b), the photo-detector spectrum is centred at DC, requiring quadrature fringe locking or additional modulation for readout.

For a homodyne interferometer, as depicted in Fig. 2.1 (b), the two beams are interfered at the same optical frequency. Following the same field analysis as above, we can write out the electric fields and resultant photo-detector signal as follows:

$$
\begin{aligned}
E_{\mathrm{sig}}(t) &= A_{\mathrm{sig}}\,e^{-i\omega t + i\phi_{\mathrm{sig}}(t)} \\
E_{\mathrm{ref}}(t) &= A_{\mathrm{ref}}\,e^{-i\omega t + i\phi_{\mathrm{ref}}(t)} \\
\mathrm{PD} &= A_{\mathrm{sig}}{}^2 + A_{\mathrm{ref}}{}^2 + 2\cos\left[\phi_{\mathrm{sig}}(t) - \phi_{\mathrm{ref}}(t)\right]
\end{aligned}
\tag{2.3}
$$

Here the interferometric phase information is spectrally centred at DC. In the absence of the heterodyne carrier, the terms inside the cosine function no longer continuously ramp over multiple fringes. Consequently, extraction of the interferometric signal requires locking to the quadrature point, where linear readout and maximum sensitivity can be achieved. Alternatively, a re-introduction of phase modulation is required to sample the interferometric fringe. This modulation needs to minimally cover both the I and Q quadratures translating to a minimal amplitude of $\pi/2$, but does not need to be continuous.

Our DEHoI modulation falls into this latter category, with four discrete phase steps evenly distributed across the full fringe (2π) for balanced and complete IQ sampling.

2.2   DEHoI Principles

The centrepiece of a digitally enhanced system is the use of a pseudo-random sequence. This sequence, or code, has desirable auto-correlation properties. By encoding the optical beam with the code and later demodulating with its delayed copy, DI enables signal selection and isolation by its optical time-of-flight. Additionally, as multiple demodulation channels can be easily parallelised in digital signal processing (DSP), DI also enables continuous, multi-point sensing without requiring any additional optical hardware [126, 127].

2.2.1   M-Sequence Auto-correlation

Traditionally, a maximum-length sequence, or m-sequence, is used for DI modulation [122]. The m-sequence [128] belongs to a family of pseudo-random codes used for multiplexing in telecommunication [129, 130, 131, 132], and was selected for DI due to its flexible length and favourable auto-correlation. An m-sequence appears statistically-random but is deterministically generated using linear feedback shift registers (LFSRs) with specific tap configurations [133, 134]. Possible arrangements of these taps are documented in [135]. The "bit depth" of a LFSR refers to the number of states it contains, and this determines the length of the m-sequence generated. For an N-bit LFSR, the m-sequence length is given by

$$L = 2^N - 1 \tag{2.4}$$

An example of a 4-bit m-sequence is shown in Fig. 2.2. Each element of the sequence is referred to as a "chip" and the speed at which this chip is encoded onto the optical phase is referred to as the "chip frequency". As the code only has 15 elements, the chips repeat after the full length, and the speed of this repetition is referred to as the "code frequency".
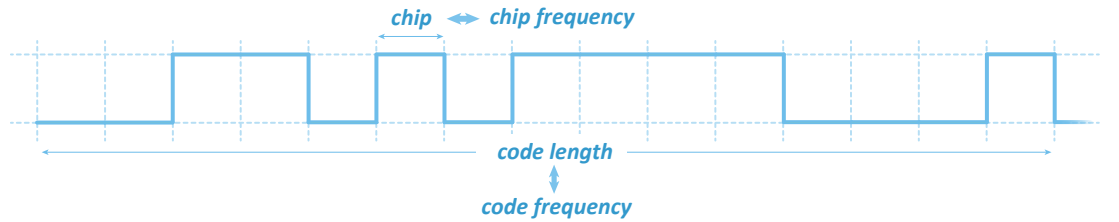
Figure 2.2: An example of a 4-bit m-sequence, identifying a chip, chip frequency, code length and code frequency.

In DI, elements of the code are mapped to certain phase steps and modulated onto one or more optical beams. The modulation typically occurs prior to optical propagation and is done using an AOM or an electro-optic modulator (EOM). Figure 2.3 shows a simple heterodyne DI setup. In this configuration, the optical phase of the signal arm is pseudo-randomly flipped by 0 or π according to the m-sequence, while the reference arm remains

unmodulated. The resulting heterodyne beat note becomes scrambled and incoherent. Its Fourier spectrum, also illustrated in the figure, consists of multiple teeth evenly spaced at the code frequency f_{code} and centred at the heterodyne frequency f_h. We also observe nulls at integer multiples of the chip frequency f_{chip}.
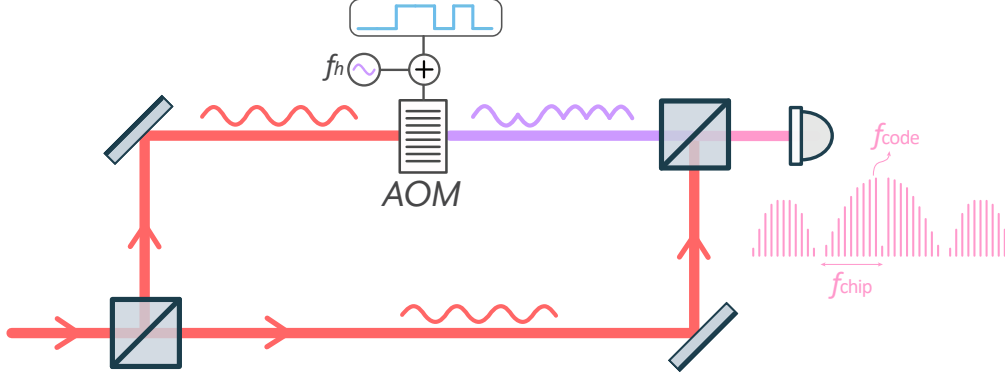
Figure 2.3: The optical compartment of a standard heterodyne DI setup consists of a frequency-shifted, encoded signal arm and an unmodulated reference arm. Upon recombination, this produces a scrambled and incoherent signal pattern. In the Fourier domain, the spectrum consists of equivdistant peaks separated by f_{code} and centred at f_h, with nulls at multiples of f_{chip}.

From the photo-detector spectrum, we see that without demodulation, the optical coherence is reduced to the physical length occupied by a single chip. This sets the limit for the ranging resolution of DI. As a convention, we define this physical length as the "chip length":

$$L_{\text{chip}} = \frac{c}{n \times f_{\text{chip}}} \tag{2.5}$$

where c is the speed of light, n the refractive index, and f_{chip} the chip frequency. For this experiment, L_{chip} lands on the order of a few meters in fibre, significantly shorter than our intended interferometer pathlength difference. This means that demodulation is required to restore coherence and reconstruct the desired interference pattern. In addition, through this process, spurious interference outside of this range gate can be suppressed. This is achieved through the auto-correlation property of the m-sequence, which is a measure of the code's similarity with respect to itself. Mathematically, the auto-correlation value as a function of relative delay can be written as

$$A(\tau) = \sum_{n=1}^{L} C(n)\, C(n - \tau) \tag{2.6}$$

where τ is the relative delay and L the code length. $C(n)$ represents a single element of the code and takes the value of 1 or -1. For an m-sequence:

$$A(\tau) = \begin{cases} L & \text{when } \tau = mL, \text{ where } m \text{ is an integer} \\ -1 & \text{all other delays} \end{cases} \tag{2.7}$$

In Fig. 2.4, we plot the auto-correlation value $A(\tau)$ of the 4-bit example m-sequence as a function of the delay τ. We can see that $A(\tau)$ reaches a peak value of 15 when the

sequence is matched with itself. For all other delays, the correlation value lies at -1. To understand how this can be utilised for ranging and gating, we consider the demodulation process as illustrated in Fig. 2.5.
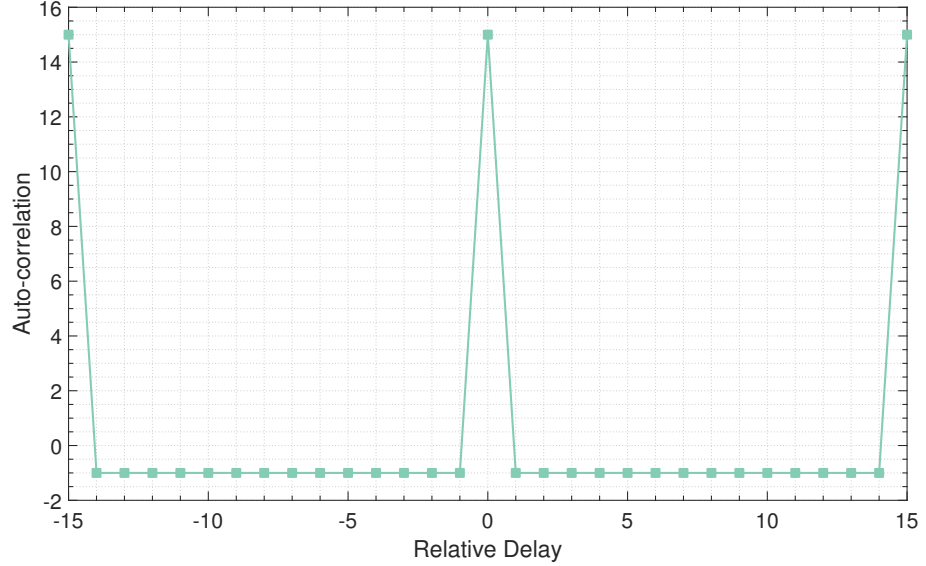
Figure 2.4: The auto-correlation value of the above a 4-bit example m-sequence is plotted against its relative delay against itself. We can see a peak value of 15 when the delay is 0 or integer multiples of the code length, and -1 at all other delays. This enables code-delay based selective interference for digital interferometry.
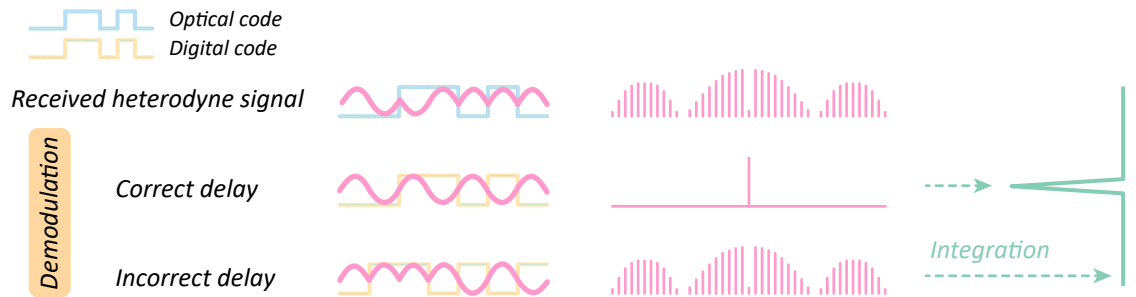
Figure 2.5: Illustration of the DI demodulation process, which is mathematically equivalent to the calculation of m-sequence auto-correlation. First, a digitally-delayed code is used to decode the received heterodyne signal. Successful recovery of the heterodyne beat note only occurs when the digital delay matches with the optical time-of-flight delay. Second, the decoded wave is summed over a full code length. Depending on the digital code delay, the interference signal is either maximally recovered or suppressed in amplitude.

In Fig. 2.5, the received heterodyne signal remains scrambled by the m-sequence modulation. In addition, the m-sequence is delayed through optical propagation between the modulator and the photo-detector. In digital signal processing, a digitally delayed copy of the code is prepared to undo the modulation. This decoding process successfully restores the original heterodyne beat note only when the digital code matches with the optical one, and this is labelled "correct delay" in the figure. Spectrally, this collapses the spread, modulated spectrum back into a single peak at the heterodyne frequency.

The second scenario, labelled "incorrect delay", decodes the received signal at a mismatched delay. In mathematical terms, this can be thought of as multiplying with a second m-sequence consisting of -1s and 1s, where -1 represents a π phase flip. When these two sequences are unmatched, the result of this multiplication is yet another m-sequence. Therefore, the incorrectly decoded signal remains scrambled and retains the same spectral shape.

The final step of DI demodulation involves integration over a full code, which completes the computation of the auto-correlation function. Referring back to Eqn. 2.6, the decoding process is mathematically represented as $C(n)C(n-\tau)$, whereas the integration sums this product over a full code length L. This links DI demodulation to a specific point on the auto-correlation function, based on the relative delay between the optical code and the digital code. As the digital delay is easily variable, this gives the ability to "range" to the desired optical time-of-flight and isolate that section of interference, while suppressing signals from other delays by a factor of L. In addition, multiple demodulation channels with individual code delays can be parallelised in DSP, allowing multiplexing of signals on a single, continuous-wave beam [136].

2.2.2   Four-level Phase Modulation

Having covered the fundamental principles of DI, we now proceed to discuss the implementation of homodyne DI. Here, we continue to leverage the auto-correlation property of the modulation code to achieve time-of-flight based signal isolation and background noise suppression. Unlike heterodyne DI, however, in homodyne we no longer have the beat frequency acting as a continuous phase scan across all quadratures. This difference can be conceptually visualised in Fig. 2.6. In replacement of the phase scan and in order to still sample both I and Q quadratures, we introduce phase-stepping for homodyne DI. Although the functionally minimal step level is two, the structural imbalance and incompleteness of that sampling architecture often result in systematic issues. Instead, we use four levels of phase stepping as shown below in "DEHoI". This allows us to cover all four quadratures of the interferometric fringe and provides a more robust phase readout.

To achieve this, we construct a four-level code from two m-sequences at different delays, as illustrated in Fig. 2.7 (a). For convention, we refer to them as I and Q codes respectively and the resultant four-level code as quadrature phase shift key (QPSK) code. We can write the generation logic for QPSK as follows:

$$
\begin{aligned}
C_I(n) \;\&\; C_Q(n) &= \pi/4 \\
\overline{C_I(n)} \;\&\; C_Q(n) &= 3\pi/4 \\
\overline{C_I(n)} \;\&\; \overline{C_Q(n)} &= -3\pi/4 \\
C_I(n) \;\&\; \overline{C_Q(n)} &= -\pi/4
\end{aligned}
\tag{2.8}
$$

where & and $\overline{C(n)}$ represent the logical AND and NOT operation respectively. Figure 2.7 (b) shows the four-level modulation on a phasor diagram. As we can see, the constellation is evenly distributed across all four quadratures. Additionally, given the statistical balance
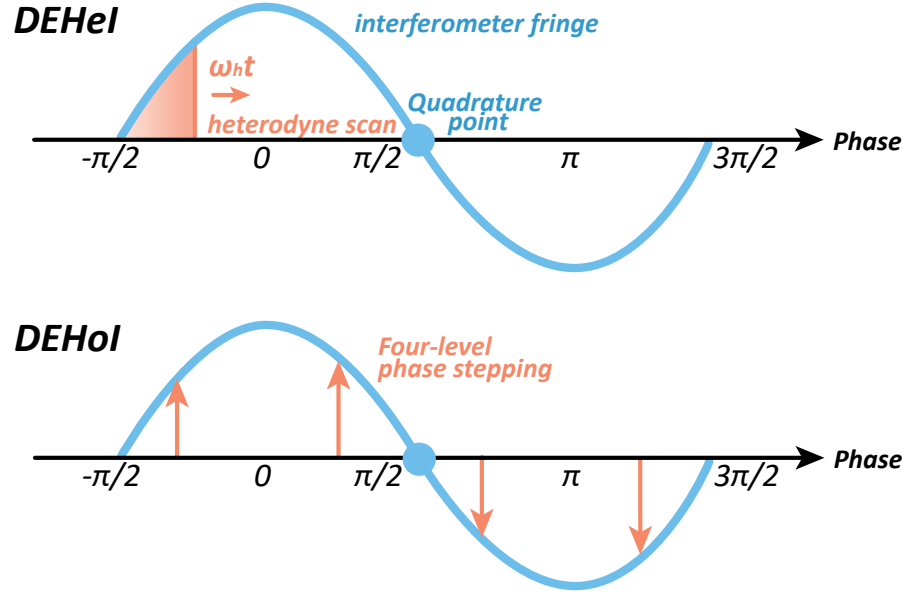
Figure 2.6: Conceptual visualisation of the difference between heterodyne and homodyne DI. In DEHeI, the beat frequency creates a continuously evolving phase which scans all quadratures of the interferometer fringe. For DEHoI, this scan is absent and replaced by a four-level code which pseudo-randomly steps and samples the quadratures.

of m-sequences, where the number of 0s and 1s differ only by 1 element regardless of code length, the amount of time spent by QPSK in each quadrature is also largely balanced.
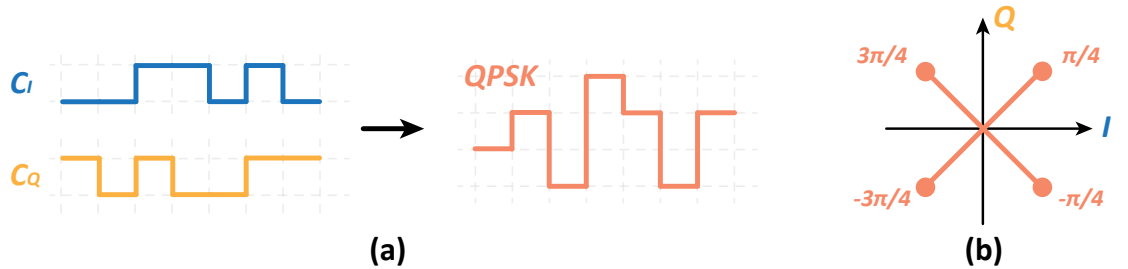
Figure 2.7: (a) Generation of QPSK from two differently delayed m-sequences. (b) Visualisation of the four-level code on a phasor diagram.

Optically, DEHoI modulation can be set up in two ways. The first configuration modulates only a single arm using an electro-optic modulator (EOM), or two AOMs with the same frequency shift placed in both arms. The second configuration moves the modulator upfront, coding both optical beams with the same QPSK sequence, as illustrated in Fig. 2.8. For this, the two beams must have different optical delays in order for them to be code-distinguishable, and demodulation directly recovers their differential phase. Our experiment adopts the second configuration to remove any active components in key interferometer paths. This provides greater thermal stability for the system with the trade-off being higher digital complexity.
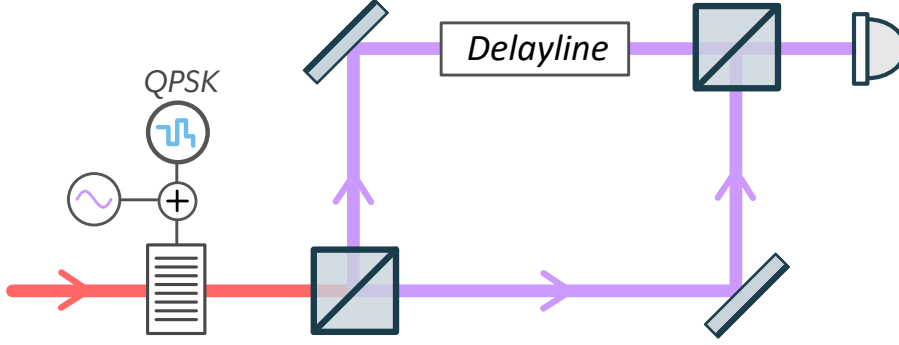
Figure 2.8: Front-modulation DEHoI optical setup. In this configuration, the QPSK modulation occurs upfront, encoding both optical beams with the same sequence. The resultant interference can be recovered through two demodulation stages each matching the optical delay of one arm.

2.2.3 Double Demodulation

Due to the presence of two QPSK codes in our optical configuration, a dual-stage demodulation architecture is implemented. This process decodes both interferometer arms at their respective delays, τ_1 and τ_2, and can be mathematically represented as:

$$
\begin{aligned}
\text{II}(t) &= \text{PD}(t) \times C_I(t - \tau_1)\, C_I(t - \tau_2) \\
\text{IQ}(t) &= \text{PD}(t) \times C_I(t - \tau_1)\, C_Q(t - \tau_2) \\
\text{QI}(t) &= \text{PD}(t) \times C_Q(t - \tau_1)\, C_I(t - \tau_2) \\
\text{QQ}(t) &= \text{PD}(t) \times C_Q(t - \tau_1)\, C_Q(t - \tau_2)
\end{aligned}
\tag{2.9}
$$

where PD refers to the received photo-detector signal, and C_I and C_Q are binary codes of ± 1s. To interpret the physical meaning of these four components, we refer to Fig. 2.9. On the left (a), the electric field in each interferometer arm is represented by its own coordinate frame, where the I and Q components form the two axes. The instantaneous phasor of the two fields are given by the multiplication of their amplitude and the unit code vectors, written in their own frame as follows:

$$
\begin{aligned}
\vec{E}_{\text{ref}}(t) &= [A_{\text{ref}} C_I(t - \tau_1),\ A_{\text{ref}} C_Q(t - \tau_1)] \\
\vec{E}_{\text{sig}}(t) &= [A_{\text{sig}} C_I(t - \tau_2),\ A_{\text{sig}} C_Q(t - \tau_2)]
\end{aligned}
\tag{2.10}
$$

In optical propagation, these two frames rotate away from their original orientation, by ϕ_{ref} and ϕ_{sig} respectively. In order to compare the two electric fields above, we need to bring them back into the same coordinate frame. For simplicity, in Fig. 2.9(b), we use the reference field coordinate frame and convert the signal field towards it, giving:

$$
\begin{aligned}
\vec{E}_{\text{ref}}(t) =\ & [A_{\text{ref}} C_I(t - \tau_1),\ A_{\text{ref}} C_Q(t - \tau_1)] \\
\vec{E}_{\text{sig}}(t) =\ & [A_{\text{sig}} C_I(t - \tau_2) \cos \Delta\phi - A_{\text{sig}} C_Q(t - \tau_2) \sin \Delta\phi, \\
& A_{\text{sig}} C_I(t - \tau_2) \sin \Delta\phi + A_{\text{sig}} C_Q(t - \tau_2) \cos \Delta\phi]
\end{aligned}
\tag{2.11}
$$

where $\Delta\phi = \phi_{\text{sig}} - \phi_{\text{ref}}$. With this in mind, we can write the optical power at the photo-detector as follows:

$$
\text{PD}(t) = \left| \vec{E}_{\text{ref}}(t) + \vec{E}_{\text{sig}}(t) \right|^2
\tag{2.12}
$$
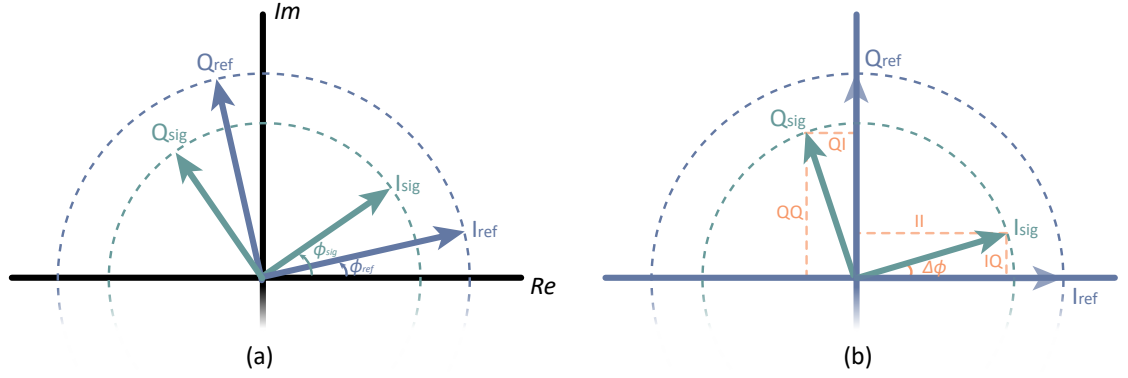
Figure 2.9: The vector coordinate of each electric field is defined in its own coordinate frame, which rotates through optical propagation. The relative angle between the reference and signal frames represents the desired interferometric phase, which is recovered through a dual-stage demodulation process, producing projections from the signal frame to the reference frame.

Expanding this out and isolating the RF terms, we have

$$\text{PD}(t) = 2A_{\text{ref}}A_{\text{sig}}[C_I(t-\tau_1)C_I(t-\tau_2)\cos\Delta\phi - C_I(t-\tau_1)C_Q(t-\tau_2)\sin\Delta\phi$$
$$+ C_Q(t-\tau_1)C_I(t-\tau_2)\sin\Delta\phi + C_Q(t-\tau_1)C_Q(t-\tau_2)\cos\Delta\phi] \tag{2.13}$$

Observing Eqn. 2.13, we see that the desired interferometric phase $\Delta\phi$ is split into in-phase and quadrature components via effectively "code interference". For example, in the first and last terms, the I or Q codes interfere with a time-shifted version of themselves, obtaining no additional phase shift and therefore representing the "in-phase" component. The middle two terms, on the other hand, interfere with their orthogonal codes with a $\pi/2$ relative phase difference, thereby representing the "quadrature" part.

We can also see now that the double decoding process, written in Eqn. 2.9, retrieves in-phase and quadrature components of interest. Noting that when delays are matched, the code cancels to 1 ($C(t-\tau) \times C(t-\tau) = 1$), we can re-write Eqn. 2.9 as:

$$\text{II}(t) = 2A_{\text{ref}}A_{\text{sig}}[\cos\Delta\phi - C_I(t-\tau_2)C_Q(t-\tau_2)\sin\Delta\phi$$
$$+ C_I(t-\tau_1)C_Q(t-\tau_1)\sin\Delta\phi + C_I(t-\tau_1)C_I(t-\tau_2)C_Q(t-\tau_1)C_Q(t-\tau_2)\cos\Delta\phi]$$
$$\text{IQ}(t) = 2A_{\text{ref}}A_{\text{sig}}[C_I(t-\tau_2)C_Q(t-\tau_2)\cos\Delta\phi - \sin\Delta\phi$$
$$+ C_I(t-\tau_1)C_I(t-\tau_2)C_Q(t-\tau_1)C_Q(t-\tau_2)\sin\Delta\phi + C_I(t-\tau_1)C_Q(t-\tau_1)\cos\Delta\phi]$$
$$\text{QI}(t) = 2A_{\text{ref}}A_{\text{sig}}[C_I(t-\tau_1)C_Q(t-\tau_1)\cos\Delta\phi$$
$$- C_I(t-\tau_1)C_I(t-\tau_2)C_Q(t-\tau_1)C_Q(t-\tau_2)\sin\Delta\phi$$
$$+ \sin\Delta\phi + C_I(t-\tau_2)C_Q(t-\tau_2)\cos\Delta\phi]$$
$$\text{QQ}(t) = 2A_{\text{ref}}A_{\text{sig}}[C_I(t-\tau_1)C_I(t-\tau_2)C_Q(t-\tau_1)C_Q(t-\tau_2)\cos\Delta\phi$$
$$- C_I(t-\tau_1)C_Q(t-\tau_1)\sin\Delta\phi + C_I(t-\tau_2)C_Q(t-\tau_2)\sin\Delta\phi + \cos\Delta\phi]$$

$$\tag{2.14}$$

Further combining relevant parts to simplify:

$$
\begin{aligned}
\mathrm{I} &= \mathrm{II} + \mathrm{QQ} \\
&= 4A_{\mathrm{ref}}A_{\mathrm{sig}}[\cos\Delta\phi + C_I(t-\tau_1)C_I(t-\tau_2)C_Q(t-\tau_1)C_Q(t-\tau_2)\cos\Delta\phi] \\
\mathrm{Q} &= \mathrm{QI} - \mathrm{IQ} \\
&= 4A_{\mathrm{ref}}A_{\mathrm{sig}}[\sin\Delta\phi - C_I(t-\tau_1)C_I(t-\tau_2)C_Q(t-\tau_1)C_Q(t-\tau_2)\sin\Delta\phi]
\end{aligned}
\tag{2.15}
$$

Upon integration over the code, the second terms in both equations become suppressed by a factor of code length relative to the first. We can therefore approximate I and Q as:

$$
\begin{aligned}
\mathrm{I} &\approx 4A_{\mathrm{ref}}A_{\mathrm{sig}}\cos\Delta\phi \\
\mathrm{Q} &\approx 4A_{\mathrm{ref}}A_{\mathrm{sig}}\sin\Delta\phi
\end{aligned}
\tag{2.16}
$$

And the interferometric phase is recovered through an arctangent operation:

$$
\Delta\phi = \arctan\left(\frac{Q}{I}\right)
\tag{2.17}
$$

FPGA Implementation

The decoding algorithm was traditionally implemented on the FPGA through a binner architecture, as detailed in [134]. This has since been upgraded to a new structure illustrated in Fig. 2.10. These two designs are functionally identical. While the new implementation strictly follows the form of Eqn. 2.9, the binner approach rotates the frame of reference by 45 degrees. This rotation merely introduces a static phase offset and is of no consequence to the intended measurement. As such, the binner approach was preferred due to its higher computational efficiency, until recently when onboard resources were no longer of concern.

As seen in Fig. 2.10, in the new decoding arrangement, I and Q codes are individually delayed by reference and signal delays. These delayed codes are then fed into six separate decoders, which flip the phase of their input as per the binary code. The decoded components, II, IQ, QI and QQ, are then further combined into the in-phase and quadrature parts of the interferometer phase.

The main advantage of the new decoding architecture over the binning design lies in a higher signal amplitude and therefore greater clearance of ADC noise floor. In the binner design, half of the ADC samples are effectively discarded as a result of the frame rotation. That is, each sample is binned entirely towards one quadrature, and the other is left empty. In contrast, in the upgraded design, every sample is projected towards both quadratures and this allows a higher total signal amplitude. While this upgrade does not improve the noise floor of this specific experiment, it could be useful for DEHoI experiments close to the quantisation limit.
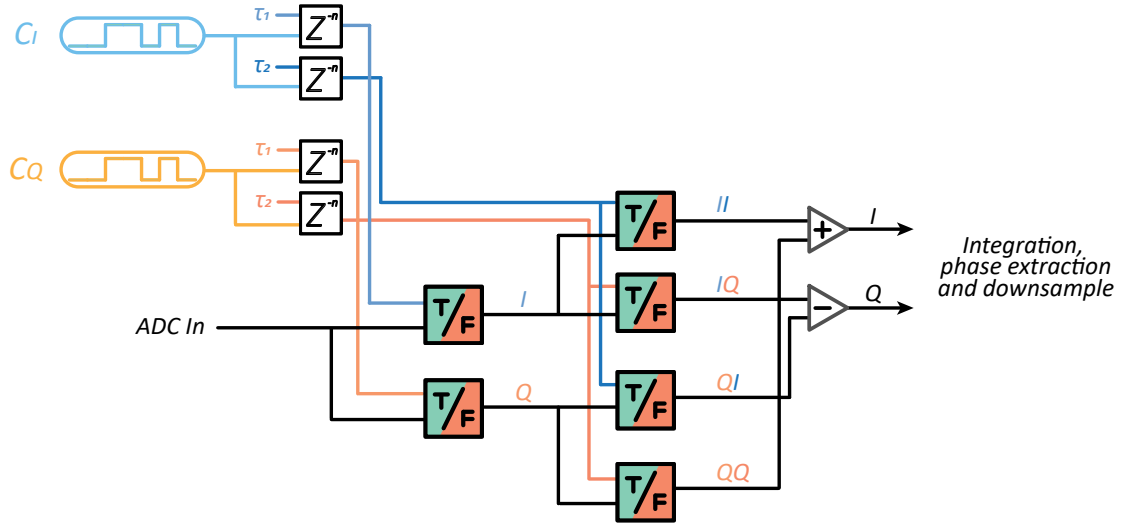
Figure 2.10: Updated FPGA implementation of dual-stage demodulation. Delayed copies of the I and Q codes are prepared to drive six individual decoders, as per the algorithms set out in Eqn. 2.9. Following this, four decoded components are combined to create the I and Q parts of the desired interferometric phase.

2.2.4   Limitations of QPSK

It is worth pointing out here that the use of QPSK, while optically simple and thermally stable, has a severe drawback. Compared to an m-sequence auto-correlation, QPSK has a modified correlation profile with a lower signal-to-noise suppression and some "ghost correlations" at non-signal delays. This is caused by each of the two m-sequences split across the I and Q quadratures. Upon integration of I and Q, only half the elements within each code gather, forming a partial, incomplete and compromised auto-correlation.

Some circumvention strategies exist for specific DEHoI designs. For example, in a single-beam modulation system, a square wave could replace one m-sequence and serve the function of quadrature switching. This square wave would then run at integer multiples of the speed of the remaining m-sequence, allowing all elements of the m-sequence to appear in both quadratures and form a full correlation. Unfortunately, this strategy cannot be readily implemented in a dual-beam modulation system as it has specific requirements on the interferometer length relative to the modulation speed.

Another potential solution involves the use of a two-layer code. This idea is commonly explored in global positioning system (GPS) and global navigation satellite system (GNSS), where a short, coarse acquisition code (C/A) and a long precision (P) code are emitted at different speeds to refine ranging precision and eliminate ambiguity [137, 138]. When applied to DEHoI, this would involve embedding a precision code within a coarse code. The latter switches between quadratures while the entire length of the former samples that quadrature. This architecture could benefit DEHoI systems tracking small and slow signals, but is unsuitable for our experiment. As data throughput is reduced to the code frequency of the top, slow code, this limits the ability to track fast dynamics. In addition, signal fidelity is also reduced as a result of the extended time between switching quadratures, to the point they may no longer be approximated as simultaneously sampled.

Perhaps the final and ultimate solution for this experiment is to pre-construct a four-level sequence with elements that form a full correlation in each quadrature. This is an ongoing field of exploration. If such a code exists, we expect at least a factor of 2 bandwidth penalty at the same modulation rate and background noise suppression. As we will discuss in Chapter 5.2, this may in fact create further complications between two opposing noise sources in this experiment.

Beyond a compromised correlation, there are two other potential limitations in DEHoI. First, as the m-sequence we use has an odd number of elements, there is a slight unbalance in the amount of sampling time in each quadrature. This effect can be seen in DEHeI too, when the heterodyne frequency is not carefully selected to compliment the DI chip frequency [123]. Second, QPSK sampling is by design sequential, not simultaneous. In the limit of a very slow chip frequency, the phase excursions in between chips may invalidate the sequentially sampled I and Q data. This issue again is not unique to homodyne, and can occur in heterodyne systems if the phase ramp is too slow. Fortunately, our experiment operates at a fast chip frequency and uses a long code with negligible imbalance, and is therefore unaffected by either of these limits.

2.3   Further Developments

Although not directly related to this experiment, some interesting DI developments have occurred over the past two years, with a strong focus on crosstalk improvement in a heterodyne arrangement. In this section we provide a brief overview of progress to date, and discuss how the improvements could be transferred into a homodyne system.

Crosstalk in a multiplexed system refers to the contamination of one channel due to signals present on the other channels. In DI, incorrectly decoded delays are suppressed by $1/L$, where L is the code length, and this sets the crosstalk limit. In practice, the achieved crosstalk suppression is often worse depending on specific DI configurations. Ref. [123] details optimal parameters for DEHeI to fully leverage the auto-correlation property of the m-sequence. To achieve suppression beyond $1/L$, however, other techniques are required.

To this end, two approaches have been considered. The first one looks for alternative codes with more desirable auto-correlation properties. For example, Golay complimentary pairs [139, 140, 141, 142] with inverted auto-correlation profiles have been demonstrated to achieve substantially increased crosstalk suppression [21]. A1 and A2 sequences have also been explored [143]. These are generated from M and have periodic zero-crossings in auto-correlation, allowing improved suppression at those delays.

While these alternative codes readily improve crosstalk suppression for DEHeI, in transferring their advantages into DEHoI, the difficulty remains in designing a QPSK that retains the desired auto-correlation property.

The second approach to improve crosstalk suppression, also developed for and demonstrated in DEHeI, seeks to engineer an ideal auto-correlation profile [124]. This works by summing two correlation profiles at different demodulation delays to cancel out most of the background. Taking the m-sequence as an example, two demodulation channels

are simultaneously operated, one at the intended signal delay, the other at a non-physical delay. In addition, the second channel is phase-flipped by 180° such that its effective correlation is inverted and cancels the -1s in first channel. This technique is named offset decoding and is similarly applicable to other binary codes. Utilising offset decoding in DEHoI remains a work in progress with the primary challenge remaining modified QPSK correlation.

To conclude the above discussions, the single bottleneck in transferring these developments into DEHoI lies still in QPSK correlation. It is clear that the current QPSK generation method is unoptimised for crosstalk suppression, and a redesigned algorithm is required. The new four-level code would consist two binary sub-codes, each capable of forming a full correlation after being "binned" and separated into I or Q quadrature by the other code. Designing such a code is a non-trivial task, but could fundamentally transform DEHoI if successful. Currently, techniques for minimising crosstalk and rejecting background noise are more mature and adaptable in DEHeI. However, due to the thermal requirement of our experiment, we are unable to use DEHeI and therefore trade noise suppression for optical stability.

Experimental Design and Implementation

This chapter focusses on the design and implementation of this experiment. In Section 3.1 to 3.4, we discuss optical, thermal and mechanical considerations that inform physical design decisions. We then shift focus to software components that interface with the optical system, and detail the FPGA implementation of DEHoI algorithms in Section 3.5.

3.1 Optical Configuration

We recall from Chapter 1.2–1.3 that a fibre frequency reference reads out source frequency noise through an armlength mismatched interferometer. This presents two design questions: what type of interferometer to use, and what differential armlength is optimal. We now discuss available options for both parameters and their suitability for this experiment.

3.1.1 Optical Topology

The optical topology is perhaps the most fundamental piece of hardware upon which the rest of experimental design is based. It also sets the ultimate achievable performance for the system. As DEHoI allows the extraction of interferometric phase between any two fibre lengths that exhibit temporal code incoherence, several optical arrangements are possible and have been demonstrated. These are depicted in Fig. 3.1, including a re-entrant fibre loop [144] (a), a Michelson interferometer [145, 146] (b) and a Mach-Zehnder interferometer [147] (c).

In all three configurations, the encoded laser light is split to interrogate two identical or near-identical interferometers. The existence of the second interferometer serves to remove common input frequency noise, and allows performance characterisation through the subtraction of the two readouts. This requires the two interferometers to exhibit thermal, acoustic and mechanical independence, such that these noises, if present, are not removed from the stability measurement. For the first topology of Fig. 3.1(a), the design is such that the backwards illuminated path can constitute the second interferometer, maximising
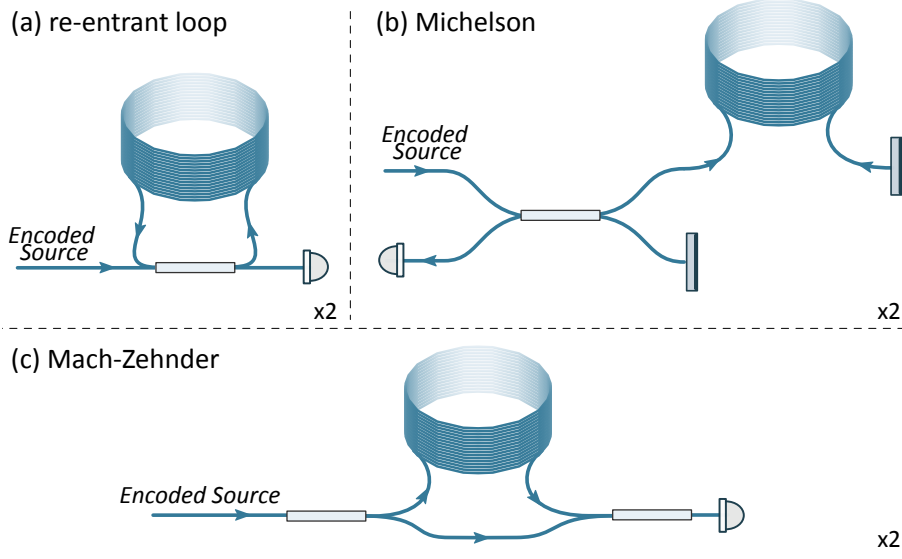
Figure 3.1: Possible optical topologies for a fibre frequency reference. (a) re-entrant fibre delayline. (b) an armlength unbalanced Michelson interferometer. (c) Mach-Zehnder interferometer representing the topology used in this experiment.

optical simplicity while ensuring the two systems are perfectly identical. However, the common fibre coil gives rise to high correlation between the twin interferometers, meaning the differential measurement does not truly measure the relative frequency stability.

In addition to intra-interferometer dependence, the optical configuration of (a) also suffers from increased crosstalk noise. As beams traversing multiple coil lengths interfere, a number of undesired interference signals are presented at the photo-detector. The proportion of these signals unsuppressed by DEHoI in demodulation then constitutes crosstalk noise (referred to as "code noise" in [148]) and requires further management. Overall, configuration (a) is not an optimal design.

The Michelson configuration (b) uses two separate interferometers that can be considered independent up to the limit of their environmental commonality. A major drawback with this design however lies in the reflective nature of the measurement, which allows the coupling of first-order Rayleigh backscattering (RBS) [149]. RBS contaminates phase measurements by introducing a statistically random error on top of the metrology field. Although DEHoI strongly suppresses RBS outside the range gate, the amount of noise contribution from within a chip still limits the low-frequency performance of this system.

Another issue associated with operating in reflection is the use of Faraday mirrors to manage polarisation wander [150]. This increases the interferometer's susceptibility to magnetic field variations and can be especially problematic for low-Earth spacecrafts or aeroplanes. In addition, Faraday mirrors are assembled from free-space optics and therefore susceptible to vibration, stress and dynamic misalignment. These effects can be further exacerbated by long lengths of fibre and significantly degrade system robustness. Therefore topology (b) is also non-ideal for a fibre frequency reference.

Finally, the Mach-Zehnder topology (c) is selected for this experiment. The single-pass, transmissive design eliminates crosstalk and first-order RBS noise seen by the previous two configurations. The two interferometers are also separately housed in individual isolation chambers to maximise independence, as will be discussed in Section 3.2. Despite these improvements, two technical challenges remain in this optical configuration.

First, the transmissive system is still susceptible to second-order RBS, where the doubly-backscattered field realigns with the forward-propagating metrology field. This second-order effect generates lower noise than its first-order counterpart, however as we will see in Chapter 5.2, it still becomes a noise floor limitation even with DEHoI code suppression. Because Rayleigh scattering is unavoidable with current fibre manufacturing technologies, it remains a fundamental challenge to any fibre interferometry.

Second, the ideal condition for relative frequency stability measurement is two independent yet identical interferometers. In pursuing independence and separating the construction and isolation of the two interferometers, minor and unintentional differences inevitably accumulate and reduce identicalness. Examples include differences in fibre coil length, splice loss and stress, and thermal profile of the isolation chambers. These result in different measurements of source frequency noise and therefore incomplete cancellation of it. We will quantify this effect later in Section 4.5.

We note that the above two challenges are not unique to the selected Mach-Zehnder configuration (c). Second-order RBS effects are difficult to mitigate, while system independence is essential for valid noise floor characterisation, despite the trade-off of potentially higher residual frequency noise. Overall, the optical topology of (c) still represents the most favourable design amongst all three possibilities.

