Phase readout for satellite interferometry

Oliver Gerberding

2014

Phase readout for satellite interferometry

Von der Fakultät für Mathematik und Physik der
Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des Grades

Doktor der Naturwissenschaften
- Dr. rer. nat. -

genehmigte Dissertation
von

Dipl.-Phys. Oliver Gerberding

geboren am 11. April 1986
in Hannover, Deutschland

2014

"Bring the phasemeter. The toaster-looking thing."
"Yeah, I know what the phasemeter is."

Thor: The Dark World, 2013

Abstract

This thesis describes the development of digital phase readout systems, so-called phasemeters, required for performing precise length measurements in and between satellites with laser interferometry at frequencies below 1 Hz. These technologies have been studied in the scope of the planned space-borne gravitational wave detector LISA (Laser Interferometer Space Antenna), and of future satellite geodesy missions such as GRACE (Gravity Recovery and Climate Experiment) Follow-On. The studies presented here were conducted between 2010 and 2013 at the Albert Einstein Institute in Hannover, Germany.

The first part of this thesis provides a comprehensive overview of the basic concepts of inter-satellite interferometry. The analogue and digital parts of the phase measurement chain are described, with a focus on the design elements that are critical for achieving $\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ performance levels under the extreme conditions of the inter-satellite link. Digital signal simulations, as well as performance tests in analogue and optical set-ups with phasemeter prototypes are presented, which were used to determine the limiting effects with realistic signals. This includes tests of a precise angular readout technique, which were performed in the scope of GRACE Follow-On breadboarding activities. The developments presented culminate in the design, implementation and testing of an elegant breadboard model of the LISA phasemeter, which has been developed in the scope of an ESA (European Space Agency) technology development activity.

The second part of this thesis describes investigations for new interferometry techniques using state-of-the-art digital signal processing. These allow the simplification of optical designs and are therefore candidates for performing the local interferometry in future satellite missions. A real time phasemeter was developed for deep phase modulation, a readout scheme that uses strong phase modulations in a homodyne/heterodyne hybrid architecture to deduce the interferometric phase. Additionally, a new interferometry technique called Digital Interferometry was experimentally investigated in collaboration with the Australian National University. This scheme allows interferometric signals to be multiplexed through a single measurement chain, allowing fundamentally new interferometer architectures to be constructed. A dedicated phasemeter system and an optical set-up have been used to study its limitations and performance. Finally, these new schemes were compared to existing technologies.

Keywords: gravitational physics, interferometry, aerospace

Kurzfassung

Diese Doktorarbeit befasst sich mit der Entwicklung von digitalen Phasenmesssystemen für präzise Längenmessungen in und zwischen Satelliten mittels Laser Interferometrie bei Frequenzen unter 1 Hz. Diese Technologien wurden untersucht im Rahmen des geplanten Weltraum basierten Gravitationswellen Detektors LISA (Laser Interferometer Space Antenna) und zukünftiger Satelliten Geodäsie Missionen wie GRACE (Gravity Recovery and Climate Experiment) Follow-On. Die Arbeit wurde von 2010 bis 2013 am Albert Einstein Institut in Hannover, Deutschland angefertigt.

Der erste Teil der Arbeit gibt einen umfassenden Überblick über Laser Interferometrie zwischen Satelliten. Die kritischen Komponenten im analogen und digitalen Teil der Phasenmessung werden genauer beschrieben. Dabei liegt der Fokus auf den Elementen die kritisch sind um Rauschlevel von $\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ unter den anspruchsvollen Bedingungen der Interferometrie zwischen Satelliten zu erreichen. Digitale, analoge und optische Tests mit Prototypen solcher Phasenmesssysteme werden präsentiert, sie wurden durchgeführt um limitierende Effekte zu identifizieren die mit realistischen Signalen auftreten. Unter anderem wurde auch ein präzises Winkelmessverfahren im Rahmen von Vorbereitungsexperimenten für GRACE Follow-On getestet. All diese Entwicklungen gipfelten in dem Design, dem Bau und dem Testen eines Prototypen des Phasenmesssystems für LISA, welches im Rahmen eines Technologie Entwicklungsprojekts der Europäischen Raumfahrtbehörde (ESA) entwickelt wurde.

Der zweite Teil dieser Arbeit beschreibt Untersuchungen an neuen Interferometrie Techniken, die sich schnelle digitale Signalverarbeitung zu nutze machen. Diese Techniken ermöglichen es den optischen Aufbau von Interferometern stark zu vereinfachen und sie sind daher Kandidaten für interferometrische Messungen in zukünftigen Satelliten Missionen. Ein Echtzeit Auslesesystem wurde entwickelt für eine Technik die starke Phasenmodulationen verwendet um die optischen Phasen auszulesen. Gemeinsam mit der Australien Nation University wurde außerdem eine neue Technik untersucht, die es erlaubt mehrere interferometrische Signal mittels einer einzigen Phasenmesskette zu detektieren, was völlig neue Interferometer Topologien ermöglicht. Abschließend wurden diese neuen mit den vorhandenen Technologien verglichen.