3.1.2   Interferometer Armlength

The second consideration in designing a fibre frequency reference is the armlength difference of the interferometer, which determines how laser frequency noise is translated into the interferometric phase readout. To illustrate this, we consider the corresponding transfer function for an unbalanced interferometer. The input and output of the system can be written as:

$$f(t) = \frac{1}{2\pi}\frac{d\phi(t)}{dt}$$
$$y(t) = \phi(t-\tau) - \phi(t) \tag{3.1}$$

where $f(t)$ is the instantaneous laser frequency, defined by the derivative of its phase $\phi(t)$. τ is the interferometer transit time corresponding to the armlength difference ($\tau = nL/c$), and $y(t)$ is the DEHoI phase readout. The transfer function of the interferometer is obtained by taking the Laplace transform of the above terms, giving:

$$H(s) = \frac{Y(s)}{F(s)} = \frac{e^{-\tau s}\Phi(s) - \Phi(s)}{\left[s\Phi(s) - \phi(0)\right]/2\pi} = 2\pi\frac{e^{-\tau s} - 1}{s} \tag{3.2}$$

Substituting in $s = 2\pi i f$, we plot the amplitude response, $|H(s)|$, for three different armlength differences ($L = 1$ km, 5 km and 20 km) in Fig. 3.2. All three traces have a flat response at low Fourier frequencies, followed by periodic nulls and a $1/f$ roll-off,

as indicated by the black dashed line. Here, we define two parameters via the armlength difference to quantitatively describe the response curves: the optical gain and the free-spectral range (FSR).
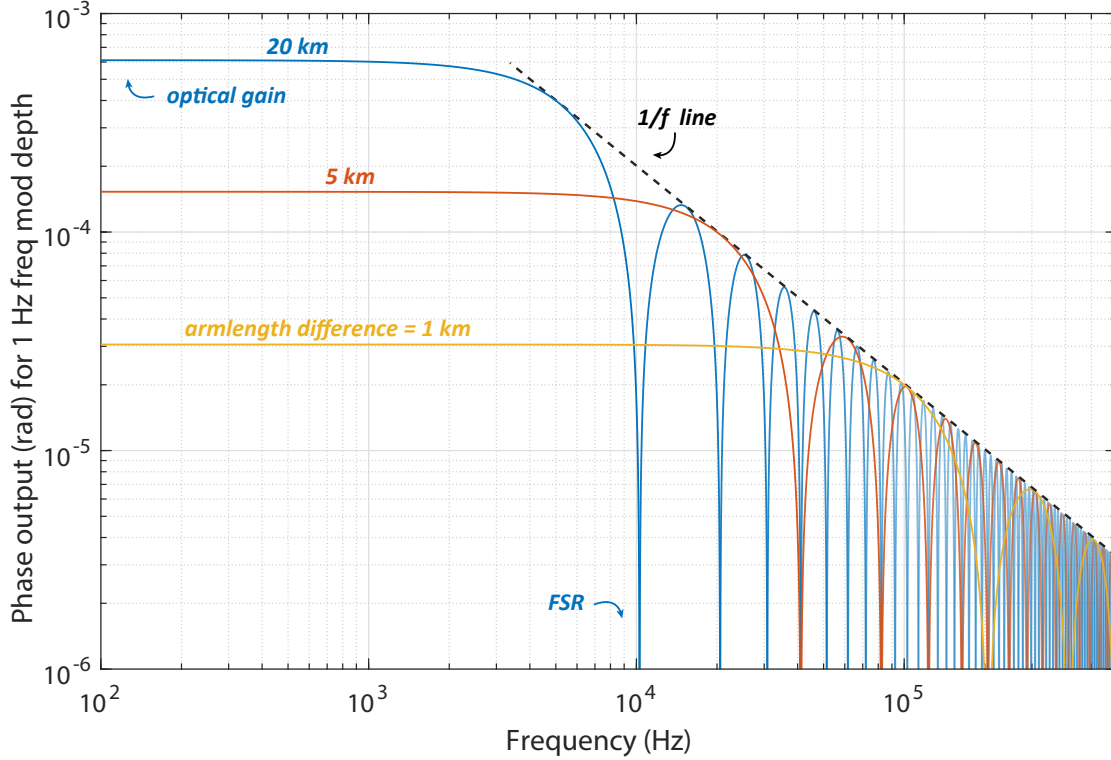
Figure 3.2: Amplitude response of three unbalanced interferometers, with armlength differences of 1 km, 5 km and 20 km respectively. The armlength difference quantitatively determines the optical gain at low Fourier frequencies and the free-spectral range (FSR) of the interferometer. All three traces intersect and roll off with a $1/f$ line.

Optical Gain

The optical gain refers to the height of the flat region in the above transfer functions, and is a measure of the interferometer's responsivity to laser frequency noise. Its numerical value represents the amount of phase excursion in radians arising from one unit Hz of frequency modulation, at a specific Fourier frequency. In practice, the Fourier frequency dependence is low when $f \ll \text{FSR}$, and the optical gain is considered constant in that region.

Comparing the three traces of Fig. 3.2, we can see that the optical gain scales proportionally with armlength difference. Indeed, equal armlength interferometers, such as a Sagnac or a balanced Michelson interferometer, have near-zero optical gain and are thus highly immune to input frequency noise. As our goal is to directly measure frequency noise, we look to the other end of the spectrum and favour a higher optical gain. This would allow larger phase signals from the same source frequency fluctuation and give a higher clearance from measurement noises. This benefit however does not scale indefinitely, as dynamic range and tracking requirements also increase with optical gain and eventually become a limitation.

Free-Spectral Range

The FSR refers to the Fourier frequency at which the first null occurs on the transfer function. Mathematically, the FSR is the inverse of the interferometer transit time τ and inversely proportional to the armlength difference L:

$$\text{FSR} = 1/\tau = c/(nL) \tag{3.3}$$

where c is the speed of light and n is the refractive index in fibre. The FSR sets the measurement cut-off frequency, beyond which laser dynamics are filtered and distorted by the interferometer. A lower FSR frequency also generates a higher amount of non-linear signal, which needs to be sufficiently resolved by DEHoI to prevent the occurrence of readout errors. Therefore for both measurement bandwidth and requirement on DEHoI operation speed, a shorter, more modest armlength difference is favourable in the optical design.

The final armlength difference selected for this experiment was 15 km. This maximises the optical gain and the resulting signal amplitude, while still providing a broad measurement bandwidth with an FSR of 13.7 kHz. This FSR was also originally designed to align with an intended DI code filter, and this was believed to enhance filter suppression and reduce aliasing. In later work, we discovered the complexity with fast, distorted laser dynamics and the necessity to track them, and consequently moved to a higher DI readout speed. In hindsight, while signal strength and noise clearance are excellent at 15 km, robust readout of laser dynamics is challenging with high gain and low FSR. Perhaps a shorter length could relax the tracking requirement and still maintain clearance of measurement noises. This however is a major experimental modification and beyond the timeframe of this project.

3.1.3   Final Optical Layout

Combining the above discussions, the final optical layout used for this experiment is shown below in Fig. 3.3. The optical source is encoded and split between the two Mach-Zehnder interferometers, each with 15 km of armlength difference. A polarisation controller is placed at the input of each interferometer to align the input field to the polarisation eigenmode of the interferometer, and balanced detection is used at the interferometer output ports. To minimise etalon effects, all optical fibre connections are fusion spliced. Dual-stage isolators are additionally spliced in upstream of attenuators and photo-detectors, and downstream of the modulator.

The long lengths of fibre provides a medium for ambient thermal, mechanical and acoustic noises to transfer into the interferometric readout. Without mitigation, these noises degrade the measurement sensitivity and obscure the signal of interest. Two isolation chambers were therefore constructed to stabilise the key interferometric paths for each interferometer. We now proceed to discuss the design of these chambers.
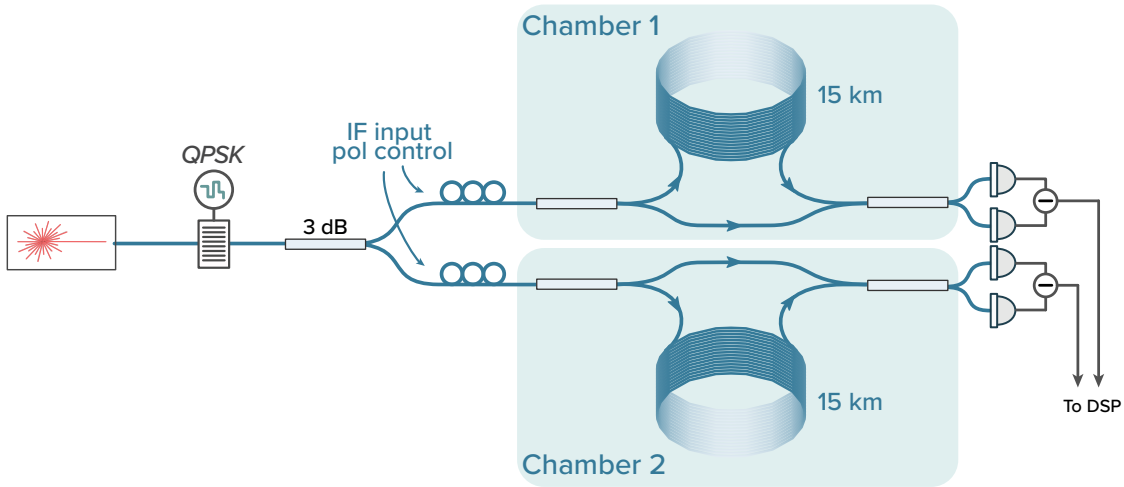
Figure 3.3: Optical layout of the experiment, showing two near-identical Mach-Zehnder interferometers with 15 km of armlength difference, simultaneously measuring the same source laser. The entire optical setup is fusion spliced with single-mode fibre, and each interferometer's input field is aligned to the interferometer's eigen-mode through an upstream polarisation controller.

3.2 Construction of Isolation Chambers

Although previous fibre frequency reference experiments also used isolation chambers, a new set was made for this project to accommodate increased fibre spool dimensions. The new chambers also incorporate a few design changes aimed to improve stability and ease of access. Fig. 3.4 illustrates the design of one isolation chamber, showing the top and side view of the system. We now detail the components that address stability concerns in each category.

3.2.1 Passive Thermal Stabilisation

The primary purpose of the chambers is to provide thermal stability for the key interferometer paths inside. This can be done either through active stabilisation or passive isolation. While the active approach potentially achieves higher accuracy and greater long-term stability, the additional sensor, heat pump and control loop constitute a more complex system and contradict our philosophy of a simple and robust design.

The passive approach, which we use, consists of two sets of nested Styrofoam and aluminium boxes. These two materials are chosen for their high thermal resistance and capacitance respectively. When cascaded, the assembly is analogous to a second-order RC filter, which blocks fast thermal dynamics and provides long-term stability for the inside environment. The thermal time constant for each chamber is estimated to be 18.5 hrs; its calculation and verification are provided in Chapter 4.3.

In fabrication, the Styrofoam panels were cut with tongue and groove joints while the aluminium panels were connected via M4 cap screws. To ensure maximum thermal
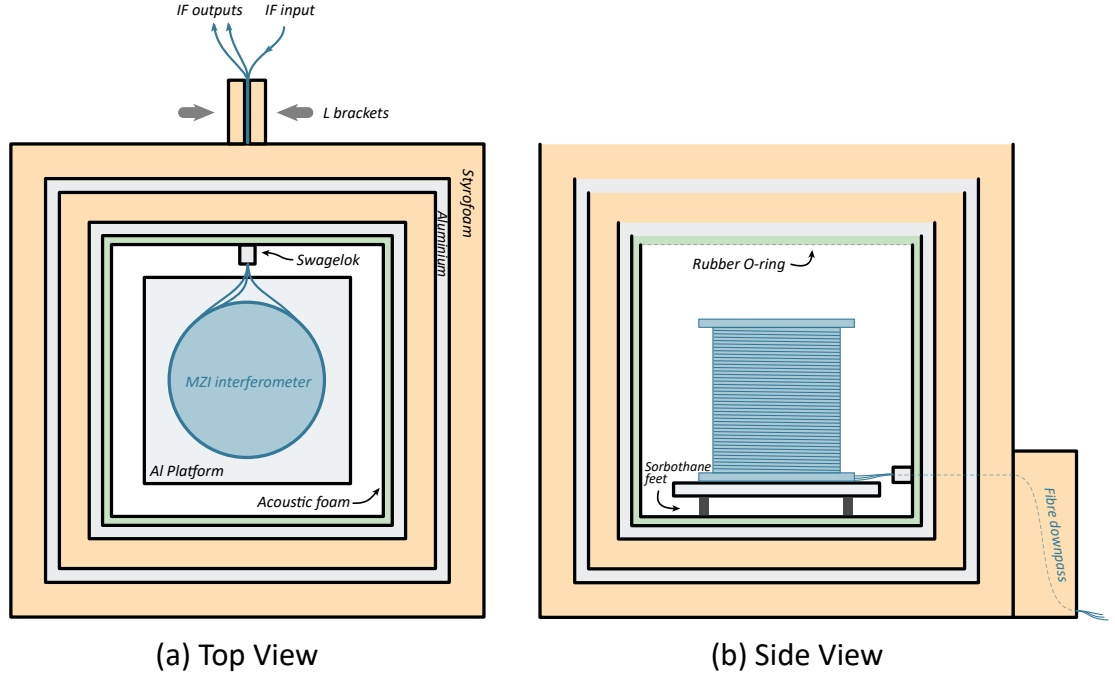
Figure 3.4: Top (a) and side (b) view of one isolation chamber designed to reduce thermal, mechanical and acoustic noise coupling for the inside interferometer.

conductivity within the aluminium boxes, thermal paste was also applied along each contact edge before screws were used to secure the assembly.

3.2.2   Mechanical Considerations

As seen in Fig.   3.4, the Mach-Zehnder interferometer is constructed on a separate aluminium platform inside the chamber.   The platform is supported by 4 pieces of cylindrical Sorbothane feet [151], which dampen mechanical vibrations from the optical table. In addition, to minimise fibre movement, the entire interferometer was secured to the side, top and bottom flange of the 15 km coil using electrical tape, as shown in Fig. 3.5. This was a crucial step conducted later in the experiment, and saw a decrease in mechanical noise from a few short sections of previously loose fibre. This concluded that interferometer rigidity is key to minimising mechanical pickup, with the optimal design being a system fully potted in a resin of similar thermal property.

Following the same rigidity principle, the input and output fibres were "sandwiched" between two pieces of Styrofoam outside the chamber, as illustrated earlier in Fig. 3.4. The Styrofoam pieces were further secured by two L-brackets mounted on the optical table.   This structure allows fibre to securely exit to the optical table with minimum allowed movement or mechanical pickup.
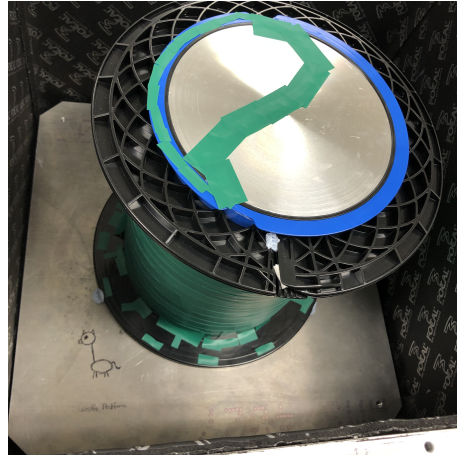
Figure 3.5: A photo of the MZI interferometer inside one of the two chambers. All optical fibre and components are securely attached to the fibre spool using electrical tape. This was shown to decrease fibre movement and the resulting mechanical noise.

3.2.3 Acoustic Insulation

To reduce audio coupling for the interferometer, a layer of sound deadening mat was applied to the inner walls of the small aluminium box, as indicated by the green area in Fig. 3.4. The mat was sourced commercially from the automotive industry [152], and consisted of a honeycombed damping material, an aluminum sheet and a viscoelastic complex (Butyl rubber). Functionally, the honeycombed damping material absorbs audio waves inside the enclosure. The aluminium sheet helps maintain rigidity of the mat for easier installation. The viscoelastic complex adheres to the inner aluminium box, and attenuates vibration and sound from the outside environment.

3.2.4 Sealing the Boxes

Perhaps the most important consideration in designing and constructing the isolation chambers was to ensure optimal sealing against the outside, unstabilised environment. Any minuscule air gaps could create an effective short-circuit, allowing ambient thermal, mechanical and audio noise to pass through. This short-circuit is most likely to occur at the fibre exit point, where a small hole needed to be drilled on every box. To circumvent this, the holes were first limited to 1 mm in diameter, the minimum amount required for three fibres of 250 µm diameter. The inner aluminium box also hosts a Swagelok connector with a Teflon ferrule insert [153]. The Teflon material was custom made into a tapered cylindrical shape and had three evenly distributed through-holes for the input and output fibres. When the Swagelok connector is tightened, the compressed Teflon ferrule forms an ultra-high-grade seal theoretically tolerating up to 1 part in 10^{12} pressure difference [154], which far exceeds the needs for this experiment.

As a precautionary measure, during construction, a thin layer of epoxy was applied outside the small aluminium box where panels are joint, sealing any potential air gaps. On the same box, a thin groove was also made on the top edges of the wall panels for a rubber o-ring insert. When installing the lid panel, the o-ring is depressed and forms the final seal for the enclosure.

3.3   Table Enclosure

While the isolation chambers provide stability for the key interferometric paths, the rest of the optics are still exposed to fluctuating laboratory conditions and susceptible to noise coupling, particularly through air currents expelled by the air conditioning unit directly above the optical table. To reduce air movement across the experiment, an optical table enclosure was also installed. Figure 3.6 shows a photo of the enclosure.

Figure 3.6: Photo of the optical table enclosure.

The enclosure was custom designed into an L shape to cover all components of this experiment. The top and side panels are made of polycarbonate plastic with two cut-outs at the back end for cable passthrough. The enclosure is supported by an aluminium frame mounted to the optical table via L-brackets. Visible gaps are sealed using a combination of foam insert and cloth tape.

While the isolation provided by the enclosure was not directly measurable on this experiment, in a separate seismic characterisation, a drop in vibration was observed when the measurement unit was placed inside. This demonstrates the enclosure is effective in reducing air movement and subsequent noise coupling through that process.

3.4   Polarisation Management

The next part of experimental design concerns polarisation management. Changes in the optical polarisation state, either from the source laser or single-mode fibre propagation, are known to induce signal fading, visibility issues and interferometer phase noise [155]. Amongst them, polarisation induced phase noise is the dominant concern for our experiment, and has been previously observed in the infrasonic frequency band in a similar fibre reference system [144]. To determine whether polarisation noise incurs an experimental limit here, a few components were added to the optical setup to enable DEHoI polarisation readout. The topology was adapted from work by McRae et al [144] and illustrated in Fig. 3.7.
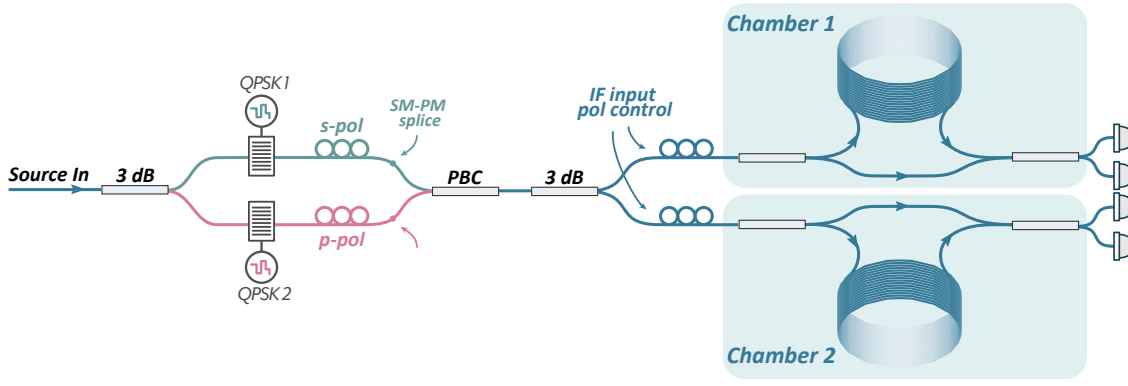
Figure 3.7: Optical layout for polarisation tracking. Two orthogonal, code-distinguishable polarisation states are prepared and launched into each interferometer. Polarisation evolution is then tracked by demodulating specific polarisation state combinations, recovering elements of the Jones matrix.

In this setup, the optical source is split into two paths that are individually coded with different QPSK sequences, before being recombined by a polarisation beam combiner (PBC). This prepares two orthogonal polarisation states, referred to as *s*- and *p*-, that are code distinguishable. The power for each polarisation state is maximised using an inline polarisation controller [156] placed prior to the SM to PM splice, ensuring optimal alignment and coupling. Two additional controllers are also inserted before the launch coupler of each interferometer to allow adjustment of the output polarisation state.

Algorithms for DEHoI Polarisation Readout

Polarisation evolution within each interferometer can be described by the Jones matrix:

$$\mathbf{E}_{\mathrm{out}} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \mathbf{E}_{\mathrm{in}} \tag{3.4}$$

where $\mathbf{E}_{\text{in}}$ and $\mathbf{E}_{\text{out}}$ represent the input and output electric fields. As s- and p- input polarisation are distinguished by code, elements of the Jones matrix can be extracted using the following demodulation logic:

$$
\begin{aligned}
a &\ni \text{QPSK1}(\tau_{\text{ref}}) \ \text{ and } \ \text{QPSK1}(\tau_{\text{sig}}) \\
b &\ni \text{QPSK1}(\tau_{\text{ref}}) \ \text{ and } \ \text{QPSK2}(\tau_{\text{sig}}) \\
c &\ni \text{QPSK2}(\tau_{\text{ref}}) \ \text{ and } \ \text{QPSK1}(\tau_{\text{sig}}) \\
d &\ni \text{QPSK2}(\tau_{\text{ref}}) \ \text{ and } \ \text{QPSK2}(\tau_{\text{sig}})
\end{aligned}
\tag{3.5}
$$

Each of these demodulation channels returns the real and imaginary components of each matrix element, allowing us to reconstruct the phase and amplitude information of that element. Here, the amplitudes are used to readout and track the interferometer polarisation state, as demonstrated in [144]. To obtain the interferometer phase, we calculate the matrix determinant through the following algorithms:

$$
\begin{aligned}
\text{I}_{\text{det}} &= \text{I}_a\text{Q}_d + \text{Q}_a\text{I}_d - \text{I}_b\text{Q}_c - \text{Q}_b\text{I}_c \\
\text{Q}_{\text{det}} &= \text{I}_a\text{I}_d + \text{Q}_b\text{Q}_c - \text{Q}_a\text{Q}_d - \text{I}_b\text{I}_c
\end{aligned}
\tag{3.6}
$$

$$
\phi_{\text{det}} = \frac{1}{2}\arctan\left(\frac{\text{Q}_{\text{det}}}{\text{I}_{\text{det}}}\right)
\tag{3.7}
$$

The determinant phase, collected from all four elements of the Jones matrix, is insensitive to polarisation variations within the system and removes polarisation noise down to the code suppression limit [144]. However, when comparing ϕ_{det} to the original readout without polarisation tracking, we did not observe a difference in noise floor. This leads to the conclusion that the experiment is not limited by polarisation noise.

As a result, measurements described in the rest of the thesis are done without polarisation tracking. This involves mostly a software change, and the optical setup from Fig. 3.7 is preserved. In all following measurements, only the s-polarisation path is driven and the p-path remains idle. The interferometer input polarisation controllers are still adjusted to maximise signal amplitude and SNR.

3.5   DSP Implementation

We now proceed to describe the software and digital signal processing code developed for this experiment. The DSP algorithms are implemented on a National Instruments (NI) Kintex 7 PXIe-7975R FPGA connected to a NI 5782R transceiver adapter module. The module offers two 14-bit analogue to digital (ADC) channels each operating at 250 MSps and two 16-bit digital-to-analogue (DAC) ports each at 500 MSps. The DAC speed can be further doubled to 1 GSps by disabling the second port. We use this single DAC mode to drive our AOM at its centre frequency of 200 MHz with DI phase modulation, at an output sample rate of 1 GHz which retains maximum DI spectral content. Configuration diagram for this feature is provided in Appendix B. The two ADC channels are used to receive the returned signal from each interferometer.

While these IO channels operate at relatively high speeds, the FPGA program runs at a slower rate for resource and timing efficiency. This is referred to as the base clock rate, and is determined by the specific component-level intellectual property (CLIP) profile used for the system. The NI 5782R supports two possible clock rates: 125 MHz in multiple sample CLIP and 250 MHz in single sample CLIP [157]. We nominally use the multiple sample CLIP as its slower time base gives more relaxed timing requirements. Synchronisation between this base clock and the IO sample clocks is configured via software and detailed in Appendix B.

An overview of the FPGA implementation is given in Fig. 3.8. We can see that the entire VI is placed in a 125 MHz single-cycle timed loop. Higher rates for the DAC and ADCs are achieved through oversampling and parallel processing of samples. Specifically, within each clock cycle, our program generates 8 samples for the DAC, each phase-offset by 1/8 of a cycle. A detailed block diagram of this is provided in Appendix A. Similarly, two input samples are provided to each ADC channel per clock, and parallel decoding maintains the 250 MHz speed by processing both samples simultaneously. The exact implementation is simply a two times parallelised version of the previously shown Fig. 2.10. After decoding, the two samples are added, returning to the base clock rate of 125 MHz for further processing.

3.5.1   Real-time Phase Readout

As part of this experiment, the demodulation processes downstream of decoding were upgraded to improve readout fidelity. This new architecture extracts DEHoI phase onboard the FPGA at the clock frequency, and is named real-time (RT) phase readout. The design and development of this work was driven by the high tracking requirement of this experiment to reduce non-linear and aliasing errors, discussed further in the following paragraphs. The software implementation was enabled by increased FPGA resources and bandwidth [158], and completed with Chathura Bandutunga in 2019. We now compare the before and after of the readout scheme, and highlight the advantages of RT readout.

Prior to the digital overhaul, traditional DEHoI readout relied on a single decimating filter to achieve two purposes, code correlation and downsampling, as shown in the top section of Fig. 3.9. The phase reconstruction process occurred in post-processing at the
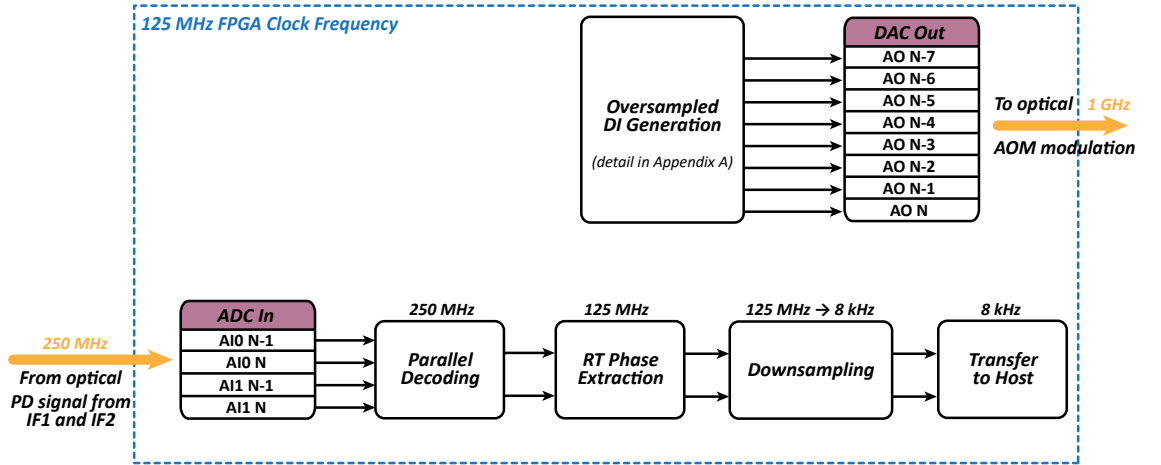
Figure 3.8: A flow diagram of the FPGA program used for this experiment. The program runs and updates at a clock frequency of 125 MHz. Higher sample rates for the DAC and ADCs are achieved by generating or processing multiple samples in parallel. The DAC output consists of a QPSK modulated sinusoid and is amplified to drive the AOM. The ADCs digitise returned signals from the two interferometers. These are then passed through decoding and phase reconstruction operations, before being downsampled and transferred to a networked host computer for recording and post-processing.

decimated sample rate, usually on the order of several kiloHertz. In contrast, the new RT system structurally separates the traditional filter into two functional blocks, code integration (a-b) and decimation (d).

The auto correlation for both I and Q is recovered using a real-time, cascaded integrator-comb (CIC) filter (a-b). The comb stage of this filter (a) incorporates a delay-line that retards samples by an entire code length. This requires a large number of samples to be stored and retrieved at full clock rate, where serialising registers becomes cumbersome and resource expensive. Instead, we allocate parts of the FPGA block RAM for this purpose, creating a dedicated memory block which we can write to and read from. As shown in Fig. 3.9(f), when accessing the memory block, we offset the read and write pointers by 1 element. For every clock cycle, both pointers have their addresses increased by 1, corresponding to a scan to the right-hand-side in the figure illustration. These addresses are returned to 0 once they reach the scan length, which is set to the DI code length. This scan loop architecture creates a robust, synchronous delay of exactly one full code. The comb stage is then followed by an integrator (b) also operating at full rate. The lack of decimation here minimises any aliasing effects prior to phase construction.

Following this, the integrated I and Q components are fed into a high-throughput rectangular-to-polar function (c) [159], returning the amplitude and phase of the desired signal. This real-time function uses the COordinate Rotation DIgital Computer (CORDIC) algorithm with pipelined operation to achieve the desired throughput rate. For this experiment, the input I and Q data both have a word length of 30 bits. The CORDIC is then configured to use a 32-stage pipeline and outputs one sample per clock cycle. The output amplitude mirrors the input word length of 30 bits, while the phase data is truncated to a precision of 16 bits as tested to be sufficient for our purposes. This

Figure 3.9: Flow diagram of traditional DEHoI readout (top) and new RT DEHoI readout (bottom). The RT system achieves full-rate phase readout as a result of splitting the correlation and decimation functions. This minimises non-linear errors within the system and provides greater flexibility of decimation rates. (a-b) Comb and integrator stages of the RT CIC code integrator. (c) RT phase extraction and unwrapping. (d) Decimation with three cascaded 2nd-order CIC filters. (e) Data packaging and transfer to Host program. (f) Detail of the variable delayline used in the comb stage of the RT code integrator. (g) Block diagram of an individual CIC downsampling filter.

reduced word length also allows more bit growth in downsampling and therefore higher dynamic range for the overall system.

There are two immediate advantages of RT phase readout. Firstly, as the arctangent algorithm is mathematically non-linear, by having it performed upstream and at full rate, we are able to minimise non-linear errors coupling into the readout. Second, the decoupling of code integration (a-b) and downsampling (d) means that we are free to choose both the architecture and decimation of the downstream filters, allowing the implementation of higher-order, anti-aliasing filters at flexible speeds.

Following the CORDIC operation, the data container is enlarged to 64 bits for both the amplitude and phase signals, giving an extra 48 bits headroom for downstream operations. The RT phase signal is further unwrapped with a custom function, ensuring continuous phase evolution. Both signals are now ready to be downsampled to a suitable recording rate.

The downsampling process, shown in Fig. 3.9(d), consists of three cascaded, 2nd-order CIC filters. Collectively, the process equates a 6th-order filter, but has the flexibility of individual decimation factor per stage. The three filters are connected through a handshaking mechanism, allowing downstream operations to occur at slower rates. The block diagram of an individual filter is provided in Fig. 3.9(g), showing the integrator, rate reduction and comb sections. The handshake is represented by the blue logic line, and calculations are enabled only in clock cycles with valid input. In this experiment, we nominally downsample by a factor of 25 per CIC filter, obtaining a final readout rate of 8 kHz. The amplitude and phase signals are then interleaved and streamed to Host at said rate using the output valid trigger from the last filter, as depicted in Fig. 3.9(e).

Bit Growth and Dynamic Range

The choice of a 6th-order downsampling structure is determined by the total filter gain and required dynamic range for the system. Although higher order filters offer better anti-aliasing properties, they also incur higher data bit growth. This bit growth would leave less headroom for actual signal dynamics in the system, therefore reducing the overall dynamic range. With container size limited to 64 bit onboard the FPGA, the only mitigation to avoid these outcomes is to perform intermediate rescaling using bitshift operations during the downsampling process. The maximum down-shift without compromising readout precision can be found by equating the expected quantisation noise at the filter output [133] with the desired system noise floor, written as follows:

$$\tilde{S}_{\phi_q}(f) = \frac{2^{-N_q}}{\sqrt{6f_s}} \times \pi = 10^{-6} \quad \mathrm{rad}/\sqrt{\mathrm{Hz}} \tag{3.8}$$

where $\tilde{S}_{\phi_q}(f)$ is the broadband quantisation noise resulting from a single interferometer RT phase readout, and we specify it to be no more than 1 μrad/$\sqrt{\mathrm{Hz}}$. An extra factor of π is multiplied due to the CORDIC output unit of π radians. This means that at $f_s = 8$ kHz, the minimum data precision required to clear quantisation noise is $N_q = 13.8$ bits. The CORDIC output, at 16 bits word length and 14 bits decimal length, is already at this limit. This means any intermediate down-shifting would result in a precision penalty, and should be avoided for our system.

Without the option to downshift and offset the bit growth in high-order filters, the design of these decimating filters becomes a balancing act between anti-aliasing properties and dynamic range. For this experiment, we aim for a dynamic range of 10^6 cycles, which is expected to cover a high-drift metrology laser over multiple days. Recalling that the container enlargement after CORDIC provides 48 bits of headroom, the remaining bit growth allowed in the filter section is:

$$48 - \log_2\left(10^6\right) = 28 \text{ bits} \tag{3.9}$$

The bit growth resulting from the filter section is calculated by [160]:

$$\text{Filter section bit growth} = \log_2\left(NR\right)^M \tag{3.10}$$

where N is the number of delays and equates 1 in our case, R is the decimating factor and M the collective filter order (order of each filter $\times$ number of filters). For our desired data throughput rate of around 8 kHz, $R = 25$ if using three cascaded filters and $R = 125$ if using two. We choose the former for its lower gain, and equate Eqn. 3.9 with Eqn. 3.10. This yields a maximum filter order of $M = 6$ to satisfy our dynamic range requirement, achieved by cascading three 2nd-order CIC filters (Fig. 3.9(d)).

From Eqn. 3.10, we see that bit growth is strongly dominated by the summed filter order M and less so by the delay N or decimating factor R. In our system, any further increase of M would quickly reduce the available dynamic range, and this sets the limit for our anti-aliasing performance. For systems with less stringent dynamic range requirement, however, M should be correspondingly increased for higher out-of-band signal suppression.

Experimental Characterisation

Having constructed the experiment, we now proceed to characterise the system's performance over both short and long timescales. The relative stability of the two interferometers is measured in Section 4.1-4.2, focussing on different Fourier regimes. From there, we quantify the thermal, mechanical and armlength differences between the two interferometers, and examine how they affect the overall noise floor over various Fourier frequency bands.

4.1 Laser Frequency Readout and Stability Characterisation

We recall from Section 3.1.2 that the single-interferometer phase readout is written as:

$$y_1(t) = \phi(t - \tau_1) - \phi(t) \tag{4.1}$$

where $\phi(t)$ is the instantaneous phase of the laser and τ_1 the interferometer transit time. The phase readout can be converted into frequency readout for timescales sufficiently longer than τ_1, such that τ_1 could be considered instantaneous:

$$\nu_1(t) = \frac{1}{2\pi} \frac{\phi(t - \tau_1) - \phi(t)}{\tau_1} = \frac{c}{2\pi n L_1} y_1(t) \tag{4.2}$$

The conversion factor, $c/(2\pi n L)$, is related to the interferometer FSR by a factor of 2π. The frequency readout is accurate for slow dynamics but starts to decline in amplitude response beyond the Fourier frequency of $\sim$FSR/2, as we will show in Section 4.5. Our readout bandwidth of 4 kHz resides within this FSR/2 limit and therefore the above conversion into frequency is justified.

The near-identical interferometer 2 obtains a second measurement of laser frequency fluctuation, $\nu_2(t)$. The relative stability of the system is then characterised by removing common laser frequency noise through subtraction:

$$\Delta\nu(t) = \nu_1(t) - \nu_2(t) = \frac{c}{2\pi n L_1} y_1(t) - \frac{c}{2\pi n L_2} y_2(t) \tag{4.3}$$

In the experimental design of Chapter 3, we have ensured readout independence between ν_1 and ν_2 by constructing two separate interferometers and isolation systems. We can therefore expect minimum common-mode noise suppression and $\Delta\nu$ to be a true relative stability measurement. $\Delta\nu$ will be used throughout the rest of thesis as our main metric of performance.

Figure 4.1 plots a 10 minute recording of the two interferometer's frequency readout, $\nu_1(t)$ and $\nu_2(t)$ (a), and their subtraction $\Delta\nu(t)$ (b). The measurement was performed with a Koheras E15 laser with a DEHoI modulation frequency of 41.7 MHz and a code length of 511 elements. We can see that over the 10 minutes, the E15's central frequency fluctuated by approximately 15 MHz, as recorded by both interferometers. The subtraction readout in (b) largely removes this laser frequency fluctuation. The residual drift, on the order of 30 kHz, is attributed to thermal differences between the two interferometers.
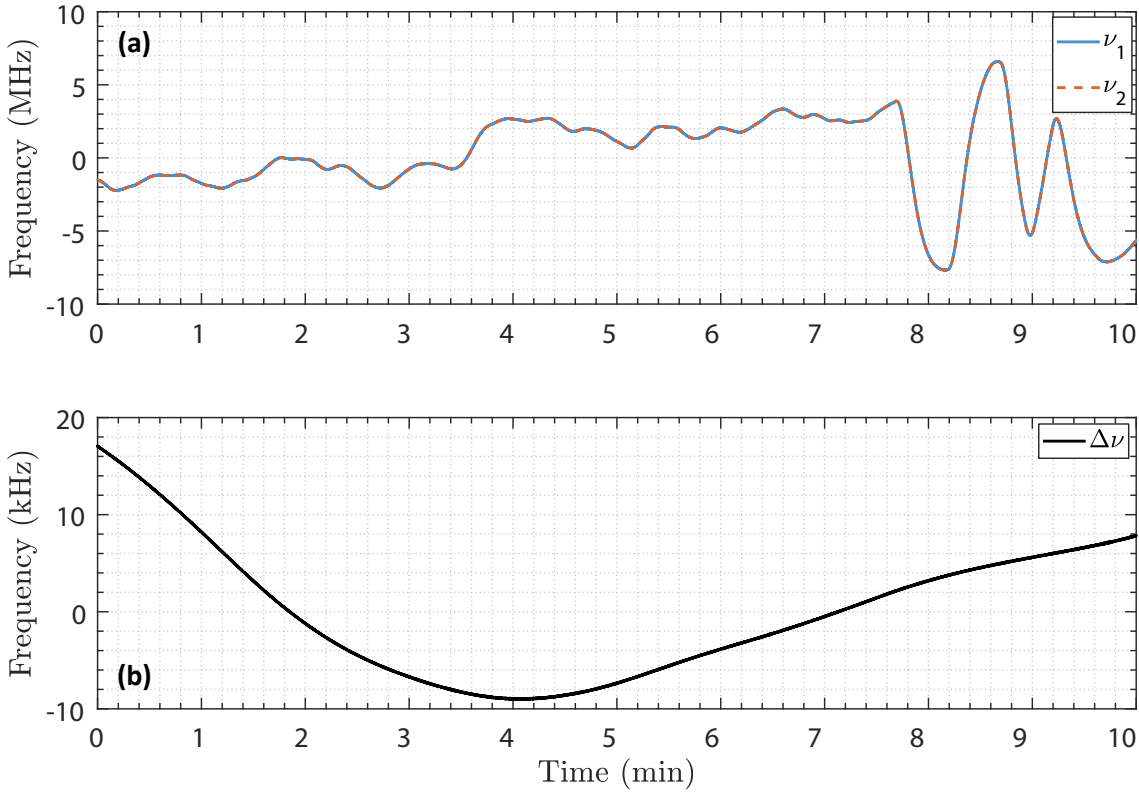
Figure 4.1: 10 minute, time-domain traces of (a) frequency readout from the two interferometers (blue and dashed red), (b) the subtraction readout (black). The free-running E15 laser fluctuated on the order of 15 MHz over the course of 10 minutes. The differential measurement shows removal of this fluctuation with the residual drift attributable to uncorrelated thermal effects.

For this particular measurement, we detrended the differential measurement by subtracting off a linear fit. This removes linear differences in thermal expansion between the two interferometers, which are too slow to be sufficiently resolved in 10 minutes and can consequently cause errors in Fourier-domain analysis. A non-detrended $\Delta\nu$ will be shown later in a longer measurement. For now, we focus on performance above 50 mHz Fourier frequency.

Figure 4.2 plots the Fourier domain, frequency spectral densities of ν_1 and $\Delta\nu$ against the left y-axis. These two traces are further divided by the central frequency of the E15

laser, and the resultant fractional stability spectral densities are indicated by the right y-axis. The spectral trace of ν_2 is left out as it shows identical features to ν_1 over the scale presented here.
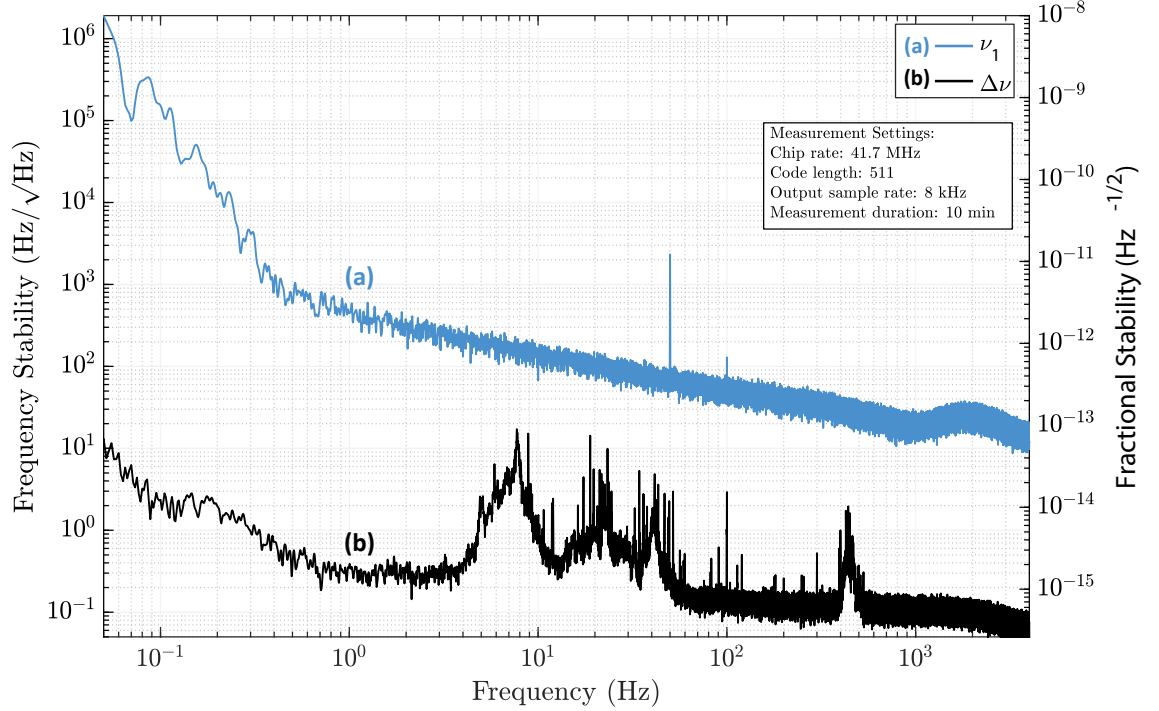
Figure 4.2: The time-domain readout traces of both a single interferometer and the subtraction are converted into Fourier-domain frequency spectral density. Trace (a) ν_1 shows the frequency characteristics of the Koheras E15 laser, with a knee frequency of approximately 500 mHz. Trace (b) $\Delta\nu$ plots the relative stability between the two interferometers. The system achieves a broadband stability of 0.1 Hz/$\sqrt{\text{Hz}}$ beyond 70 Hz, and is limited by different noise sources across several frequency regimes.