Schlagworte: Gravitationsphysik, Interferometrie, Raumfahrt

Contents

Acronyms

AAF	anti-aliasing filter
AEI	Albert Einstein Institute
AC	alternating current
ADPLL	all-digital phase-locked loop
ADC	analogue-to-digital converter
AFE	analogue front-end
ANU	Australian National University
AOM	acusto-optic modulator
BPSK	binary phase shift key
CIC	cascaded integrator-comb
CPU	central processing unit
DAC	digital-to-analogue converter
DC	direct current
DDR	double data rate
DEHeI	digitally enhanced heterodyne interferometry
DEHoI	digitally enhanced homodyne interferometry
DFACS	drag-free and attitude control system
DI	digital interferometry
DLL	delay-locked loop
DPM	deep phase modulation
DSP	digital signal processing
DSS	digital signal simulator
DTU	Danish Technical University
DWS	differential wavefront sensing
EBB	elegant breadboard
eLISA	evolved LISA
EOAM	electro-optic amplitude modulator
EOM	electro-optic modulator
ESA	European Space Agency
FDS	frequency distribution subsystem
FFT	fast Fourier-transform
FIR	finite impulse response
FMC	FPGA mezzanine card
FPGA	field-programmable gate array
GFO	GRACE Follow-On

Contents

GOCE	Gravity Field and Steady-State Ocean Circulation Explorer
GRACE	Gravity Recovery and Climate Experiment
GRS	gravitational reference sensor
GW	gravitational wave
IFO	interferometer
IQ	in-phase and quadrature phase
IIR	infinite impulse response
InGaAs	indium gallium arsenide
IP	internet protocol
LISA	Laser Interferometer Space Antenna
LFSR	linear feedback shift register
LO	local oscillator
LRI	Laser Ranging Interferometer
LUT	look-up table
NASA	National Aeronautics and Space Administration
NCO	numerically controlled oscillator
NGO	New Gravitational wave Observatory
NPRO	non-planar ring oscillator
NSF	noise shape function
NTC	negative temperature coefficient thermistor
OB	optical bench
PA	phase accumulator
PAAM	point-ahead angle mechanism
PC	personal computer
PCB	printed circuit board
PL	pipeline
PLL	phase-locked loop
PMS	phase measurement system
PD	photo diode
PI	proportional-integral
PID	proportional-integral-differential
PIR	phase increment register
PT	pilot tone
QPD	quadrant photo diode
QPSK	quadrature phase shift key
PIR	phase increment register
PR	photo receiver
PRN	pseudo-random noise
RF	radio frequency
RIN	relative intensity noise
RMS	root mean square
RTEMS	Real-Time Executive for Multiprocessor Systems
SB	side band

SC	spacecraft
SER	serialiser
SG	signal generator
SIM	simulator
SMA	SubMiniature version A connector
SNR	signal-to-noise ratio
TA	thermoelectric assembly
TCP	transmission control protocol
TDI	time-delay interferometry
TIA	trans-impedance amplifier
TMA	triple-mirror assembly
TS	temperature sensor
UDP	user datagram protocol
VGA	variable gain amplifier
VHDL	very high speed integrated circuit hardware description language

List of Figures

Chapter 1

Introduction

Laser interferometry allows ultra precise measurements of length variations over long distances. It is therefore the tool of choice to track the relative motion of satellites or test masses within them. Precise knowledge of these motions makes it possible to determine the underlying forces of the system with high accuracy and can therefore be used to study them. This is especially useful for measuring small effects of gravity at low frequencies, which are not accessible with ground based devices. The development of the different components of such laser interferometers is an on going international effort. The two main applications for which these techniques are currently studied, are space based gravitational wave detection and satellite geodesy. This thesis describes investigations for the readout of such interferometers, that are applicable to both the inter-satellite interferometry, as well as the local one, that is used, for example, to measure the motion of test masses inside the satellites. Digital phase measurement systems, also called phasemeters, are used to track the interferometric phase signals in real time and with high precision. The development and testing of such phasemeters for different interferometer techniques and applications is the core subject of this thesis.

In the following the two applications of space based interferometers are introduced in some more detail. The first one is space based gravitational wave detection, addressed by the mission concept of the Laser Interferometer Space Antenna (LISA). The second application is satellite based geodesy, addressed by missions like GRACE Follow-on. The last section of this chapter presents the structure of this thesis.
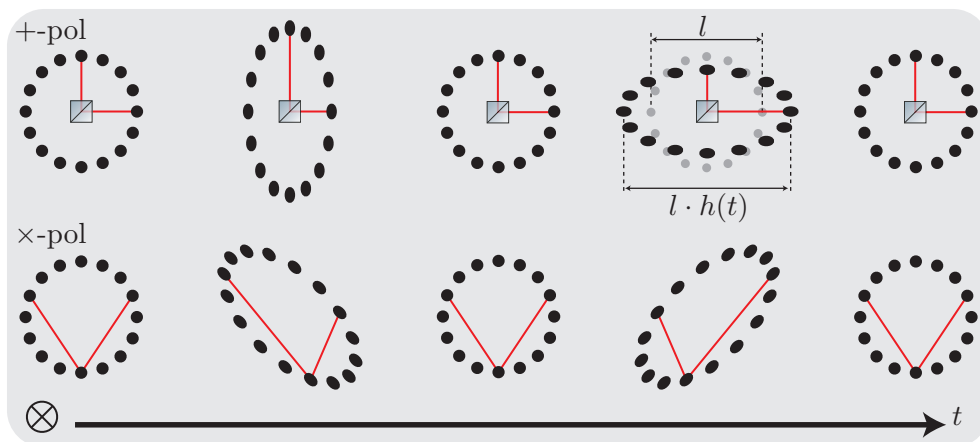
Figure 1.1.: Illustration of the two polarisations of gravitational waves by their effect on free test masses. Both waves are shown over one oscillation cycle. The induced relative length change is determined by the time dependent wave strain $h(t)$ [1].

1.1. Gravitational waves and the Laser Interferometer Space Antenna

Einstein's General Theory of Relativity [2] predicts the existence of gravitational waves (GWs). These ripples in space-time are generated by accelerated masses and they propagate almost freely through the universe [1]. They stretch and compress space-time perpendicular to their propagation direction and they exist in two polarisation states, as shown in Figure 1.1. The amplitude of the waves h, also called strain, is a dimensionless quantity that describes the induced relative length changes. Typical values of strain are on the order of $h = 10^{-21}$, which makes a direct detection of these deformations of space-time very demanding. However, gravitational waves have already been detected indirectly. Russell Hulse and Joseph H. Taylor were the first to observe a binary pulsar system which showed the exact energy loss that was predicted by the emission of gravitational waves [3].

A direct detection of gravitational waves would allow astronomers to deduce detailed information about their sources, which is encoded in the amplitude, phase and frequency of each wave. Access to these parameters makes it possible to use GWs as new astronomical messengers, which would complement the current observations that rely on electro-magnetic radiation and neutrinos. Since GWs can be generated by all massive objects, including electro-magnetically dark ones, new, unknown objects could also be revealed. The prospect of such a gravitational wave astronomy has led to the development of various detector types, each one sensitive for a different range of frequencies and therefore poten-
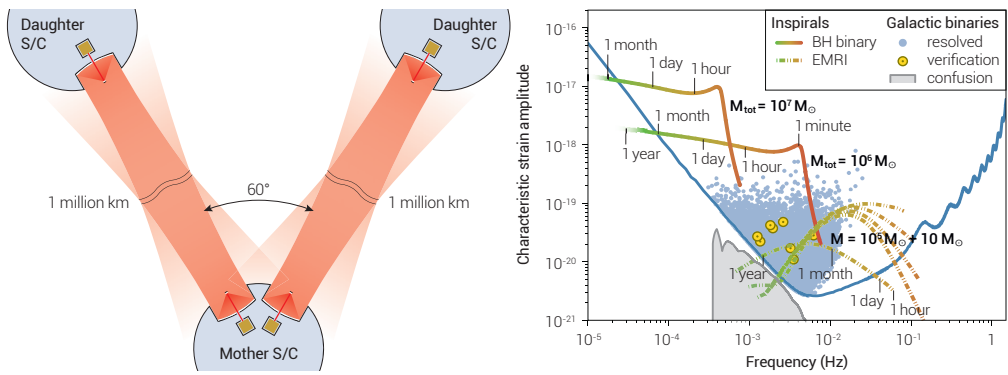
Figure 1.2.: The left picture shows the measurement scheme of eLISA. Laser Interferometry is used to determine the length changes in the arms between free floating test masses, separated by several million kilometres [9]. The right side shows the unit less characteristic strain amplitudes for eLISA source in comparison to its sensitivity. Some potential sources for eLISA are already well known from electro-magnetic observations, the so-called verification binaries.

tial sources. Besides space-based detection, which is described in the following paragraph, two detector classes are most relevant. The first are large ground-based laser interferometers, which aim to detect GWs in the audio frequency band (10 Hz to 10 kHz) [4]. Several of these detectors exist all around the globe and some of them are currently updated to reach higher sensitivities that will potentially allow the first direct detection [5, 6, 7]. The second promising detector class are pulsar-timing arrays [8], they track electro-magnetic radiation from known pulsars to deduce gravitational waves present at ultra low frequencies ($\approx 10^{-8}$–10^{-9} Hz).

Gravitational wave detection with space based instruments makes it possible to measure signals in a frequency range of 0.1 mHz to 1 Hz. A variety of GWs from interesting sources are expected in this band, including galactic binary systems, massive black hole binaries and extreme mass-ratio in-spirals [9]. The basic working principle of such detectors is to measure the relative motion of free floating test masses, hosted inside satellites, that have a separation of several million kilometres. Gravitational waves will induce length changes in these baselines that are read out via laser interferometry. A measurement noise floor of $\mathrm{pm}/\sqrt{\mathrm{Hz}}$ makes it possible to determine the relative arm length changes, and therefore the effective gravitational wave strain, with levels of $10^{-21}/\sqrt{\mathrm{Hz}}$. The existing detector designs utilise at least two such interferometer arms. This is in order to be able to reduce the influence of laser frequency noise, which would otherwise spoil the achievable performance. Different variants of this type of mission have been studied. A design for a joint mission between the European

Space Agency (ESA) and the National Aeronautics and Space Administration (NASA) was pursued until 2011 [10]. This specific concept is often referred to as the original LISA. In 2011 a re-scoped mission design was developed, called the New Gravitational wave Observatory (NGO) [11] or evolved LISA (eLISA). In 2013 the European Space Agency chose the science theme "The Gravitational Universe" for its third large mission in the Cosmic Vision program, which is aimed to be launched in 2034 (see Figure 1.2). This science case is based on the sensitivities achieved with the eLISA design, which makes it the likely candidate for the corresponding mission.

A technology-demonstrator mission for LISA, called LISA Pathfinder, will be launched in 2015. It will test some of the key technologies required for LISA, including drag-free control of free floating test masses and laser interferometric position readout with $\mathrm{pm}/\sqrt{\mathrm{Hz}}$ precision [12, 13].

The developments for inter-satellite interferometry presented in this thesis were mainly focused on the original LISA design. They can, however, also be applied for eLISA or future, even further evolved, variants of the mission. Amongst other things, this thesis includes the results of a technology development activity of the LISA metrology system, funded by the European Space Agency. The new techniques investigated in this thesis for local interferometry could also be utilised in future designs of eLISA or similar missions.