Trace (a) details the E15 laser dynamics across various timescales. A $1/\sqrt{f}$ type frequency noise roll-off can be observed from approximately 500 mHz to 1 kHz. A small feature, centred at 2 kHz, protrudes beyond this $1/\sqrt{f}$ trend, likely due to noise peaking of a control loop within the laser, while at 400 mHz there is an abrupt transition to a $1/f^3$ dependence that dominates low frequency stability.

In the relative stability FSD of trace (b), the noise floor can be categorised into four regimes. Above 70 Hz, a frequency stability of 0.1 Hz/$\sqrt{\text{Hz}}$ is achieved. This broadband noise floor is later found to be due to second-order RBS effects, discussed in Section 5.2. The slight roll-off above 2 kHz is due to an anti-aliasing filter prior to data decimation. Between 5 Hz and 70 Hz, we see a few distinct mechanical and acoustic modes, rising up to two orders of magnitude above the background frequency noise. At infrasonic frequencies, the system stability lands close to the fundamental fibre thermo-mechanical noise between 400 mHz and 2 Hz, modelled in Section 5.1. Below this frequency band, we see the onset of fibre temperature drift, which is examined further in Section 4.2 - 4.3.

To determine the optimal operation timescale of our system, an Allan deviation plot is also generated for the fractional stability measurement, $\Delta\nu/\nu_c$. The computation was done using the Modified Allan function in Stable32 [161], with a confidence factor of 0.683. In Fig. 4.3, we can see the system reaches an optimal fractional stability of 1.29×10^{-14} at 2 ms integration time, while a stability below the 10^{-13} mark is maintained up till 0.3 s.
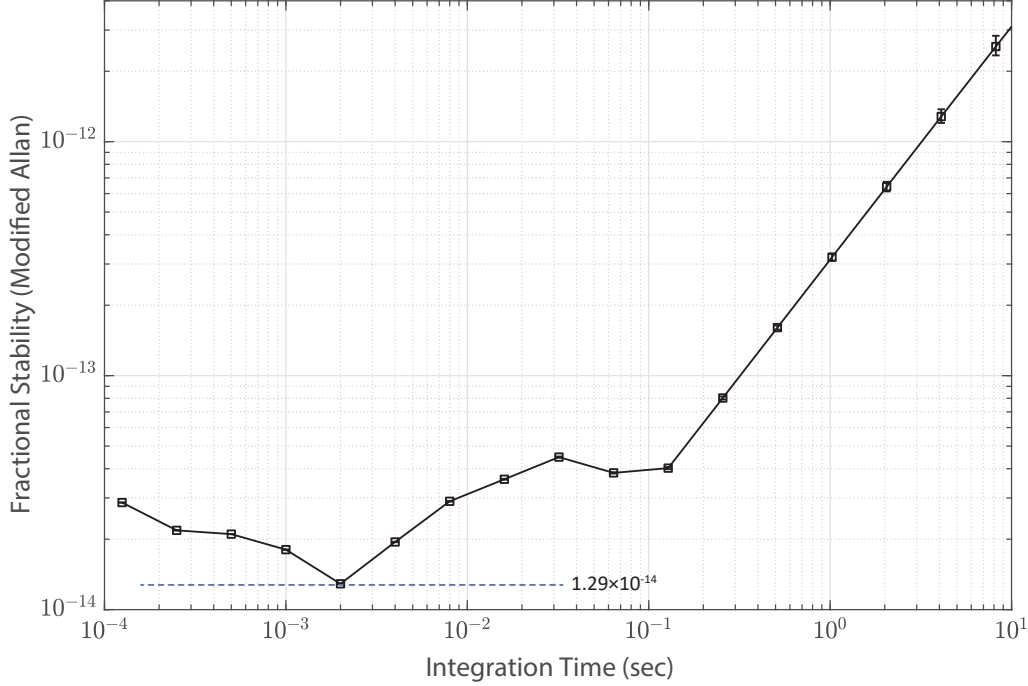
Figure 4.3: Allan deviation plot showing the fractional frequency stability of the fibre frequency reference as a function of averaging time. The plot was generated using Stable32. An optimal stability of 1.29×10^{-14} is achieved at 2 ms, while better than 10^{-13} stability is maintained up to 0.3 s.

The optimal integration time here is shorter than that reported in the previous generation FFR, which was 0.1 - 1 sec [134]. This can be explained by the current system achieving lower noise floor at high Fourier frequencies, translating to better stability below 10^{-2} sec integration time. The worsened stability between 10^{-2} - 0.25 sec can be attributed to mechanical and acoustic pickup, as evident in the Fourier band of 4 - 100 Hz in Fig. 4.2. Finally, at above 1 sec integration time, the fractional stability of both systems is degraded by the onset of temperature drift.

4.2   Long-term Stability Characterisation Using an Optical Frequency Comb

Having characterised the high frequency stability of our system, we now shift focus to the performance over long timescales. We further compare with an independent reference, an optical frequency comb (OFC), to provide a measure of any correlated noise sources between the two interferometers due to their shared laboratory environment.

4.2.1   OFC Operation Principle

The Fourier domain frequency comb is generated by a mode-locked, ultra-fast pulse laser with fixed phase relationship between successive emissions. The repetition rate of the pulses (f_r) determines the spectral separation of each comb tooth, while the offset frequency (f_o) specifies the starting frequency for the comb. These two parameters collectively set the absolute optical frequency of each comb tooth, and we write that of the ith as follows:

$$\nu_i = f_o + i \times f_r \tag{4.4}$$

Here, f_o is related to the carrier envelop phase ϕ_{ce}, which evolves between successive pulses due to the difference in the phase (v_p) and group velocities (v_g) of the resonator field inside the mode-locked laser. The increment of ϕ_{ce} per pulse is written as [162]:

$$\Delta\phi_{ce} = \left(\frac{1}{v_g} - \frac{1}{v_p}\right) l_c\,\omega_c \tag{4.5}$$

where l_c is the round-trip length of the laser cavity, and ω_c is the carrier frequency. The resultant frequency-domain shift from this phase increment is expressed as [162]:

$$f_o = \frac{1}{2\pi} f_r \Delta\phi_{ce} \tag{4.6}$$

From Eqn. 4.4, we see that the frequency stability of each OFC tooth is linked to that of f_o and f_r. As both of them fall into the RF regime, direct detection is possible and stabilisation solutions are readily available. The pulse repetition rate, f_r, is detected by heterodyning two adjacent teeth. To isolate the offset frequency, f_o, a second harmonic generator (SHG) is used to double the optical frequency of one tooth (ν_1). This is then compared with another tooth (ν_2) an octave above the original one, generating a heterodyne beat note of

$$2\nu_1 - \nu_2 = 2\left(f_o + if_r\right) - \left(f_o + 2if_r\right) = f_o \tag{4.7}$$

Once f_o and f_r have been measured, they are typically stabilised to an atomic reference and inherit the stability of that reference. The OFC used in this experiment is a commercial unit by Menlo Systems [163] stabilised to a Stanford Research System Rubidium Frequency Standard [97]. The rubidium standard has superior stability at low Fourier frequencies and therefore is an ideal candidate for characterising the thermal stability and independence of our system.

4.2.2 Experimental Integration of OFC

In the experimental setup of Fig. 4.4, the two fibre references were identically configured as previous. About 50% of the laser power was tapped off to interfere with the OFC, and the input attenuation within the fibre references was correspondingly adjusted to maintain the same incident photo-detector power.
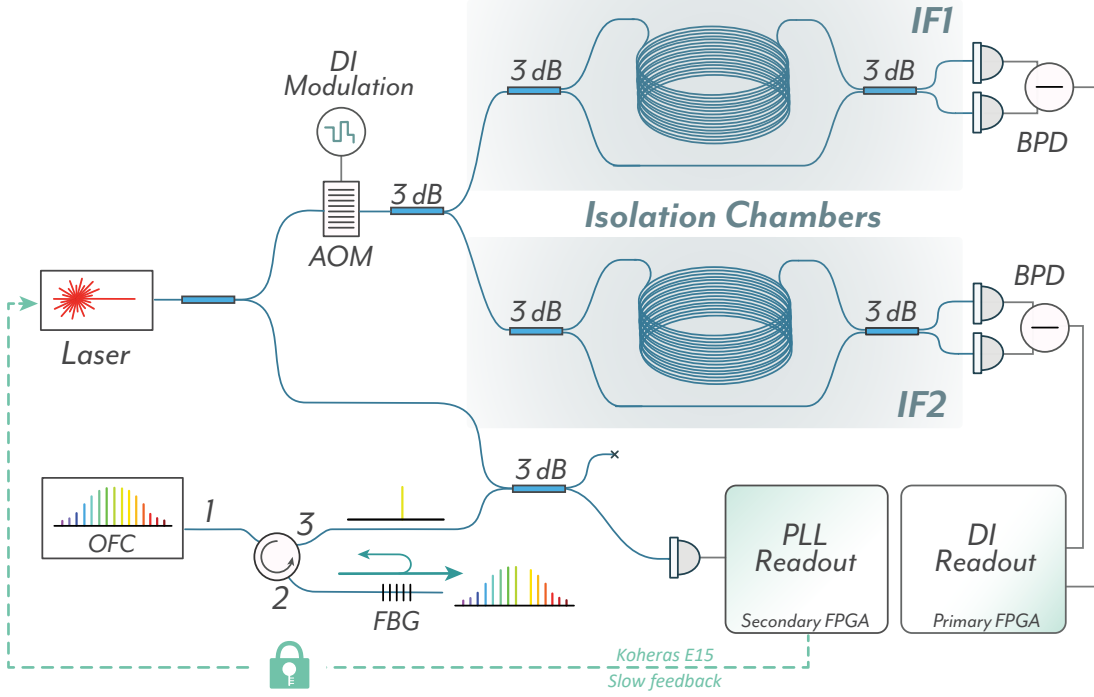
Figure 4.4: Experimental layout for the simultaneous, three-system measurement. The fibre references were identically configured as previous, with the exception of a 50% tap-off from the source laser. This tap-off was interfered with the OFC, and a single beat note within 125 MHz was selected using a FBG and a 125 MHz photo-detector. Following digitisation, this beat frequency was tracked and recorded using a phase-locked loop (PLL). A slow feedback was implemented for the Koheras E15 laser to keep its frequency within the PLL dynamic range of ∼50 MHz.

To maximise the useful optical power for the comb-laser interference, we used a fibre Bragg grating (FBG) with a pass band of 0.1 nm [164] to select the nearest neighbour comb teeth to our test laser. By removing all but a few comb teeth, the shot-noise-limited sensitivity of the measured beat note was improved.

Given the spectral separation of the OFC at 250 MHz, the comb tooth nearest the laser wavelength would result in a beat frequency less than 125 MHz. We therefore used a 125 MHz Newport 1811 photodetector to further isolate this single beat note. The photo-detector output was digitised using a NI 5782 IO Module connected to a NI 7975R Kintex 7 FPGA, a second, identical system to the one used for the fibre reference readout.

The two FPGA systems were clock synchronised and simultaneously called using a common host program running on an integrated PC, allowing concurrent readout and recording of all three references. The readout of the comb beat note was handled by a phase-locked loop (PLL) running on the secondary FPGA. The PLL architecture, illustrated in Fig.

4.5, was adapted from the original development by Roberts [133], with improvements in bandwidth and resolution thanks to increased hardware resources.
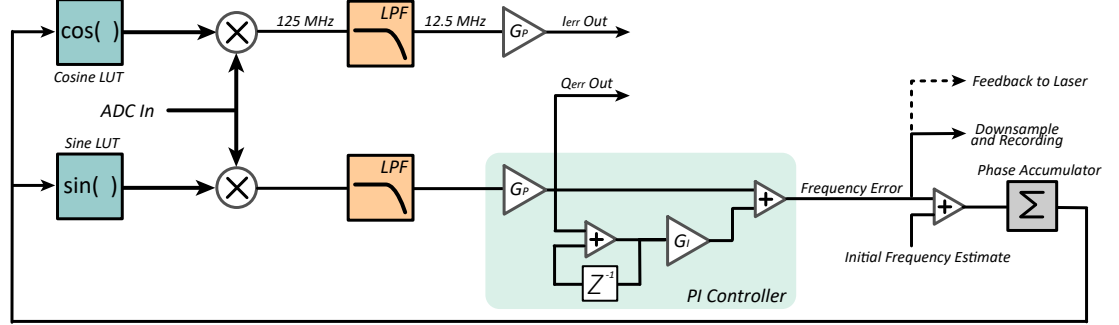
Figure 4.5: FPGA block diagram of the phase-locked loop. The digitised beat note is mixed with an initial frequency estimate to generate an error in both I and Q. The up-converted frequency content in this error is removed through a low-pass filter. This is then passed through a PI controller and converted into frequency error in FPGA counts. The PLL loop is closed by adding the frequency error onto the initial estimate and updating the mixing frequency. The beat note dynamics are recorded from the frequency error readout, which is also used for laser pre-stabilisation if the free-running drift exceeds the PLL dynamic range.

In the block diagram in Fig. 4.5, the digitised beat note is first mixed with sine and cosine numerically controlled oscillators (NCOs) at an initial frequency given by the frequency estimate parameter. Following low pass filtering this yields the IQ projections of the beat note relative to the NCO. This differential frequency is subsequently passed through a PI controller with tunable proportional (G_P) and integral (G_I) gains, generating an output which we call the frequency error. This numerical value represents the instantaneous error in our initial frequency estimate. The final part of the architecture adds this frequency error as a phase step into the phase accumulator, updating the mixing frequency within the two look-up tables.

There are two key parameters that define the performance of the PLL: bandwidth and dynamic range. Bandwidth represents the fastest dynamics that the PLL can track, and is inversely proportional to the decimation factor of the PLLs and the total logic delay within the loop. For metrology lasers with low high-frequency noise, fast dynamics in the beat note are dominated by noise from the comb laser. In our testing, we found a LPF update rate of 5 MHz to generally yield sufficient bandwidth for the OFC laser, although a higher rate of 12.5 MHz was used in the below measurement for additional noise tolerance.

The theoretical tracking range of the PLL equates the FPGA Nyquist frequency of 62.5 MHz, but is in practice reduced to about 50 MHz due to degraded signal fidelity towards Nyquist. To maximise the usage of this dynamic range, the beat frequency was centred at 25 MHz prior to the commencement of our measurements, by fine-tuning the repetition rate f_r on the OFC. However, through several trail measurements, we found that the free-running Koheras E15 drift frequently exceeded the 50 MHz dynamic range. This meant that we could not use the PLL as an open-loop readout, and a pre-stabilisation scheme was developed.

For this, the output "Frequency Error" indicated in Fig. 4.5 was used as an error signal in addition to being downsampled and recorded. The feedback path used a separate PI

controller, and was output via the DAC to drive the piezo actuator modulating the E15 fibre cavity. As only the slow laser frequency drift was of concern here, we electronically low-pass filtered the piezo drive signal down to 0.1 Hz using a Stanford Research Systems SR560 pre-amplifier. Given the low gain and bandwidth of the loop, we obtained on the order of 4 hours of data prior to the lock disengaging.

At the time of this experiment, another metrology laser, the Orbits Lightwave, was also available for characterisation. However, this laser suffered from a mode-hop issue, with discrete frequency jumps on the order of 15-20 MHz occurring up to three times per hour, subject to laboratory temperature stability. Brief attempts to re-calibrate the laser operation point and to pre-stabilise it both proved unsuccessful. These frequent and fast frequency steps created significant challenge for both the PLL tracking bandwidth and its dynamic range, and we were unable to obtain a long-term measurement with this laser.

4.2.3 Measurement Results

Fig. 4.6 below plots the time-domain readout for all three references and their subtractions over a proportion of the 4 hr measurement time. Unlike the previous short-term measurement, here the traces were not detrended as the longer timescale affords sufficient low-frequency resolution. This allows us to see significant differences between the two fibre references, ν_1 (a) and ν_2 (b), with their subtraction $\Delta\nu$ (d) drifting by 4 MHz over the course of the measurement.
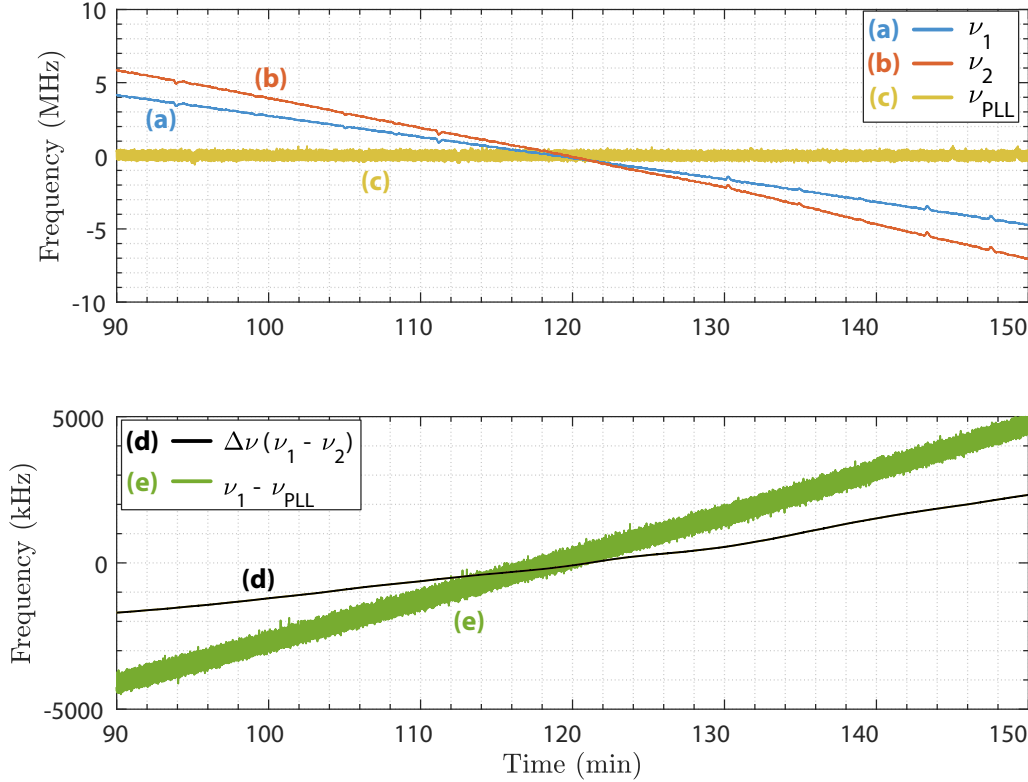
Figure 4.6: Time-domain readout of the three-reference measurement using the Koheras E15 laser. (a) readout by IF1. (b) readout by IF2. (c) beat note readout by the PLL. (d) subtraction between IF1 and IF2. (e) subtraction between IF1 and PLL.

Because the laser was stabilised to the rubidium referenced OFC, a highly stable standard at long timescales, the drift in trace (a) and (b) can be attributed to fibre length change, which at these timescales corresponds to temperature drift. We can further quantify this temperature drift by subtracting one interferometer readout (ν_1, trace (a)) with that of the PLL (ν_{PLL}, trace (c)), removing any residual laser frequency noise in the process. The result is plotted in trace (e). By comparing it with the direct interferometer subtraction, $\Delta\nu$ (d), we can determine the amount of correlated drift between the two interferometers and thus estimate their temperature independence.

Quantifying the temperature independence this way is easiest done in the frequency domain. In Fig. 4.7, we transform the above time-domain readout traces and plot their corresponding frequency spectral density. From both the IF1 readout (a) and the PLL readout (c), we see improved E15 laser frequency noise below 0.2 Hz compared to its free-running spectrum (b) obtained in a separate measurement. Beyond 0.2 Hz, the laser performance is degraded with many additional harmonic peaks, likely due to the control and electronic noise within the pre-stabilisation feedback loop.
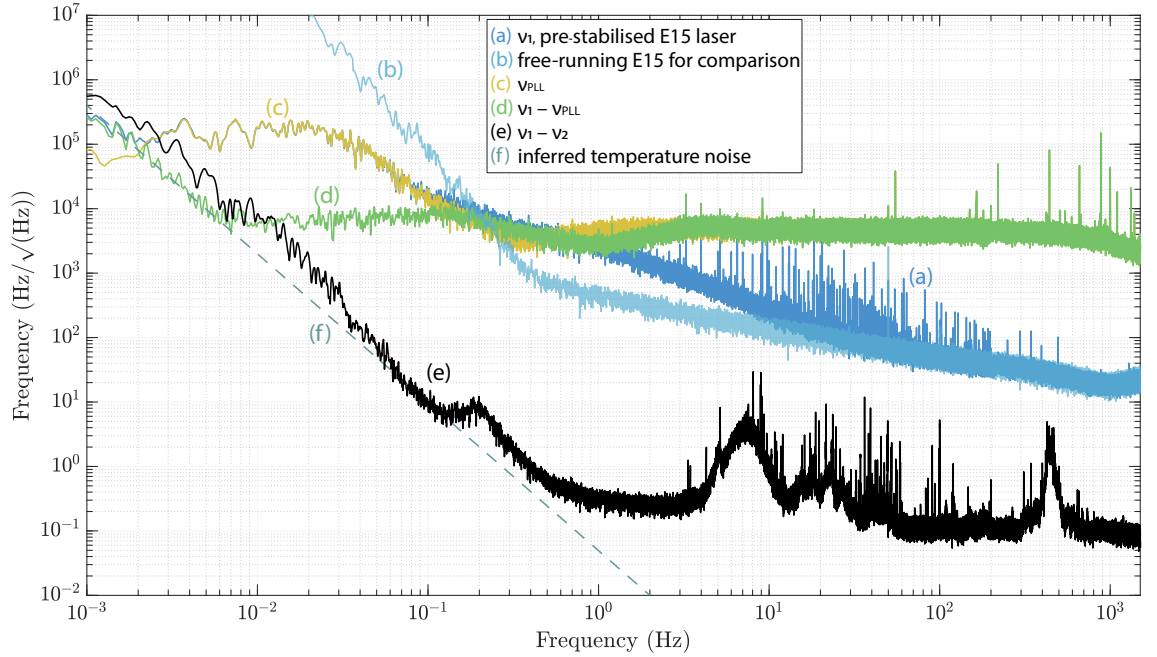
Figure 4.7: Frequency spectral densities of the three-reference measurement using the Koheras E15 laser. The E15 was pre-stabilised to ensure its drift remained within the PLL tracking range. (a) Readout by IF1. (b) IF1 readout in a separate measurement where the laser was free-running. (c) Readout by the PLL loop. (d) Subtraction between the single interferometer and PLL readout. (e) Subtraction of the readout of the two fibre interferometers. (f) A $1/f^{2.5}$ type temperature noise can be inferred from trace (d).

Comparing the IF1 readout (a) with that of the PLL (c), we see an overlap between 2.5 mHz - 0.5 Hz. Below 2.5 mHz, temperature-driven fibre drift exceeds the residual laser frequency noise, confirming our observation from Fig. 4.6. Above 0.5 Hz, the stability of the E15 laser exceeds that of the OFC, as the IF1 spectrum continues to drop while the PLL remains at a flat noise floor of 5×10^3 Hz/$\sqrt{\mathrm{Hz}}$ throughout the measurement band.

Trace (d) maps out the temperature noise of a single fibre interferometer between 1 - 6 mHz, beyond which it is limited by the OFC noise floor. A $1/f^{2.5}$ slope (f) was fitted to the trace (d) as the inferred, single interferometer temperature noise. This was compared to the subtraction spectrum between the two interferometers (e), and we can see that trace (e) lies above $\sqrt{2}$ of trace (e) across the spectrum, apart from the Fourier range of 50 - 120 mHz. This means that for most frequencies, incoherent temperature noise addition applies to the two fibre interferometers, while between 50 - 120 mHz a small amount of correlation exists and this results in slight common-mode suppression. From this we can conclude that although located in the same laboratory, due to their independent thermal housing, the two interferometers are largely temperature independent. The differences in the thermal response of the two chambers are discussed in Section 4.3.

4.3   Thermal Chamber Stability Characterisation

We recall from Chapter 3.2 the construction of the isolation chambers, which consist of two nested sets of Styrofoam and aluminium boxes. In this section, we model the thermal response of this architecture, and experimentally measure the temperature profile for both chambers.

4.3.1   Estimating Thermal Time Constant

The top view of the nested box structure is repeated in Fig. 4.8(a). For both layers, the Styrofoam boxes have a thickness of 5 cm and the aluminium ones 1 cm. The equivalent thermal circuit of this system is provided in Fig. 4.8(b), where the index denotes the isolation layer from the outside inwards.
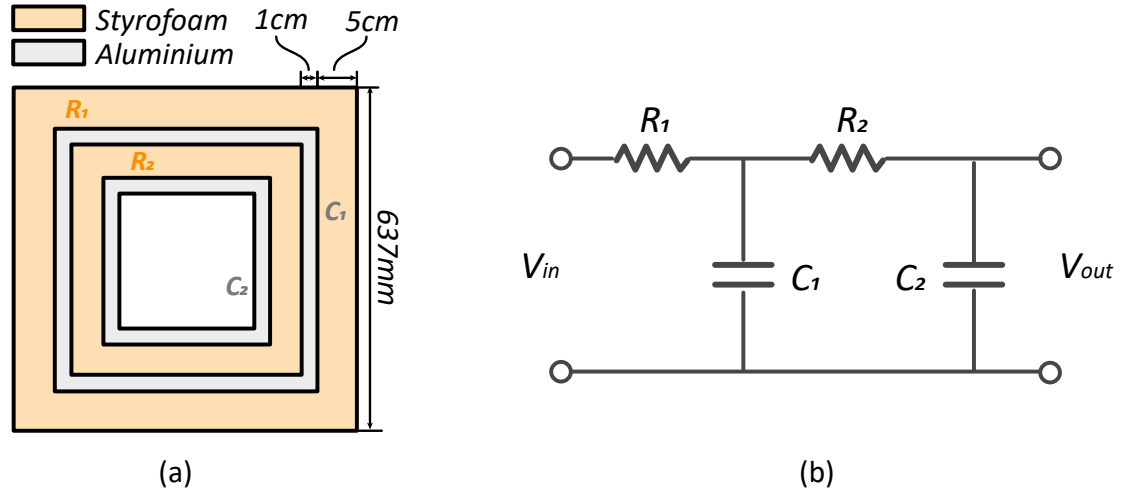
Figure 4.8: (a) Top view of the dual-stage thermal isolation chamber. Two Styrofoam boxes and two aluminium boxes are cascaded together, forming a passive, second-order thermal low-pass filter. The Styrofoam layers are 5 cm thick while the aluminium pieces are 1 cm thick. (b) The thermal circuit of the system, where R_1 and R_2 correspond to the outside and inside Styrofoam boxes, and C_1 and C_2 the aluminium boxes.

For a single-order RC circuit, the time constant is simply the multiplication of R and C. A second-order RC filter however differs from two independent, cascaded first-order filters. This is especially true when the second stage has a similar impedance and interacts with the first stage. In the case of these thermal chambers, the inside boxes interact with the outside ones by reducing their thermal load. Therefore the thermal constant of the dual-stage chamber must be determined through its transfer function and the 3 dB cut-off frequency.

The transfer function of a second-order RC filter is given by [165]

$$G = \frac{V_{\text{out}}}{V_{\text{in}}} = \frac{1}{1 + (j\omega)\left[C_1 R_1 + C_2(R_1 + R_2)\right] + (j\omega)^2 C_1 C_2 R_1 R_2} \tag{4.8}$$

The thermal resistance of Styrofoam can be calculated as

$$R = \left(\frac{I}{\kappa}\right) / A \qquad (4.9)$$

where I is the thickness of the material, A the surface area, and κ the thermal conductivity. For Styrofoam, $\kappa = 0.03$ W/m/K. All Styrofoam layers used in the chambers are $I = 5$ cm thick, therefore surface area A is the cause for different thermal resistance values across inner and outer layers. Based on box dimensions, we can estimate $R_1 = 0.91$ K/W and $R_2 = 1.55$ K/W.

We can then calculate the thermal capacitance which is dominated by the aluminium layers, with the below expression [166]:

$$C = c \times m \qquad (4.10)$$

where c is the specific heat of a material and is equal to 900 J/kg/°C for aluminium. m is the mass and can be estimated by the volume and density of material used. This results in a thermal capacitance of $C_1 = 34705$ J/K and $C_2 = 18675$ J/K for the outside and inside aluminium boxes respectively.

We can now substitute the required values into Eqn. 4.8 and plot the frequency-domain amplitude response of the thermal isolation chamber. This is shown in Fig. 4.9. We find a 3 dB cutoff frequency of $f_c = 2.38$ μHz. Given that $f_c = 1/(2\pi RC)$, the effective time constant is estimated to be $\tau = 1/(2\pi f_c) = 18.5$ hrs . We now proceed to compare this estimate with direct experimental measurement of chamber temperature.
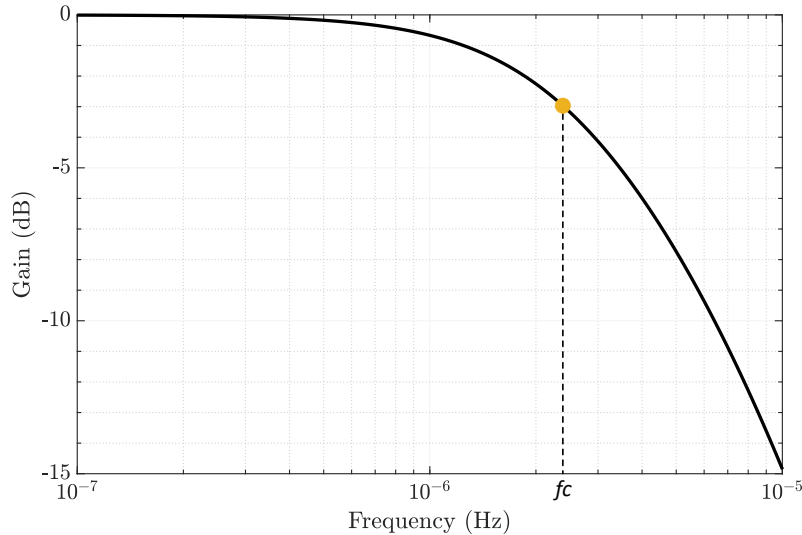
Figure 4.9: Modelled transfer function of the two-stage thermal isolation chambers, plotting gain as a function of frequency. We find a 3 dB cut-off frequency of $f_c = 2.38$ μHz, corresponding to an estimated time constant of 18.5 hrs.

4.3.2   Experimental Measurement of Thermal Stability

During the commissioning process, three temperature sensors were installed across the experiment to characterise the thermal performance of the isolation chambers. These sensors were sourced from Texas Instruments (model number TMP117), and have an accuracy of ± 0.1°C. Temperature readout is digital and obtained through I²C communication via a LabJack U3-HV unit. Figure 4.10 below illustrates the circuit layout for this measurement.
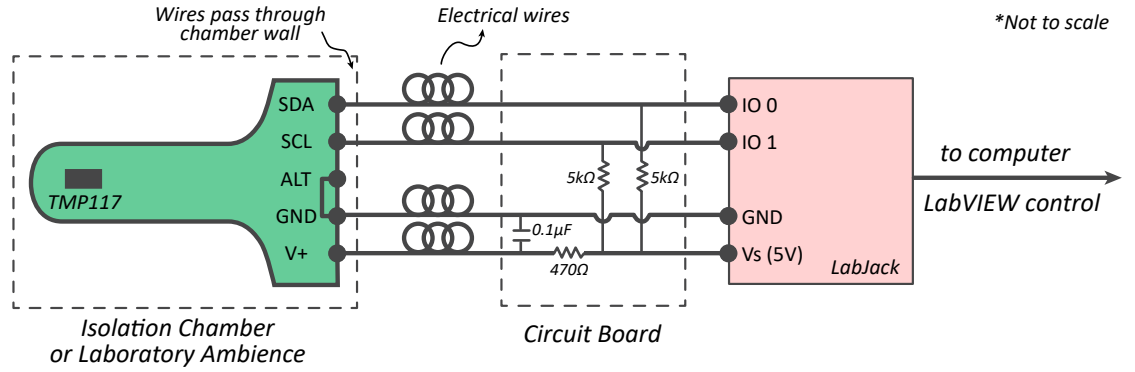
Figure 4.10: Circuit layout for long-timescale temperature measurement. We use the breakout board of the TMP117EVM evaluation module, where the sensing unit TMP117 is located at the tip of the board. This was placed either inside an isolation chamber or outside in ambient laboratory air. The breakout board has 5 pins connected to a circuit board through electrical wires. Two pullup resistors of 5 kΩ were provided to SDA and SCL pins, while a bypass capacitor of 0.1 μF was connected between V_+ and GND. An adjustment resistor of 470 Ω was also inserted in the voltage line. The alert (ALT) pin was not in use and therefore grounded. These electrical components were mounted on the circuit board and further connected to corresponding LabJack channels. The LabJack unit was then connected to a computer and data acquisition was handled by LabVIEW.

Amongst the three sensors, two were placed inside each isolation chamber, with the thermal pad suspended mid-way between the inner wall and the fibre coil. As these sensors operate on low power, the self heating effect within the chambers is expected to be minimal. Each sensor is soldered onto four electrical wires, which pass through chamber walls and connect to the control circuit outside. The third sensor was placed outside the chambers to measure ambient laboratory temperature.

Following the installation and setup of the sensors, a measurement was carried out over 94 days, retrieving and recording temperature simultaneously from all three locations once per second. The resultant time-domain readout is plotted in Fig. 4.11. While the ambient laboratory temperature (trace a) experiences daily and sometimes hourly fluctuations of up to 4 degrees, temperatures inside the isolation chambers (trace b and c) change at a much slower rate. This confirms the low-pass filter effect of thermal chambers as discussed in the section above.

The simultaneous measurement also allows for the calculation of transfer functions for both isolation chambers. This is done by computing the amplitude spectral density (ASD) of all three traces in Fig. 4.11, and dividing the frequency response of an isolation chamber by that of the ambient air. The transfer functions of both chambers are plotted in Fig.
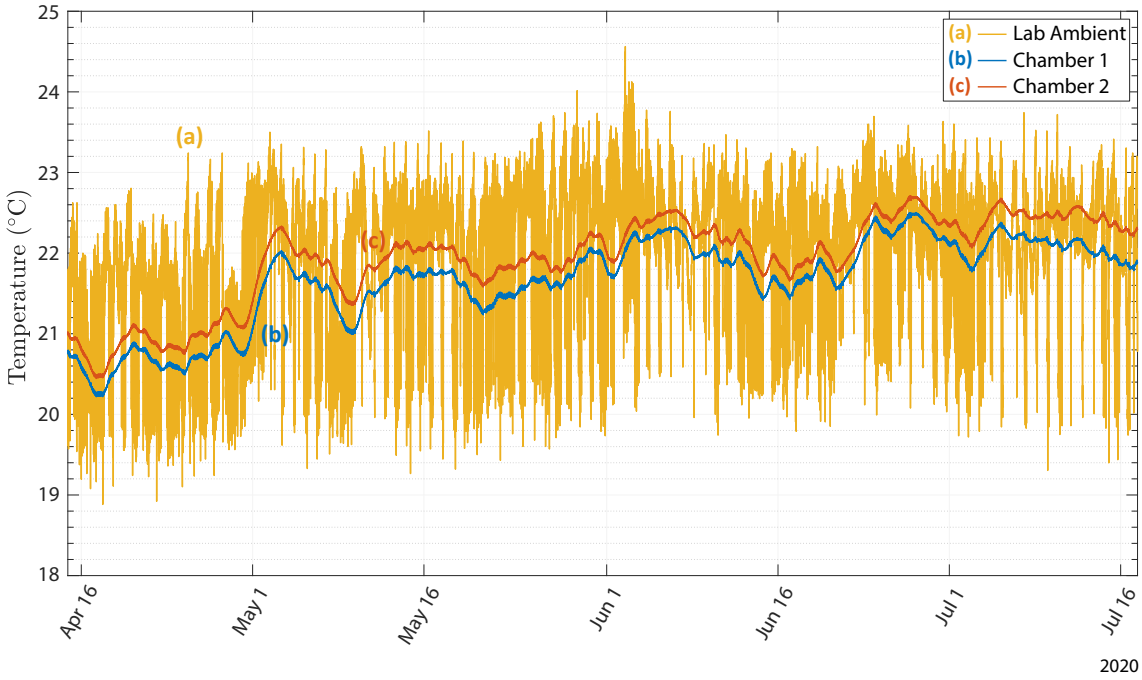
Figure 4.11: Time domain readout of the 94 day temperature measurement, showing ambient laboratory temperature in yellow (a), and chamber 1 and 2 in blue (b) and red (c) respectively. We see that whilst the laboratory temperature fluctuates over hourly and daily timescales, temperature inside the chambers changes at a significantly slower pace and over a reduced range.

4.12. The high frequency range of the x-axis is limited to 3 mHz; beyond 3 mHz, chamber temperature measurement is limited by the TMP117 noise floor of 4.5 mK/$\sqrt{\text{Hz}}$.

From Fig. 4.12, we see that both chambers have a rippled frequency response. This is likely due to the electrical wires that connect the temperature sensors to the control board, which pass through chamber walls, creating an effective thermal short-circuit between the inside and outside environment. Although the passthrough punch holes are minimised in their diameter (3 mm for aluminium and 5 mm for Styrofoam), given the high thermal conductivity of electrical wires, they can still induce significant thermal oscillation and corrupt the stability afforded by the chambers. The actual equivalent thermal circuit for these chambers is likely different to that described in Fig. 4.8(b), and therefore these transfer functions are not directly comparable to the modelled transfer function of Fig. 4.9.

Fortunately, we can still approximate the time constant for both chambers using a thermal interferometer analogy. Here, we consider two paths for laboratory temperature fluctuation to flow into the thermal chamber. The electrical wire path with its low thermal resistance is analogous to the short arm of the interferometer, while the original intended path through the chamber is the long arm with a transit time of τ, where τ is the chamber time constant. When these two paths of the thermal interferometer are combined, the resultant readout transfer function displays nulls at integer multiples the FSR, where FSR = $1/\tau$. This allows us to estimate the time constant by simply computing the inverse of the first null frequency in Fig. 4.12.
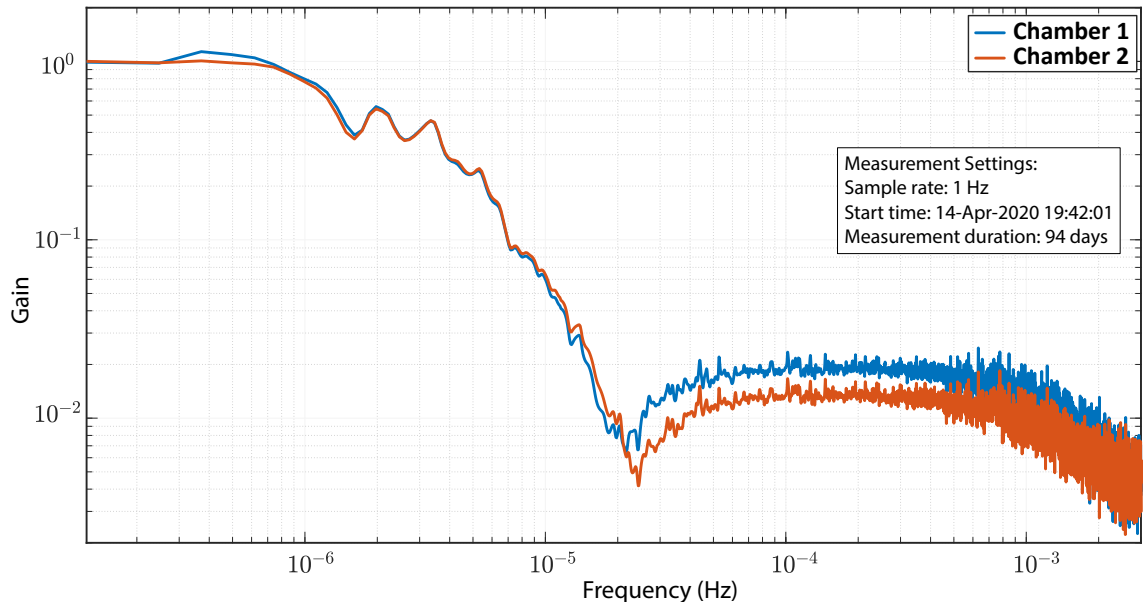
Figure 4.12: Transfer function of the two isolation chambers, calculated by dividing the Fourier-domain ASD of inside temperature measurement by the ambient laboratory measurement. Both chambers have ripples in their transfer function likely due to a thermal short-circuit, which alters the thermal dynamics between the inside and outside environment. Fortunately, we can still apply the thermal interferometer analogy to approximate the time constant for both chambers. The nulls in their frequency response correspond to an estimated time constant of 13.2 hrs and 11.4 hrs for chamber 1 and 2 respectively.

Using this analogy, we obtain a time constant of 13.2 hrs for box 1 and 11.4 hrs for box 2. Both are lower than the theoretical estimate of 18.5 hrs, likely due to minuscule gaps between various box surfaces. The difference in the number and distribution of these gaps across the two chambers then results in different time constants for them. This also explains the lack of thermal correlation between our interferometers at very low frequencies.

Upon completion of this measurement, the temperature probes were removed from the experiment, and the chambers were resealed to eliminate the short-circuit.