1.2. Satellite geodesy and GRACE Follow-On

The Earth's gravity field is influenced by multiple effects, including tides, climate, earth quakes, water reservoirs and more. The field is often displayed as a geoid, an equipotential surface map of the Earth (an example is shown in Figure 1.3). A measurement of the geoid, and its variation over time, can be used to deduce the underlying mass distribution and its variations, which are often inaccessible otherwise. This information can be applied in earth sciences like geophysics, geology, physical geography and hydrology. One way to determine the gravity field is by probing the forces acting on satellites or test masses orbiting the planet. This principle is called satellite geodesy [14], and the relevant method to measure these forces is to track the relative motion between two or more masses in "free flight" around the earth that are ideally only influenced by gravity. This concept has already been utilised in different missions: The Gravity Recovery and Climate Experiment (GRACE) monitored the relative motion of two satellites, acting as test masses, with a microwave ranging instrument [15]. The Gravity Field and Steady-State Ocean Circulation Explorer (GOCE) utilised several test masses located within a single satellite [16, 17]. Their motion was read out in three dimensions via capacitive sensors.

The large success of these missions has led to an international effort to develop technologies that will allow such gravity measurements to be continued
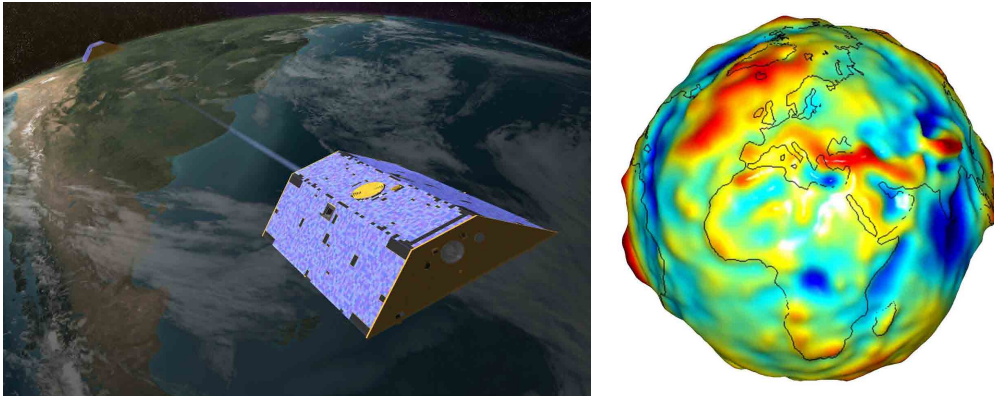
Figure 1.3.: The left picture shows an artists impression of the two GRACE satellites orbiting Earth. The right side shows the first geoid calculated from GRACE data. (source: http://science.nasa.gov/missions/grace/)

with even higher precision [18, 19, 20]. To achieve this the sensitivity of the utilised instruments, as well as the data processing and modelling of the gravity fields has to be improved. The aspect that is relevant for this thesis is the improvement of the motion sensing with the use of precision laser interferometry. Designs of future missions might utilise inter-satellite laser interferometry and local interferometry for test mass readout, or both of these technologies.

A fast continuation of the GRACE mission called GRACE Follow-On (GFO) is currently under development. The main goal of GFO is to continue the measurements of the geoid time variations with unchanged overall sensitivity. In addition to the original microwave instrument, GFO will include a new inter-satellite ranging system, the Laser Ranging Interferometer (LRI), which will be used to demonstrate the feasibility and performance of this technology for future missions. This instrument will potentially be the first working implementation of an inter-satellite laser interferometer. Phase readout systems utilised in interferometer breadboarding activities for GFO, conducted at the Albert Einstein Institute, were developed as part of this thesis.

1.3. Thesis structure

The first part of this thesis focuses on the readout of inter-satellite interferometers. The second part discusses investigations of new and alternative techniques that can be utilised in future missions to perform the local interferometry, like the readout of test mass positions.

Part I begins with the introduction of inter-satellite laser interferometry and its basic concepts in Chapter 2. This also includes an overview of the specific implementations for LISA and GRACE Follow-On, as well as an introduction

to the phase readout principle. The details of the phase measurement chain are described in the next two chapters. Chapter 3 focuses on the analogue components. Chapter 4 discusses the digital signal processing algorithms. Principle methods of testing phase readout systems are discussed in Chapter 5. The results of simulations to investigate the performance limits of the phase readout digital signal processing algorithms are presented in Chapter 6. The implementation of a phasemeter system for laboratory experiment for GRACE Follow-On and some results from optical experiments are shown in Chapter 7. A number of smaller experimental investigations was carried out to prepare the development of a phasemeter for LISA. These are summarised in Chapter 8, which also introduces some implementation details as well as first results of advanced phase measurement test set-ups using analogue and optical signals. Chapter 9 describes the elegant breadboard model of the LISA phasemeter that was developed as part of an ESA technology development activity. The system is presented with a focus on critical elements together with the achieved test results. Part I concludes with an outlook on future tests and aspects that still have to be investigated to validate and prepare the phase measurement systems for the actual missions in Chapter 10.

Part II first describes an interferometer scheme using strong phase modulations in Chapter 11. A real-time readout system was developed for this technique, which is a potential alternative to the heterodyne systems used in LISA and LISA Pathfinder. A new type of interferometry, that uses digitally generated phase modulations to multiplex interferometric signals, was also investigated. Chapter 12 gives an overview about this multiplexing scheme, which makes it possible to construct fundamentally simpler interferometers, and presents results from an optical experiment. A comparison between these new interferometer schemes and classic heterodyne techniques is drawn in Chapter 13, with the readout of a test mass in two directions as an example case for demonstrating the advantages of multiplexed interferometry.

Finally, the presented results and findings of this thesis are summarised in Chapter 14.

Part I.

Readout of inter-satellite laser interferometers

Chapter 2

Inter-satellite laser interferometry

This chapter introduces the fundamentals of distance measurements between satellites using laser interferometry. This thesis focuses on the architecture that was initially developed for LISA and parts of which have later been adopted for GRACE Follow-On (GFO). Other conceptual designs are also currently studied [21, 22], but they are not discussed here. In general the architecture presented here has been studied for missions with large satellite distances ($> 100\,\text{km}$), large relative speeds ($> 10\,\text{m/s}$) and a measurement band below $1\,\text{Hz}$.

First, a description of the optical signals is given, introducing concepts and vocabulary that is used in this thesis and the referenced literature. This is followed by the basic concept of a heterodyne Mach-Zehnder interferometer that is, for example, used in laboratory experiments. This allows an easy access to the topic and it sets the stage for the comparison with other interferometry techniques investigated in the second part of this thesis. The readout principle of such an interferometer (IFO) is described in the third section and it is relevant throughout the full first part of the thesis. Inter-satellite interferometry is then finally introduced in the fourth section and the required techniques are described together with the most relevant noise sources. The laser ranging instrument developed for GRACE Follow-On is then presented as the first example of a specific implementations. Finally, the LISA interferometry, the central application aimed at in this thesis, is described and additionally utilised techniques are introduced that are necessary for the complex LISA metrology concept.

2.1. Optical signals

There are two main reasons for utilising laser beams for interferometric distance measurements. First, their small wavelength contributes a huge amplification factor of distance variations into respective phase changes that allows length fluctuations to be measured with high precision. Second, the high achievable stabilities of laser beams and sources make them also the best available rulers for distance variation measurements.

In the following the properties of laser beams are summarised. For achieving the measurement precision desired for gravitational wave detection, either in space or on ground, all of these parameters have to be taken into account

- Wavelength: Monochromatic laser light at a wavelength $\lambda = 1064\,\text{nm}$ is assumed in this thesis and is the current baseline for GFO and LISA. This is due to the fact that the laser sources at this wavelength (especially non-planar ring oscillator (NPRO) lasers [23, 24]) offer high stabilities and sufficient optical power with an overall acceptable power consumption. Equally suitable sources with different wavelengths are currently under research and might become an option in future studies. For $1064\,\text{nm}$ a measured phase φ can be converted into a distance variation l with $l = \frac{\varphi}{2\pi} 1064\,\text{nm}$.

- Frequency and phase: Corresponding to the wavelength a laser beam at $1064\,\text{nm}$ has a carrier frequency of $f = c/\lambda \approx 282\,\text{THz}$. Even though the light is considered monochromatic, the frequency and the corresponding angular frequency ω_x of any real laser beam are not constant. Their variations are what is often denoted as laser frequency noise or laser line width. Therefore each beam also carries a phase $\varphi_x(t)$ that depends on its frequency noise and its propagation. Frequency and phase can both be used to describe the phase of a laser beam and their relations are:

$$\varphi_x(t) = \int_0^\infty \omega_x(t)\ dt = 2\pi \int_0^\infty f_x(t)\ dt, \tag{2.1}$$

$$f_x(t) = \frac{\omega_x(t)}{2\pi} = \frac{1}{2\pi}\frac{d\varphi_x}{dt}. \tag{2.2}$$

- Propagation direction: Each beam has a propagation direction defined in three dimensions. The beam direction and its wavelength are also encoded in the wave vector $\vec{k}(x, y, z)$.

- Beam parameters: Gaussian optics is the fundamental theory used to describe laser beams in interferometers. In addition to its direction each beam has a waist size w_0 and waist position z_0 along its propagation.

- Transverse modes: Two types of transverse beam shapes are most relevant for inter-satellite interferometry and this thesis, fundamental Gaussian modes and flat top beams. For laboratory experiments and beams on a local spacecraft purely Gaussian modes are assumed, delivered either by a laser source or an optical fibre. Flat top beams are assumed for light being received from a distant satellite. The aperture at a receiving spacecraft (SC) will cut out a small part of the almost flat wave front that was generated by the divergence due to the long beam propagation. The most relevant property of flat top beams for this thesis is, that their interference with a local Gaussian beam is not ideal and will reduce the effective interference contrast.

- Optical power: Each laser beam has an oscillating electro-magnetic field with an amplitude E_x and a corresponding optical power $P_x \propto |E_x|^2$. Photo diodes are too slow to detect 100s of THz and only sense $\langle |E_x|^2 \rangle$ averaged over many periods.

- Polarisation: The direction of the oscillation of the electric field of the beam is defined as its polarisation. It is always perpendicular to the propagation direction and only components of the same polarisation can interfere with each other.

The above description of the geometrical and optical aspects of inter-satellite interferometry are very brief. Their full complexity and influence on mission design is a topic of elaborate research [25, 26, 27, 28, 29, 30] with an influence at least as relevant as the *readout* aspects focused on in this thesis. The apparent separation into two topics is done here purely for simplification and readability, it is not representative of the full analysis and planning required for a precision interferometer.

2.2. Heterodyne Mach-Zehnder interferometry

One of the fundamental types of interferometers is a Mach-Zehnder IFO. A basic set-up using a heterodyne scheme is depicted in Figure 2.1. Two laser beams with frequencies ω_1 and ω_2 are split into two interferometers. The laser frequencies can either be generated by using two lasers (like shown in Figure 2.1) or, for example, by using acousto-optic-modulators, as done in LISA Pathfinder [31]. The two beams are once interfered directly, to generate a reference phase φ_r, and they are interfered a second time, after one beam has propagated through a system under test, where it collected a measurement phase φ_m, that typically encodes the path length variations to be measured.