4.4 Mechanical Resonance Investigations

From the relative stability measurement shown in Fig. 4.2, we saw coupling of mechanical resonances in the readout noise floor, with distinct peaks at 8 Hz, 23 Hz, 43 Hz and 430 Hz. In order to identify the source of these mechanical features, a seismic measurement was conducted. The goal of this experiment was to map out the vibration profile for both the laboratory and the optical table of use, and to identify any mechanical correlation between these systems.

The device used for this measurement was a Silicon Audio 203-60 optical seismometer [167], with three mutually perpendicular sensors in X, Y and Z directions. An interferometric measurement was made on each axis onboard the device, producing an electronic voltage proportional to the acceleration along that axis. This voltage can be accessed through the output terminals of the SA 203-60. For this experiment, we combined all voltage terminals and the power supply ports onto a 9-pin Sub-D connector and secured it onto the body of the seismometer. A Sub-D extension cable then wired these connections to the desired location, and a subsequent breakout board allowed connection to a 15 V power supply and three ADC channels.

The ADC used for this measurement was an NI PXIe-5764 digitiser card with four 16-bit ADC channels operating at 1GSps, of which three were used for the corresponding seismometer outputs. The conversion factor of the seismometer was specified to be 60 V/g, where g is the gravitational acceleration. The high gain of this instrument meant that large voltage signals could be produced in excess of the ADC damage threshold. We therefore constructed three electrical voltage limitors and ensured the measurements were conducted at a quiet time.

The seismometer was first placed inside the polycarbonate enclosure on the optical table, shielded from strong air current. The X axis was aligned with the short side (1.2 m) of the optical table while the Y axis measured acceleration along the long side (2.4 m). The Z axis corresponded to vertical movements. The PXIe-5764 digitised all three voltages at 1GSps and down-sampled them to 8 kHz using a 6th order CIC filter. The resultant readout was recorded to disk for 10 minutes, and calibrated to acceleration. In the blue, red and yellow trace of Fig. 4.13, we plot the amplitude spectral density for the X, Y and Z axis readout.

In Fig. 4.13, we see some mechanical features centred at 8 Hz on the X and Y axes, 23 Hz on Z, and 43 Hz on all three axes. These features correlate well with the peaks seen in our optical noise floor, indicating that those resonances likely originate from the optical table. Between 200 Hz and 4 kHz, all axes have strong mechanical and acoustic pickup, with the dominant resonance occurring at 2.7 kHz. From this, we can speculate that the 430 Hz feature seen in our experiment noise floor could be a resonance of one of the optical spools, driven by the seismic pickup by the optical table around that frequency band.

To ensure the validity of the above measurement, a null measurement was also carried out on the 5764, with all three ADC channels open and unterminated. The resultant amplitude spectral density is plotted in the black trace in Fig. 4.13, showing the ADC noise floor of the digitiser. As the seismic movements were small, the output voltage from the seismometer was on the order of 10 mVpp for all three axes. This meant that only a
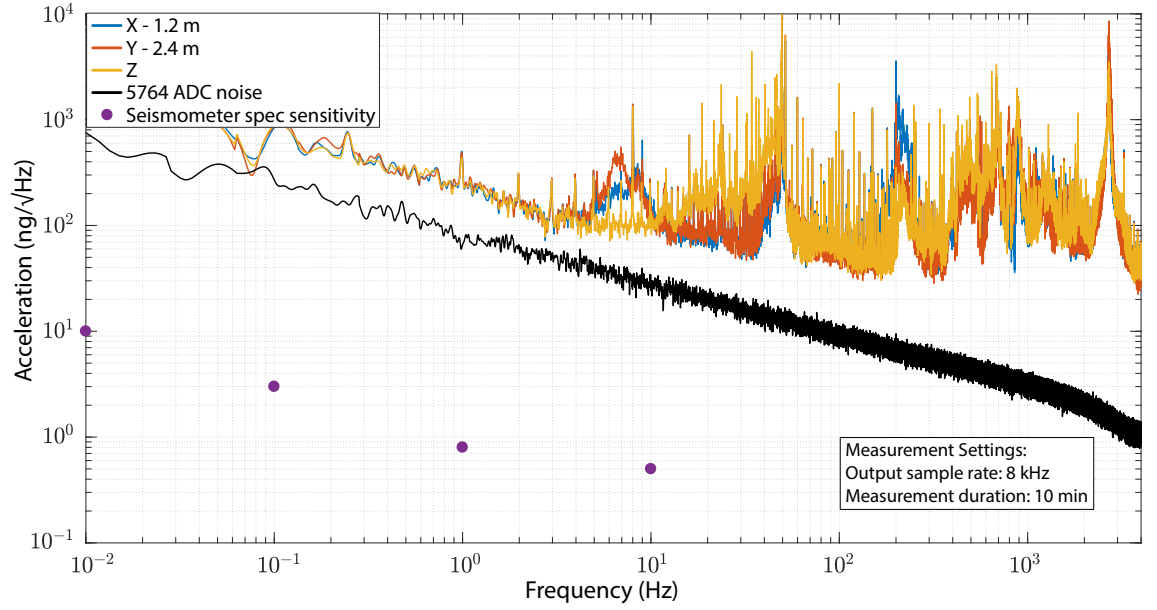
Figure 4.13: Seismic profile of the optical table in frequency domain and measurement noise floor. Blue: acceleration in X axis, parallel to the short side of the table. Red: seismic movement in Y axis, aligned with the long side of table. Yellow: acceleration profile for the vertical axis. Black: null measurement characterising ADC noise floor. Purple: specified sensitivity for the seismometer.

small proportion of the ADC dynamic range was utilised, and therefore the measurement was close to being limited by ADC front end noise below 10 Hz. However, we can still be confident in the 8 Hz features seen in the X and Y axes, which protrude significantly above the ADC noise floor. This null measurement also validated the previous measurement above 10 Hz Fourier frequency, where ADC noise rolls off as $1/\sqrt{f}$ and no longer imposes a limitation.

Figure 4.13 also plots the noise floor specification of the seismometer in purple dots [167], which resides below our ADC noise floor across the entire measurement bandwidth. The seismometer in addition has a pass band of 5 mHz - 1.5 kHz. Overall, we do not see the sensitivity or frequency response of the seismometer as a limitation for these measurements.

4.4.1   Measuring the Optical Table Transfer Function

To further characterise the mechanical performance of the optical table, we proceeded to measure its mechanical transfer function. For this, the seismometer was moved onto the laboratory floor, preserving its axis orientation. A wooden enclosure was placed on top of the seismometer to reduce mechanical and acoustic coupling from air pressure. The measurement was conducted in a similar fashion over 10 minutes, and the resultant amplitude spectral density was computed for all three axes. The transfer function of the optical table was then obtained by dividing the table response by the ground response for each axis respectively, the result of which is plotted in Fig 4.14.

We can see that the 8 Hz feature remains present in the horizontal axes and the 43 Hz in the vertical direction. These two frequencies are therefore likely resonances of the optical
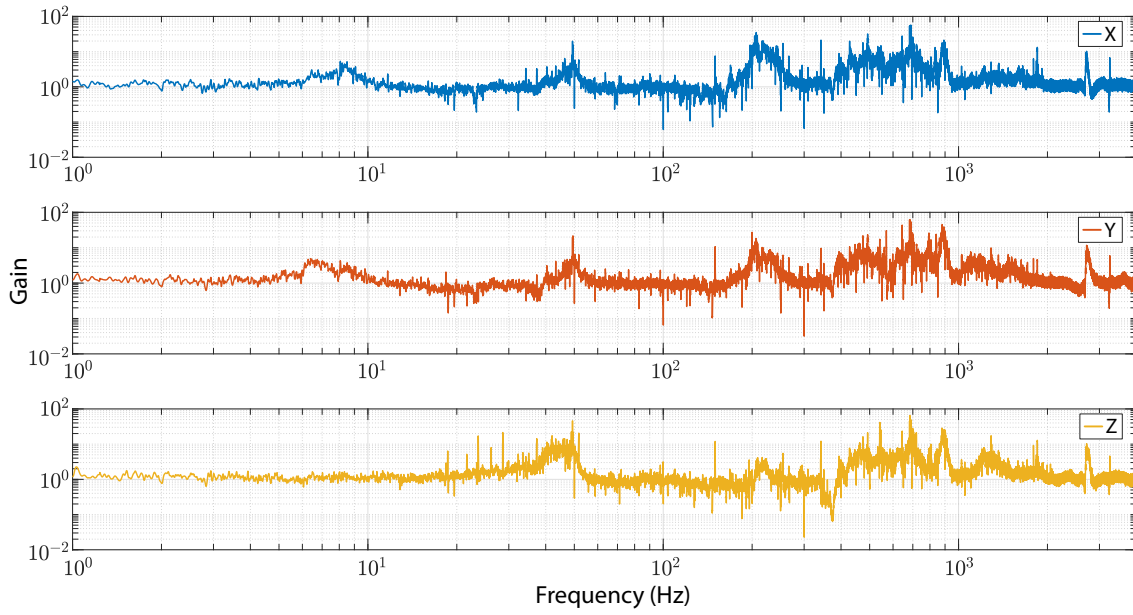
Figure 4.14: Mechanical transfer function of the optical table used for this experiment in X (blue), Y (red) and Z (yellow) directions. 8 Hz and 43 Hz features can be attributed to table resonances.

table. Specifically, the resonance in Y appears to be at a slightly lower frequency, possibly due to longer table dimension in that direction. At 23 Hz, seismic noise can be found in the amplitude spectral density of the floor measurement. This means that the table does not resonate at 23 Hz but transfers movement of the building with minimal suppression.

At high Fourier frequencies, the optical table has more mechanical pickup, most noticeably around 220 Hz, between 350 Hz and 1 kHz and at 2.7 kHz. These peaks fall in the typical "resonance region" of an optical table, and can be effectively reduced via active vibration damping [168]. Alternatively, without treating the optical table, future work could look to increase the mechanical stability of the fibre interferometer itself, for example by using an epoxy-potted fibre coil.

4.5   Residual Laser Frequency Noise

The stability measurement for our experiment is dependent on the complete cancellation of laser frequency noise between the two interferometers. This however can be hampered by the length difference of the two systems resulting in different measurements of the laser frequency. In this section, we quantify the amount of residual frequency noise in our stability measurement, using three different subtraction methods. We also confirm for each method whether residual laser frequency noise imposes an experimental limit.

4.5.1   Interferometer Response Differences

We recall from Section 3.1.2 the transfer function from laser frequency fluctuation to interferometer phase readout is:

$$|H(s)| = \left| \frac{e^{-2\pi i f \tau} - 1}{if} \right| \tag{4.11}$$

where $\tau = L/nc$ is the optical transit time through the fibre coil. In this experiment, the coils used for the two interferometers differ in length by about 25 m, which corresponds with 0.17% difference in their total length. This results in small differences in both the optical gain and the FSR between the two interferometers, as seen below in Fig. 4.15, and gives rise to incomplete cancellation of laser frequency noise in the subtraction readout.
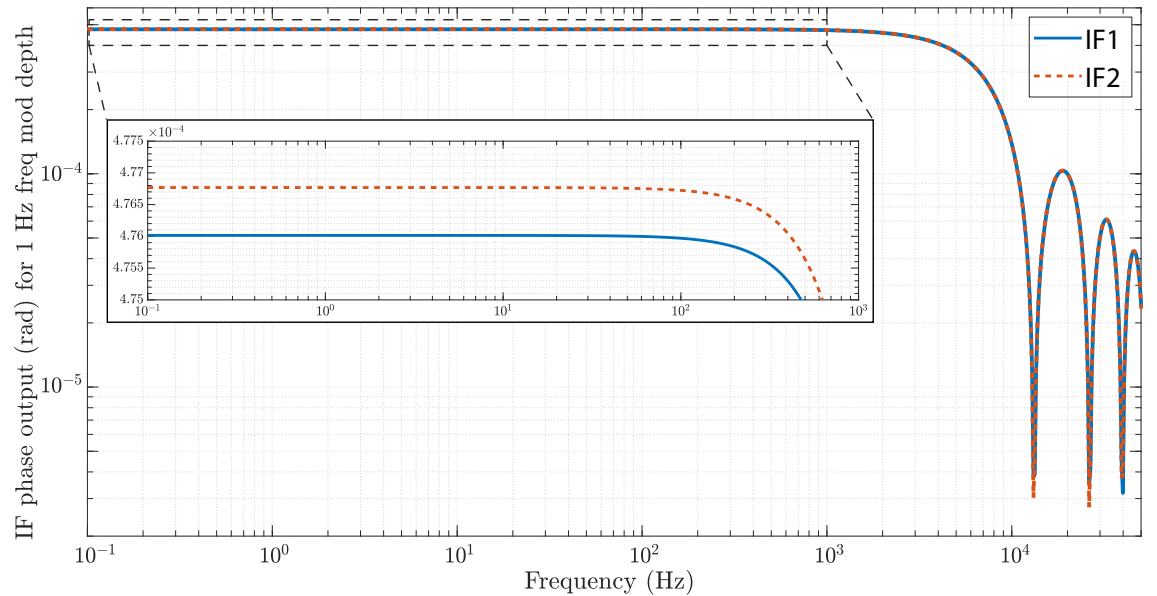
Figure 4.15: Modelled amplitude response of the two interferometers used in this experiment. The 0.17% armlength difference between the two systems results in slight differences in both the optical gain, shown in the inset, and the cut-off FSR frequency.

For the following analysis, we define a laser frequency "suppression ratio" as the ratio between the amplitude response of a single interferometer and that of the subtraction. This allows us to determine whether residual frequency noise is a limitation by comparing it with the experimentally measured $\nu_1/\Delta\nu$ ratio.

4.5.2 Determining the Suppression Ratio

Direct phase subtraction

We first characterise the achievable suppression ratio for a direct phase subtraction between the two interferometers, without the conversion into frequency. This is done by computing the transfer function from laser frequency fluctuation to differential phase readout. The input and output of this system are written as follows:

$$
\begin{aligned}
\text{input: } f(t) &= \frac{1}{2\pi}\frac{d\phi(t)}{dt} \\
\text{output: } \Delta y(t) &= y_1(t) - y_2(t) \\
&= [\phi(t) - \phi(t - \tau_1)] - [\phi(t) - \phi(t - \tau_2)]
\end{aligned}
\tag{4.12}
$$

where $f(t)$ is the instantaneous laser frequency and equates the derivative of its phase $\phi(t)$. $y_1(t)$ and $y_2(t)$ are the respective phase readout of the two interferometers, and τ_1 and τ_2 correspond to their individual transit time. Taking the Laplace transform, we arrive at the following transfer function:

$$
\Delta H_\phi(s) = 2\pi \frac{e^{-s\tau_1} - e^{-s\tau_2}}{s}
\tag{4.13}
$$

In trace (b) of Fig. 4.16, we plot the absolute value of $\Delta H_\phi(s)$ in dashed blue against the left y-axis. The amplitude response of interferometer 1, given by Eqn. 4.11, is also included in trace (a) for comparison. From this, we see a laser frequency suppression ratio of ~ 630 at all Fourier frequencies below the interferometer FSR. In an experimental measurement, this translates to a residual laser frequency noise floor that is a factor of 630 below the measured laser frequency level. This clearance is insufficient for us to reach the desired sensitivity, particularly at low Fourier frequencies where the laser drift is high. Therefore a more sophisticated subtraction method is required to calibrate for the gain and transit time differences between the two interferometers.

FSR calibration

The phase-to-frequency readout conversion described in Eqn. 4.2 allows partial correction for the interferometer length difference, as individual FSR values are multiplied to the original phase readout. For convention, this process will be referred to as FSR calibration. To quantify the suppression ratio using FSR calibration, we compute the transfer function for both a single interferometer frequency readout and for the frequency subtraction. We write out the time-domain input and outputs as follows:

$$
\begin{aligned}
\text{input: } f(t) &= \frac{1}{2\pi}\frac{d\phi(t)}{dt} \\
\text{output for single IF: } \nu_1(t) &= [\phi(t - \tau_1) - \phi(t)] \times \frac{1}{2\pi\tau_1} \\
\text{output for subtraction: } \Delta\nu(t) &= \nu_1(t) - \nu_2(t)
\end{aligned}
\tag{4.14}
$$
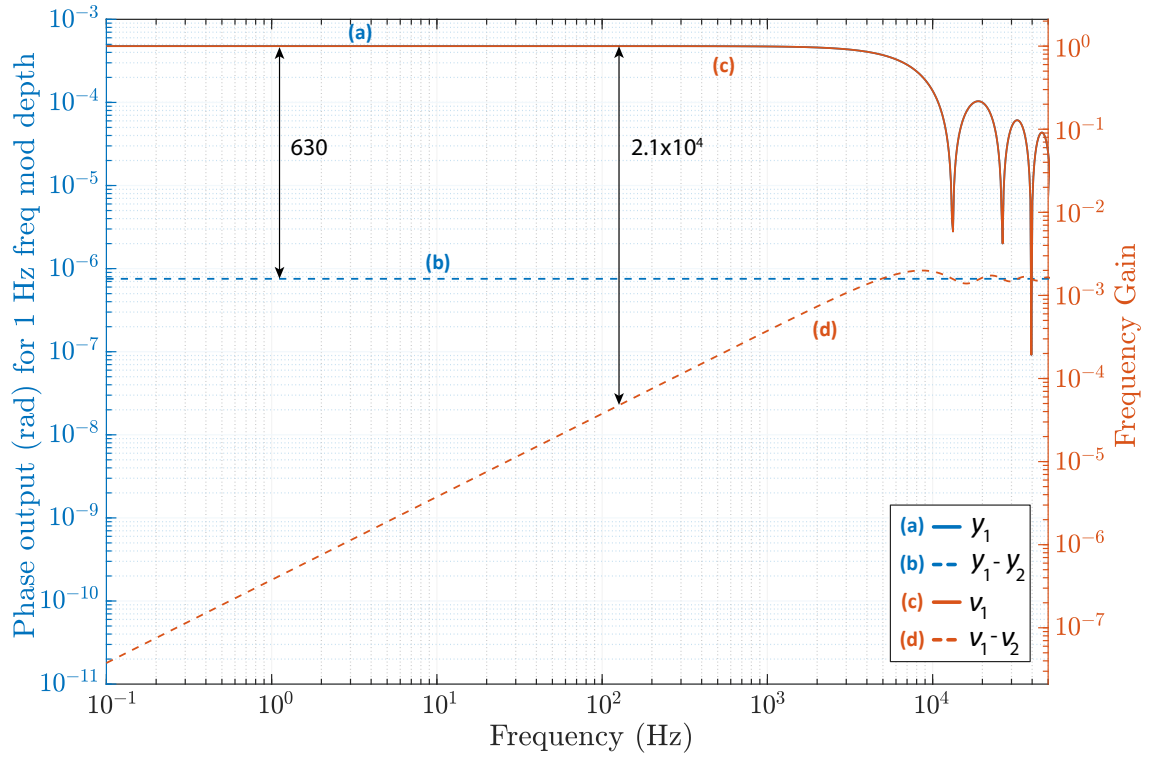
Figure 4.16: Comparison between two readout subtraction methods. Solid blue (overlapped with solid red): interferometer 1 phase response. Dashed blue: direct phase subtraction response. Solid red: interferometer 1 frequency response. Dashed red: frequency subtraction response. We see that FSR calibrated frequency subtraction (dashed red) achieves higher input frequency suppression over most frequency bands. However, the method is not guaranteed to sufficiently remove input frequency noise at kilohertz Fourier frequencies, due to the calibration's inability to correct for actual interferometer FSR mismatch.

Applying Laplace transform again, we find the following transfer functions:

$$H_{\nu_1}(s) = \frac{e^{-s\tau_1} - 1}{s\tau_1}$$
$$\Delta H_\nu(s) = \frac{e^{-s\tau_1} - 1}{s\tau_1} - \frac{e^{-s\tau_2} - 1}{s\tau_2} \tag{4.15}$$

The amplitude of $H_{\nu_1}(s)$ is plotted in trace (c) in Fig. 4.16 against the right y-axis. We can see that the response follows the same shape as the phase response in trace (a), but the gain is calibrated to 1 for $f \ll$ FSR. In addition, at these low Fourier frequencies, the FSR calibration reduces the difference between the two interferometers response that was observed earlier in Fig. 4.15.

The amplitude response of the frequency subtraction, $|\Delta H_\nu(s)|$, is plotted in trace (d) of Fig. 4.16 against the right y-axis. This response follows a f dependence up to approximately half the average FSR frequency, beyond which we see the onset of a few ripples. Comparing with the direct phase subtraction response of trace (b), the FSR calibration achieves higher suppression ratio up to ~5 kHz. The improvement is particularly noticeable at low Fourier frequencies, and this allows suppression of laser frequency noise below other experimental noise sources. Above 5 kHz, the FSR calibration

demonstrates no further advantage to direct phase subtraction. This is because the calibration merely rescales the interferometer phase outputs and does not adjust for the differential transit time. Thus, it is unable to correct for the mismatch towards individual interferometer nulls.

As can be inferred from the earlier spectral plot of Fig. 4.2, most metrology lasers we develop for will have a typical "knee frequency" between 0.1 and 1 Hz. This knee frequency divides the spectral performance. High drift is usually observed below the knee frequency, while above it, the laser noise rolls off as a polynomial of f and the suppression requirement becomes more relaxed. Still, depending on the specific dynamics of the laser in use, the available clearance in the kilohertz band for both subtraction methods may still be insufficient, and residual laser frequency noise could still become a limitation in this frequency range.

In order to explicitly determine whether the experiment was residual frequency noise limited in any Fourier regime, we now use Time-Delay Interferometry (TDI) to correct for both the gain and transit time differences between the two interferometers, and compare the resulting sensing sensitivity.

4.5.3 Time Delay Interferometry

Time delay interferometry (TDI) was initially proposed by Tinto et al [169] to allow unequal armlength interferometers to achieve sensitivity well below the frequency noise of the interrogation laser. The technique has since been extensively studied and modelled for the Laser Interferometric Space Antenna (LISA) mission [170, 38, 171, 172], in which three satellites orbiting the solar system form multiple armlength unbalanced interferometers, and laser frequency stabilisation alone becomes insufficient for the desired sensitivity.

The operation of TDI allows the circumvention of laser frequency noise while preserving the desired sensing signal. This is done by combining the readout of at least two interferometers and their respectively time-delayed versions, synthesising a Sagnac-type response. The resulting construct, given precise knowledge of the interferometer pathlength, can be free of laser frequency noise. This presents a computationally simple, time-domain solution without measurement time or bandwidth constraints, as opposed to its frequency-domain counterpart [173]. However, it is worth noting that despite theoretical complete laser noise removal, the actual suppression achieved experimentally is often limited by pathlength fluctuations and the ability to accurately track and promptly adjust corresponding delays. These technical details are examined in detail in several LISA studies [174, 175, 38, 176].

Adapting LISA TDI to our system, the required time-domain construct to cancel source frequency noise is

$$x\left(t\right) = y_1\left(t - \tau_2\right) - y_1\left(t\right) - y_2\left(t - \tau_1\right) + y_2\left(t\right) \tag{4.16}$$

where $y_1(t)$ and $y_2(t)$ are the respective phase readout of the two interferometers, and are subtracted with a copy of themselves delayed by the other's transit time. If we substitute in Eqn. 4.1, we see that the laser frequency terms $\phi(t)$ are completely cancelled out.

TDI modulation function

Clearly, the frequency response of $x(t)$ would differ from that of $\Delta y(t) = y_1(t) - y_2(t)$ as calculated earlier. We now use Laplace transform again to quantify this difference and determine how $x(t)$ alters the spectrum of our noise floor. We know that in addition to residual laser frequency noise, the subtraction readout $\Delta y(t)$ also contains uncommon noise sources between the two interferometers, written as

$$\Delta n(t) = n_1(t) - n_2(t) \tag{4.17}$$

For the TDI construct $x(t)$, no residual frequency noise is present, and therefore its readout could be simplified to

$$x(t) = n_1(t - \tau_2) - n_1(t) - n_2(t - \tau_1) + n_2(t) \tag{4.18}$$

Taking the Laplace transform of $\Delta n(t)$ and $x(t)$, we arrive at the transfer function of TDI computation

$$H_{\text{TDI}}(s) = \frac{X(s)}{\Delta N(s)} = \frac{\left(e^{-\tau_2 s} - 1\right) N_1(s) - \left(e^{-\tau_1 s} - 1\right) N_2(s)}{N_1(s) - N_2(s)} \tag{4.19}$$

We further approximate that $\tau_1 = \tau_2$ and simplify the above to

$$H_{\text{TDI}}(s) = e^{-\tau_{\text{avg}} s} - 1 \tag{4.20}$$

where $\tau_{\text{avg}} = (\tau_1 + \tau_2)/2$. The amplitude of Eqn. 4.20, $|H_{\text{TDI}}(s)|$, is plotted in Fig. 4.17, showing the behaviour of uncorrelated noise in the TDI construct. We can see that up to half the average interferometer FSR, the modulation function exhibits a f dependency with a peak gain of 2 and a unity gain frequency of 2.2 kHz.

Knowing the TDI modulation function allows us to recover our original interferometer noise floor by applying the inverse transfer function. For frequencies well below the interferometer FSR, this could be simplified to a multiplication of f_0/f in the frequency domain, where f_0 is the unity gain frequency. The resulting spectrum is then directly comparable to that of the two subtraction readouts discussed earlier, allowing us to observe any noise floor difference from the removal of residual frequency noise.

Although we have only computed the Laplace transform for noise terms, this analysis also applies for other interferometric measurands. In the example of LISA spacecrafts, gravitational wave signals are another component of the interferometer readout $y(t)$. Following the same calculation process, we see that these signals are preserved in the TDI readout and modulated by the same function. Therefore, recovery of these signals to their original amplitude simply requires the application of the inverse modulation function.

FPGA implementation

There are two methods to experimentally implement the TDI delays seen in Eqn. 4.16: in real-time on the FPGA and in post-processing. For LISA, post-processed TDI is
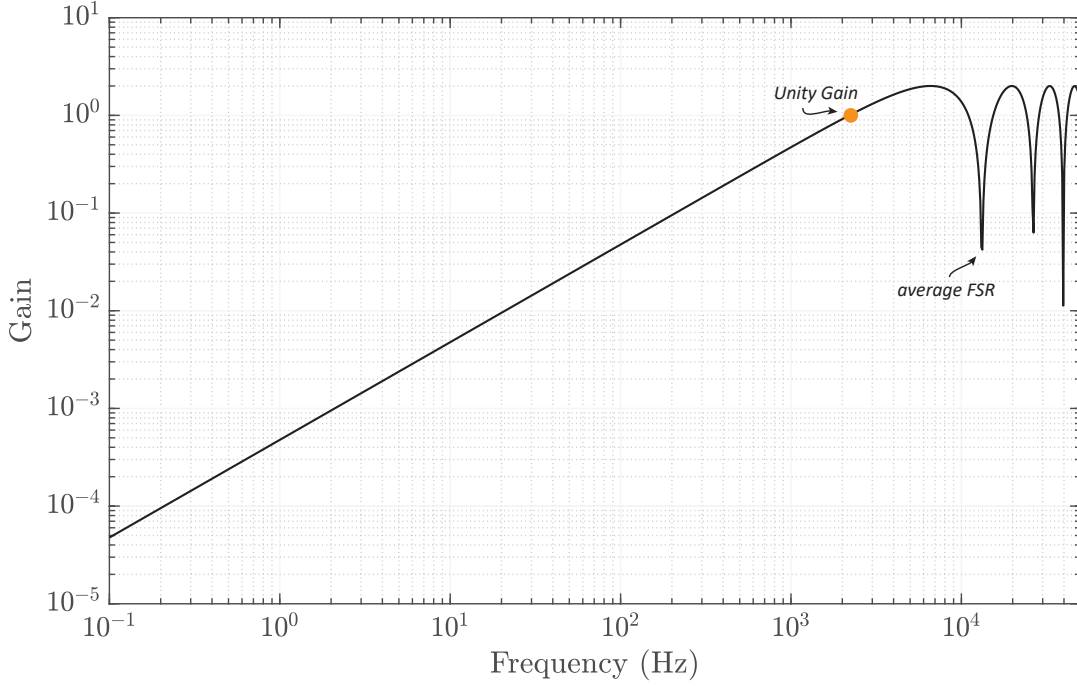
Figure 4.17: The interferometer noise floor is modulated by a frequency dependent function when applying TDI. This function increases in gain as a function of f up until half of the interferometer FSR, crossing unity gain at 2.2 kHz for our interferometers. Nulls can be observed at integer multiples of the FSR.

strongly favoured for the reduction of onboard calculation and chance of error, relaxed data telemetry requirement and increased algorithm flexibility on Earth [177]. The implementation of post-processed TDI often requires a fractional delay filter to obtain sub-sample resolution. While fractional delay filters are highly effective at interpolating and recovering downsampled data, they are however susceptible to interpolation errors [178, 177].

In comparison, our experiment operates on ground and is controlled locally. This, in addition to sufficient FPGA resources, favours the real-time approach for our purposes. When implemented on our FPGA platform, we are able to custom delay the RT phase readout of each interferometer, with an adjustment resolution set by the FPGA clock cycle of 8 ns. The block diagram of this procedure is shown in Fig. 4.18 for interferometer 1. A memory delay block, similar to that implemented for the RT code filter (Fig. 3.9(f)), takes in the CORDIC phase output of interferometer 1 and delays it by the transit time of interferometer 2. Following this, both the prompt $(y_1(t))$ and delayed $(y_1(t - \tau_2))$ outputs are downsampled and streamed to host at 8 kHz. This block diagram is duplicated for interferometer 2, and the final linear combination of the TDI construct, $x(t)$, is computed in post-processing.

Following the FPGA implementation, delays for each interferometer were varied within their approximate range until maximum laser frequency suppression was achieved [179]. The final clock delays were fixed to 9485 and 9470 cycles for 1550 nm corresponding to τ_2 and τ_1 respectively.
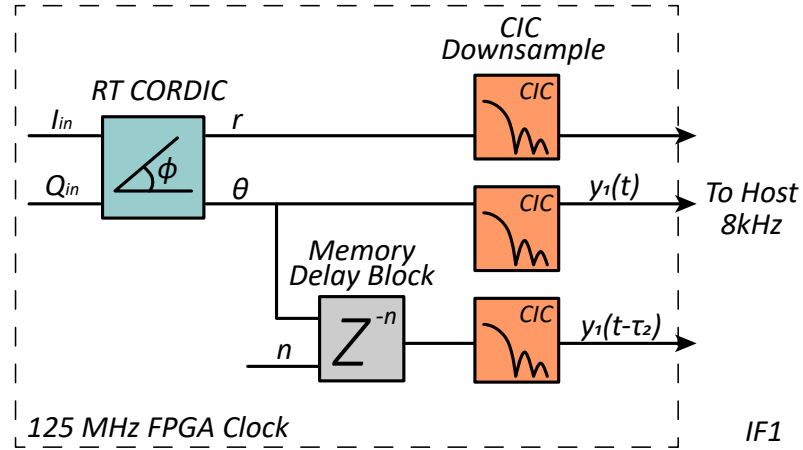
Figure 4.18: FPGA implementation of real-time TDI. Each interferometer's CORDIC phase output is delayed by the transit time of the other interferometer in integer multiples of 8 ns, before being filtered to 8 kHz for host monitoring and recording. Real-time TDI is the preferred approach for our experiment as it allows precise control of relative delays while avoiding loss of data or interpolation errors due to downsampling.

4.5.4 Experimental Verification

Having modelled the achievable laser frequency noise suppression for three subtraction methods, we now proceed to experimentally verify these ratios. For this, a sinusoidal frequency modulation (FM) tone was added through the encoding AOM, simulating laser frequency noise at the FM frequency. We subsequently quantify the ratio between this tone as measured by a single interferometer and by the three subtraction methods. The results are summarised in Fig. 4.19, where theoretical suppression, as given in Fig. 4.16, are repeated in trace (a) and (b) for direct phase subtraction and FSR calibration respectively. Experimentally measured values are plotted in dots at specific FM frequencies, blue for direct phase subtraction (c), red for FSR calibration (d), and green for unmodulated TDI (e, where the f dependency has been removed).

These FM tests only occurred between Fourier frequencies of 60 Hz to 1.9 kHz. This is because to obtain a valid suppression measurement, residual FM tone for all three subtraction methods must protrude beyond the experimental noise floor. This becomes difficult at low Fourier frequencies due to raised mechanical and temperature noise floor, as well as the FSR method's increased suppression resulting in no protrusion despite high modulation depth. At high frequencies, the ability to track fast frequency changes becomes a challenge, forcing smaller modulation depth and resulting in insufficient clearance to characterise TDI suppression. Nonetheless, the frequency range tested here does contain the kHz regime of concern, where FSR calibration may not sufficiently reject source frequency noise. These tests also enable comparison between the three methods and earlier models.

For direct phase subtraction and FSR calibration, experimental data show good agreement with their analytical estimates. The FSR method achieves substantially higher suppression at low Fourier frequencies but deteriorates in performance towards the interferometer FSR frequency. While the TDI readout in theory has no residual frequency noise, the outcome
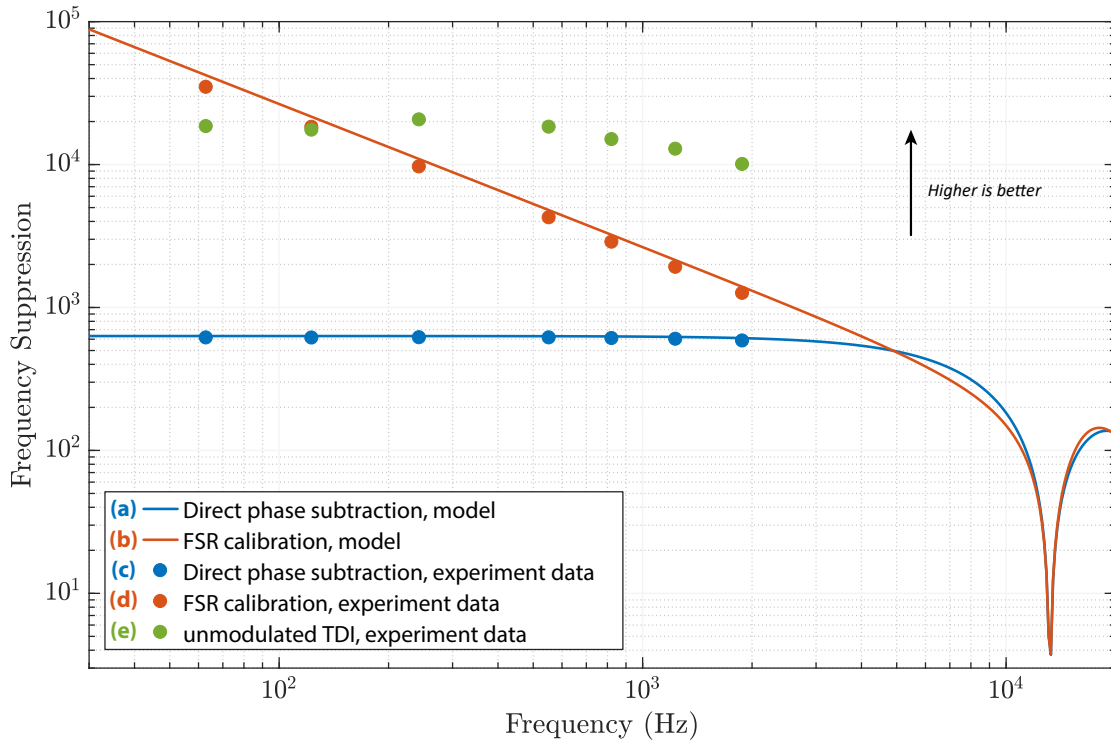
Figure 4.19: Summary of frequency modulation test results. The suppression ratio was determined for each subtraction method at each modulation frequency, with higher values representing greater suppression. (a) analytical model for suppression obtained by direct phase subtraction. (b) model for suppression by FSR calibration. (c) suppression by direct phase subtraction at several test FM frequencies. (d) experimentally measured suppression for FSR calibration. (e) experimental data for TDI readout with f modulation removed. The suppression ratio achieved by TDI (e) in the kHz regime exceeded the experimental clearance requirement for residual frequency noise but resulted in no noise floor improvement. From this we concluded that the current sensitivity was not limited by residual frequency noise, and FSR calibration alone could provide sufficient suppression at all Fourier frequencies.

is dependent on the ability to precisely delay by the required amount, which for us is limited to the FPGA clock resolution. This limitation also explains poorer suppression using the TDI readout below ∼150 Hz.

At high Fourier frequencies, however, TDI has a clear advantage over the other two methods, maintaining over four orders of magnitude frequency suppression. For the frequency range tested here, this ratio is unambiguously beyond the experimentally achieved clearance by a minimum factor of 10. Examining the frequency-domain spectral density of the unmodulated TDI readout, however, we find good overlap between its noise floor and that of direct phase subtraction and FSR calibration, for all audio frequencies except the test FM tone. Despite higher and sufficient frequency suppression in the kHz band, TDI readout does not improve the experimental noise floor. This concludes that the current experimental sensitivity is not limited by residual laser frequency noise, and FSR calibration alone should provide sufficient suppression across both infrasonic and audio frequency bands.

System Noise Budget

In previous chapters, we have characterised the temperature stability of our fibre interferometers, and attributed mechanical features in the 4 - 70 Hz Fourier regime to optical table resonances. To complete the experimental characterisation, in this chapter, we model and quantify limiting noise sources for the remaining Fourier bands. In Section 5.1, we examine the effects of fibre thermal noise. This is followed by Section 5.2 - 5.3 where we discuss broadband noise limitations including double Rayleigh backscattering (DRS) and DEHoI tracking. Noise contributions from additive sources are calculated in Section 5.4. Finally, these elements are summarised in Section 5.5 where we present a complete noise budget for our system.

5.1   Fibre Thermal Noise

The long lengths of optical fibre used in this experiment act as an effective thermal amplifier, increasing the phase error induced by both temperature fluctuation and intrinsic thermal noise. The differentiating factor between the two thermal noise sources is that the former is dependent on environmental temperature changes, whereas the latter is primarily dependent on the absolute temperature at which the material's internal degrees-of-freedom fluctuate. As such, temperature noise can be effectively reduced via passive isolation or active stabilisation, whereas mitigation of intrinsic thermal noise requires substantial cooling, an often impractical task for our intended usage scenarios. In addition, for glass-like materials, reduction of temperature often results in increased loss angle and, as we will see in this section, limits the reduction of thermal noise. For these two reasons, intrinsic thermal noise represents a fundamental limit for our system, and will be the focus for this section.

Early research in fibre thermal noise has drawn analogy between it and mirror noise seen in LIGO interferometers, where spontaneous, localised temperature fluctuations induce microscopic changes in interferometer pathlength through displacement of the coating material [180]. In the case of a fibre interferometer, these minuscule temperature variations result in the expansion and contraction of the glass as well as changes in its refractive index. These two effects are studied and modelled independently by Wanser [181] and Foster [182], both predicting a frequency independent fibre thermal noise floor.

Both theories demonstrate good experimental agreement in the high frequency regime [183, 184, 185, 186, 187]. At infrasonic frequencies, however, experimental measurements deviate

from their predictions, as a $1/f$ type noise in excess of laser phase noise is widely observed [188, 117, 189, 190]. This additional noise is believed to originate from thermo-mechanical effects, where random expansion and contraction of the fibre occur due to internal friction or structural damping [191]. Following these experimental observations and theoretical proposals, Duan sought to bring both types of thermal effects under a universal framework through the fluctuation-dissipation theorem [11]. Our modelling in this section is based on this prior work by Duan, in which two mechanisms of fibre thermal noise are considered: 1) thermo-dynamic noise, akin to the derivations presented by Wanser and Foster, and 2) thermo-mechanical noise, which fixes the low-frequency experimental discrepancy in earlier models [11, 190, 192].

5.1.1   Thermo-mechanical Noise

The thermo-mechanical noise estimate is derived through a normal mode expansion analysis detailed in [191] and [193], where each normal mode is treated as a harmonic oscillator to find the overall spectral density for fibre length change from structural damping. An important simplification of this calculation occurs in Eqn. 25 of [191]. This assumes the frequency range of concern is significantly below the first longitudinal resonance of the fibre, and is not entirely justified for our system with 15 km of fibre corresponding to a first longitudinal mechanical resonance of approximately 0.24 Hz [190]. An exact calculation for our experiment would therefore require the numerical evaluation of an infinite number of higher order modes, as per Eqn. 24 in [191].

Fortunately, as the mode order increases, its impact on noise spectral density decreases exponentially. As we computed the analytical sum of 20 modes, this in fact converged to the value given by the aforementioned assumption and simplification. We further observe that the assumption still holds true for us at very low frequencies, and the $1/f$ trend continues through the mechanical resonances at higher frequencies, as shown in [191]. In addition, these resonances have not been observed experimentally, and could be caused by oversimplifications in the model such as assuming a constant, frequency independent loss angle ϕ_0. If we assume no distinct mechanical peaks and a continuation of the $1/f$ trend, we can justify applying the simplified thermo-mechanical noise estimate to all Fourier frequencies. The simplified phase spectral density for thermo-mechanical noise is [11]:

$$\tilde{S}_\phi(f) = \sqrt{\frac{4\pi L}{\lambda^2}\frac{2k_B T n^2 \phi_0}{3E_0 A}\frac{1}{f}} \;\; \text{rad}/\sqrt{\text{Hz}} \tag{5.1}$$

We can further convert this into frequency spectral density by multiplying with the interferometer FSR. We also add another factor of $\sqrt{2}$ for the incoherent addition between the two interferometers:

$$\tilde{S}_\nu(f) = \sqrt{2}\,\frac{c}{2\pi n L}\,\tilde{S}_\phi(f) \;\; \text{Hz}/\sqrt{\text{Hz}} \tag{5.2}$$

A glossary of the constants is provided in Tab 5.1. Substituting in these values, Eqn. 5.2 becomes

$$\tilde{S}_\nu(f) = \frac{26.48}{\sqrt{Lf}} \;\; \text{Hz}/\sqrt{\text{Hz}} \tag{5.3}$$

Name	Symbol	Value
Boltzmann constant	k_B	1.380649×10^{-23}
Temperature	T	298 K
Mechanical loss angle	ϕ_0	0.01
Young's modulus without loss	E_0	19 GPa
Cross-sectional area of fibre	A	$\pi(125\,\mu\text{m})^2$
Optical wavelength	λ	1550 nm
Refractive index	n	1.47
Thermo-optic coefficient	$\frac{dn}{dT}$	9.488×10^{-6} K^{-1}
Linear expansion coefficient	α_L	5×10^{-7} K^{-1}
Thermal conductivity	κ	1.37 W/m/K
Power profile radius	r_0	$r_0{}^2 = a^2/2$
Mode field radius	a	5.2×10^{-6} m
Thermal diffusivity	D	8.2×10^{-7} m^2/s

Table 5.1: Glossary of thermal constants.