For our purposes we simplify the full description of a propagating electro-
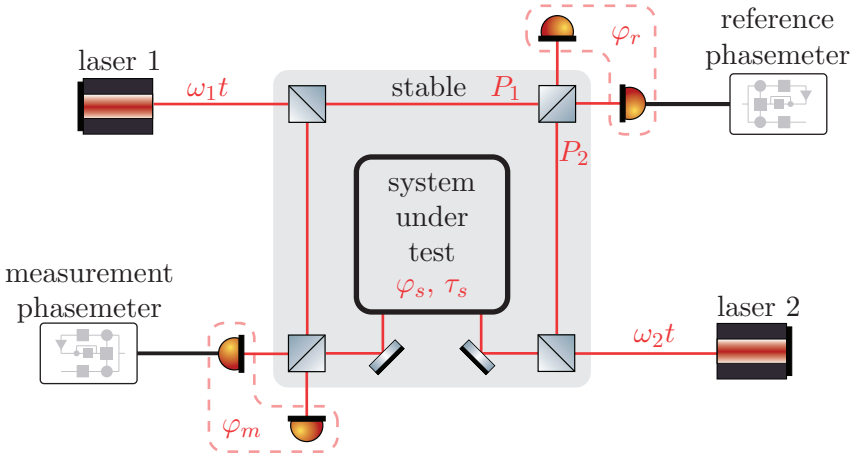
Figure 2.1.: Sketch of a general heterodyne Mach-Zehnder interferometer for measuring length changes in a system under test. The set-up requires two interferometers, one for the actual measurement (φ_{m}) and one for the reference measurement (φ_{r}). The shown example uses two lasers and two phasemeters for the phase readout. In principle only a single IFO could be used, but in practice the phase stability of the light entering the IFO is not sufficient and therefore a local reference is required. For achieving high sensitivities, all beam paths, after the splitting into the two interferometers, have to be sufficiently stable and noise free.

magnetic wave for beam $\vec{E}_1$ to

$$\vec{E}_1(x, y, z, t) = E_1 \, e^{i(\varphi_1(t))} \, \vec{g}_1(x, y, z). \tag{2.3}$$

This description includes the strength of the electric field E_1, an oscillating term that includes the phase and a geometrical term $\vec{g}_1$, a normalised electric field vector that contains the geometrical and optical properties of the beam. This term is complex, since it also contains phase information that becomes relevant later on.

The conversion of the phase difference between two beams into an electronic signal is now done in a two step process, explained by the example of the reference interferometer. First, the beams are combined using a beam splitter, which splits each beams power depending on the power reflectivity r and introduces a 90° phase shift for the transmitted beam [4]. The field at one output of the beam splitter can then be written as

$$\vec{E}_{\mathrm{r}}(x, y, z, t) = i \sqrt{t} \, \vec{E}_1(x, y, z, t) + \sqrt{r} \, \vec{E}_2(x, y, z, t) \tag{2.4}$$

Now, in the second step, the signal is detected by a large single element photo diode (PD), which converts the intensity of the electric fields into a photo cur-

rent. The intensity for square-law detection and for a power reflectivity of $r = 0.5$ is given as

$$I_r(x, y, z, t) = \frac{c\varepsilon_0 n}{2} \left| i \sqrt{\frac{1}{2}} \vec{E}_1(x, y, z, t) + \sqrt{\frac{1}{2}} \vec{E}_2(x, y, z, t) \right|^2. \qquad (2.5)$$

Here c is the speed of light, ε_0 is the vacuum permittivity and n the refractive index of the propagation medium. The photo diode integrates this intensity over its active area [25], this leads to an effective received optical power of

$$P_r(t) = \frac{c\varepsilon_0 n}{2} \left(\frac{E_1^2}{2} + \frac{E_2^2}{2} + g_{1,2} E_1 E_2 \cos(\varphi_1 - \varphi_2 + \varphi_{g_{1,2}}) \right) \qquad (2.6)$$

Here the integrated geometrical factors contributed to the amplitude of the interference pattern $(g_{1,2})$ as well as to the phase difference $(\varphi_{g_{1,2}})$. They do however not contribute to the non-phase dependent terms, which is valid as long as the active area of the PD is large compared to the beam size. The geometrical contributions to the phase are neglected in the following. Finally we can compute the generated photo current of the reference interferometer and write it in a more common format

$$i_r(t) = R_{\mathrm{PD}} \left(\frac{P_1 + P_2}{2} + \sqrt{\eta_e P_1 P_2} \cos\left[\varphi_r(t) \right] \right). \qquad (2.7)$$

Here the electric fields were replaced with the respective beam power impending on the interference beam splitter. The geometry influence was rewritten as heterodyne efficiency η_e and the current was computed by multiplying the received power with the photo diode responsivity R_{PD}, which describes the efficiency of the diode to convert photons into electrons ($[R_{\mathrm{PD}}] = \mathrm{A/W}$). The detected phase difference is now denoted as the reference phase $\varphi_r = \varphi_1 - \varphi_2$.

For everyday use this formula is often simplified by denoting the constant photo current on the photo diode i^{DC} and the contrast c, which determined the amount of signal that is influenced by fluctuations of the signal phase. Together this leads to a photo current

$$i_r(t) = i_r^{\mathrm{DC}} \left(1 + c \cos\left[\varphi_r(t) \right] \right) = i_r^{\mathrm{DC}} + i_r^{\mathrm{AC}}(t). \qquad (2.8)$$

The notation also introduces the convention to split the signal into two parts, one unmodulated by phase changes (i^{DC}), the DC signal and one modulated (i^{AC}), the AC signal or interferometric beat note. For the interferometers described in this thesis the AC signals are at frequencies in the MHz range, allowing a clear distinction between AC and DC parts. Though the reader should be reminded, that this partition is arbitrary. For different interferometer configurations, namely homodyne schemes, the above distinction between AC and DC signals is not suitable.