In trace (b) of Fig. 5.1 below, we plot the thermo-mechanical noise estimate for our interferometer length of $L = 15$ km. Comparing this with our experimental noise floor of trace (a), identical to trace (b) shown earlier in Fig. 4.2, we can identify a thermo-mechanical noise limited region of 400 mHz - 2 Hz, indicated by the shaded area.
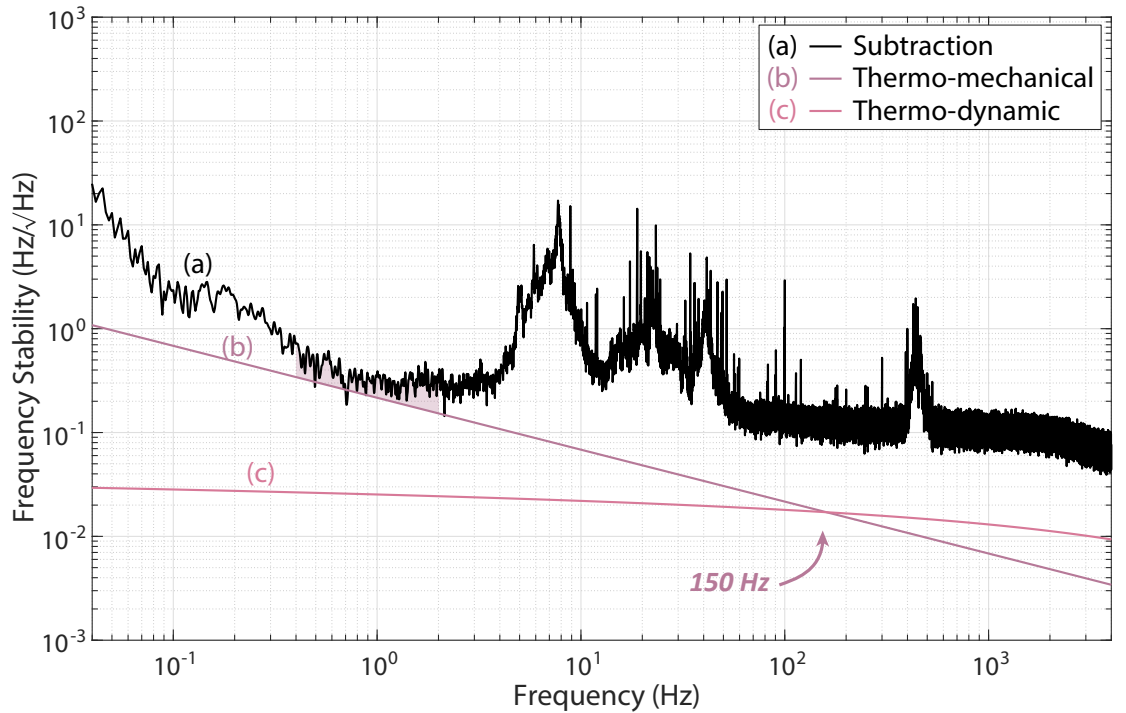
Figure 5.1: The experimental noise floor (trace a) is compared with two types of fibre thermal noises. Thermo-mechanical noise (trace b) presents a $1/\sqrt{f}$ dependency, and resides within a factor of two of the experimental sensitivity between 400 mHz to 2 Hz (shaded region). In contrast, thermo-dynamic noise (trace c) is a broadband noise source and does not currently impose an experimental limitation. These two thermal noise sources always intersect at the same Fourier frequency regardless of fibre interferometer length, at approximately 150 Hz for our system.

5.1.2   Thermo-dynamic Noise

Thermo-dynamic noise in optical fibre consists of spontaneous fibre length change and refractive index variation. The two processes are jointly related to localised temperature fluctuations, and their combined noise power spectral density $S_\phi(\omega)$ can be written as a function of the temperature power spectral density $S_{\delta T}(\omega)$ [191]:

$$S_\phi(\omega) = \frac{4\pi^2 L^2}{\lambda^2} \left(\frac{dn}{dT} + n\alpha_L \right)^2 S_{\delta T}(\omega) \tag{5.4}$$

Further, $S_{\delta T}(\omega)$ is derived using the fluctuation-dissipation theorem to be [191]:

$$S_{\delta T}(\omega) = \frac{k_B T^2}{4\pi^2 L \kappa} \text{Re} \left[\exp\left(\frac{i\omega r_0^2}{2D} \right) E_1 \left(\frac{i\omega r_0^2}{2D} \right) \right] \tag{5.5}$$

where $E_1(x)$ is the special function of the exponential integral, and the other constants are specified in the glossary of Tab. 5.1. Substituting $S_{\delta T}(\omega)$ from Eqn. 5.5 into Eqn. 5.4 then allows us to compute the power spectral density of thermo-dynamic noise. When converting to single-sided phase spectral density, an extra factor of $\sqrt{4\pi}$ applies:

$$\tilde{S}_\phi(f) = \sqrt{4\pi} \sqrt{S_\phi(2\pi f)} \tag{5.6}$$

The same process is followed to convert this phase spectral density into frequency spectral density, and an extra factor of $\sqrt{2}$ is added to account for the second interferometer. The resultant thermo-dynamic noise estimate for $L = 15$ km is plotted in trace (c) of Fig. 5.1, intersecting with thermo-mechanical noise (trace (b)) at approximately 150 Hz. Currently, thermo-dynamic noise does not impose a limitation for our experiment in any frequency regime. It does however lie within a factor of 10 of our experimental noise floor above 150 Hz, and could become a dominant noise source if higher noise sources are addressed.

One interesting observation from the modelling of fibre thermal noise is that the two mechanisms always intersect at the same Fourier frequency regardless of fibre length. This can be verified by equating Eqn. 5.1 and Eqn. 5.6, and we find that L is cancelled out on both sides. For a transmissive, SMF-28e fibre interferometer with thermal constants assumed in Tab. 5.1, the cross-over frequency occurs at $\sim$150 Hz. Therefore in practice, these two noise sources always dominate different parts of the spectrum regardless of the optical configuration, and their relative scale remains constant.

5.2   Double Rayleigh Backscattering

Scattering in optical fibre occurs as a result of many microscopic defects "frozen in" the glass during the manufacturing process. These inhomogeneities allow light to scatter in various directions, a proportion of which is re-captured and guided backwards, constituting Rayleigh backscattering (RBS). Measurement noise arises when the RBS field interacts with a main metrology field, adding a small, rotating phase error onto the original phasor of interest.

The dynamics of RBS induced phase noise correlates with the source frequency noise [194, 195, 107]. A low-noise, metrology-grade laser drives a lesser phase error, while a broadband source allows for a short coherence length in which interference can occur. This phenomenon has led to the use of both high and low coherence sources in an optical fibre gyroscope to reduce RBS noise [196, 197, 198, 2].

As our system operates in transmission, we are immune to first-order RBS noise but still susceptible to its second-order effect, double Rayleigh backscattering (DRS). DRS occurs when the back-reflected field undergoes a second scattering event. This reverses the scattering field direction of travel, realigning it with the metrology field [199]. The dynamics of DRS is similar to that of RBS. Here we adapt the analytical model developed for RBS [200, 201, 134], and estimate DRS induced phase error for our system.

5.2.1   DRS Field Analysis

The two metrology fields in our Mach-Zehnder interferometer along with the total DRS field can be written as:

$$
\begin{aligned}
E_{\mathrm{ref}}(t) &= \sqrt{P_{\mathrm{ref}}}e^{i\omega(t)t}, \\
E_{\mathrm{sig}}(t) &= \underbrace{\sqrt{P_{\mathrm{sig}}}e^{i\omega(t)t+i\phi(t)}}_{\text{Signal field}} + \underbrace{\sqrt{P_{\mathrm{sig}}\langle\gamma^2\rangle}e^{i\omega(t)t+i\theta(t)}}_{\text{total DRS field}},
\end{aligned}
\tag{5.7}
$$

where P_{ref} and P_{sig} denote the optical powers in the respective interferometer arms, $\phi(t)$ the interferometric phase and $\theta(t)$ the DRS phase. As the reference path is short, we assume no DRS on that path and only add it to the signal field. γ represents the proportion of the total optical power that is backscattered, and is dependent on four factors: the fibre length l in which coherent scattering can occur, scattering coefficient α_S, attenuation coefficient α and capturing coefficient $\mathcal{S}$ [202]. The expectation value of γ in a transmissive interferometer is given by [149]

$$
\langle\gamma\rangle = \alpha_S \mathcal{S} \int_0^l e^{-\alpha z}dz = \alpha_S \mathcal{S} \, \frac{1-e^{-\alpha l}}{\alpha}
\tag{5.8}
$$

For SMF-28e optical fiber operating at 1550 nm, scattering is the dominant source of loss. We can therefore further approximate that $\alpha_S = \alpha$, giving

$$
\langle\gamma\rangle = \mathcal{S}\left(1 - e^{-\alpha l}\right)
\tag{5.9}
$$

When considering DRS, γ is squared due to each DRS event consisting of two scattering events. The phasor diagram of Fig. 5.2 illustrates the DRS noise mechanism. The main measurand of the MZI interferometer, the phase difference between the two arms, is given by $\phi(t)$ of amplitude $\beta\sqrt{P_{\text{ref}}P_{\text{sig}}}$ in dark green, where β is the collective gain from the photo-detector and DSP demodulation algorithms. As β is common to all signals analysed here, we remove it from the following calculations. In Fig. 5.2, the smaller, orange phasor added on top of the main MZI phase vector represents DRS interfering with another metrology field, P_{met}. This interference is offset from $\phi(t)$ by $\theta(t)$, and has an amplitude of $\sqrt{P_{\text{sig}}\langle\gamma^2\rangle P_{\text{met}}}$. The resultant readout error for the $\phi(t)$ can be calculated by:

$$\Delta\phi_{\text{err}}(t) = \arctan\left(\frac{\sqrt{P_{\text{sig}}\langle\gamma^2\rangle P_{\text{met}}}\sin\theta(t)}{\sqrt{P_{\text{ref}}P_{\text{sig}}}}\right) \tag{5.10}$$

which with small angle approximation is equivalent to:

$$\Delta\phi_{\text{err}}(t) = \frac{\sqrt{P_{\text{sig}}\langle\gamma^2\rangle P_{\text{met}}}}{\sqrt{P_{\text{ref}}P_{\text{sig}}}}\sin\theta(t) \tag{5.11}$$
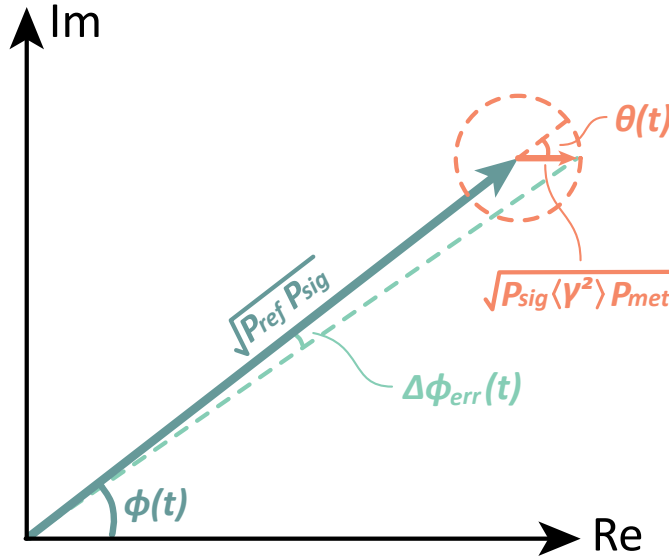
Figure 5.2: An illustration of the interaction between DRS and the metrology fields. Dark green: interference between the reference and signal paths without DRS. This represents the intended interferometric measurand $\phi(t)$ of amplitude $\sqrt{P_{\text{ref}}P_{\text{sig}}}$. Orange: phase vector representing DRS interfering with either the reference or signal field ($\sqrt{P_{\text{met}}}$). This vector contains the ensemble average of DRS phase $\theta(t)$, which is driven by the source phase noise and capped at $\sqrt{P_{\text{sig}}\langle\gamma^2\rangle P_{\text{met}}}$. Light green: resultant measurement error $\Delta\phi_{\text{err}}(t)$.

5.2.2   DEHoI Gating and DRS Coupling Mechanisms

In order to determine $\Delta\phi_{\text{err}}(t)$, we now consider the scenarios by which DRS contributes noise in a DEHoI system. We recall from earlier that, in the case of a highly coherent laser

source, the optical coherence length with DEHoI modulation is equivalent to the physical length of a DEHoI symbol, or chip length:

$$L_c = L_{\text{chip}} = \frac{c}{n \times f_{\text{chip}}}. \tag{5.12}$$

And the total physical length of the modulation code is

$$L_{\text{code}} = L_{\text{chip}} \times N \tag{5.13}$$

where N equals the number of code elements.

For this analysis, we can classify DRS effects into two main categories: code-coherent and code-incoherent. Code-coherent DRS fields have the same DEHoI delay as either the reference or signal field. They are coherently demodulated by DEHoI and indistinguishable from the two metrology fields. Within our operating parameters, the physical length of the optical coil well exceeds that of the code. Therefore there are multiple codes cascaded in the delayline, allowing code-coherent DRS to occur not just within a single chip (L_c), but also integer codes ($\eta \times L_{\text{code}}$) upstream.

The second category, code-incoherent DRS, refers to any DRS field whose optical time-of-flight corresponds to different code delays than the reference and signal fields. As a result, they are incoherently demodulated and subsequently suppressed in the DEHoI readout. The residual, unsuppressed noise is expected to constitute a small proportion of broadband noise.

We now develop analytical expressions for $\Delta\phi_{\text{err}}$ for each category.

5.2.3   Code-Coherent DRS

The two scenarios of code-coherent DRS are illustrated in Fig. 5.3. In both scenarios, a proportion of the interferometer long arm is drawn with code-coherent sections highlighted. The first scenario also depicts microscopic scatterers randomly distributed along the fibre, which reverse the direction of travel for a small percentage of the incident metrology field, forming DRS.

DRS within a single chip

In the first instance, DRS occurs within the physical length of a single chip, accumulating no more than L_c extra travel length than the signal field. This means the first backscattering event can occur anywhere along the fibre, while the second must occur within $L_c/2$ of the first. We can therefore adapt Eqn. 5.9 to write $\langle \gamma_1{}^2 \rangle$ as:

$$\langle \gamma_1{}^2 \rangle = \mathcal{S}^2 \left(1 - e^{-\alpha L} \right) \left(1 - e^{-\alpha L_c/2} \right) \tag{5.14}$$
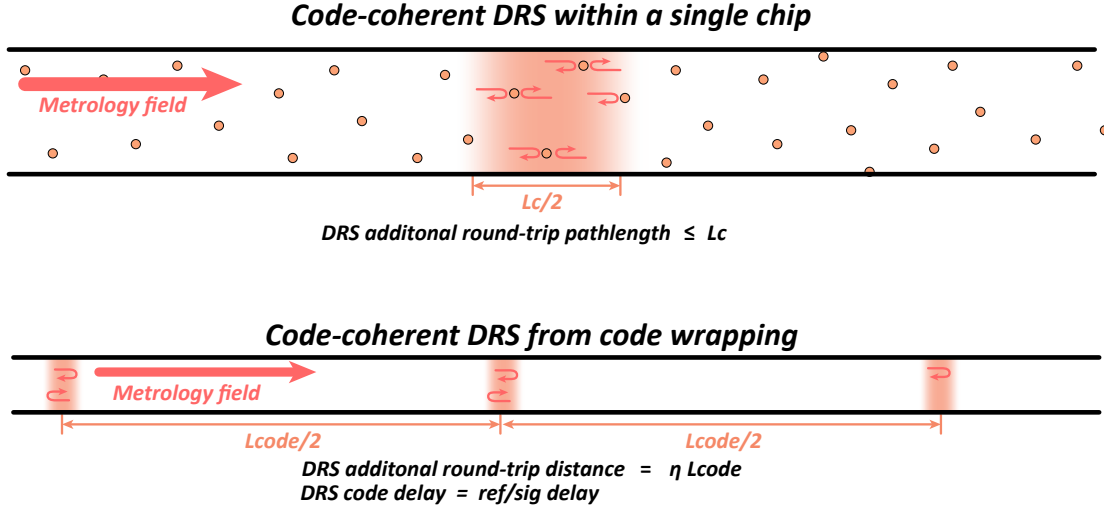
Figure 5.3: Code-coherent DRS can occur in two ways. 1) Code-coherent DRS within a single chip: two consecutive backscattering events occur within a physical distance of $L_c/2$ (highlighted region), resulting in an additional delay of less than a chip for the DRS field, maintaining coherence. 2) Code-coherent DRS from code wrapping: as there are more than one physical code within the fibre coil, the second scattering event can also occur at integer half codes upstream (highlighted regions). The resultant round-trip delay for the DRS field in this scenario equates multiple complete codes, allowing it to re-align with either metrology field, restoring code coherence and constituting noise.

where $L = 15$ km is the length of the fibre coil. The resultant DRS field in this scenario has the same code delay as the signal field, and is demodulated to coherently interfere with the reference field, giving

$$\Delta\phi_{\text{err1}}(t) = \frac{\sqrt{P_{\text{sig}}\langle\gamma_1{}^2\rangle P_{\text{ref}}}}{\sqrt{P_{\text{ref}}P_{\text{sig}}}}\sin\theta(t) = \sqrt{\langle\gamma_1{}^2\rangle}\sin\theta(t) \tag{5.15}$$

As DRS phase $\theta(t)$ is driven by laser phase noise, in the event of rapid phase fluctuation well in excess of 2π, $\Delta\phi_{\text{err1}}(t)$ is manifested as a cyclic error. We can consider this error to be spectrally white and estimate its spectral density by the RMS value of $\sin\theta(t)$:

$$\tilde{S}_{\phi_{\text{err1}}} = \sqrt{\frac{\langle\gamma_1{}^2\rangle}{2}} \tag{5.16}$$

From a previous study, the onset frequency of $\tilde{S}_{\phi_{\text{err1}}}$ was determined to be around 0.1 Hz when measuring a metrology grade laser [149]. The Fourier region below 0.1 Hz provides sufficient input frequency noise to drive the DRS phasor $\theta(t)$, validating the RMS spectral density estimate. Conversely, above 0.1 Hz, the frequency noise of these metrology lasers is reduced, and $\theta(t)$ can no longer rotate and average over 2π within the timescales of concern. The actual single-chip DRS spectrum at these high Fourier frequencies is expected to decline rapidly. $\tilde{S}_{\phi_{\text{err1}}}$ therefore is an estimate for the shelf level of the spectrum, with a corner frequency around 0.1 Hz.

DRS re-aligning with signal delay

In the second scenario depicted in Fig. 5.3, one backscattering event occurs within the total length of the delay coil, followed by a second event integer half codes upstream. Through this process, an additional round-trip delay of integer complete codes is obtained. We can write out the $\langle \gamma_2{}^2 \rangle$ value for this scenario as follows, noting that the proportion, or likelihood, of a second scattering event is dependent on the position at which the first occurs:

$$\langle \gamma_2{}^2 \rangle = \mathcal{S}^2 \sum_{\eta=1}^{\text{Floor}(2M)} \frac{2M - \eta}{2M} e^{-\alpha \eta L_{\text{code}}} \underbrace{(1 - e^{-\alpha L})}_{\text{1st scatter}} \underbrace{(1 - e^{-\alpha L_c/2})}_{\text{2nd scatter}}, \qquad (5.17)$$

where M represents the number of code wraps within the fibre coil and is determined by

$$M = \frac{L}{L_{\text{code}}} \qquad (5.18)$$

Eqn. 5.17 consists of four components. Firstly, $\mathcal{S}(1 - e^{-\alpha L})$ represents the first scattering event, free to occur anywhere along the 15 km fibre delayline. $\mathcal{S}(1 - e^{-\alpha L_c/2})$ calculates the second scattering event, which must occur within the coherence length of a single chip at integer codes upstream. The additional attenuation accrued by the scatter field is accounted for by the term $e^{-\alpha \eta L_{\text{code}}}$. Lastly, the summation is weighted by the factor $(2M - \eta)/2M$. This represents the fraction of second scattering events that fit into the physical length of the fibre as a function of η.

To illustrate this weighting factor, we consider the scenario where the first scattering event occurs within $L_{\text{code}}/2$ from the start of the fibre. In this case, there is not enough back travel length for the second event to occur. Conversely, when the first scatter event occurs near the end of the fibre, sufficient length allows the second reflection to occur a single or many half codes upstream. This number of half codes is determined by η. We can see that smaller η values onset earlier in the fibre coil, and therefore gain a greater proportion in the weight factor. We also note that M here does not need to be an integer, and the weighting factor remains exact for different η values.

Having travelled an additional ηL_{code} distance, DRS field remains at the same code delay as the signal field, and is demodulated to interfere with the reference field. The resultant phase error for this scenario can then be written as

$$\Delta \phi_{\text{err2}}(t) = \frac{\sqrt{P_{\text{sig}} \langle \gamma_2{}^2 \rangle P_{\text{ref}}}}{\sqrt{P_{\text{ref}} P_{\text{sig}}}} \sin \theta(t) = \sqrt{\langle \gamma_2{}^2 \rangle} \sin \theta(t) \qquad (5.19)$$

Again, the spectral characteristics of $\Delta \phi_{\text{err2}}(t)$ is dependent on the DRS phasor $\theta(t)$. We however draw distinction from the first scenario, as the DRS fields here have traversed significantly longer distance, between 1 to 12 extra code lengths with respect to the signal field when using a modulation frequency of $f_{\text{chip}} = 41.7$ MHz and a code length of $N = 511$ elements. This corresponds to time delays of between 12 and 144 microseconds, allowing substantially higher coupling of laser frequency noise. These extensive yet distributed delays effectively provide high optical gain for $\theta(t)$, allowing the cyclic error depicted in Fig. 5.2 to occur even at high Fourier frequencies or with low input frequency noise. The

result of this is broadband DRS noise spread evenly across our measurement spectrum, approximated again by the RMS value of $\sin\theta(t)$:

$$\tilde{S}_{\phi_{\text{err2}}} = \sqrt{\frac{\langle\gamma_2{}^2\rangle}{2}} \tag{5.20}$$

DRS re-aligning with reference delay

Apart from re-synchronising with the signal code delay, DRS fields can also traverse extra distances to re-align with the reference field and, in DEHoI demodulation, coherently interfere with the signal field. The exact calculation for $\langle\gamma_3{}^2\rangle$ depends on the relative delay between the reference and signal fields, which adds a static offset, or attenuation, into Eqn. 5.17. As the level of this offset varies by operation parameters but the fundamental process is similar, we can make the approximation that

$$\langle\gamma_3{}^2\rangle \approx \langle\gamma_2{}^2\rangle \tag{5.21}$$

The interference between these DRS fields and the signal field therefore produces a phase error of

$$\Delta\phi_{\text{err3}}(t) = \frac{\sqrt{P_{\text{sig}}\langle\gamma_3{}^2\rangle P_{\text{sig}}}}{\sqrt{P_{\text{ref}}P_{\text{sig}}}}\sin\theta(t) = \sqrt{e^{-\alpha L}\langle\gamma_3{}^2\rangle}\sin\theta(t) \tag{5.22}$$

And its corresponding spectral density, broadband across the measurement bandwidth, can be estimated as

$$\tilde{S}_{\phi_{\text{err3}}} = \sqrt{\frac{e^{-\alpha L}\langle\gamma_3{}^2\rangle}{2}} \tag{5.23}$$

5.2.4   Code-Incoherent DRS

Finally, the majority of DRS events within our system are code-incoherent. These DRS fields do not synchronise with the code delay of either metrology field, and are therefore suppressed by DEHoI demodulation. To calculate the proportion of code-incoherent DRS, we use the original integration shown in Eqn. 5.8. This is because while the first backscattering event can occur anywhere along the fibre coil, the second event has a range dependency on the first, and therefore the two integrations can no longer be decoupled as we have done previously.

$$\langle\gamma_4{}^2\rangle = \alpha_S{}^2 \mathcal{S}^2 \int_0^L e^{-\alpha z_1} dz_1 \int_0^{z_1} e^{-\alpha z_2} dz_2 \tag{5.24}$$

Strictly speaking, the calculation of $\langle\gamma_4{}^2\rangle$ also includes all code-coherent DRS scenarios as discussed above. The contribution from them however is minimal, as coherence only occurs across two chips per complete code. We can therefore consider the entirety of $\langle\gamma_4{}^2\rangle$ as code-incoherent DRS effects. These DRS fields are free to interact with either

metrology field to produce a code-suppressed phase error. Here we take the average of the two interactions, giving

$$\Delta\phi_{\mathrm{err4}}(t) = \text{Suppression} \times \frac{1}{2}\left(\frac{\sqrt{P_{\mathrm{sig}}\langle\gamma_4{}^2\rangle P_{\mathrm{ref}}}}{\sqrt{P_{\mathrm{ref}}P_{\mathrm{sig}}}} + \frac{\sqrt{P_{\mathrm{sig}}\langle\gamma_4{}^2\rangle P_{\mathrm{sig}}}}{\sqrt{P_{\mathrm{ref}}P_{\mathrm{sig}}}}\right)\sin\theta(t) \qquad (5.25)$$

$$= \text{Suppression} \times \frac{1 + \sqrt{e^{-\alpha L}}}{2}\sqrt{\langle\gamma_4{}^2\rangle}\sin\theta(t) \qquad (5.26)$$

And the resultant phase spectral density, broadened both by extensive travel lengths of DRS fields and by incoherent DEHoI demodulation, is considered spectrally white across our signal band:

$$\tilde{S}_{\phi_{\mathrm{err4}}} = \text{Suppression} \times \frac{1 + \sqrt{e^{-\alpha L}}}{2}\sqrt{\frac{\langle\gamma_4{}^2\rangle}{2}} \qquad (5.27)$$

Determining DEHoI code suppression

As DEHoI modulation uses one PRN code to pseudo-randomly toggle between the I and Q quadratures, data throughput for each quadrature is halved compared to heterodyne operation. This results in compromised DEHoI code correlation, and we simulate and show this effect in Fig. 5.4. Here we examine the normalised demodulated signal amplitude as a function of code delays selected for the reference and signal fields. The code length used here is $N = 511$, allowing a total of $N \times N$ code delay combinations. For clarity, we zoom in around a few specific reference delays and sweep the entirety of signal delays.

We see that the QPSK correlation spectrum is far from a single, clean peak with background suppression of $1/N$, as is the case with heterodyne operation. Instead, the ideal N peak is collapsed across two dimensions, and the otherwise concentrated energy is aliased into non-signal delays, resulting in "ghost correlation" spread across the delay spectrum. These spurious peaks reduce the isolation afforded by DEHoI, with higher peaks producing stronger detrimental effects. Table 5.2 summarises the distribution of these ghost peaks.

Correlation Value	$N/2$	$N/4$	$N/2\sqrt{2}$	1
Number of peaks	$N - 1$	$6N - 18$	4	$N^2 - 7N + 15$

Table 5.2: DEHoI code correlation values listed against the number of correlations at that value.

We note that out of the $N - 1$ maximum height correlations, only one is the desired signal and the rest are artefacts. At those delays, suppression is poorest at 1. Similarly, there are $6N - 18$ delays where background noise is only suppressed by $1/2$. Using Tab. 5.2, we can quantify DEHoI code suppression by the weighted average of these correlation peaks, where the total number of non-signal delays is $N^2 - 1$:

Suppression

$$= 1 \times \frac{N - 2}{N^2 - 1} + \frac{1}{2} \times \frac{6N - 18}{N^2 - 1} + \frac{1}{\sqrt{2}} \times \frac{4}{N^2 - 1} + \frac{2}{N} \times \frac{N^2 - 7N + 15}{N^2 - 1} \qquad (5.28)$$
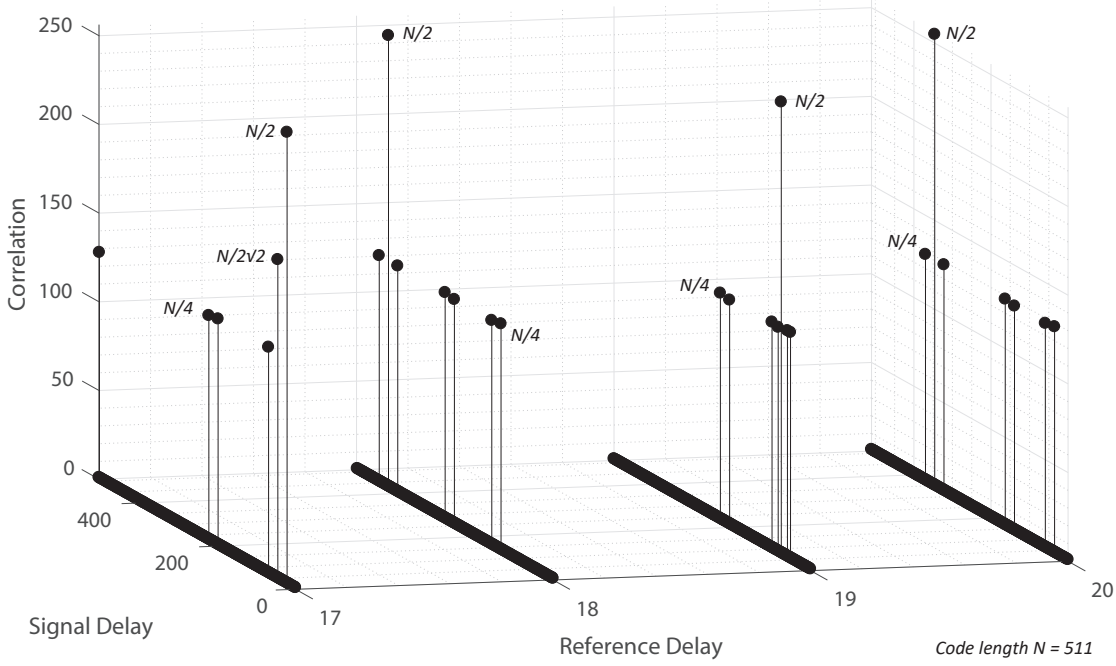
Figure 5.4: An illustration of the QPSK correlation spectrum, simulated using a code length of $N = 511$ elements. The maximum signal-to-noise ratio is halved to $N/2$. This reduced energy is aliased into other non-signal delays, producing further, smaller ghost peaks that degrade the suppression of DEHoI. The overall distribution of the correlation spectrum can be considered as a two-dimensional collapse of the ideal, single correlation peak of N, and is a result of the sequential sampling of each quadrature.

Equation 5.28 represents the actual amount of DEHoI suppression in amplitude for signals outside the selected demodulation delay pair. We assume uniformity of this suppression value across our 15 km delay coils, and apply it to our code-incoherent DRS calculation.

5.2.5   Total Numerical Estimates

Having examined both categories of DRS noise coupling, we can now compute the total DRS noise floor for our system. As each scenario discussed above can be considered an individual, non-correlated process, the total broadband DRS noise is the quadrature summation of $\tilde{S}_{\phi_{\mathrm{err2}}}$, $\tilde{S}_{\phi_{\mathrm{err3}}}$ and $\tilde{S}_{\phi_{\mathrm{err4}}}$:

$$\tilde{S}_{\phi_{\mathrm{broadband}}} = \sqrt{\tilde{S}_{\phi_{\mathrm{err2}}}^2 + \tilde{S}_{\phi_{\mathrm{err3}}}^2 + \tilde{S}_{\phi_{\mathrm{err4}}}^2} \qquad (5.29)$$

The DRS shelf spectral density, onsetting at around 0.1 Hz, is the incoherent summation of the above and $\tilde{S}_{\phi_{\mathrm{err1}}}$:

$$\tilde{S}_{\phi_{\mathrm{shelf}}} = \sqrt{\tilde{S}_{\phi_{\mathrm{err1}}}^2 + \tilde{S}_{\phi_{\mathrm{err2}}}^2 + \tilde{S}_{\phi_{\mathrm{err3}}}^2 + \tilde{S}_{\phi_{\mathrm{err4}}}^2} \qquad (5.30)$$

The above can be further converted into frequency spectral density by multiplying with the corresponding interferometer FSR. We also take into account the incoherent addition

of two interferometers, adding another factor of $\sqrt{2}$ to the total noise spectral density estimate:

$$\tilde{S}_{\nu_{\text{broadband}}} = \sqrt{2}\,\tilde{S}_{\phi_{\text{broadband}}} \times \frac{c}{2\pi nL} \tag{5.31}$$

$$\tilde{S}_{\nu_{\text{shelf}}} = \sqrt{2}\,\tilde{S}_{\phi_{\text{shelf}}} \times \frac{c}{2\pi nL} \tag{5.32}$$

Table 5.3 below lists the numerical evaluation for $\tilde{S}_{\nu_{\text{broadband}}}$ and $\tilde{S}_{\nu_{\text{shelf}}}$ at a few common DEHoI operation settings, with related constants specified in the left column. While the shelf DRS level resides below thermo-mechanical noise for all evaluated DEHoI parameters, broadband DRS imposes a noise floor limitation for our system. At the modulation frequency of 41.7 MHz and code length of 511 elements, same parameters used for the measurement of Fig. 5.1(a), we project a broadband DRS noise floor of 0.12 Hz/$\sqrt{\text{Hz}}$ overlapping with the experimental sensitivity above 70 Hz. The reduction or mitigation of broadband DRS noise is therefore critical to improving the overall system sensitivity, particularly towards the thermo-dynamic limit which currently resides under the DRS noise floor.

Constants	DEHoI Operating Paremeters	$\tilde{S}_{\nu_{\text{broadband}}}$	$\tilde{S}_{\nu_{\text{shelf}}}$
$\mathcal{S} = \frac{3}{8}\frac{\text{NA}^2}{n^2}$	$f_{\text{chip}} = 41.7$ MHz, $N = 511$	0.12 Hz/$\sqrt{\text{Hz}}$	0.13 Hz/$\sqrt{\text{Hz}}$
NA $= 0.14$ $n = 1.46$	$f_{\text{chip}} = 25$ MHz, $N = 511$	0.12 Hz/$\sqrt{\text{Hz}}$	0.14 Hz/$\sqrt{\text{Hz}}$
$\alpha = \alpha_S =$ 0.18 dB/km	$f_{\text{chip}} = 41.7$ MHz, $N = 1023$	0.08 Hz/$\sqrt{\text{Hz}}$	0.10 Hz/$\sqrt{\text{Hz}}$
$L = 15$ km	$f_{\text{chip}} = 25$ MHz, $N = 1023$	0.07 Hz/$\sqrt{\text{Hz}}$	0.10 Hz/$\sqrt{\text{Hz}}$
	$f_{\text{chip}} = 41.7$ MHz, $N = 2047$	0.05 Hz/$\sqrt{\text{Hz}}$	0.07 Hz/$\sqrt{\text{Hz}}$
	$f_{\text{chip}} = 25$ MHz, $N = 2047$	0.04 Hz/$\sqrt{\text{Hz}}$	0.08 Hz/$\sqrt{\text{Hz}}$

Table 5.3: Numerical estimate for both broadband and shelf DRS spectral density at a few DEHoI operating parameters. Constants used for these evaluations are specified in the left column.

5.2.6   Broadband DRS Dependency and Mitigation Discussions

We see from the above analysis that the amount of broadband DRS noise present in the system, $\tilde{S}_{\nu_{\text{broadband}}}$, is dependent on a number of factors: DEHoI modulation frequency f_{chip}, code length N and finally fibre coil length L.

Amongst the three, f_{chip} plays a relatively minor role in determining the DRS noise floor, as confirmed by the comparison between 41.7 MHz and 25 MHz modulation frequencies in Tab. 5.3. This is because the extra physical length occupied by the code afforded by slower modulation is largely offset by the increased coherence length and therefore coupling of DRS at each demodulation code delay.

Alternatively, lengthening the DEHoI code can effectively suppress DRS noise by reducing the number of code wraps within the fibre coil. This however results in a further complication involving DEHoI operation bandwidth and tracking capability. We will discuss this further in the next section, noting for now that tracking challenges increase with lengthened code (N) or lowered code update rate (f_{code}). Depending on severity, the manifestation of tracking difficulties ranges from increased broadband phase noise to frequent cycle slips. Therefore, in optimising interferometer sensitivity, we are forced to maximise f_{chip} and reduce code length until tracking errors and cycle slipping are eliminated. This occurs at $f_{\mathrm{chip}} = 41.7$ MHz and $N = 511$ symbols when using the Koheras E15 and X15 lasers, and at $f_{\mathrm{chip}} = 41.7$ MHz and $N = 1023$ for the Orbits Lightwave laser which has reduced high frequency noise. None of these metrology grade lasers however allows any further lengthening of the DEHoI code (e.g. to $N = 2047$) to suppress DRS without first incurring tracking limitations. Therefore with the 15 km interferometers, the lowest level of broadband DRS resides at 0.07 Hz/$\sqrt{\mathrm{Hz}}$, close to our experimental noise floor.

Finally, we consider the last parameter on which DRS noise floor is dependent, the interferometer armlength difference L. Figure 5.5 plots the estimated broadband DRS noise floor as a function of interferometer length L, for three different code lengths (trace (b)-(d)) at $f_{\mathrm{chip}} = 41.7$ MHz. As the next noise source below DRS at high Fourier frequencies is the thermo-dynamic limit, its estimate at 300 Hz is also plotted in trace (a) as a function of L.

We see that for all three code lengths, broadband DRS noise is dominated by code-coherent effects ($\Delta\phi_{\mathrm{err2}}$ and $\Delta\phi_{\mathrm{err3}}$) with a weak dependence on L. As L increases, the noise level is slightly enhanced with the onset of every additional code wrap, resulting in ripples in trace (b) - (d). When L is shortened to less than half of the physical code length L_{code}, however, code-coherent DRS is eliminated. This results in a sharp drop of total DRS noise down to the background level of code-incoherent effect, crossing the thermo-dynamic limit in the process. As fibre thermo-dynamic noise cannot be mitigated, there is little to be gained from further suppression of DRS beyond that limit. The intersections between DRS and thermo-dynamic (orange dots) therefore represent the theoretical optimal fibre lengths for broadband sensitivity for each code length.

Apparently, our current interferometer length of $L = 15$ km resides high above these predicted optimal lengths. However, this does not immediately conclude that short interferometers are beneficial. For one, Eqn. 5.3 predicts thermo-mechanical noise to scale as $1/\sqrt{L}$. As we are already thermo-mechanical noise limited between 400 mHz - 2 Hz, the stability in this region will further degrade as L decreases. This presents an explicit trade-off between low and high frequency performance, with longer L favouring low frequency stability and shorter L reducing high frequency DRS noise.

Additionally, the aforementioned tracking errors could still occur with a short interferometer and undermine DRS noise reduction. We now explore this in detail and show that, regardless of L, experimental sensitivity at the thermo-dynamic limit (orange dots in Fig. 5.5) may not be achievable in this current architecture.
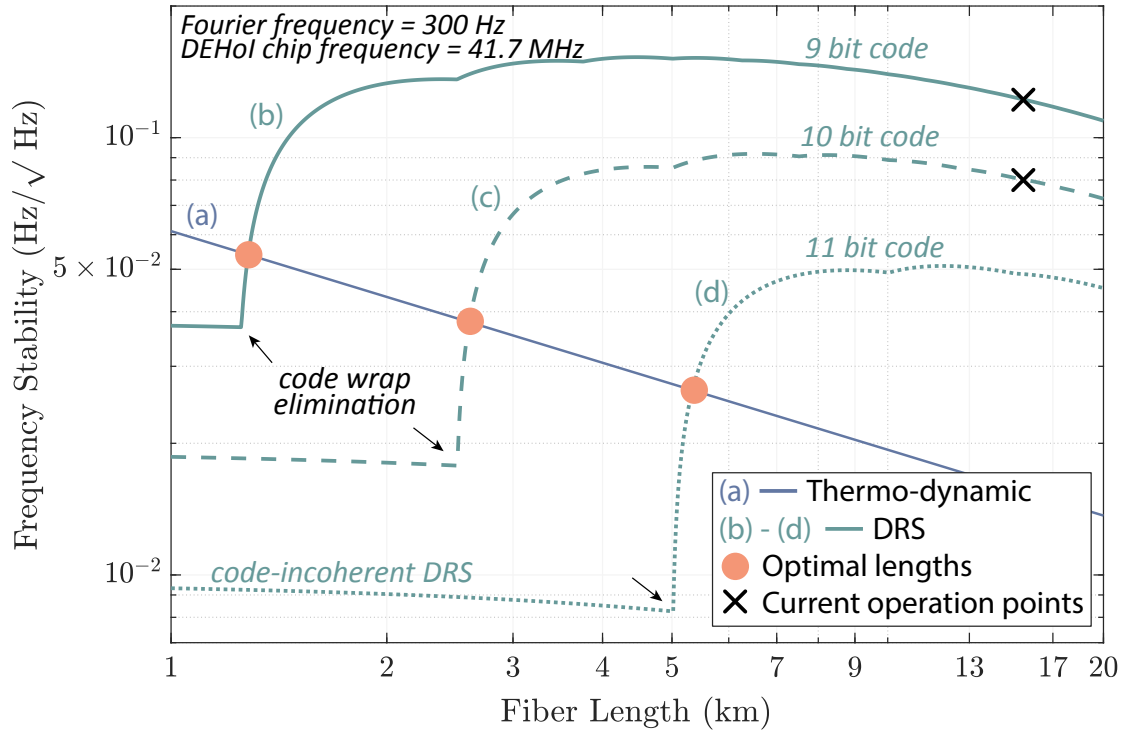
Figure 5.5: Analytical predictions for the optimum interferometer length that achieves best high-frequency sensitivity. Broadband DRS noise estimates at 41.7 MHz DEHoI chip frequency are plotted against fibre length for 9, 10 and 11 bit codes in trace (b) - (d). Modelled level for thermo-dynamic noise at 300 Hz is also plotted over the same range in trace (a). When fibre length reduces to eliminate code-wrapping, DRS noise levels drop sharply and cross over thermo-dynamic noise estimate, revealing an optimum fibre length (orange dots) for each code length. Our current operation points are highlighted in black crosses, at 15 km and 9 bit for the Koheras lasers and 10 bit for Orbits Lightwave.

5.3 DEHoI Phase Tracking Limitations

At the core of DEHoI phase tracking is the code update rate f_{code}, also referred to as the DEHoI readout bandwidth. Dynamics beyond f_{code} are filtered before phase reconstruction could occur, and this results in discrepancy between the DEHoI readout phase and the original interferometric phase. Figure 5.6 conceptually visualises this process. We see that the full interferometer response (blue) is incompletely tracked, with information beyond f_{code} (shaded area) lost. The amount of tracking error arising from this depends on the position of f_{code} relative to the interferometer response, and the input dynamics from the laser.
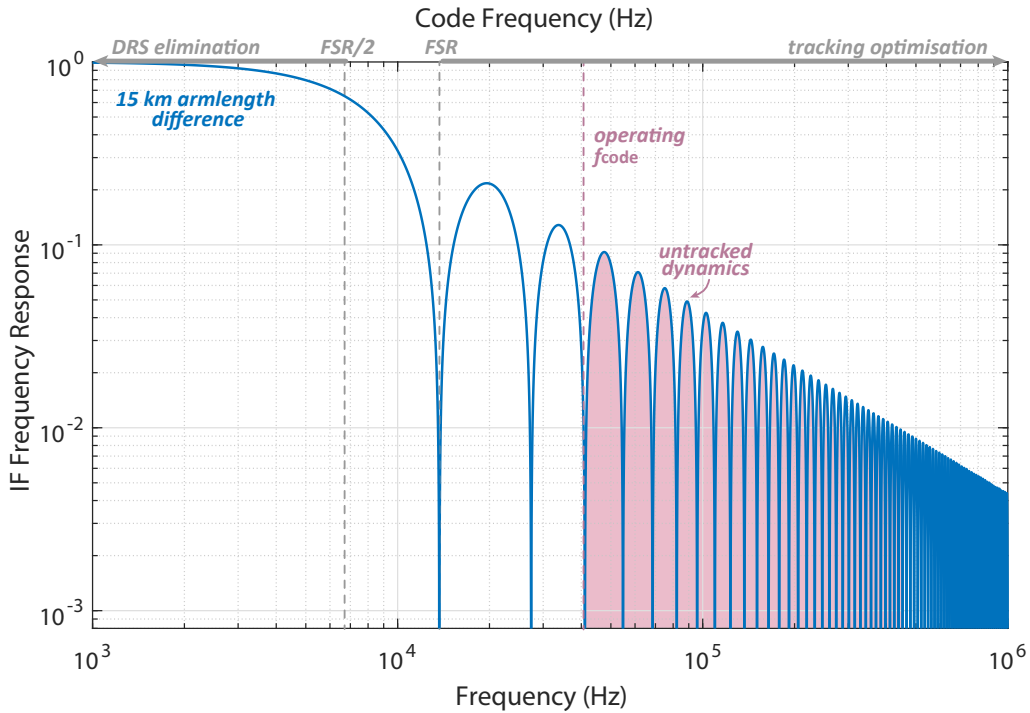
Figure 5.6: Conceptual illustration of tracking error formation. First, input laser frequency noise is transformed into interferometric phase according to the transfer function plotted in blue. This phase signal contains fast dynamics beyond the code refresh rate, which is lost in the code averaging filter and converted into tracking error (shaded red region). Higher code update rate is therefore favourable to reduce tracking noise, and a minimum $f_{\text{code}} = \text{FSR}$ is required in this experiment to prevent cycle slips (top right grey arrow). This forms a trade-off with code-coherent DRS elimination, which requires a maximum $f_{\text{code}} \leq \text{FSR}/2$ (top left grey arrow).

The position of f_{code} on the interferometer transfer function determines the ratio between tracked and untracked dynamics. Intuitively, the resultant relative error decreases as f_{code} is increased and integration time shortened. In our 15 km experiment, we have observed a minimum $f_{\text{code}} \geq \text{FSR}$ to prevent cycle slipping, and a consistent drop of broadband tracking error with increasing f_{code}, particularly at multiples of the FSR ($f_{\text{code}} = m \times \text{FSR}$).

For shorter interferometers, f_{code} must increase to maintain the same proportion of untracked dynamics and consequently the same tracking error. In the meantime, effective DRS suppression requires code wrapping to be eliminated, $f_{\text{code}} \leq \text{FSR}/2$. These two considerations push f_{code} in opposite directions and form an explicit trade-off for the

broadband noise floor, illustrated by the grey arrows in Fig. 5.6. Regardless of L, one noise source always prevails and prevents any further improvement of the experimental sensitivity.

In this respect, the use of DEHoI in a long armlength mismatched interferometer appears to be a double edged sword. Without DEHoI code gating, all DRS events are coherent and interfere with both metrology fields, resulting in a significantly higher overall noise floor. With DEHoI, the integration time required to achieve DRS suppression reduces the ability to reliably track high frequency dynamics within the system, resulting in non-linear readout errors and again broadband noise.

We now consider the second variable that determines the amount of tracking error, input laser frequency noise. Here, we are especially concerned with high-speed laser dynamics beyond f_{code}, as this is the proportion that directly contributes to tracking noise. Low-speed frequency drift on the other hand should not impose any tracking limitation as long as the dynamic range requirement is met. Therefore in selecting the most suitable laser for this experiment, their high-frequency profiles outweigh the low-frequency performance.

To illustrate this, in Fig. 5.7, we provide a comparison between three metrology lasers that were available to this experiment. The solid lines on the left represent experimental measurements of their respective frequency noise, and dashed lines at high frequencies are noise estimates extrapolated from the laser specification where available. We note that the FFR measurements on the left include the combined transfer functions from the interferometer, code integration and downsampling filters. This results in a gradual roll-off towards the interferometer FSR (13.7 kHz), f_{code} (nominally 40.7 kHz or 81.5 kHz) and the sampling Nyquist frequency (4 kHz). The specification traces carry no transfer functions and therefore can appear slightly disconnected with experimental measurements.

Comparing the three lasers, the best long-term stability is achieved by the Koheras X15 laser [203], while the E15 [204] has the highest drift at roughly 50 MHz peak-to-peak over half a day. This equates to a dynamic range requirement of approximately 3650 cycles, which is well within the system design described in Section 3.5.1.

At the high frequency end, both Koheras lasers experience a spectral hump in the kiloHertz regime. The X15 in particular transitions to a white noise floor above ~10 kHz, leading to increased tracking difficulty. In comparison, the Orbits Lightwave laser [205] switches to a $1/f^{-1/2}$ type roll-off at 2 kHz, resulting in substantially lower noise towards and beyond the code frequency. This difference can also be observed experimentally. Using the same chip frequency, we are able to extend the code length by a factor of 2 ($N = 1023$) when using the Orbits laser, whereas for both Koheras lasers, maximum code length remains at $N = 511$ to avoid cycle slips.

Following this thought process, a potential solution to the tracking versus DRS dilemma is to use a non-planar ring oscillator (NPRO) laser [206]. The structural design of NPROs allows them to have a typical frequency noise spectrum of 10 kHz/$\sqrt{\text{Hz}}$/f [207, 208], plotted in the dashed green line in Fig. 5.7. This $1/f$ dependence allows us to trade low-frequency drift for high frequency stability, which is highly beneficial provided dynamic range requirement is met. We can see from Fig. 5.7 that at 1 kHz, the NPRO exceeds
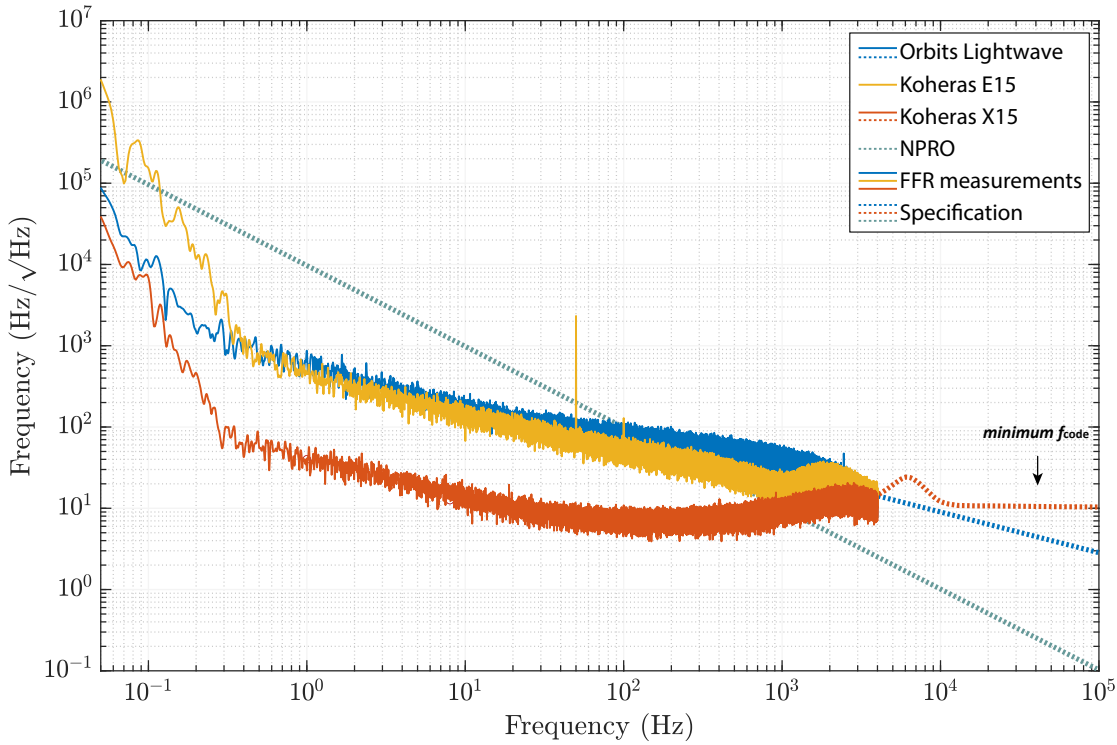
Figure 5.7: Input laser frequency noise beyond the code update rate is the second variable to determine DEHoI tracking noise. In this plot we compare the spectral performance of three metrology lasers: Orbits Lightwave (blue), Koheras E15 (yellow) and Koheras X15 (red). Solid traces represent experimental measurements of laser frequency noise, while dashed lines are extrapolated from specification documents or models. Amongst the metrology lasers, the Orbits exhibits the lowest high frequency noise, in agreement with our experimental observation of lower tracking error. We further compare this with the expected performance of a NPRO laser at 1064 nm, which has the potential to further ease the minimum f_{code} to maintain tracking capability.

the stability of all three metrology lasers tested in this experiment. The main advantage however is seen at even higher frequencies above f_{code}, where the NPRO's performance becomes at least an order of magnitude better. This allows the reduction of untracked laser dynamics, and therefore the reduction of tracking error, by directly reducing the input laser noise. In doing so we remove the need to increase f_{code}, and avoid improving tracking at the cost of worsening DRS.

It is worth noting that the 10 kHz/$\sqrt{\text{Hz}}$/f performance discussed here applies only to NPRO lasers using a Nd:YAG osillator at 1064 nm, and a similar architecture is not available for our current experimental wavelength of 1550 nm. Transitioning to this type of optical source therefore requires major experimental reconfiguration and remains future work.

5.4   Additive Noise

We now consider and calculate additive noise sources in our system, including noise contributions from the analogue front end electronics prior to ADC digitisation, and shot noise and dark noise from the photo-detector. We note here that the use of RT DEHoI readout modifies the noise level of these additive sources, and we quantify this modification via a discriminant calculation. The exact mechanism for their propagation and coupling in a DEHoI system remains under investigation and outside the scope of this thesis.

5.4.1   ADC Front End Noise

ADC front end noise arises from the analogue electronics of the ADC module, as opposed to quantisation noise which occurs during the digitisation process. The amount of ADC front end noise in our phase readout can be quantified by setting up an electronic bypass test, as illustrated in Fig. 5.8. We synthesised a null DEHoI interferometic signal on the FPGA by first delaying the modulation code by the respective optical transit time of the two arms. These two digitally simulated interferometer arms were then differentiated and used to generate the photo-detector output. The resultant signal was output via the DAC and looped back to the ADC and subsequently the demodulation signal processing chain. An analogue electronic attenuator was used to reduce to our nominal ADC input voltage of 800 mV. The output of the demodulation signal processing chain, using a standard RT architecture, was recorded for analysis.
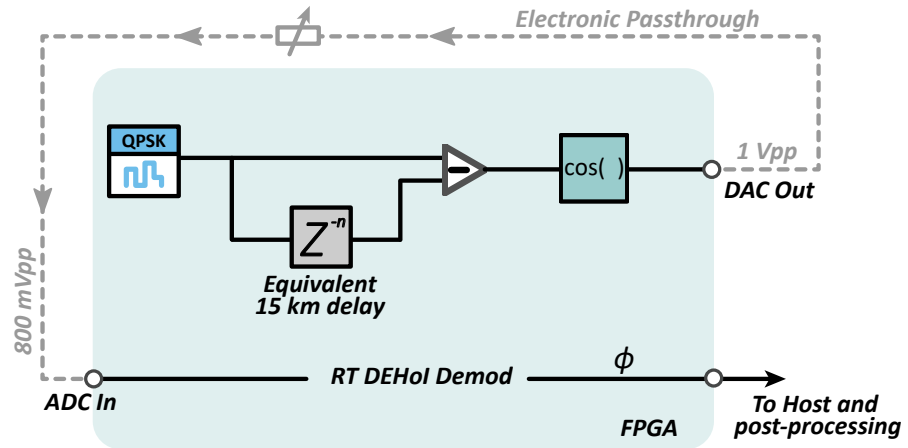
Figure 5.8: Illustration of the ADC noise measurement process. On the FPGA, a QPSK code was delayed by the number of clock cycles corresponding to the interferometer delay. This was subsequently subtracted from a prompt code copy representing the short optical path. The resultant phase difference was wrapped in a cosine function, producing a synthesised DEHoI signal. This signal was routed through the electronic converters of the IO module and attenuated to our nominal input voltage of 800 mV. Once re-digitised, we followed the same demodulation process as used for optical signals. The resultant phase measurement characterises the amount of ADC front end noise in our experiment.

The measurement was conducted at our nominal chip frequency of 41.7 MHz, code length of 1023 elements and decimated sample rate of 8 kHz. The resultant phase spectral density

is shown in Fig. 5.9, showing a flat high-frequency ADC noise floor of 5.5×10^{-7} rad/$\sqrt{\text{Hz}}$. This converts to an equivalent frequency spectral density of 1.2×10^{-3} Hz/$\sqrt{\text{Hz}}$ for a 15 km armlength mismatched interferometer, and 1.7×10^{-3} Hz/$\sqrt{\text{Hz}}$ when considering the incoherent subtraction of two interferometers.
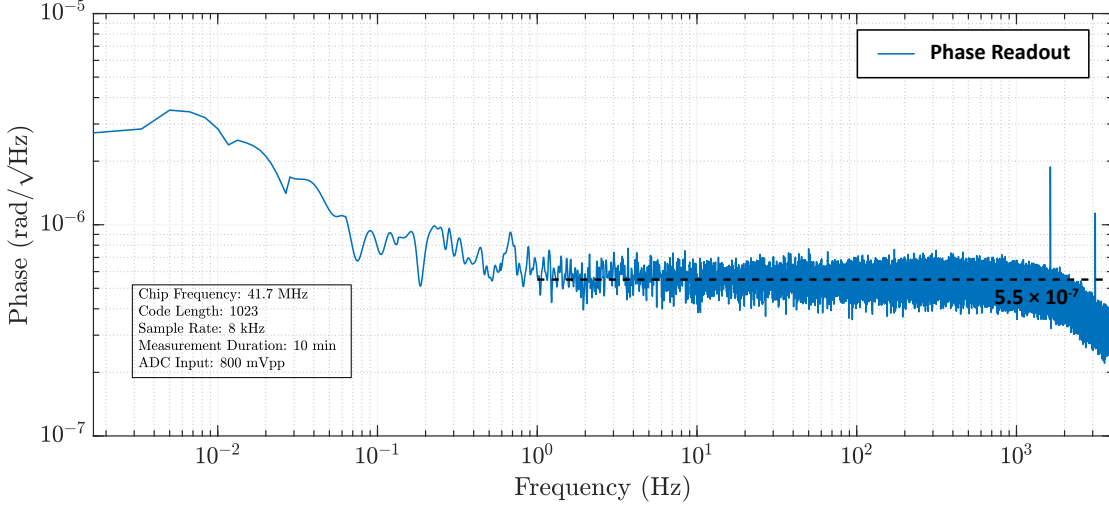
Figure 5.9: Phase spectral density of ADC front end noise, measured through RT DEHoI readout at nominal input voltage and DI parameters. Above 1 Hz, the noise floor resides at 5.5×10^{-7} rad/$\sqrt{\text{Hz}}$.

Additive Noise Discriminant

To determine the coupling coefficient for additive noise, from its raw form into our readout phase noise, we compare the result above with a separate measurement of raw ADC input voltage fluctuation. This measurement was made over a time window of $\sim$0.13 ms at the ADC sample rate of 250 MHz, with the duration limited by the memory buffer size at that measurement speed. The tested ADC channel was terminated, and the recording was converted into voltage spectral density shown in Fig. 5.10.

From Fig. 5.10, we observe a flat noise spectrum of 9 nV/$\sqrt{\text{Hz}}$ above 10 MHz Fourier frequency and a gradual roll-up towards lower frequencies. In RT DEHoI demodulation, both I and Q quadratures receive an equal amount of voltage error directly from the ADC input. However, the low-frequency (below f_{chip}) dynamics of Fig. 5.10 are commonly inherited by the two quadratures and thus become highly suppressed in the phase output. Dynamics faster than f_{chip}, conversely, can be considered uncorrelated and couple into the final phase readout. At our nominal chip frequency of $f_{\text{chip}} = 41.7$ MHz, the raw ADC voltage noise of concern is represented by the shaded area above, at 9 nV/$\sqrt{\text{Hz}}$.

We can now define a discriminant to describe the coupling of raw voltage fluctuation into RT DEHoI phase readout error:

$$\tilde{S}_\phi = \text{Discriminant}\left(f_{\text{chip}}, N\right) \times \frac{\tilde{S}_V \; \left[\text{V}/\sqrt{\text{Hz}}\right]}{A \; [\text{V}]} \quad \left[\text{rad}/\sqrt{\text{Hz}}\right] \tag{5.33}$$
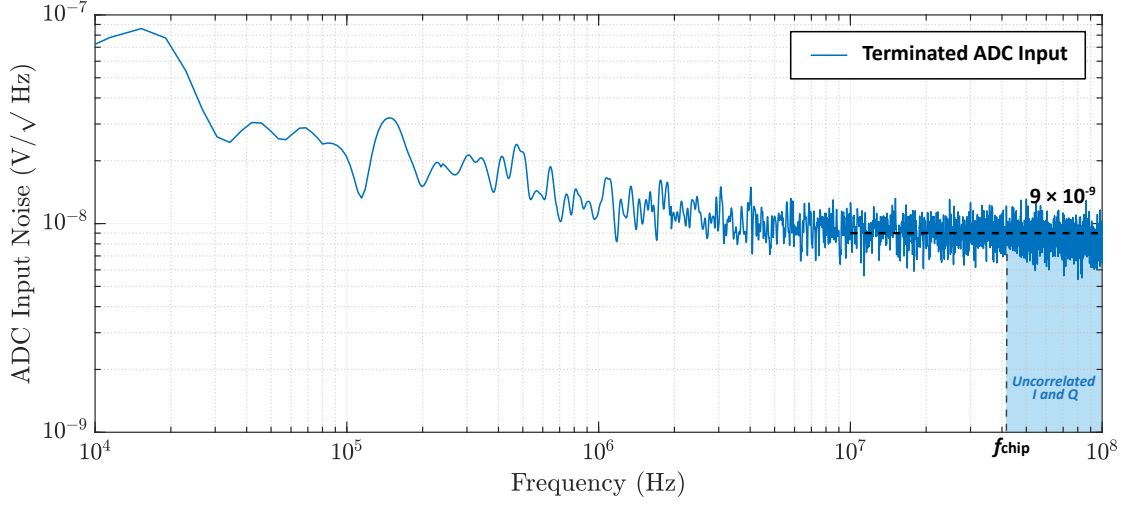
Figure 5.10: Input voltage spectral density of the NI 5782 ADC, measured at 250 MHz with the ADC channel terminated. The shaded area represents uncorrelated additive noise received by I and Q in demodulation. This occurs at Fourier frequencies higher than the DI chip frequency and contributes to phase error in the readout. The raw voltage noise in this frequency regime is measured to be 9 nV/$\sqrt{\text{Hz}}$.

where $\tilde{S}_V$ is the voltage spectral density of the additive noise observed at ADC input, A is the amplitude of the DEHoI signal and is equal to the ADC input voltage V_{in} when only one signal is present (i.e. not multiplexing), and the discriminant is a unitless coefficient dependent on specific DI parameters. Calculating for our experiment where $f_{\text{chip}} = 41.7$ MHz, $N = 1023$, $\tilde{S}_V = 9$ nV/$\sqrt{\text{Hz}}$, $A = V_{\text{in}} = 800$ mV and $\tilde{S}_\phi = 5.5 \times 10^{-7}$ rad/$\sqrt{\text{Hz}}$, the discriminant value is:

$$\text{Discriminant} \, (41.7 \, \text{MHz}, 1023) = 48.89 \tag{5.34}$$

5.4.2   Shot Noise

Also referred to as photon-counting error, shot noise is the dominate quantum-mechanical noise in an interferometer [209]. It arises from uncertainties in the number of photons absorbed by the photo-detector, generating a white, photo-current spectral noise of [210]:

$$\tilde{S}_{I_{\text{shot}}} = \sqrt{2e\rho P_{\text{det}}} \quad \left[\text{A}/\sqrt{\text{Hz}} \right] \tag{5.35}$$

where e is the elementary charge, ρ the photo-detector responsivity ($\rho \approx 1$ A/W at 1550 nm), and P_{det} the total incident optical power. The current spectral density $\tilde{S}_{I_{\text{shot}}}$ can be converted into voltage spectral density via the transimpedance gain T_G of the photo-detector:

$$\tilde{S}_{V_{\text{shot}}} = \tilde{S}_{I_{\text{shot}}} \times T_G \quad \left[\text{V}/\sqrt{\text{Hz}} \right] \tag{5.36}$$

For the Insight BPD-1 photo-detector [211] used in this experiment, its transimpedance gain can be calculated from the RF output conversion gain by:

$$\text{RF Conversion Gain} = T_G \times \rho = 5000 \quad [\text{V/W}] \tag{5.37}$$

Combining the above equations and substituting in $P_{\text{det}} = 60\ \mu\text{W}$, which represents the combined optical power on both ports of the balanced detector, the raw shot noise voltage spectral density incident on the ADC is:

$$\tilde{S}_{V_{\text{shot}}} = 2.19 \times 10^{-8} \quad \left[\text{V}/\sqrt{\text{Hz}}\right] \tag{5.38}$$

Converting this into readout phase noise via the above derived discriminant:

$$\tilde{S}_{\phi_{\text{shot}}} = \text{Discriminant}\,(41.7\ \text{MHz}, 1023) \times \frac{\tilde{S}_{V_{\text{shot}}}}{A} = 1.34 \times 10^{-6} \quad \left[\text{rad}/\sqrt{\text{Hz}}\right] \tag{5.39}$$

Equation 5.39 can then be calibrated into frequency spectral density, with an additional factor of $\sqrt{2}$ representing the incoherent subtraction of two interferometers. This gives the final shot noise estimate for our system:

$$\tilde{S}_{\nu_{\text{shot}}} = \frac{\sqrt{2}\,c}{2\pi n L}\,\tilde{S}_{\phi_{\text{shot}}} = 4.1 \times 10^{-3} \quad \left[\text{Hz}/\sqrt{\text{Hz}}\right] \tag{5.40}$$

5.4.3   Dark Noise

Dark noise is another broadband noise source relating to the output voltage fluctuation of the photo-detector. Unlike shot noise, its value is independent of incident optical power and present in the absence of illumination. Manufacturers typically characterise this noise source by the amount of optical power required to reach unity with the dark noise floor, the noise equivalent power (NEP). The Insight detector's NEP in low-gain mode is specified as [211]:

$$\text{NEP} = 5 \times 10^{-12} \quad \left[\text{W}/\sqrt{\text{Hz}}\right] \tag{5.41}$$

which is converted into voltage spectral density via the RF conversion gain:

$$\tilde{S}_{V_{\text{dark}}} = \text{NEP} \times \text{RF Conversion Gain} = 2.5 \times 10^{-8} \quad \left[\text{V}/\sqrt{\text{Hz}}\right] \tag{5.42}$$

Using the same DI discriminant, the resultant phase error spectral density is:

$$\tilde{S}_{\phi_{\text{dark}}} = \text{Discriminant}\,(41.7\ \text{MHz}, 1023) \times \frac{\tilde{S}_{V_{\text{dark}}}}{A} = 1.53 \times 10^{-6} \quad \left[\text{rad}/\sqrt{\text{Hz}}\right] \tag{5.43}$$

which is similarly calibrated into two-interferometer frequency noise floor:

$$\tilde{S}_{\nu_{\text{dark}}} = \frac{\sqrt{2}\,c}{2\pi n L}\,\tilde{S}_{\phi_{\text{dark}}} = 4.7 \times 10^{-3} \quad \left[\text{Hz}/\sqrt{\text{Hz}}\right] \tag{5.44}$$

5.5   Total System Noise Budget

Combining the modelling and measurements of previous sections, we can now summarise the total system noise budget below in Fig. 5.11. We see that the experimental noise floor (trace a) is limited by DRS noise (d) above 70 Hz and thermo-mechanical noise (b) between 0.4 - 2 Hz. This, in combination with temperature and mechanical investigations discussed in previous chapters, allows complete identification of noise limitations in all Fourier regimes.
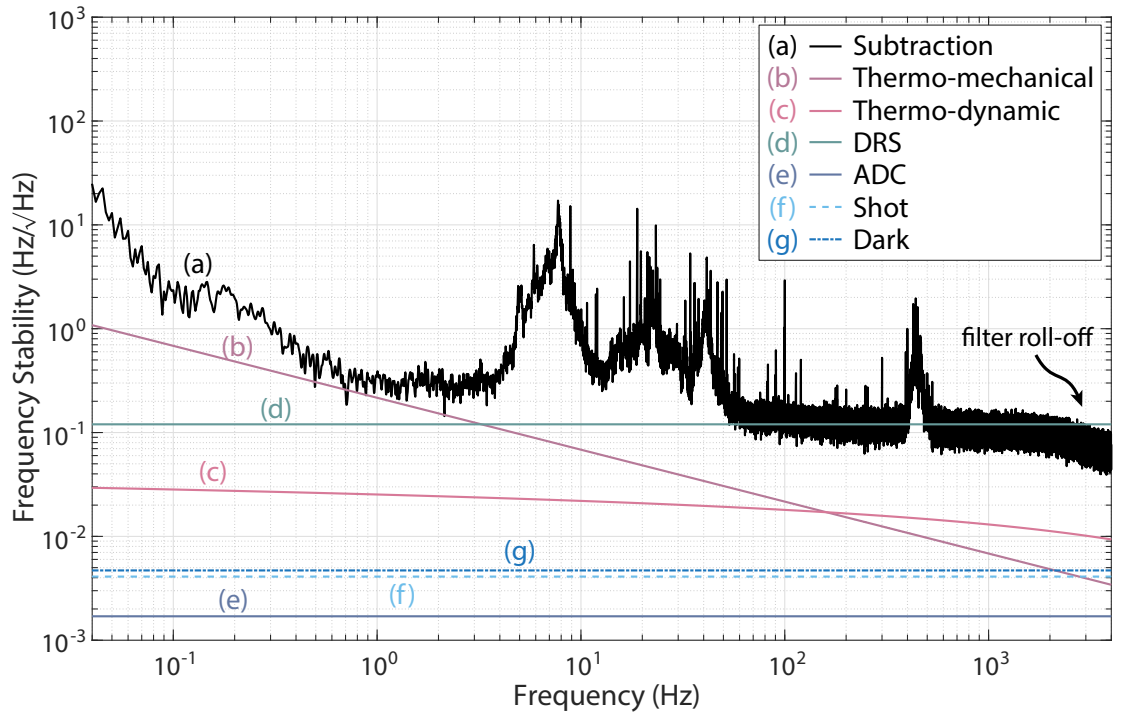
Figure 5.11: Total noise budget of our experiment. (a) experimental noise floor given by the FSR subtraction of the two interferometers. (b) modelled thermo-mechanical noise, which limits the experimental sensitivity over the 400 mHz - 2 Hz Fourier frequency band. (c) thermo-dynamic noise intercepts thermo-mechanical noise at approximately 150 Hz and is not a current limitation. (d) modelled DRS noise level, currently limiting the experiment at all frequencies beyond 70 Hz. (e) measured frequency equivalent noise from the front-end electronics of the ADC module. (f) estimated photo-detector shot noise in the current system, using RT DEHoI readout. (g) estimated photo-detector dark noise for the current system using the same readout parameters.

Although not a current limitation, the modelling and evaluation of thermo-dynamic noise (trace c) provides an important tool for the design and optimisation of future systems. Specifically, the relationship between thermo-dynamic and DRS noise will inform the optimal fibre length difference, while the numerical value of thermo-dynamic noise will set the ultimate limit for broadband noise floor and the trade-off optimisation.

Lastly, additive noise sources (trace e - g) form the lower end of the noise budget and do not currently impose any limitation. However, as these noise sources scale proportionally to $1/L$, as opposed to the $1/\sqrt{L}$ dependence of thermo-dynamic noise, their collective contribution could surpass that of thermo-dynamic noise in a system with shorter coil length.

Conclusion and Outlook

In this thesis, we have developed a passive, all optical fibre frequency reference. Using two near-identical interferometers, we achieved a state-of-the-art relative frequency stability of 0.1 Hz/$\sqrt{\text{Hz}}$ above 70 Hz Fourier frequency, limited by broadband DRS noise. In other Fourier regimes, the system performance was limited by mechanical resonances between 4 - 70 Hz, thermo-mechanical noise between 0.4 - 2 Hz, and temperature fluctuations below 0.4 Hz, all of which were modelled or measured in previous chapters. The experimental characterisation was further completed by an examination of other non-limiting noise sources, including polarisation drift, residual laser frequency noise, thermo-dynamic noise and additive noise.

6.1 Summary of Isolation Chambers and System Mechanics

The thermal stability of the isolation chambers were characterised in Chapter 4.3. Using a circuit analogy, we computed the thermal time constant of these chambers and experimentally verified it with a three-month temperature measurement. From this measurement, we found a 14% difference between the time constant of the two chambers. The difference was in agreement with our previous finding of high temperature independence between the two interferometers, which was characterised via a three-reference measurement including an optical frequency comb in Chapter 4.2.

In Chapter 4.4, we measured the seismic profile of the laboratory building and the optical table. Most mechanical resonances seen in our experimental noise floor were attributable to the building and table, with the exception of one higher-frequency feature identified as an interferometer resonance. To surpass these mechanical limitations, a future system will need to incorporate vibration control strategies. One simple, proven solution is to curate the entire fibre interferometer in an epoxy resin of similar thermal property [212, 213].

6.2 Fibre Thermal Noise Summary

In Chapter 5.1, we applied Duan's thermal noise model [214, 191, 11] to our fibre frequency reference, and demonstrated a thermo-mechanical noise limited sensitivity between 0.4 - 2 Hz. This was the first experimental observation of fibre thermal noise in a long armlength

interferometer which, in combination with observations at shorter lengths [11, 111, 190, 192, 215], formed complete verification for Duan's theory.

Building on Duan's work, we also numerically evaluated the thermo-dynamic noise level for our system, which was the second dominating broadband noise source below DRS. The thermo-dynamic model, along with that of DRS, were then combined to inform the optimal interferometer length to achieve the highest broadband stability.

6.3   DRS Summary

In Chapter 5.2, we developed a comprehensive model to predict noise contributions from Rayleigh scattering in a digitally-enhanced, transmissive fibre interferometer. Numerical evaluation of this model produced a broadband DRS noise estimate of 0.1 Hz/$\sqrt{\text{Hz}}$ at our DEHoI parameters, in agreement with the experimentally observed noise floor above 70 Hz. Although we focussed on characterising DRS with DEHoI implementation, this model is also applicable to general fibre interferometry by modifying the scattering length of consideration.

The range gating property of DEHoI in this experiment provided significant DRS suppression, which would otherwise contribute a broadband noise floor on the order of 2 Hz/$\sqrt{\text{Hz}}$ for our interferometer length [1]. However, the averaging time required to achieve this suppression resulted in limitations in the readout speed and consequently the ability to track fast laser dynamics. This formed an explicit broadband trade-off between DRS and DEHoI tracking noise, discussed in Section 5.3. For a future fibre reference, this trade-off is expected to remain, and optimisation of high-frequency performance will involve navigating these two parameters. The DRS model developed in this work will therefore contribute half of the solution, while more work is required to quantitatively estimate tracking noise.

6.4   Towards a Future FFR - the 1 µm Transition

The work in this thesis has demonstrated a fibre frequency reference with state-of-the-art sensitivity and investigated all limiting noise sources. With these technologies and scientific insight, we believe the system is now ready to transition to the 1064 nm optical wavelength and align with current and future space applications [216, 217, 218, 219]. The popularity of 1064 nm is explained by the availability of NPRO lasers using a Nd:YAG oscillator. These offer superior stability at high Fourier frequencies with a $1/f$ type noise roll-off. For us, a NPRO laser will emit less frequency fluctuation at high Fourier frequencies, potentially easing the tracking challenge and its trade-off with DRS.

We now consider system noise level changes as a result of the wavelength transition. First, DRS noise level is dependent on the scattering coefficient [220] and propagation loss [221], both of which increase at 1064 nm. Although a higher scattering coefficient would appear to induce more noise, its effect is in fact offset by the higher attenuation. As we propagate

---

[1]Calculated from $\tilde{S}_{\phi_{\text{err4}}}$ in Eqn. 5.27 without the suppression factor.

these new coefficients through our DRS model, we arrive at an estimated noise floor of $0.05 \, \text{Hz}/\sqrt{\text{Hz}}$ for a 15 km interferometer operating at the same DEHoI parameters ($f_{\text{chip}} = 41.7$ MHz and $N = 1023$).

Second, both types of thermal noise are inversely proportional to the optical wavelength, and will therefore increase by a factor of $\sim$1.46 at 1064 nm. This will likely broaden the thermo-mechanical limited region, with the exact range dependent on temperature and mechanical stability. Further, thermo-dynamic noise will rise closer to the DRS limit, at $0.027 \, \text{Hz}/\sqrt{\text{Hz}}$ at 200 Hz, potentially offering a probe opportunity if DRS noise could be lowered.

The forecast noise budget for a future FFR is summarised below in Tab. 6.1. We estimate the system to have an increased frequency band to characterise thermo-mechanical noise, up to 40 Hz. The broadband improvement will come from lowered DRS noise at 1064 nm, while improved laser stability will reduce DEHoI tracking limitation and potentially allow probing of thermo-dynamic noise.

	Noise Source	Value ($\text{Hz}/\sqrt{\text{Hz}}$)
$L = 15$ km	Thermo-mechanical	$0.32/\sqrt{f}$
$\lambda = 1064$ nm	DRS/tracking optimisation	0.05
	Cross-over frequency	40 Hz
	Thermo-dynamic	$0.027 \, \text{Hz}/\sqrt{\text{Hz}}$ at 200 Hz

Table 6.1: Estimated noise budget for a future FFR system operating at 1064 nm with an armlength difference of 15 km.

Oversampled Signal Generation

To utilise the full DAC output rate of 1 GSps, we use oversampled signal generation on the FPGA, and present eight phase-offset samples every clock cycle. This allows the FPGA program to run at a slower rate of 125 MHz for greater timing and resource efficiency. The block diagram of our signal generation is shown in Fig. A.1. On the top left, a phase accumulator (shaded blue) steps the phase value each clock cycle. The accumulator is set to wrap upon data overflow. The input value to this accumulator can be calibrated to the generated signal frequency (f_{sig}) via the following equation:

$$f_{\text{sig}} = \text{Phase Step Value} \times f_{\text{clock}} \times 2^{-n} \tag{A.1}$$

where n is the bit depth of the accumulator, and f_{clock} is the FPGA clock frequency.
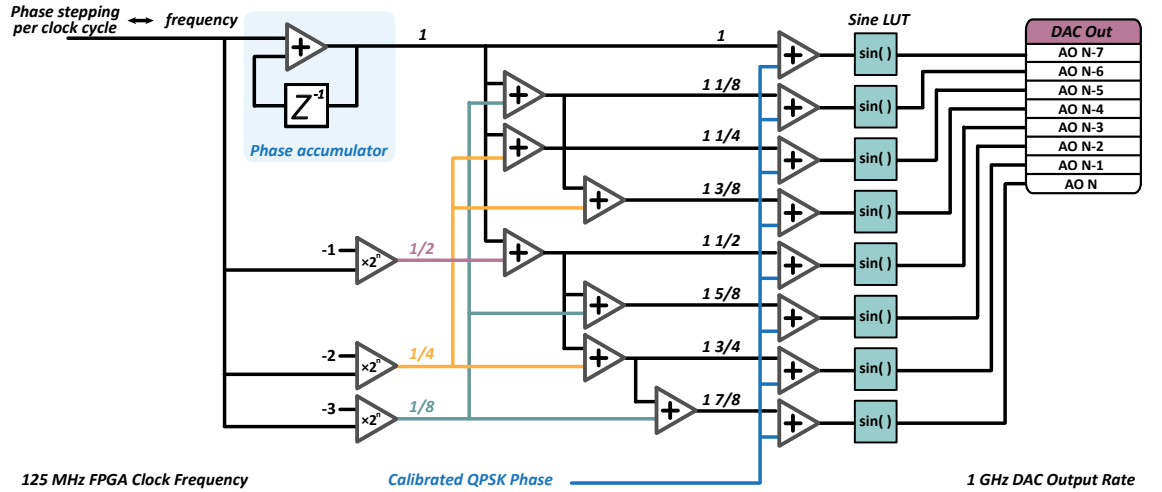
Figure A.1: Generation of the DAC signal for this experiment. A phase accumulator (shaded blue) creates a linear phase ramp, wrapping over the bit depth of the ramp which corresponds to a modulo 2π (1 cycle) phase. Eight copies of the output ramp are generated, each carrying 1/8 extra phase advance than the previous. These samples are then modulated by a calibrated QPSK signal and sent into their respective LUTs. The final data points are presented to the DAC and output sequentially at 8 times the FPGA clock frequency.

The oversampling is achieved by tapping off half (purple), quarter (yellow) and eighth (green) of the input phase step and adding them to the accumulated output. Each row is offset from the previous one by 1/8 of a cycle, labelled as 1 to 1 7/8 in Fig. A.1. We then add in QPSK phase modulation (blue), which is scaled to the input bit depth

of the ensuing look-up tables (LUTs) such that the spacing between each QPSK level corresponds to $\pi/2$. These modulated, oversampled phase values are then fed into their respective LUTs and passed onto the DAC.

As the AOM we use in this experiment has a rise time of 10 ns corresponding to a maximum modulation frequency of 100 MHz [222], there is no improvement from oversampling the QPSK generation. We therefore use the same QPSK value for all 8 samples generated per clock cycle. However, when moving towards an EOM based system with faster DI modulation, oversampled codes will become beneficial. For m-sequences, this can be done by performing XOR operation on the current tap combination as well as N future combinations, where N is the oversampling factor. Following this, the LFSR is advanced by N bits and prepared for the next clock cycle with N new outputs. Oversampled m-sequences have been demonstrated in several DI applications requiring high ranging precision [223, 224].

NI 5782R Configurations

This appendix documents two specific CLIP configurations for the NI 5782R adapter module, clock synchronisation and single DAC operation. These settings are achieved through an initialisation sequence on the Host program, shown in Fig. B.1. Upon opening a reference to the FPGA program, the Host performs standard status checks and waits till command line is idle, before moving on to the CLIP configurations.
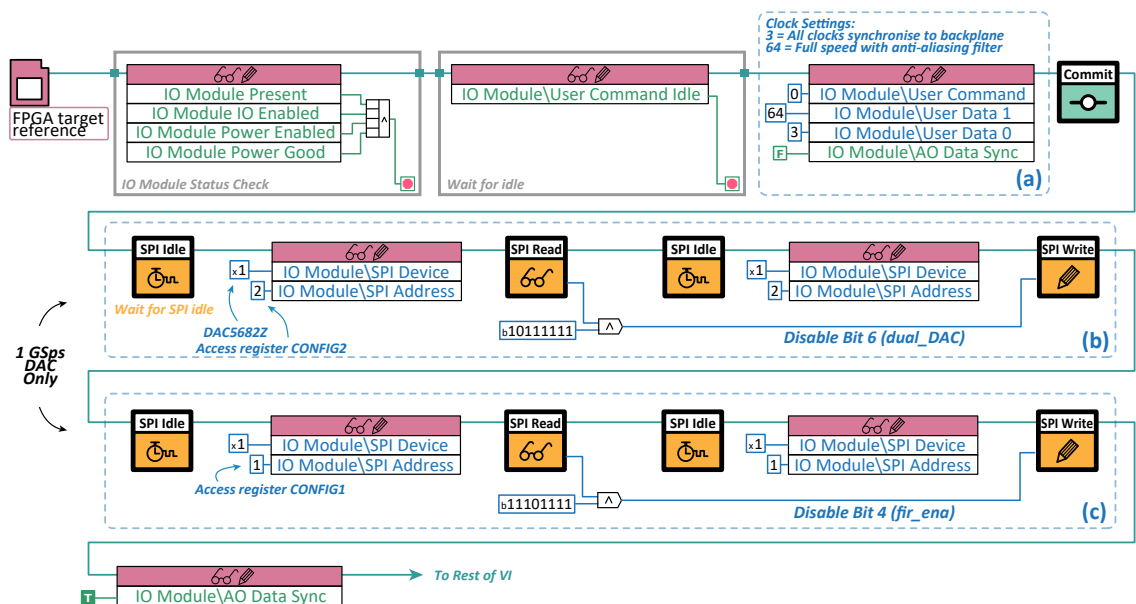
Figure B.1: Initialisation protocol for the NI 5782R adapter module. After initial status checks, the system proceeds to define (a) clock synchronisation options and DAC behaviour. For single DAC operation maximising available bandwidth, two additional configurations are required: (b) disable output from the second DAC; (c) disable the interpolating FIR filter.

B.1   Clock synchronisation

In the clock setting block of Fig. B.1(a), we load option "3" to "User Data 0" [225]. This allows the internal voltage-controlled oscillator (VCO) of the IO module to lock to the PXI_CLK10 10 MHz system reference clock, distributed across the PXIe-1092 chassis and

accessed through the backplane of the PXIe-7975R FPGA [157, 226, 227]. This clock is subsequently distributed to the ADC and DAC parts of the IO module.

The reason for using the chassis synchronisation clock, instead of the internal reference of the IO module, is twofold. First, the chassis clock is stabilised to an OCXO (Oven Controlled Crystal Oscillator) with higher stability, at ±80 ppb compared to ±1 ppm of the TCXO (Temperature Compensated Crystal Oscillator) onboard the IO module [157, 226]. Second, for simultaneous operation with multiple transceiver cards, as demonstrated in Chapter 4.2, the chassis clock becomes a universal reference to which the full system is synchronised.

B.2   1 GSps DAC

To maximise DAC bandwidth, most of this experiment uses single DAC operation at 1 GSps output rate, with polarisation readout measurements being the only exception as they required both independent DAC outputs. The additional bandwidth contains more spectral information of DI modulation, producing cleaner QPSK codes with sharper transitions and reduced oscillation. This also reduces the difference between the analogue-filtered modulation code and the ideal digital demodulation code, and leads to lowered demodulation error and subsequent noise floor.