The photo current is now converted into a voltage by the use of a trans-impedance amplifier (TIA). Further amplification, which is often separated for DC and AC signals, might also be necessary. In general we can include this into our computation by assuming a single complex impedance Z, that defines an amplitude and phase response

$$v_r(t) = \mathcal{L}^{-1}(Z(s) \cdot i_r(s)) \tag{2.9}$$

For now we assume a constant real impedance $Z = Z_0$. Later on we will further investigate the properties of the impedance by using its Laplace transform $Z(s)$, allowing us to describe its amplitude and, more importantly, phase dependency of the beat note frequency.

The phase information, our signal of interest, is decoded in the AC signals coming from the photo diode. The extraction of the phase term from 2.9 is described in the next section. With no loss of generality we can define the reference phase to depend only on the two laser frequencies. While we assume ω_1 to be perfectly stable, we define ω_2 as time-dependent.

$$\varphi_r(t) = [\omega_1 - \omega_2(t)]\, t = \omega_0(t)t. \tag{2.10}$$

By doing this and otherwise assuming the same interference condition at the measurement interferometer its detected phase can be written as

$$\varphi_m(t) = [\omega_1 - \omega_2(t + \tau_s)]\, t - \varphi_s(t) = \omega_0(t)t - \varphi_s(t). \tag{2.11}$$

Both phase signals contain a constant phase ramp at the frequency difference ω_0, which gives rise to the oscillation of the optical signals. The phase generated in the system under test can now be determined by computing the phase difference $\varphi_d(t)$ between both signals,

$$\varphi_d(t) = \varphi_r(t) - \varphi_m(t) = \varphi_s(t) + [\omega_2(t + \tau_s) - \omega_2(t)]\, t. \tag{2.12}$$

If one assumes a constant $\omega_2(t)$ (basically no laser frequency noise or initial phase noise on the laser), the determined phase φ_d equals the phase signal φ_s and that allows the phase changes in the system under test to be tracked.

The case for heterodyne interferometer

Heterodyne schemes are used in inter-satellite interferometry, and in general, because they allow interferometers to measure very large phase signals with $\varphi \gg 2\pi$. Interferometer schemes using homodyne detection, like ground based gravitational wave detectors [32], can achieve much lower noise levels, using strong light beams and large, quiet mirrors, but they can only operate in a range much smaller than a single fringe ($\varphi \ll 2\pi$).
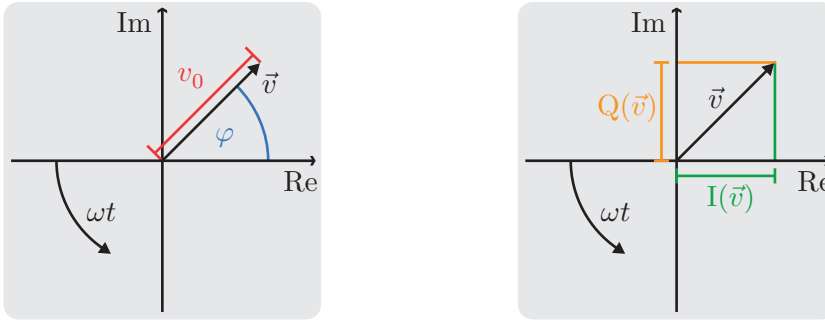
Figure 2.2.: Phasor diagram for signal 2.13. The signal can be projected onto the real and imaginary axis and thereby the phase and amplitude can be deduced. The phasor is always shown in the frame of the beat note frequency ω.

2.3. Readout principle of heterodyne interferometers

The phase of the beat note generated by the TIA can be extracted by digitising it and then using an in-phase, quadrature-phase (IQ) demodulation technique. At this point we discard the digitisation for simplicity and describe the process with a generalised input signal with constant frequency ω,

$$v(t) = v_0 \cos\left[\varphi(t) + \omega t\right]. \tag{2.13}$$

Since the input signal is defined by an amplitude v_0 and a phase φ, one can picture it as a complex number or vector $\vec{v}$ in a coordinate system rotating at the signal frequency, as shown in Figure 2.2. If we now project $\vec{v}$ onto the real and imaginary axis we can determine the in-phase component $I_\omega(\vec{v})$ and the quadrature component $Q_\omega(\vec{v})$,

$$\begin{aligned} I_\omega(v(t)) &= v_0 \cos\left[\varphi(t)\right], \\ Q_\omega(v(t)) &= v_0 \sin\left[\varphi(t)\right]. \end{aligned} \tag{2.14}$$

From these quadrature components one can then again calculate the signal amplitude and phase

$$\begin{aligned} v_0 &= \sqrt{I_\omega^2 + Q_\omega^2}, \\ \varphi(t) &= \arctan\frac{Q_\omega}{I_\omega}. \end{aligned} \tag{2.15}$$