The NI 5782R contains a Serial Peripheral Interface (SPI) engine that interfaces with the DAC, allowing the FPGA access to all registers that define the DAC's operating modes. As the DAC defaults to dual channel mode with $2\times$ interpolation, we change the two corresponding taps, *dual_DAC* and *fir_ena* as per the register map given in [228]. This process is illustrated in Fig. B.1(b) and (c). Upon selecting the DAC device number and interested register address, the target tap is disabled via an AND operation with that bit set to false. The change is then pushed using the *SPI_Write* command.

For dual DAC operation, block (b) and (c) are removed from the initialisation sequence, and DAC behaviour is directly set by the "User Data 1" command shown in block (a). Here, option "64" splits the full 1 GSps throughput rate across two DAC channels and enables anti-aliasing filters to reduce imaging around the Nyquist frequency [225].

Bibliography

[1] Craig Cahillane, Georgia L. Mansell, Georgia L. Mansell, and Daniel Sigg. Laser frequency noise in next generation gravitational-wave detectors. *Optics Express*, 29(25):42144–42161, December 2021.

[2] Jacob Chamoun. *A Laser-driven Fiber Optic Gyroscope For Inertial Navigation Of Aircraft*. PhD thesis, Stanford University, December 2016.

[3] Benjamin Biedermann. *Menlo Systems ORS1500 Optical Reference System: Design and Performance*, July 2013.

[4] Y. Shevy, D. Shevy, R. Lee, and D. Provenzano. Slow light laser oscillator. In *2010 Conference on Optical Fiber Communication (OFC/NFOEC), Collocated National Fiber Optic Engineers Conference*, pages 1–3, March 2010.

[5] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye. A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity. *Nature Photonics*, 6(10):687–692, October 2012.

[6] J. Millo, D. V. Magalhães, C. Mandache, Y. Le Coq, E. M. L. English, P. G. Westergaard, J. Lodewyck, S. Bize, P. Lemonde, and G. Santarelli. Ultrastable lasers based on vibration insensitive cavities. *Physical Review A*, 79(5):053829, May 2009.

[7] Y. Y. Jiang, A. D. Ludlow, N. D. Lemke, R. W. Fox, J. A. Sherman, L.-S. Ma, and C. W. Oates. Making optical atomic clocks more stable with 10^{-16}-level laser stabilization. *Nature Photonics*, 5(3):158–161, March 2011.

[8] NKT Photonics. *Laser Spectral Linewidth*, 2013.

[9] Nicolas Von Bandel, Mikhaël Myara, Mohamed Sellahi, Tahar Souici, Rémi Dardaillon, and Philippe Signoret. Time-dependent laser linewidth: Beat-note digital acquisition and numerical analysis. *Optics Express*, 24(24):27961, November 2016.

[10] Mo Chen, Zhou Meng, Jianfei Wang, and Wei Chen. Ultra-narrow linewidth measurement based on Voigt profile fitting. *Optics Express*, 23(5):6803, March 2015.

[11] Lingze Duan. Thermal Noise-Limited Fiber-Optic Sensing at Infrasonic Frequencies. *IEEE Journal of Quantum Electronics*, 51(2):1–6, February 2015.

[12] D. Halford, J.H. Shoaf, and A.S. Risley. Spectral Density Analysis: Frequency Domain Specification and Measurement of Signal Stability. In *27th Annual Symposium on Frequency Control*, pages 421–431, Cherry Hill , New Jersey, USA, 1973. IEEE.

[13] F.J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1):51–83, January 1978.

[14] D.W. Allan. Statistics of atomic frequency standards. *Proceedings of the IEEE*, 54(2):221–230, 1966.

[15] D.W. Allan and J.A. Barnes. A Modified "Allan Variance" with Increased Oscillator Characterization Ability. In *Thirty Fifth Annual Frequency Control Symposium*, pages 470–475, May 1981.

[16] W J Riley. Handbook of Frequency Stability Analysis. *National Institute of Standards and Technology Special Publication 1065*, page 136, July 2008.

[17] Newport. *Types of Lasers*.

[18] Stuart D. Jackson. Towards high-power mid-infrared emission from a fibre laser. *Nature Photonics*, 6(7):423–431, July 2012. Number: 7 Publisher: Nature Publishing Group.

[19] Michalis N. Zervas and Christophe A. Codemard. High Power Fiber Lasers: A Review. *IEEE Journal of Selected Topics in Quantum Electronics*, 20(5):219–241, September 2014.

[20] *The Worldwide Market for Lasers: Trends and Five-Year Forecast (2019 – 2025)*.

[21] James T. Spollard, Lyle E. Roberts, Callum S. Sambridge, Kirk McKenzie, and Daniel A. Shaddock. Mitigation of phase noise and Doppler-induced frequency offsets in coherent random amplitude modulated continuous-wave LiDAR. *Optics Express*, 29(6):9060–9083, March 2021.

[22] Nadav Arbel, Maya Yanko, and Avishay Eyal. The sparsity of the response of a quasi-distributed fiber optic sensing system allows 'overclocking' its interrogation. *Optical Fiber Technology*, 61:102373, January 2021.

[23] Davide Grassani, Eirini Tagkoudi, Hairun Guo, Clemens Herkommer, Fan Yang, Tobias J. Kippenberg, and Camille-Sophie Brès. Mid infrared gas spectroscopy using efficient fiber laser driven photonic chip-based supercontinuum. *Nature Communications*, 10(1):1553, April 2019.

[24] A. L. Schawlow and C. H. Townes. Infrared and Optical Masers. *Physical Review*, 112(6):1940–1949, December 1958.

[25] Minh A. Tran, Duanni Huang, and John E. Bowers. Tutorial on narrow linewidth tunable semiconductor lasers using Si/III-V heterogeneous integration. *APL Photonics*, 4(11):111101, November 2019.

[26] NKT Photonics. *Koheras Single-Frequency Fiber Lasers*.

[27] iXblue. *IXC-CLFO @ 1.5 μm, Single-Frequency Fiber Laser*.

[28] Orbits Lightwave, Inc. *Ethernal SlowLight Lasers*.

[29] M. Bailes, B. K. Berger, P. R. Brady, M. Branchesi, K. Danzmann, M. Evans, K. Holley-Bockelmann, B. R. Iyer, T. Kajita, S. Katsanevas, M. Kramer, A. Lazzarini, L. Lehner, G. Losurdo, H. Lück, D. E. McClelland, M. A. McLaughlin, M. Punturo, S. Ransom, S. Raychaudhury, D. H. Reitze, F. Ricci, S. Rowan, Y. Saito, G. H. Sanders, B. S. Sathyaprakash, B. F. Schutz, A. Sesana, H. Shinkai, X. Siemens, D. H. Shoemaker, J. Thorpe, J. F. J. van den Brand, and S. Vitale. Gravitational-wave physics and astronomy in the 2020s and 2030s. *Nature Reviews Physics*, 3(5):344–366, May 2021.

[30] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. Adya, C. Affeldt, N. Aggarwal, O. D. Aguiar, A. Ain, P. Ajith, A. Alemic, B. Allen, D. Amariutei, S. B. Anderson, W. G. Anderson, K. Arai, M. C. Araya, C. Arceneaux, J. S. Areeda, G. Ashton, S. Ast, S. M. Aston, P. Aufmuth, C. Aulbert, B. E. Aylott, S. Babak, P. T. Baker, S. W. Ballmer, J. C. Barayoga, M. Barbet, S. Barclay, B. C. Barish, D. Barker, B. Barr, L. Barsotti, J. Bartlett, M. A. Barton, I. Bartos, R. Bassiri, J. C. Batch, C. Baune, B. Behnke, A. S. Bell, C. Bell, M. Benacquista, J. Bergman, G. Bergmann, C. P. L. Berry, J. Betzwieser, S. Bhagwat, R. Bhandare, I. A. Bilenko, G. Billingsley, J. Birch, S. Biscans, C. Biwer, J. K. Blackburn, L. Blackburn, C. D. Blair, D. Blair, O. Bock, T. P. Bodiya, P. Bojtos, C. Bond, R. Bork, M. Born, Sukanta Bose, P. R. Brady, V. B. Braginsky, J. E. Brau, D. O. Bridges, M. Brinkmann, A. F. Brooks, D. A. Brown, D. D. Brown, N. M. Brown, S. Buchman, A. Buikema, A. Buonanno, L. Cadonati, J. Calderón Bustillo, J. B. Camp, K. C. Cannon, J. Cao, C. D. Capano, S. Caride, S. Caudill, M. Cavaglià, C. Cepeda, R. Chakraborty, T. Chalermsongsak, S. J. Chamberlin, S. Chao, P. Charlton, Y. Chen, H. S. Cho, M. Cho, J. H. Chow, N. Christensen, Q. Chu, S. Chung, G. Ciani, F. Clara, J. A. Clark, C. Collette, L. Cominsky, M. Constancio, D. Cook, T. R. Corbitt, N. Cornish, A. Corsi, C. A. Costa, M. W. Coughlin, S. Countryman, P. Couvares, D. M. Coward, M. J. Cowart, D. C. Coyne, R. Coyne, K. Craig, J. D. E. Creighton, T. D. Creighton, J. Cripe, S. G. Crowder, A. Cumming, L. Cunningham, C. Cutler, K. Dahl, T. Dal Canton, M. Damjanic, S. L. Danilishin, K. Danzmann, L. Dartez, I. Dave, H. Daveloza, G. S. Davies, E. J. Daw, D. DeBra, W. Del Pozzo, T. Denker, T. Dent, V. Dergachev, R. T. DeRosa, R. DeSalvo, S. Dhurandhar, M. D´ıaz, I. Di Palma, G. Dojcinoski, E. Dominguez, F. Donovan, K. L. Dooley, S. Doravari, R. Douglas, T. P. Downes, J. C. Driggers, Z. Du, S. Dwyer, T. Eberle, T. Edo, M. Edwards, M. Edwards, A. Effler, H.-B. Eggenstein, P. Ehrens, J. Eichholz, S. S. Eikenberry, R. Essick, T. Etzel, M. Evans, T. Evans, M. Factourovich, S. Fairhurst, X. Fan, Q. Fang, B. Farr, W. M. Farr, M. Favata, M. Fays, H. Fehrmann, M. M. Fejer, D. Feldbaum, E. C. Ferreira, R. P. Fisher, Z. Frei, A. Freise, R. Frey, T. T. Fricke, P. Fritschel, V. V. Frolov, S. Fuentes-Tapia, P. Fulda, M. Fyffe, J. R. Gair, S. Gaonkar, N. Gehrels, L. Á Gergely´, J. A. Giaime, K. D. Giardina, J. Gleason, E. Goetz, R. Goetz, L. Gondan, G. González, N. Gordon, M. L. Gorodetsky, S. Gossan, S. Goßler, C. Gräf, P. B. Graff, A. Grant, S. Gras, C. Gray, R. J. S. Greenhalgh, A. M. Gretarsson, H. Grote, S. Grunewald, C. J. Guido, X. Guo, K. Gushwa, E. K. Gustafson, R. Gustafson, J. Hacker, E. D. Hall, G. Hammond, M. Hanke, J. Hanks, C. Hanna, M. D. Hannam, J. Hanson, T. Hardwick, G. M. Harry, I. W. Harry, M. Hart, M. T. Hartman, C.-J. Haster, K. Haughian, S. Hee, M. Heintze, G. Heinzel, M. Hendry, I. S. Heng, A. W. Heptonstall, M. Heurs, M. Hewitson, S. Hild, D. Hoak, K. A. Hodge, S. E. Hollitt,

K. Holt, P. Hopkins, D. J. Hosken, J. Hough, E. Houston, E. J. Howell, Y. M. Hu, E. Huerta, B. Hughey, S. Husa, S. H. Huttner, M. Huynh, T. Huynh-Dinh, A. Idrisy, N. Indik, D. R. Ingram, R. Inta, G. Islas, J. C. Isler, T. Isogai, B. R. Iyer, K. Izumi, M. Jacobson, H. Jang, S. Jawahar, Y. Ji, F. Jiménez-Forteza, W. W. Johnson, D. I. Jones, R. Jones, L. Ju, K. Haris, V. Kalogera, S. Kandhasamy, G. Kang, J. B. Kanner, E. Katsavounidis, W. Katzman, H. Kaufer, S. Kaufer, T. Kaur, K. Kawabe, F. Kawazoe, G. M. Keiser, D. Keitel, D. B. Kelley, W. Kells, D. G. Keppel, J. S. Key, A. Khalaidovski, F. Y. Khalili, E. A. Khazanov, C. Kim, K. Kim, N. G. Kim, N. Kim, Y.-M. Kim, E. J. King, P. J. King, D. L. Kinzel, J. S. Kissel, S. Klimenko, J. Kline, S. Koehlenbeck, K. Kokeyama, V. Kondrashov, M. Korobko, W. Z. Korth, D. B. Kozak, V. Kringel, B. Krishnan, C. Krueger, G. Kuehn, A. Kumar, P. Kumar, L. Kuo, M. Landry, B. Lantz, S. Larson, P. D. Lasky, A. Lazzarini, C. Lazzaro, J. Le, P. Leaci, S. Leavey, E. O. Lebigot, C. H. Lee, H. K. Lee, H. M. Lee, J. R. Leong, Y. Levin, B. Levine, J. Lewis, T. G. F. Li, K. Libbrecht, A. Libson, A. C. Lin, T. B. Littenberg, N. A. Lockerbie, V. Lockett, J. Logue, A. L. Lombardi, M. Lormand, J. Lough, M. J. Lubinski, H. Lück, A. P. Lundgren, R. Lynch, Y. Ma, J. Macarthur, T. MacDonald, B. Machenschalk, M. MacInnis, D. M. Macleod, F. Magaña-Sandoval, R. Magee, M. Mageswaran, C. Maglione, K. Mailand, I. Mandel, V. Mandic, V. Mangano, G. L. Mansell, S. Márka, Z. Márka, A. Markosyan, E. Maros, I. W. Martin, R. M. Martin, D. Martynov, J. N. Marx, K. Mason, T. J. Massinger, F. Matichard, L. Matone, N. Mavalvala, N. Mazumder, G. Mazzolo, R. McCarthy, D. E. McClelland, S. McCormick, S. C. McGuire, G. McIntyre, J. McIver, K. McLin, S. McWilliams, G. D. Meadors, M. Meinders, A. Melatos, G. Mendell, R. A. Mercer, S. Meshkov, C. Messenger, P. M. Meyers, H. Miao, H. Middleton, E. E. Mikhailov, A. Miller, J. Miller, M. Millhouse, J. Ming, S. Mirshekari, C. Mishra, S. Mitra, V. P. Mitrofanov, G. Mitselmakher, R. Mittleman, B. Moe, S. D. Mohanty, S. R. P. Mohapatra, B. Moore, D. Moraru, G. Moreno, S. R. Morriss, K. Mossavi, C. M. Mow-Lowry, C. L. Mueller, G. Mueller, S. Mukherjee, A. Mullavey, J. Munch, D. Murphy, P. G. Murray, A. Mytidis, T. Nash, R. K. Nayak, V. Necula, K. Nedkova, G. Newton, T. Nguyen, A. B. Nielsen, S. Nissanke, A. H. Nitz, D. Nolting, M. E. N. Normandin, L. K. Nuttall, E. Ochsner, J. O'Dell, E. Oelker, G. H. Ogin, J. J. Oh, S. H. Oh, F. Ohme, P. Oppermann, R. Oram, B. O'Reilly, W. Ortega, R. O'Shaughnessy, C. Osthelder, C. D. Ott, D. J. Ottaway, R. S. Ottens, H. Overmier, B. J. Owen, C. Padilla, A. Pai, S. Pai, O. Palashov, A. Pal-Singh, H. Pan, C. Pankow, F. Pannarale, B. C. Pant, M. A. Papa, H. Paris, Z. Patrick, M. Pedraza, L. Pekowsky, A. Pele, S. Penn, A. Perreca, M. Phelps, V. Pierro, I. M. Pinto, M. Pitkin, J. Poeld, A. Post, A. Poteomkin, J. Powell, J. Prasad, V. Predoi, S. Premachandra, T. Prestegard, L. R. Price, M. Principe, S. Privitera, R. Prix, L. Prokhorov, O. Puncken, M. Pürrer, J. Qin, V. Quetschke, E. Quintero, G. Quiroga, R. Quitzow-James, F. J. Raab, D. S. Rabeling, H. Radkins, P. Raffai, S. Raja, G. Rajalakshmi, M. Rakhmanov, K. Ramirez, V. Raymond, C. M. Reed, S. Reid, D. H. Reitze, O. Reula, K. Riles, N. A. Robertson, R. Robie, J. G. Rollins, V. Roma, J. D. Romano, G. Romanov, J. H. Romie, S. Rowan, A. Rüdiger, K. Ryan, S. Sachdev, T. Sadecki, L. Sadeghian, M. Saleem, F. Salemi, L. Sammut, V. Sandberg, J. R. Sanders, V. Sannibale, I. Santiago-Prieto, B. S. Sathyaprakash, P. R. Saulson, R. Savage, A. Sawadsky, J. Scheuer, R. Schilling, P. Schmidt, R. Schnabel, R. M. S. Schofield, E. Schreiber, D. Schuette, B. F. Schutz, J. Scott, S. M. Scott, D. Sellers, A. S. Sengupta,

A. Sergeev, G. Serna, A. Sevigny, D. A. Shaddock, M. S. Shahriar, M. Shaltev, Z. Shao, B. Shapiro, P. Shawhan, D. H. Shoemaker, T. L. Sidery, X. Siemens, D. Sigg, A. D. Silva, D. Simakov, A. Singer, L. Singer, R. Singh, A. M. Sintes, B. J. J. Slagmolen, J. R. Smith, M. R. Smith, R. J. E. Smith, N. D. Smith-Lefebvre, E. J. Son, B. Sorazu, T. Souradeep, A. Staley, J. Stebbins, M. Steinke, J. Steinlechner, S. Steinlechner, D. Steinmeyer, B. C. Stephens, S. Steplewski, S. Stevenson, R. Stone, K. A. Strain, S. Strigin, R. Sturani, A. L. Stuver, T. Z. Summerscales, P. J. Sutton, M. Szczepanczyk, G. Szeifert, D. Talukder, D. B. Tanner, M. Tápai, S. P. Tarabrin, A. Taracchini, R. Taylor, G. Tellez, T. Theeg, M. P. Thirugnanasambandam, M. Thomas, P. Thomas, K. A. Thorne, K. S. Thorne, E. Thrane, V. Tiwari, C. Tomlinson, C. V. Torres, C. I. Torrie, G. Traylor, M. Tse, D. Tshilumba, D. Ugolini, C. S. Unnikrishnan, A. L. Urban, S. A. Usman, H. Vahlbruch, G. Vajente, G. Valdes, M. Vallisneri, A. A. van Veggel, S. Vass, R. Vaulin, A. Vecchio, J. Veitch, P. J. Veitch, K. Venkateswara, R. Vincent-Finley, S. Vitale, T. Vo, C. Vorvick, W. D. Vousden, S. P. Vyatchanin, A. R. Wade, L. Wade, M. Wade, M. Walker, L. Wallace, S. Walsh, H. Wang, M. Wang, X. Wang, R. L. Ward, J. Warner, M. Was, B. Weaver, M. Weinert, A. J. Weinstein, R. Weiss, T. Welborn, L. Wen, P. Wessels, T. Westphal, K. Wette, J. T. Whelan, S. E. Whitcomb, D. J. White, B. F. Whiting, C. Wilkinson, L. Williams, R. Williams, A. R. Williamson, J. L. Willis, B. Willke, M. Wimmer, W. Winkler, C. C. Wipf, H. Wittel, G. Woan, J. Worden, S. Xie, J. Yablon, I. Yakushin, W. Yam, H. Yamamoto, C. C. Yancey, Q. Yang, M. Zanolin, Fan Zhang, L. Zhang, M. Zhang, Y. Zhang, C. Zhao, M. Zhou, X. J. Zhu, M. E. Zucker, S. Zuraw, and J. Zweizig. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, March 2015.

[31] D V Martynov, E D Hall, B P Abbott, R Abbott, T D Abbott, C Adams, R X Adhikari, R A Anderson, S B Anderson, K Arai, M A Arain, S M Aston, L Austin, S W Ballmer, M Barbet, D Barker, B Barr, L Barsotti, J Bartlett, M A Barton, I Bartos, J C Batch, A S Bell, I Belopolski, J Bergman, J Betzwieser, G Billingsley, J Birch, S Biscans, C Biwer, E Black, C D Blair, C Bogan, C Bond, R Bork, D O Bridges, A F Brooks, D D Brown, L Carbone, C Celerier, G Ciani, F Clara, D Cook, S T Countryman, M J Cowart, D C Coyne, A Cumming, L Cunningham, M Damjanic, R Dannenberg, K Danzmann, R DeSalvo, K L Dooley, S Doravari, J C Driggers, S E Dwyer, A Effler, T Etzel, M Evans, T M Evans, M Factourovich, H Fair, D Feldbaum, R P Fisher, S Foley, M Frede, A Freise, P Fritschel, V V Frolov, P Fulda, M Fyffe, V Galdi, J A Giaime, K D Giardina, J R Gleason, R Goetz, S Gras, C Gray, R J S Greenhalgh, H Grote, C J Guido, K E Gushwa, E K Gustafson, R Gustafson, G Hammond, J Hanks, J Hanson, T Hardwick, G M Harry, K Haughian, J Heefner, M C Heintze, A W Heptonstall, D Hoak, J Hough, A Ivanov, K Izumi, M Jacobson, E James, R Jones, S Kandhasamy, S Karki, M Kasprzack, S Kaufer, K Kawabe, W Kells, N Kijbunchoo, E J King, P J King, D L Kinzel, J S Kissel, K Kokeyama, W Z Korth, G Kuehn, P Kwee, M Landry, B Lantz, A Le Roux, B M Levine, J B Lewis, V Lhuillier, N A Lockerbie, M Lormand, M J Lubinski, A P Lundgren, T MacDonald, M MacInnis, D M Macleod, M Mageswaran, K Mailand, S Marka, Z Marka, A S Markosyan, E Maros, I W Martin, R M Martin, J N Marx, K Mason, T J Massinger, F Matichard, N Mavalvala, R McCarthy, D E McClelland, S McCormick, G McIntyre, J McIver, E L Merilh, M S Meyer, P M Meyers, J Miller, R Mittleman, G Moreno, C L Mueller, G Mueller, A Mullavey,

J Munch, P G Murray, L K Nuttall, J Oberling, J O'Dell, P Oppermann, Richard J Oram, B O'Reilly, C Osthelder, D J Ottaway, H Overmier, J R Palamos, H R Paris, W Parker, Z Patrick, A Pele, S Penn, M Phelps, M Pickenpack, V Piero, I Pinto, J Poeld, M Principe, L Prokhorov, O Puncken, V Quetschke, E A Quintero, F J Raab, H Radkins, P Raffai, C R Ramet, C M Reed, S Reid, D H Reitze, N A Robertson, J G Rollins, V J Roma, J H Romie, S Rowan, K Ryan, T Sadecki, E J Sanchez, V Sandberg, V Sannibale, R L Savage, R M S Schofield, B Schultz, P Schwinberg, D Sellers, A Sevigny, D A Shaddock, Z Shao, B Shapiro, P Shawhan, D H Shoemaker, D Sigg, B J J Slagmolen, J R Smith, M R Smith, N D Smith-Lefebvre, B Sorazu, A Staley, A J Stein, A Stochino, K A Strain, R Taylor, M Thomas, P Thomas, K A Thorne, E Thrane, K V Tokmakov, C I Torrie, G Traylor, G Vajente, G Valdes, A A van Veggel, M Vargas, A Vecchio, P J Veitch, K Venkateswara, T Vo, C Vorvick, S J Waldman, M Walker, R L Ward, J Warner, B Weaver, R Weiss, T Welborn, P Weßels, C Wilkinson, P A Willems, L Williams, B Willke, I Wilmut, L Winkelmann, C C Wipf, J Worden, G Wu, H Yamamoto, C C Yancey, H Yu, L Zhang, M E Zucker, and J Zweizig. The Sensitivity of the Advanced LIGO Detectors at the Beginning of Gravitational Wave Astronomy. page 20, April 2016.

[32] Rana X. Adhikari. Gravitational Radiation Detection with Laser Interferometry. *Reviews of Modern Physics*, 86(1):121–151, February 2014.

[33] B P Abbott, R Abbott, R Adhikari, P Ajith, B Allen, G Allen, R S Amin, S B Anderson, W G Anderson, M A Arain, M Araya, H Armandula, P Armor, Y Aso, S Aston, P Aufmuth, C Aulbert, S Babak, P Baker, S Ballmer, C Barker, D Barker, B Barr, P Barriga, L Barsotti, M A Barton, I Bartos, R Bassiri, M Bastarrika, B Behnke, M Benacquista, J Betzwieser, P T Beyersdorf, I A Bilenko, G Billingsley, R Biswas, E Black, J K Blackburn, L Blackburn, D Blair, B Bland, T P Bodiya, L Bogue, R Bork, V Boschi, S Bose, P R Brady, V B Braginsky, J E Brau, D O Bridges, M Brinkmann, A F Brooks, D A Brown, A Brummit, G Brunet, A Bullington, A Buonanno, O Burmeister, R L Byer, L Cadonati, J B Camp, J Cannizzo, K C Cannon, J Cao, L Cardenas, S Caride, G Castaldi, S Caudill, M Cavaglià, C Cepeda, T Chalermsongsak, E Chalkley, P Charlton, S Chatterji, S Chelkowski, Y Chen, N Christensen, C T Y Chung, D Clark, J Clark, J H Clayton, T Cokelaer, C N Colacino, R Conte, D Cook, T R C Corbitt, N Cornish, D Coward, D C Coyne, J D E Creighton, T D Creighton, A M Cruise, R M Culter, A Cumming, L Cunningham, S L Danilishin, K Danzmann, B Daudert, G Davies, E J Daw, D DeBra, J Degallaix, V Dergachev, S Desai, R DeSalvo, S Dhurandhar, M Díaz, A Dietz, F Donovan, K L Dooley, E E Doomes, R W P Drever, J Dueck, I Duke, J-C Dumas, J G Dwyer, C Echols, M Edgar, A Effler, P Ehrens, E Espinoza, T Etzel, M Evans, T Evans, S Fairhurst, Y Faltas, Y Fan, D Fazi, H Fehrmenn, L S Finn, K Flasch, S Foley, C Forrest, N Fotopoulos, A Franzen, M Frede, M Frei, Z Frei, A Freise, R Frey, T Fricke, P Fritschel, V V Frolov, M Fyffe, V Galdi, J A Garofoli, I Gholami, J A Giaime, S Giampanis, K D Giardina, K Goda, E Goetz, L M Goggin, G González, M L Gorodetsky, S Goßler, R Gouaty, A Grant, S Gras, C Gray, M Gray, R J S Greenhalgh, A M Gretarsson, F Grimaldi, R Grosso, H Grote, S Grunewald, M Guenther, E K Gustafson, R Gustafson, B Hage, J M Hallam, D Hammer, G D Hammond, C Hanna, J Hanson, J Harms, G M Harry, I W Harry, E D Harstad, K Haughian, K Hayama, J Heefner, I S Heng,

A Heptonstall, M Hewitson, S Hild, E Hirose, D Hoak, K A Hodge, K Holt, D J Hosken, J Hough, D Hoyland, B Hughey, S H Huttner, D R Ingram, T Isogai, M Ito, A Ivanov, B Johnson, W W Johnson, D I Jones, G Jones, R Jones, L Ju, P Kalmus, V Kalogera, S Kandhasamy, J Kanner, D Kasprzyk, E Katsavounidis, K Kawabe, S Kawamura, F Kawazoe, W Kells, D G Keppel, A Khalaidovski, F Y Khalili, R Khan, E Khazanov, P King, J S Kissel, S Klimenko, K Kokeyama, V Kondrashov, R Kopparapu, S Koranda, D Kozak, B Krishnan, R Kumar, P Kwee, P K Lam, M Landry, B Lantz, A Lazzarini, H Lei, M Lei, N Leindecker, I Leonor, C Li, H Lin, P E Lindquist, T B Littenberg, N A Lockerbie, D Lodhia, M Longo, M Lormand, P Lu, M Lubinski, A Lucianetti, H Lück, B Machenschalk, M MacInnis, M Mageswaran, K Mailand, I Mandel, V Mandic, S Márka, Z Márka, A Markosyan, J Markowitz, E Maros, I W Martin, R M Martin, J N Marx, K Mason, F Matichard, L Matone, R A Matzner, N Mavalvala, R McCarthy, D E McClelland, S C McGuire, M McHugh, G McIntyre, D J A McKechan, K McKenzie, M Mehmet, A Melatos, A C Melissinos, D F Menéndez, G Mendell, R A Mercer, S Meshkov, C Messenger, M S Meyer, J Miller, J Minelli, Y Mino, V P Mitrofanov, G Mitselmakher, R Mittleman, O Miyakawa, B Moe, S D Mohanty, S R P Mohapatra, G Moreno, T Morioka, K Mors, K Mossavi, C MowLowry, G Mueller, H Müller-Ebhardt, D Muhammad, S Mukherjee, H Mukhopadhyay, A Mullavey, J Munch, P G Murray, E Myers, J Myers, T Nash, J Nelson, G Newton, A Nishizawa, K Numata, J O'Dell, B O'Reilly, R O'Shaughnessy, E Ochsner, G H Ogin, D J Ottaway, R S Ottens, H Overmier, B J Owen, Y Pan, C Pankow, M A Papa, V Parameshwaraiah, P Patel, M Pedraza, S Penn, A Perraca, V Pierro, I M Pinto, M Pitkin, H J Pletsch, M V Plissi, F Postiglione, M Principe, R Prix, L Prokhorov, O Punken, V Quetschke, F J Raab, D S Rabeling, H Radkins, P Raffai, Z Raics, N Rainer, M Rakhmanov, V Raymond, C M Reed, T Reed, H Rehbein, S Reid, D H Reitze, R Riesen, K Riles, B Rivera, P Roberts, N A Robertson, C Robinson, E L Robinson, S Roddy, C Röver, J Rollins, J D Romano, J H Romie, S Rowan, A Rüdiger, P Russell, K Ryan, S Sakata, L Sancho de la Jordana, V Sandberg, V Sannibale, L Santamaría, S Saraf, P Sarin, B S Sathyaprakash, S Sato, M Satterthwaite, P R Saulson, R Savage, P Savov, M Scanlan, R Schilling, R Schnabel, R Schofield, B Schulz, B F Schutz, P Schwinberg, J Scott, S M Scott, A C Searle, B Sears, F Seifert, D Sellers, A S Sengupta, A Sergeev, B Shapiro, P Shawhan, D H Shoemaker, A Sibley, X Siemens, D Sigg, S Sinha, A M Sintes, B J J Slagmolen, J Slutsky, J R Smith, M R Smith, N D Smith, K Somiya, B Sorazu, A Stein, L C Stein, S Steplewski, A Stochino, R Stone, K A Strain, S Strigin, A Stroeer, A L Stuver, T Z Summerscales, K-X Sun, M Sung, P J Sutton, G P Szokoly, D Talukder, L Tang, D B Tanner, S P Tarabrin, J R Taylor, R Taylor, J Thacker, K A Thorne, A Thüring, K V Tokmakov, C Torres, C Torrie, G Traylor, M Trias, D Ugolini, J Ulmen, K Urbanek, H Vahlbruch, M Vallisneri, C Van Den Broeck, M V van der Sluys, A A van Veggel, S Vass, R Vaulin, A Vecchio, J Veitch, P Veitch, C Veltkamp, A Villar, C Vorvick, S P Vyachanin, S J Waldman, L Wallace, R L Ward, A Weidner, M Weinert, A J Weinstein, R Weiss, L Wen, S Wen, K Wette, J T Whelan, S E Whitcomb, B F Whiting, C Wilkinson, P A Willems, H R Williams, L Williams, B Willke, I Wilmut, L Winkelmann, W Winkler, C C Wipf, A G Wiseman, G Woan, R Wooley, J Worden, W Wu, I Yakushin, H Yamamoto, Z Yan, S Yoshida, M Zanolin, J Zhang, L Zhang, C Zhao, N Zotov, M E Zucker, H zur Mühlen, and J Zweizig. LIGO: The Laser

Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 72(7):076901, July 2009.

[34] Alexa Staley. *Locking the Advanced LIGO Gravitational Wave Detector: With a Focus on the Arm Length Stabilization Technique.* PhD thesis, Columbia University, 2015.

[35] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, K. Ackley, C. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, C. Affeldt, N. Aggarwal, O. D. Aguiar, A. Ain, P. Ajith, B. Allen, P. A. Altin, S. B. Anderson, W. G. Anderson, K. Arai, M. C. Araya, C. C. Arceneaux, J. S. Areeda, K. G. Arun, G. Ashton, M. Ast, S. M. Aston, P. Aufmuth, C. Aulbert, S. Babak, P. T. Baker, S. W. Ballmer, J. C. Barayoga, S. E. Barclay, B. C. Barish, C. P.L. Berry, C. Bond, W. Del Pozzo, W. M. Farr, A. Freise, A. C. Green, C. J. Haster, H. Middleton, C. M. Mow-Lowry, D. Töyrä, A. Vecchio, J. Veitch, S. Vinciguerra, W. D. Vousden, H. Wang, M. Wang, and LIGO Scientific Collaboration. Exploring the sensitivity of next generation gravitational wave detectors. *Classical and Quantum Gravity*, 34(4), January 2017.

[36] Pau Amaro-Seoane, Heather Audley, Stanislav Babak, John Baker, Enrico Barausse, Peter Bender, Emanuele Berti, Pierre Binetruy, Michael Born, Daniele Bortoluzzi, Jordan Camp, Chiara Caprini, Vitor Cardoso, Monica Colpi, John Conklin, Neil Cornish, Curt Cutler, Karsten Danzmann, Rita Dolesi, Luigi Ferraioli, Valerio Ferroni, Ewan Fitzsimons, Jonathan Gair, Lluis Gesa Bote, Domenico Giardini, Ferran Gibert, Catia Grimani, Hubert Halloin, Gerhard Heinzel, Thomas Hertog, Martin Hewitson, Kelly Holley-Bockelmann, Daniel Hollington, Mauro Hueller, Henri Inchauspe, Philippe Jetzer, Nikos Karnesis, Christian Killow, Antoine Klein, Bill Klipstein, Natalia Korsakova, Shane L. Larson, Jeffrey Livas, Ivan Lloro, Nary Man, Davor Mance, Joseph Martino, Ignacio Mateos, Kirk McKenzie, Sean T. McWilliams, Cole Miller, Guido Mueller, Germano Nardini, Gijs Nelemans, Miquel Nofrarias, Antoine Petiteau, Paolo Pivato, Eric Plagnol, Ed Porter, Jens Reiche, David Robertson, Norna Robertson, Elena Rossi, Giuliana Russano, Bernard Schutz, Alberto Sesana, David Shoemaker, Jacob Slutsky, Carlos F. Sopuerta, Tim Sumner, Nicola Tamanini, Ira Thorpe, Michael Troebs, Michele Vallisneri, Alberto Vecchio, Daniele Vetrugno, Stefano Vitale, Marta Volonteri, Gudrun Wanner, Harry Ward, Peter Wass, William Weber, John Ziemer, and Peter Zweifel. Laser Interferometer Space Antenna. *arXiv:1702.00786 [astro-ph]*, February 2017.

[37] Karsten Danzmann and the LISA study team. LISA: Laser interferometer space antenna for gravitational wave measurements. *Classical and Quantum Gravity*, 13(11A):A247–A250, November 1996.

[38] Michele Vallisneri. Synthetic LISA: Simulating time delay interferometry in a model LISA. *Physical Review D: Particles and Fields*, 71(2):022001, January 2005.

[39] Daniel A. Shaddock. Operating LISA as a Sagnac interferometer. *Physical Review D*, 69(2):022001, January 2004.

[40] Richard P. Kornfeld, Bradford W. Arnold, Michael A. Gross, Neil T. Dahya, William M. Klipstein, Peter F. Gath, and Srinivas Bettadpur. GRACE-FO: The Gravity Recovery and Climate Experiment Follow-On Mission. *Journal of Spacecraft and Rockets*, 56(3):931–951, 2019.

[41] James Ira Thorpe and Kirk McKenzie. Arm locking with the GRACE follow-on laser ranging interferometer. *Physical Review D*, 93(4):042003, February 2016.

[42] B. S. Sheard, G. Heinzel, K. Danzmann, D. A. Shaddock, W. M. Klipstein, and W. M. Folkner. Intersatellite laser ranging instrument for the GRACE follow-on mission. *Journal of Geodesy*, 86(12):1083–1095, December 2012.

[43] MenloSystems. *ORS Ultrastable Lasers.*

[44] F Acernese, M Agathos, K Agatsuma, D Aisa, N Allemandou, A Allocca, J Amarni, P Astone, G Balestri, G Ballardin, F Barone, J-P Baronick, M Barsuglia, A Basti, F Basti, Th S Bauer, V Bavigadda, M Bejger, M G Beker, C Belczynski, D Bersanetti, A Bertolini, M Bitossi, M A Bizouard, S Bloemen, M Blom, M Boer, G Bogaert, D Bondi, F Bondu, L Bonelli, R Bonnand, V Boschi, L Bosi, T Bouedo, C Bradaschia, M Branchesi, T Briant, A Brillet, V Brisson, T Bulik, H J Bulten, D Buskulic, C Buy, G Cagnoli, E Calloni, C Campeggi, B Canuel, F Carbognani, F Cavalier, R Cavalieri, G Cella, E Cesarini, E Chassande Mottin, A Chincarini, A Chiummo, S Chua, F Cleva, E Coccia, P-F Cohadon, A Colla, M Colombini, A Conte, J-P Coulon, E Cuoco, A Dalmaz, S D'Antonio, V Dattilo, M Davier, R Day, G Debreczeni, J Degallaix, S Deléglise, W Del Pozzo, H Dereli, R De Rosa, L Di Fiore, A Di Lieto, A Di Virgilio, M Doets, V Dolique, M Drago, M Ducrot, G Endrőczi, V Fafone, S Farinon, I Ferrante, F Ferrini, F Fidecaro, I Fiori, R Flaminio, J-D Fournier, S Franco, S Frasca, F Frasconi, L Gammaitoni, F Garufi, M Gaspard, A Gatto, G Gemme, B Gendre, E Genin, A Gennai, S Ghosh, L Giacobone, A Giazotto, R Gouaty, M Granata, G Greco, P Groot, G M Guidi, J Harms, A Heidmann, H Heitmann, P Hello, G Hemming, E Hennes, D Hofman, P Jaranowski, R J G Jonker, M Kasprzack, F Kéfélian, I Kowalska, M Kraan, A Królak, A Kutynia, C Lazzaro, M Leonardi, N Leroy, N Letendre, T G F Li, B Lieunard, M Lorenzini, V Loriette, G Losurdo, C Magazzù, E Majorana, I Maksimovic, V Malvezzi, N Man, V Mangano, M Mantovani, F Marchesoni, F Marion, J Marque, F Martelli, L Martellini, A Masserot, D Meacher, J Meidam, F Mezzani, C Michel, L Milano, Y Minenkov, A Moggi, M Mohan, M Montani, N Morgado, B Mours, F Mul, M F Nagy, I Nardecchia, L Naticchioni, G Nelemans, I Neri, M Neri, F Nocera, E Pacaud, C Palomba, F Paoletti, A Paoli, A Pasqualetti, R Passaquieti, D Passuello, M Perciballi, S Petit, M Pichot, F Piergiovanni, G Pillant, A Piluso, L Pinard, R Poggiani, M Prijatelj, G A Prodi, M Punturo, P Puppo, D S Rabeling, I Rácz, P Rapagnani, M Razzano, V Re, T Regimbau, F Ricci, F Robinet, A Rocchi, L Rolland, R Romano, D Rosińska, P Ruggi, E Saracco, B Sassolas, F Schimmel, D Sentenac, V Sequino, S Shah, K Siellez, N Straniero, B Swinkels, M Tacca, M Tonelli, F Travasso, M Turconi, G Vajente, N van Bakel, M van Beuzekom, J F J van den Brand, C Van Den Broeck, M V van der Sluys, J van Heijningen, M Vasúth, G Vedovato, J Veitch, D Verkindt, F Vetrano, A Viceré, J-Y Vinet, G Visser, H Vocca, R Ward, M Was, L-W Wei, M Yvert, A Zadro żny, and J-P Zendri. Advanced Virgo: A second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, January 2015.

[45] Patrick Jr Meras, Ilya Y. Poberezhskiy, Daniel H. Chang, Jason Levin, and Gary D. Spiers. Laser frequency stabilization for coherent lidar applications using novel

all-fiber gas reference cell fabrication technique. *24th International Laser Radar Conference (ILRC)*, June 2008.

[46] E. Anne Curtis, Thomas Bradley, Geoffrey P. Barwood, Christopher S. Edwards, Natalie V. Wheeler, Richard Phelan, David J. Richardson, Marco N. Petrovich, and Patrick Gill. Laser frequency stabilization and spectroscopy at 2051 nm using a compact CO_2-filled Kagome hollow core fiber gas-cell system. *Optics Express*, 26(22):28621, October 2018.

[47] Yu-Hung Lai, Myoung-Gyun Suh, Yu-Kun Lu, Boqiang Shen, Qi-Fan Yang, Heming Wang, Jiang Li, Seung Hoon Lee, Ki Youl Yang, and Kerry Vahala. Earth rotation measured by a chip-scale ring laser gyroscope. *Nature Photonics*, 14(6):345–349, June 2020.

[48] Liquid Instruments. *Laser Frequency Offset Locking*, Sept 2021.

[49] Michel Lintz, Duy-Ha Phung, J.-P. Coulon, Benoit Faure, and Thomas Leveque. Dual, feed-forward + feed-back laser frequency control for efficient and convenient diode laser line narrowing. In *International Conference on Space Optics — ICSO 2016*, volume 10562, pages 1588–1595. SPIE, September 2017.

[50] T. Sala, D. Gatti, A. Gambetta, N. Coluccelli, G. Galzerano, P. Laporta, and M. Marangoni. Wide-bandwidth phase lock between a CW laser and a frequency comb based on a feed-forward configuration. *Optics Letters*, 37(13):2592–2594, July 2012.

[51] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31(2):97–105, June 1983.

[52] Eric D. Black. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, January 2001.

[53] D. A. Shaddock, M. B. Gray, and D. E. McClelland. Frequency locking a laser to an optical cavity by use of spatial mode interference. *Optics Letters*, 24(21):1499–1501, November 1999.

[54] Daniel Anthony Shaddock. *Advanced Interferometry for Gravitational Wave Detection*. PhD thesis, November 2000.

[55] Namisha Chabbra, Andrew R. Wade, Emily Rose Rees, Andrew J. Sutton, Alberto Stochino, Robert L. Ward, Daniel A. Shaddock, and Kirk McKenzie. High stability laser locking to an optical cavity using tilt locking. *Optics Letters*, 46(13):3199–3202, 2021.

[56] Bonnie L. Schmittberger and David R. Scherer. A Review of Contemporary Atomic Frequency Standards. *arXiv preprint arXiv:2004.09987*, April 2020.

[57] Fritz Riehle. Optical clock networks. *Nature Photonics*, 11(1):25–31, January 2017.

[58] N. Poli, C. W. Oates, P. Gill, and G. M. Tino. Optical atomic clocks. *La Rivista del Nuovo Cimento*, 36(12):555–624, December 2013.

[59] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. *Reviews of Modern Physics*, 87(2):637–701, June 2015.

[60] Fritz Riehle. Towards a redefinition of the second based on optical atomic clocks. *Comptes Rendus Physique*, 16(5):506–515, June 2015.

[61] Patrick Gill. When should we change the definition of the second? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1953):4109–4130, October 2011.

[62] N. Huntemann, C. Sanner, B. Lipphardt, Chr. Tamm, and E. Peik. Single-ion atomic clock with 3×10^{-18} systematic uncertainty. *Physical Review Letters*, 116(6):063001, February 2016.

[63] C. Chou, D. Hume, J. Koelemeij, D. Wineland, and T. Rosenband. Frequency comparison of two high-accuracy al$^+$ optical clocks. *Physical Review Letters*, 104(7):070802, February 2010.

[64] Ali Al-Masoudi, Sören Dörscher, Sebastian Häfner, Uwe Sterr, and Christian Lisdat. Noise and instability of an optical lattice clock. *Physical Review A*, 92(6):063814, December 2015.

[65] Martin M. Boyd, Andrew D. Ludlow, Sebastian Blatt, Seth M. Foreman, Tetsuya Ido, Tanya Zelevinsky, and Jun Ye. ^{87}Sr lattice clock with inaccuracy below 10^{-15}. *Physical Review Letters*, 98(8):083002, February 2007.

[66] Gretchen K. Campbell, Andrew D. Ludlow, Sebastian Blatt, Jan W. Thomsen, Michael J. Martin, Marcio H. G. de Miranda, Tanya Zelevinsky, Martin M. Boyd, Jun Ye, Scott A. Diddams, Thomas P. Heavner, Thomas E. Parker, and Steven R. Jefferts. The absolute frequency of the ^{87}sr optical clock transition. *Metrologia*, 45(5):539–548, September 2008.

[67] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye. An optical lattice clock with accuracy and stability at the 10^{-18} level. *Nature*, 506(7486):71–75, February 2014.

[68] T. L. Nicholson, S. L. Campbell, R. B. Hutson, G. E. Marti, B. J. Bloom, R. L. McNally, W. Zhang, M. D. Barrett, M. S. Safronova, G. F. Strouse, W. L. Tew, and J. Ye. Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty. *Nature Communications*, 6(1):6896, April 2015.

[69] X. Baillard, M. Fouché, R. Le Targat, P. G. Westergaard, A. Lecallier, F. Chapelet, M. Abgrall, G. D. Rovera, P. Laurent, P. Rosenbusch, S. Bize, G. Santarelli, A. Clairon, P. Lemonde, G. Grosche, B. Lipphardt, and H. Schnatz. An optical lattice clock with spin-polarized ^{87}sr atoms. *The European Physical Journal D*, 48(1):11, December 2007.

[70] F.-L. Hong, M. Musha, M. Takamoto, H. Inaba, S. Yanagimachi, A. Takamizawa, K. Watabe, T. Ikegami, M. Imae, Y. Fujii, M. Amemiya, K. Nakagawa, K. Ueda, and H. Katori. Measuring the frequency of a Sr optical lattice clock using a 120 km coherent optical transfer. *Optics Letters*, 34(5):692–694, March 2009.

[71] Ichiro Ushijima, Masao Takamoto, Manoj Das, Takuya Ohkubo, and Hidetoshi Katori. Cryogenic optical lattice clocks. *Nature Photonics*, 9(3):185–189, March 2015.

[72] N. Hinkley, J. A. Sherman, N. B. Phillips, M. Schioppo, N. D. Lemke, K. Beloy, M. Pizzocaro, C. W. Oates, and A. D. Ludlow. An Atomic Clock with 10^{-18} Instability. *Science*, 341(6151):1215–1218, September 2013.

[73] Klaus Döringshoff, Franz B. Gutsch, Vladimir Schkolnik, Christian Kürbis, Markus Oswald, Benjamin Pröbster, Evgeny V. Kovalchuk, Ahmad Bawamia, Robert Smol, Thilo Schuldt, Matthias Lezius, Ronald Holzwarth, Andreas Wicht, Claus Braxmaier, Markus Krutzik, and Achim Peters. Iodine Frequency Reference on a Sounding Rocket. *Physical Review Applied*, 11(5):054068, May 2019.

[74] Wataru Kokuyama, Kenji Numata, and Jordan Camp. Simple iodine reference at 1064 nm for absolute laser frequency determination in space applications. *Applied Optics*, 49(32):6264–6267, November 2010.

[75] Thilo Schuldt, Klaus Döringshoff, Evgeny V. Kovalchuk, Anja Keetman, Julia Pahl, Achim Peters, and Claus Braxmaier. Development of a compact optical absolute frequency reference for space with 10^{-15} instability. *Applied Optics*, 56(4):1101–1106, February 2017.

[76] Er Jun Zang, Jian Ping Cao, Ye Li, Cheng Yang Li, Yong Kai Deng, and Chun Qing Gao. Realization of Four-Pass $_2$ Absorption Cell in 532-nm Optical Frequency Standard. *IEEE Transactions on Instrumentation and Measurement*, 56(2):673–676, April 2007.

[77] Masato Yoshida, Kazuki Yoshida, Keisuke Kasai, and Masataka Nakazawa. 1.55 um hydrogen cyanide optical frequency-stabilized and 10 GHz repetition-rate-stabilized mode-locked fiber laser. *Optics Express*, 24(21):24287–24296, October 2016.

[78] David J. Jones, Scott A. Diddams, Jinendra K. Ranka, Andrew Stentz, Robert S. Windeler, John L. Hall, and Steven T. Cundiff. Carrier-Envelope Phase Control of Femtosecond Mode-Locked Lasers and Direct Optical Frequency Synthesis. *Science*, 288(5466):635–639, April 2000.

[79] Scott A. Diddams, David J. Jones, Jun Ye, Steven T. Cundiff, John L. Hall, Jinendra K. Ranka, Robert S. Windeler, Ronald Holzwarth, Thomas Udem, and T. W. Hänsch. Direct Link between Microwave and Optical Frequencies with a 300 THz Femtosecond Laser Comb. *Physical Review Letters*, 84(22):5102–5105, May 2000.

[80] A. I. Ferguson, J. N. Eckstein, and T. W. Hänsch. A subpicosecond dye laser directly pumped by a mode-locked argon laser. *Journal of Applied Physics*, 49(11):5389–5391, November 1978.

[81] D. E. Spence, P. N. Kean, and W. Sibbett. 60-fsec pulse generation from a self-mode-locked Ti:sapphire laser. *Optics Letters*, 16(1):42–44, January 1991.

[82] Jinendra K. Ranka, Robert S. Windeler, and Andrew J. Stentz. Visible continuum generation in air–silica microstructure optical fibers with anomalous dispersion at 800 nm. *Optics Letters*, 25(1):25–27, January 2000.

[83] T. A. Birks, W. J. Wadsworth, and P. St J. Russell. Supercontinuum generation in tapered fibers. *Optics Letters*, 25(19):1415–1417, October 2000.

[84] J. N. Eckstein, A. I. Ferguson, and T. W. Hänsch. High-Resolution Two-Photon Spectroscopy with Picosecond Light Pulses. *Physical Review Letters*, 40(13):847–850, March 1978.

[85] J.L. Hall. Optical frequency measurement: 40 years of technology revolutions. *IEEE Journal of Selected Topics in Quantum Electronics*, 6(6):1136–1144, November 2000.

[86] Th Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch. Accurate measurement of large optical frequency differences with a mode-locked laser. *Optics Letters*, 24(13):881–883, July 1999.

[87] Jun Ye, Jin-Long Peng, R. Jason Jones, Kevin W. Holman, John L. Hall, David J. Jones, Scott A. Diddams, John Kitching, Sebastien Bize, James C. Bergquist, Leo W. Hollberg, Lennart Robertsson, and Long-Sheng Ma. Delivery of high-stability optical and microwave frequency standards over an optical fiber network. *JOSA B*, 20(7):1459–1467, July 2003.

[88] J. Reichert, R. Holzwarth, Th. Udem, and T. W. Hänsch. Measuring the frequency of light with mode-locked lasers. *Optics Communications*, 172(1):59–68, December 1999.

[89] M. Niering, R. Holzwarth, J. Reichert, P. Pokasov, Th. Udem, M. Weitz, T. W. Hänsch, P. Lemonde, G. Santarelli, M. Abgrall, P. Laurent, C. Salomon, and A. Clairon. Measurement of the hydrogen $1S$- $2S$ transition frequency by phase coherent comparison with a microwave cesium fountain clock. *Physical Review Letters*, 84(24):5496–5499, June 2000.

[90] Jun Ye, Tai Hyun Yoon, John L. Hall, Alan A. Madej, John E. Bernard, Klaus J. Siemsen, Louis Marmet, Jean-Marie Chartier, and Annick Chartier. Accuracy Comparison of Absolute Optical Frequency Measurement between Harmonic-Generation Synthesis and a Frequency-Division Femtosecond Comb. *Physical Review Letters*, 85(18):3797–3800, October 2000.

[91] R. Holzwarth, Th. Udem, T. W. Hänsch, J. C. Knight, W. J. Wadsworth, and P. St. J. Russell. Optical Frequency Synthesizer for Precision Spectroscopy. *Physical Review Letters*, 85(11):2264–2267, September 2000.

[92] Birgitta Bernhardt, Akira Ozawa, Patrick Jacquet, Marion Jacquey, Yohei Kobayashi, Thomas Udem, Ronald Holzwarth, Guy Guelachvili, Theodor W. Hänsch, and Nathalie Picqué. Cavity-enhanced dual-comb spectroscopy. *Nature Photonics*, 4(1):55–57, January 2010.

[93] Guy Millot, Stéphane Pitois, Ming Yan, Tatevik Hovhannisyan, Abdelkrim Bendahmane, Theodor W. Hänsch, and Nathalie Picqué. Frequency-agile dual-comb spectroscopy. *Nature Photonics*, 10(1):27–30, January 2016.

[94] N. Kuse and M. E. Fermann. Frequency-modulated comb LIDAR. *APL Photonics*, 4(10):106105, October 2019.

[95] Jacob Nürnberg, Benjamin Willenberg, Christopher R. Phillips, and Ursula Keller. Dual-comb ranging with frequency combs from single cavity free-running laser oscillators. *Optics Express*, 29(16):24910–24918, August 2021.

[96] Hollie Wright, Jinghua Sun, Jinghua Sun, David McKendrick, Nick Weston, and Derryck T. Reid. Two-photon dual-comb LiDAR. *Optics Express*, 29(23):37037–37047, November 2021.

[97] Stanford Research Systems. *FS725 Rubidium Frequency Standard Operation and Service Manual*.

[98] Thomas P. Heavner, Elizabeth A. Donley, Filippo Levi, Giovanni Costanzo, Thomas E. Parker, Jon H. Shirley, Neil Ashby, Stephan Barlow, and S. R. Jefferts. First accuracy evaluation of NIST-F2. *Metrologia*, 51(3):174–182, May 2014.

[99] X. Y. Zeng, Y. X. Ye, X. H. Shi, Z. Y. Wang, K. Deng, J. Zhang, and Z. H. Lu. Thermal-noise-limited higher-order mode locking of a reference cavity. *Optics Letters*, 43(8):1690–1693, April 2018.

[100] A. D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye. Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15}. *Optics Letters*, 32(6):641–643, March 2007.

[101] S. A. Webster, M. Oxborrow, S. Pugla, J. Millo, and P. Gill. Thermal-noise-limited optical cavity. *Physical Review A: Atomic, Molecular, and Optical Physics*, 77(3):033847, March 2008.

[102] Sebastian Häfner, Stephan Falke, Christian Grebing, Stefan Vogt, Thomas Legero, Mikko Merimaa, Christian Lisdat, and Uwe Sterr. 8×10^{-17} fractional laser frequency instability with a long room-temperature cavity. *Optics Letters*, 40(9):2112–2115, May 2015.

[103] Y. Y. Jiang, A. D. Ludlow, N. D. Lemke, J. A. Sherman, J. Von Stecher, R. W. Fox, L. S. Ma, A. M. Rey, and C. W. Oates. Improving the stability and accuracy of the Yb optical lattice clock. In *Proceedings of the IEEE International Frequency Control Symposium and Exposition*, pages 8–10, 2011.

[104] EP Ippen and RH Stolen. Stimulated Brillouin scattering in optical fibers. *Applied Physics Letters*, 21(11):539–541, 1972.

[105] R. G. Smith. Optical power handling capacity of low loss optical fibers as determined by stimulated raman and brillouin scattering. *Applied Optics*, 11(11):2489–2494, November 1972.

[106] R. H. Stolen and Chinlon Lin. Self-phase-modulation in silica optical fibers. *Physical Review A: Atomic, Molecular, and Optical Physics*, 17(4):1448–1453, April 1978.

[107] S. W. Lloyd, M. J. F. Digonnet, and S. Fan. Modeling coherent backscattering errors in fiber optic gyroscopes for sources of arbitrary line width. *Journal of Lightwave Technology*, 31(13):2070–2078, July 2013.

[108] Jong H. Chow, Benjamin S. Sheard, David E. McClelland, Malcolm B. Gray, and Ian C. M. Littler. Photothermal effects in passive fiber Bragg grating resonators. *Opt. Lett., OL*, 30(7):708–710, April 2005.

[109] Ian C. M. Littler, Thomas Grujic, and Benjamin J. Eggleton. Photothermal effects in fiber Bragg gratings. *Appl. Opt., AO*, 45(19):4679–4685, July 2006.

[110] Fabien Kéfélian, Haifeng Jiang, Pierre Lemonde, and Giorgio Santarelli. Ultralow-frequency-noise stabilization of a laser by locking to an optical fiber-delay line. *Optics Letters*, 34(7):914, April 2009.

[111] Jing Dong, Yongqi Hu, Junchao Huang, Meifeng Ye, Qiuzhi Qu, Tang Li, and Liang Liu. Subhertz linewidth laser by locking to a fiber delay line. *Applied Optics*, 54(5):1152–1156, February 2015.

[112] Nicola Chiodo, Khelifa Djerroud, Ouali Acef, André Clairon, and Peter Wolf. Lasers for coherent optical satellite links with large dynamics. *Applied Optics*, 52(30):7342–7351, October 2013.

[113] Haifeng Jiang, Fabien Kéfélian, Pierre Lemonde, André Clairon, and Giorgio Santarelli. An agile laser with ultra-low frequency noise and high sweep linearity. *Optics Express*, 18(4):3284–3297, February 2010.

[114] Ashby P. Hilton, Philip S. Light, Lauris Talbot, and Andre N. Luiten. Optimal design for spectral narrowing and fast frequency sweep of an interferometer-stabilized laser. *Optics Letters*, 45(1):45, January 2020.

[115] Benjamin S. Sheard, Malcolm B. Gray, and David E. McClelland. High-bandwidth laser frequency stabilization to a fiber-optic delay line. *Applied Optics*, 45(33):8491–8499, November 2006.

[116] Nicolas Bourbeau Hébert, Ashby P. Hilton, Philip S. Light, and Andre N. Luiten. Hertz-level frequency comparisons between diverse color lasers without a frequency comb. *Optics Letters*, 45(15):4196, August 2020.

[117] G. A. Cranch. Frequency noise reduction in erbium-doped fiber distributed-feedback lasers by electronic feedback. *Opt. Lett., OL*, 27(13):1114–1116, July 2002.

[118] K Takahashi, M Ando, and K Tsubono. Stabilization of laser intensity and frequency using optical fiber. *Journal of Physics: Conference Series*, 122:012016, July 2008.

[119] Radek Šmíd, Martin Čížek, Břetislav Mikel, and Ondřej Číp. Frequency noise suppression of a single mode laser with an unbalanced fiber interferometer for subnanometer interferometry. *Sensors*, 15(1):1342–1353, 2015.

[120] Masato Wada, Feng-Lei Hong, and Hajime Inaba. Frequency noise measurement and its uncertainty estimation of an optical frequency comb using a delay line interferometer. *Measurement Science and Technology*, 31(12):125012, October 2020.

[121] Andrew J. Sutton, Oliver Gerberding, Gerhard Heinzel, and Daniel A. Shaddock. Digitally enhanced homodyne interferometry. *Optics Express*, 20(20):22195, 2012.

[122] Daniel A. Shaddock. Digitally enhanced heterodyne interferometry. *Optics Letters*, 32(22):3355–3357, November 2007.

[123] Paul G. Sibley, Robert L. Ward, Lyle E. Roberts, Samuel P. Francis, and Daniel A. Shaddock. Crosstalk reduction for multi-channel optical phase metrology. *Optics Express*, 28(7):10400–10424, March 2020.

[124] Anneshwa Dey, Ya Zhang, Justin Wong, Paul Sibley, Chathura Bandutunga, Malcolm Gray, and Jong Chow. Algebraic cancellation of inter-channel crosstalk in multiplexed heterodyne interferometry. *Optics Letters*, October 2021.

[125] James T. Spollard, Lyle E. Roberts, Paul G. Sibley, Callum S. Sambridge, Daniel A. Shaddock, and Kirk McKenzie. Improved cross-talk suppression for digitally enhanced interferometry using Golay complementary pairs. *Optics Letters*, 47(7):1570–1573, April 2022.

[126] Danielle M. R. Wuchenich, Timothy T.-Y. Lam, Jong H. Chow, David E. McClelland, and Daniel A. Shaddock. Laser frequency noise immunity in multiplexed displacement sensing. *Optics Letters*, 36(5):672, March 2011.

[127] Glenn de Vine, David S. Rabeling, Bram J. J. Slagmolen, Timothy T.-Y. Lam, Sheon Chua, Danielle M. Wuchenich, David E. McClelland, and Daniel A. Shaddock. Picometer level displacement metrology with digitally enhanced heterodyne interferometry. *Optics Express*, 17(2):828–837, January 2009.

[128] S. Golomb. On the classification of balanced binary sequences of period $2^n - 1$(Corresp.). *IEEE Transactions on Information Theory*, 26(6):730–732, November 1980.

[129] R. H. Barker. Group Synchronizing of Binary Digital Sequences. *Communication Theory*, 1953.

[130] Robert Gold. Optimal Binary Sequences for Spread Spectrum Multiplexing. *IEEE Trans. Information Theory*, pages 619–621, 1967.

[131] T. Kasami. Some lower bounds on the minimum weight of cyclic codes of composite length. *IEEE Transactions on Information Theory*, 14(6):814–818, November 1968.

[132] M. Golay. Complementary series. *IRE Transactions on Information Theory*, 7(2):82–87, April 1961.

[133] Lyle Edward Roberts. *Internally Sensed Optical Phased Arrays*. PhD thesis, Australian National University, July 2016.

[134] Chathura Priyankara Bandutunga. *Digitally Enhanced Interferometry for Precision Metrology*. PhD thesis, Australian National University, 2020.

[135] Roy Ward and Timothy C.A. Molteno. *Table of Linear Feedback Shift Registers*, 2012.

[136] Ya Zhang, Chathura P. Bandutunga, Malcolm B. Gray, and Jong H. Chow. Multi-target CW interferometric acoustic measurements on a single optical beam. *Optics Express*, 27(13):18477–18483, June 2019.

[137] PNT Capability Area Integration, Portfolio Architect Space and Missile Systems Center – LAAFB. *NAVSTAR GPS Space Segment/ User Segment L1C Interfaces*, May 2021.

[138] PNT Capability Area Integration, Portfolio Architect, Space and Missile Systems Center - LAAFB. *NAVSTAR GPS Space Segment/User Segment L5 Interfaces*, May 2021.

[139] Ali Pezeshki, Robert Calderbank, Stephen D. Howard, and William Moran. Doppler Resilient Golay Complementary Pairs for Radar. In *2007 IEEE/SP 14th Workshop on Statistical Signal Processing*, pages 483–487, August 2007.

[140] Matthew G. Parker, Kenneth G. Paterson, and Chintha Tellambura. Golay Complementary Sequences. *Wiley Encyclopedia of Telecommunications*, page eot367, April 2003.

[141] P. Borwein and R. Ferguson. A complete description of Golay pairs for lengths up to 100. *Mathematics of Computation*, 73(246):967–985, 2004.

[142] Elana Kalashnikov. An Introduction to Golay Complementary Sequences. *Eureka*, 4(1):40–48, July 2014.

[143] Chikao Nagasawa, Makoto Abo, Hideki Yamamoto, and Osamu Uchino. Random modulation cw lidar using new random sequence. *Applied Optics*, 29(10):1466–1470, April 1990.

[144] Terry G. McRae, Silvie Ngo, Timothy T.-Y. Lam, David E. McClelland, Jong H. Chow, Daniel A. Shaddock, and Malcolm B. Gray. Algebraic cancellation of polarisation noise in fibre interferometers. *Optics Express*, 24(10):10486–10494, May 2016.

[145] Terry G. McRae, Silvie Ngo, Daniel A. Shaddock, Magnus T. L. Hsu, and Malcolm B. Gray. Digitally enhanced optical fiber frequency reference. *Optics Letters*, 39(7):1752–1755, April 2014.

[146] Chathura P. Bandutunga, Terry G. McRae, Ya Zhang, Malcolm B. Gray, and Jong H. Chow. Infrasonic performance of a passively stabilized, all-fiber, optical frequency reference. *Optics Express*, 28(7):9280–9287, March 2020.

[147] Ya Zhang, Chathura P. Bandutunga, Terry G. McRae, Malcolm B. Gray, and Jong H. Chow. Double Rayleigh scattering in a digitally enhanced, all-fiber optical frequency reference. *Optics Express*, 29(17):26319–26331, August 2021.

[148] Silvie Ngo, Daniel A. Shaddock, Terry G. McRae, Timothy T-Y. Lam, Jong H. Chow, and Malcolm B. Gray. Suppressing Rayleigh backscatter and code noise from all-fiber digital interferometers. *Optics Letters*, 41(1):84–87, January 2016.

[149] Chathura P. Bandutunga, Ya Zhang, Terry G. Mcrae, Malcolm B. Gray, and Jong H. Chow. Coherent rayleigh backscatter phase noise in digitally enhanced fiber interferometers. *Journal of Lightwave Technology*, 39(8):2625–2630, 2021.

[150] A. D. Kersey, M. J. Marrone, and M. A. Davis. Polarisation-insensitive fibre optic Michelson interferometer. *Electronics Letters*, 27(6):518–520, March 1991.

[151] *Thorlabs - AV3 Ø18.0 Mm Sorbothane Feet, Adhesive Mounting Surface, 4 Pieces.*

[152] Focal. *Focal B.A.M. XXXL.*

[153] Swagelok. *Gaugeable Tube Fittings and Adapter Fittings.*

[154] Eric R.I. Abraham and Eric A. Cornell. Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems. *Applied Optics*, 37(10):1762, April 1998.

[155] Clay K. Kirkendall and Anthony Dandridge. Overview of high performance fibre-optic sensing. *Journal of Physics D: Applied Physics*, 37(18):R197–R216, September 2004.

[156] *Thorlabs - CPC250 In-Line Fiber Optic Polarization Controller for Ø250 Mm Bare Fiber.*

[157] *NI 5782R User Manual and Specifications - National Instruments.*

[158] Xilinx. *FPGAs & 3D ICs.*

[159] *High Throughput Rectangular to Polar Function - LabVIEW 2018 FPGA Module Help - National Instruments.*

[160] Richard Lyons. Understanding cascaded integrator-comb filters. *Embedded Systems Programming Magazine*, January 2005.

[161] Bill Riley. Stable32 User Manual. page 348.

[162] Jun Ye and Steven T. Cundiff, editors. *Femtosecond Optical Frequency Comb: Principle, Operation, and Applications.* Kluwer Academic Publishers, Boston, 2005.

[163] MenloSystems. *User Manual – FC1500 Optical Frequency Synthesizer*, August 2013.

[164] TechnicaSA. *T10 Customized FBG Sensor.*

[165] *Second-Order RC Networks.*

[166] *Thermal Time Constant and Specific Heat - John Dunn, Consultant, Ambertec P.E., P.C.*

[167] Silicon Audio. *Low-Noise Optical Seismic Sensor.*

[168] Newport. *Active Vibration Damping.*

[169] Massimo Tinto and J. W. Armstrong. Cancellation of laser noise in an unequal-arm interferometer detector of gravitational radiation. *Physical Review D: Particles and Fields*, 59(10):102003, April 1999.

[170] J. W. Armstrong, F. B. Estabrook, and Massimo Tinto. Time-delay interferometry for space based gravitational wave searches. *ApJ*, 527(2):814–826, December 1999.

[171] Massimo Tinto and Sanjeev V. Dhurandhar. Time-Delay Interferometry. *Living Reviews in Relativity*, 17(1):6, August 2014.

[172] Jean-Baptiste Bayle, Marc Lilley, Antoine Petiteau, and Hubert Halloin. Effect of filters on the time-delay interferometry residual laser noise for LISA. *Physical Review D: Particles and Fields*, 99(8):084023, April 2019.

[173] Giacomo Giampieri, Ronald W. Hellings, Massimo Tinto, and James E. Faller. Algorithms for unequal-arm Michelson interferometers. *Optics Communications*, 123(4):669–678, 1996.

[174] Massimo Tinto, F. B. Estabrook, and J. W. Armstrong. Time-delay interferometry for LISA. *Physical Review D: Particles and Fields*, 65(8):082003, April 2002.

[175] Neil J Cornish and Ronald W Hellings. The effects of orbital motion on LISA time delay interferometry. *Class. Quantum Grav.*, 20(22):4851–4860, November 2003.

[176] Gang Wang and Wei-Tou Ni. Numerical simulation of time delay interferometry for TAIJI and new LISA. *Research in Astronomy and Astrophysics*, 19(4):058, April 2019.

[177] D. A. Shaddock, B. Ware, R. E. Spero, and M. Vallisneri. Postprocessed time-delay interferometry for LISA. *Physical Review D: Particles and Fields*, 70(8):081101, October 2004.

[178] Massimo Tinto, Daniel A. Shaddock, Julien Sylvestre, and J. W. Armstrong. Implementation of time-delay interferometry for LISA. *Physical Review D: Particles and Fields*, 67(12):122003, June 2003.

[179] Samuel P Francis, Daniel A Shaddock, Andrew J Sutton, Glenn De Vine, Brent Ware, Robert E Spero, William M Klipstein, and Kirk McKenzie. Tone-assisted time delay interferometry on GRACE Follow-On. *Physical Review D*, 92(1):012005, 2015.

[180] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae. *Physics Letters A*, 264(1):1–10, December 1999.

[181] K. H. Wanser. Fundamental phase noise limit in optical fibres due to temperature fluctuations. *Electronics Letters*, 28(1):53–54, January 1992.

[182] S. Foster, A. Tikhomirov, and M. Milnes. Fundamental Thermal Noise in Distributed Feedback Fiber Lasers. *IEEE Journal of Quantum Electronics*, 43(5):378–384, May 2007.

[183] S. Knudsen, A.B. Tveten, and A. Dandridge. Measurements of fundamental thermal induced phase fluctuations in the fiber of a Sagnac interferometer. *IEEE Photonics Technology Letters*, 7(1):90–92, January 1995.

[184] K. Krakenes and K. Blotekjaer. Comparison of fiber-optic Sagnac and Mach-Zehnder interferometers with respect to thermal processes in the fiber. *Journal of Lightwave Technology*, 13(4):682–686, April 1995.

[185] Alan D. Kersey. A Review of Recent Developments in Fiber Optic Sensor Technology. *Optical Fiber Technology*, 2(3):291–317, July 1996.

[186] V. Annovazzi-Lodi, S. Donati, and S. Merlo. Thermodynamic phase noise in fibre interferometers. *Optical and Quantum Electronics*, 28(1):43–49, January 1996.

[187] R. P. Moeller and W. K. Burns. Observation of thermal noise in a dynamically biased fiber-optic gyro. *Optics Letters*, 21(3):171, February 1996.

[188] Erlend Rønnekleiv. Frequency and Intensity Noise of Single Frequency Fiber Bragg Grating Lasers. *Optical Fiber Technology*, 7(3):206–235, July 2001.

[189] Scott Foster, Geoffrey A. Cranch, and Alexei Tikhomirov. Experimental evidence for the thermal origin of $1/f$ frequency noise in erbium-doped fiber lasers. *Physical Review A: Atomic, Molecular, and Optical Physics*, 79(5):053802, May 2009.

[190] R. E. Bartolo, A. B. Tveten, and A. Dandridge. Thermal Phase Noise Measurements in Optical Fiber Interferometers. *IEEE Journal of Quantum Electronics*, 48(5):720–727, May 2012.

[191] Lingze Duan. General treatment of the thermal noises in optical fibers. *Physical Review A: Atomic, Molecular, and Optical Physics*, 86(2):023817, August 2012.

[192] Jing Dong, Junchao Huang, Tang Li, and Liang Liu. Observation of fundamental thermal noise in optical fibers down to infrasonic frequencies. *Applied Physics Letters*, 108(2):021108, January 2016.

[193] Kazuhiro Yamamoto, Masaki Ando, Keita Kawabe, and Kimio Tsubono. Theoretical approach to thermal noise caused by an inhomogeneously distributed loss: Physical insight by the advanced modal expansion. *Physical Review D: Particles and Fields*, 75(8):082002, April 2007.

[194] W. Burns and R. Moeller. Rayleigh backscattering in a fiber gyroscope with limited coherence sources. *Journal of Lightwave Technology*, 1(2):381–386, 1983.

[195] Kazumasa Takada. Calculation of Rayleigh backscattering noise in fiber-optic gyroscopes. *J. Opt. Soc. Am. A, JOSAA*, 2(6):872–877, June 1985.

[196] W. Burns, Chin-Lin Chen, and R. Moeller. Fiber-optic gyroscopes with broad-band sources. *Journal of Lightwave Technology*, 1(1):98–105, March 1983.

[197] Seth W. Lloyd, Vinayak Dangui, Michel J. F. Digonnet, Shanhui Fan, and G. S. Kino. Measurement of reduced backscattering noise in laser-driven fiber optic gyroscopes. *Optics Letters*, 35(2):121, 2010.

[198] Seth W. Lloyd, Shanhui Fan, and Michel J. F. Digonnet. Experimental Observation of Low Noise and Low Drift in a Laser-Driven Fiber Optic Gyroscope. *J. Lightwave Technol., JLT*, 31(13):2079–2085, July 2013.

[199] Michael Fleyer, Seth Heerschap, Geoffrey A. Cranch, and Moshe Horowitz. Noise induced in optical fibers by double Rayleigh scattering of a laser with a $1/f^\nu$ frequency noise. *Optics Letters*, 41(6):1265–1268, March 2016.

[200] P. Gysel and R. K. Staubli. Statistical properties of Rayleigh backscattering in single-mode fibers. *Journal of Lightwave Technology*, 8(4):561–567, 1990.

[201] Michael Fleyer, James P. Cahill, Moshe Horowitz, Curtis R. Menyuk, and Olukayode Okusaga. Comprehensive model for studying noise induced by self-homodyne detection of backward Rayleigh scattering in optical fibers. *Optics Express*, 23(20):25635–25652, October 2015.

[202] E. Brinkmeyer. Backscattering in single-mode fibres. *Electronics Letters*, 16(9):329–330, April 1980.

[203] NKT Photonics. *Koheras BASIK Ultra-low Phase Noise OEM Fiber Laser.*

[204] NKT Photonics. *Koheras ADJUSTIK Low-noise, Single-Frequency Fiber Laser.*

[205] Orbits Lightwave, Inc. *Ethernal™ SlowLight™ Laser Modules.*

[206] Thomas J. Kane and Robert L. Byer. Monolithic, unidirectional single-mode Nd:YAG ring laser. *Optics Letters*, 10(2):65, February 1985.

[207] Lumentum. *Continuous-Wave (CW) Single-Frequency IR Laser NPRO® 125/126 Series.*

[208] Coherent. *Ultra-Narrow Linewidth CW DPSS Laser.*

[209] Carlton M. Caves. Quantum-mechanical noise in an interferometer. *Physical Review D*, 23(8):1693–1708, April 1981.

[210] Andrew J. Stevenson, Malcolm B. Gray, Hans-A. Bachor, and David E. McClelland. Quantum-noise-limited interferometric phase measurements. *Applied Optics*, 32(19):3481–3493, July 1993.

[211] Insight Photonic Solutions, Inc. *High Performance Balanced Photodetector (BPD-1).*

[212] Coherent, Nufern. *Fiber Optic Gyroscope Coils.*

[213] Advanced Navigation. *BOREAS D90*, September 2021.

[214] L.Z. Duan. Intrinsic thermal noise of optical fibres due to mechanical dissipation. *Electronics Letters*, 46(22):1515, 2010.

[215] George Skolianos, Arushi Arora, Martin Bernier, and Michel J. F. Digonnet. Photonics sensing at the thermodynamic limit. *Optics Letters*, 42(10):2018–2021, May 2017.

[216] Roger Haagmans, Christian Siemes, Luca Massotti, Olivier Carraz, and Pierluigi Silvestrin. ESA's next-generation gravity mission concepts. *Rendiconti Lincei. Scienze Fisiche e Naturali*, 31(1):15–25, October 2020.

[217] W M Folkner, K McKenzie, D Shaddock, R Spero, R Thompson, D Wuchenich, N Yu, M Stephens, J Leitch, M Davis, R Pierce, and Ball Aerospace. Laser Frequency Stabilization for GRACE-II. page 8, June 2010.

[218] Shuichi Sato, Seiji Kawamura, Masaki Ando, Takashi Nakamura, Kimio Tsubono, Akito Araya, Ikkoh Funaki, Kunihito Ioka, Nobuyuki Kanda, Shigenori Moriwaki, Mitsuru Musha, Kazuhiro Nakazawa, Kenji Numata, Shin-ichiro Sakai, Naoki Seto, Takeshi Takashima, Takahiro Tanaka, Kazuhiro Agatsuma, Koh-suke Aoyanagi, Koji Arai, Hideki Asada, Yoichi Aso, Takeshi Chiba, Toshikazu Ebisuzaki, Yumiko Ejiri, Motohiro Enoki, Yoshiharu Eriguchi, Masa-Katsu Fujimoto, Ryuichi Fujita, Mitsuhiro Fukushima, Toshifumi Futamase, Katsuhiko Ganzu, Tomohiro Harada, Tatsuaki Hashimoto, Kazuhiro Hayama, Wataru Hikida, Yoshiaki Himemoto, Hisashi Hirabayashi, Takashi Hiramatsu, Feng-Lei Hong, Hideyuki Horisawa, Mizuhiko Hosokawa, Kiyotomo Ichiki, Takeshi Ikegami, Kaiki T. Inoue, Koji Ishidoshiro, Hideki Ishihara, Takehiko Ishikawa, Hideharu Ishizaki, Hiroyuki Ito, Yousuke Itoh, Nobuki Kawashima, Fumiko Kawazoe, Naoko Kishimoto, Kenta Kiuchi, Shiho Kobayashi, Kazunori Kohri, Hiroyuki Koizumi, Yasufumi Kojima, Keiko Kokeyama, Wataru Kokuyama, Kei Kotake, Yoshihide Kozai, Hideaki Kudoh, Hiroo Kunimori, Hitoshi Kuninaka, Kazuaki Kuroda, Kei-ichi Maeda, Hideo Matsuhara, Yasushi Mino, Osamu Miyakawa, Shinji Miyoki, Mutsuko Y.

Morimoto, Tomoko Morioka, Toshiyuki Morisawa, Shinji Mukohyama, Shigeo Nagano, Isao Naito, Kouji Nakamura, Hiroyuki Nakano, Kenichi Nakao, Shinichi Nakasuka, Yoshinori Nakayama, Erina Nishida, Kazutaka Nishiyama, Atsushi Nishizawa, Yoshito Niwa, Taiga Noumi, Yoshiyuki Obuchi, Masatake Ohashi, Naoko Ohishi, Masashi Ohkawa, Norio Okada, Kouji Onozato, Kenichi Oohara, Norichika Sago, Motoyuki Saijo, Masaaki Sakagami, Shihori Sakata, Misao Sasaki, Takashi Sato, Masaru Shibata, Hisaaki Shinkai, Kentaro Somiya, Hajime Sotani, Naoshi Sugiyama, Yudai Suwa, Rieko Suzuki, Hideyuki Tagoshi, Fuminobu Takahashi, Kakeru Takahashi, Keitaro Takahashi, Ryutaro Takahashi, Ryuichi Takahashi, Tadayuki Takahashi, Hirotaka Takahashi, Takamori Akiteru, Tadashi Takano, Keisuke Taniguchi, Atsushi Taruya, Hiroyuki Tashiro, Yasuo Torii, Morio Toyoshima, Shinji Tsujikawa, Yoshiki Tsunesada, Akitoshi Ueda, Ken-ichi Ueda, Masayoshi Utashima, Yaka Wakabayashi, Hiroshi Yamakawa, Kazuhiro Yamamoto, Toshitaka Yamazaki, Jun'ichi Yokoyama, Chul-Moon Yoo, Shijun Yoshida, and Taizoh Yoshino. DECIGO: The Japanese space gravitational wave antenna. *Journal of Physics: Conference Series*, 154:012040, March 2009.

[219] Seiji Kawamura, Masaki Ando, Naoki Seto, Shuichi Sato, Takashi Nakamura, Kimio Tsubono, Nobuyuki Kanda, Takahiro Tanaka, Jun'ichi Yokoyama, Ikkoh Funaki, Kenji Numata, Kunihito Ioka, Takeshi Takashima, Kazuhiro Agatsuma, Tomotada Akutsu, Koh-suke Aoyanagi, Koji Arai, Akito Araya, Hideki Asada, Yoichi Aso, Dan Chen, Takeshi Chiba, Toshikazu Ebisuzaki, Yumiko Ejiri, Motohiro Enoki, Yoshiharu Eriguchi, Masa-Katsu Fujimoto, Ryuichi Fujita, Mitsuhiro Fukushima, Toshifumi Futamase, Tomohiro Harada, Tatsuaki Hashimoto, Kazuhiro Hayama, Wataru Hikida, Yoshiaki Himemoto, Hisashi Hirabayashi, Takashi Hiramatsu, Feng-Lei Hong, Hideyuki Horisawa, Mizuhiko Hosokawa, Kiyotomo Ichiki, Takeshi Ikegami, Kaiki T. Inoue, Koji Ishidoshiro, Hideki Ishihara, Takehiko Ishikawa, Hideharu Ishizaki, Hiroyuki Ito, Yousuke Itoh, Kiwamu Izumi, Isao Kawano, Nobuki Kawashima, Fumiko Kawazoe, Naoko Kishimoto, Kenta Kiuchi, Shiho Kobayashi, Kazunori Kohri, Hiroyuki Koizumi, Yasufumi Kojima, Keiko Kokeyama, Wataru Kokuyama, Kei Kotake, Yoshihide Kozai, Hiroo Kunimori, Hitoshi Kuninaka, Kazuaki Kuroda, Sachiko Kuroyanagi, Kei-ichi Maeda, Hideo Matsuhara, Nobuyuki Matsumoto, Yuta Michimura, Osamu Miyakawa, Umpei Miyamoto, Shinji Miyoki, Mutsuko Y. Morimoto, Toshiyuki Morisawa, Shigenori Moriwaki, Shinji Mukohyama, Mitsuru Musha, Shigeo Nagano, Isao Naito, Kouji Nakamura, Hiroyuki Nakano, Kenichi Nakao, Shinichi Nakasuka, Yoshinori Nakayama, Kazuhiro Nakazawa, Erina Nishida, Kazutaka Nishiyama, Atsushi Nishizawa, Yoshito Niwa, Taiga Noumi, Yoshiyuki Obuchi, Masatake Ohashi, Naoko Ohishi, Masashi Ohkawa, Kenshi Okada, Norio Okada, Kenichi Oohara, Norichika Sago, Motoyuki Saijo, Ryo Saito, Masaaki Sakagami, Shin-ichiro Sakai, Shihori Sakata, Misao Sasaki, Takashi Sato, Masaru Shibata, Hisaaki Shinkai, Ayaka Shoda, Kentaro Somiya, Hajime Sotani, Naoshi Sugiyama, Yudai Suwa, Rieko Suzuki, Hideyuki Tagoshi, Fuminobu Takahashi, Kakeru Takahashi, Keitaro Takahashi, Ryutaro Takahashi, Ryuichi Takahashi, Tadayuki Takahashi, Hirotaka Takahashi, Takamori Akiteru, Tadashi Takano, Nobuyuki Tanaka, Keisuke Taniguchi, Atsushi Taruya, Hiroyuki Tashiro, Yasuo Torii, Morio Toyoshima, Shinji Tsujikawa, Yoshiki Tsunesada, Akitoshi Ueda, Ken-ichi Ueda, Masayoshi Utashima, Yaka Wakabayashi, Kent Yagi, Hiroshi Yamakawa, Kazuhiro Yamamoto, Toshitaka Yamazaki, Chul-Moon Yoo,

Shijun Yoshida, Taizoh Yoshino, and Ke-Xun Sun. The Japanese space gravitational wave antenna: DECIGO. *Classical and Quantum Gravity*, 28(9):094011, April 2011.

[220] Hesham Sakr, Yong Chen, Gregory T. Jasion, Thomas D. Bradley, John R. Hayes, Hans Christian H. Mulvad, Ian A. Davidson, Eric Numkam Fokoua, and Francesco Poletti. Hollow core optical fibres with comparable attenuation to silica fibres between 600 and 1100 nm. *Nature Communications*, 11(1):6030, November 2020.

[221] Corning. *Corning® HI 1060 & RC HI 1060 Specialty Optical Fibers*.

[222] Gooch & Housego. *Fiber-Q 1550 nm Fiber Coupled Acousto-Optic Modulator*, December 2016.

[223] Andrew John Heaton Sutton. *Digital Techniques for Coherent Optical Metrology*. PhD thesis, Australian National University, March 2013.

[224] Ya Zhang, Yuwei Zhao, Chathura P. Bandutunga, Malcolm B. Gray, and Jong H. Chow. Stereoscopic audio-band vibrometry with source triangulation and interferometric sensitivity. In *2020 Conference on Lasers and Electro-Optics Pacific Rim (CLEO-PR)*, pages 1–2, August 2020.

[225] National Instruments. *NI-5782 Multi-Sample CLIP I/O Reference*.

[226] National Instruments. *PXIe-1092 Specifications*.

[227] National Instruments. *NI PXIe-7975R Specifications*.

[228] Texas Instruments. *DAC5682Z 16-Bit, 1.0 GSPS 2x-4x Interpolating Dual-Channel Digital-to-Analog Converter (DAC)*.