The projection of the vector onto the real and imaginary axis is done by multiplying with a sine and a cosine oscillating at ω and by a subsequent filtering of terms at 2ω, as shown on the left side of Figure 2.3. Depending on the dynamics of the input signals this readout scheme is already sufficient and it is,
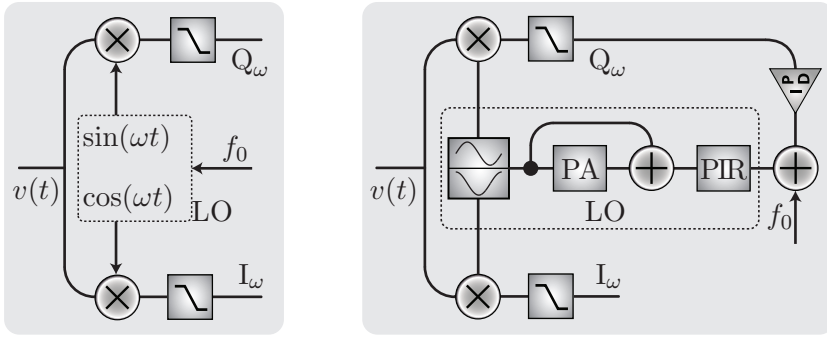
Figure 2.3.: The left side shows the IQ demodulation principle used for the readout of heterodyne interferometers with fixed signal frequencies like LISA Pathfinder. Sine and cosine generated from a local oscillator (LO) are mixed with the incoming signal and the product is then filtered to exclude the second harmonic terms. The right side shows the extension to a full phase-locked loop, in which the quadrature signal is amplified in a servo controller and a resulting actuator signal is fed into the local oscillator. The current frequency of the LO in a digital implementation is stored in the phase increment register (PIR) and its current phase in the phase accumulator (PA) [36].

for example, used in LISA Pathfinder [33] and other experiments [34]. For input signals with a large dynamic range, like the distance measurement between moving satellites, this readout style is not sufficient. The demodulation and signal frequency will deviate from each other and therefore introduce a phase error for marginal detuning [35] and a complete signal loss for large detuning.

The IQ demodulation is therefore extended to a phase-locked loop (PLL), a feedback loop that controls the demodulation frequency and phase to track the incoming signal [37]. An overview of such a PLL is shown on the right side of Figure 2.3. The feedback loop actuates the local oscillator frequency to keep the quadrature signal small, which minimises the phase difference between the signal and the LO ($\varphi \ll 1$). Therefore the phase information is transferred from Q_ω to the actuator signal, which in this case is the phase of the local oscillator. While the PLL is locked I_ω also yields a value directly proportional to the input signal amplitude, as can be seen from Equation 2.14. A detailed model of the PLL is given in Chapter 4. The device implementing the beat note digitisation and the phase readout algorithm is called a phase measurement system or phasemeter.

2.4. Inter-satellite interferometry

For implementing an interferometric distance measurement between satellites a master-transponder scheme [38, 20] is the current baseline architecture for LISA and GFO [9, 39]. Due the heritage of this scheme from the developments for
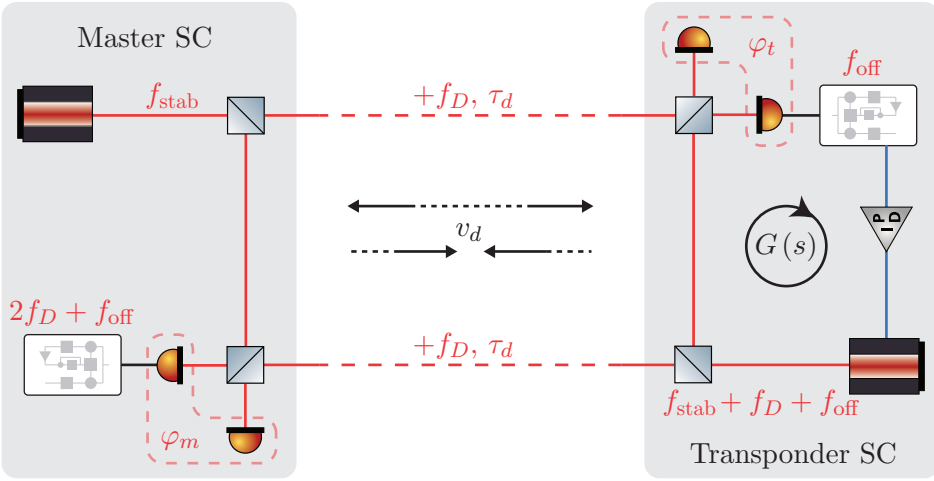
Figure 2.4.: Sketch of the LISA-arm interferometry concept. The shown optical layout is simplified (LISA, for example, uses a single optical axes for beam transmission and receiving). The critical parameters are the satellite distance (here as light travel time τ_d) and the relative satellite speeds v_d, which determines the Doppler shifts. The laser in the master satellite is usually pre-stabilised (not shown).

LISA it also sometime referred to as *LISA arm interferometry* [10].

The interferometric link between two satellites is established by using two light sources, one on each spacecraft as shown in Figure 2.4. This allows operating the link even for large distances, which cause a strong reduction of the received power due to beam divergence. Using only a single light source, together, with a passive reflection is not feasible, since the received laser power after a full round trip would be to small to establish a continuous link and to achieve the desired sensitivities.

One of the two laser sources acts as a master, while the other source on the distant spacecraft acts as a transponder. This is done by implementing a phase lock between the received light and the transponder laser with a frequency offset f_{off} that allows all signal frequencies to be kept within the measurement band. For both LISA and GFO the interferometric beat note frequencies are in a range between 2 MHz and 25 MHz, depending on their precise orbits. This range is mostly driven by the large frequency Doppler-shifts f_D induced by relative satellite speeds v_d.

$$f_D = \frac{v_d}{c} f_0 = \frac{v_d}{\lambda_0}$$

$$f_D^{1064\,\text{nm}} \approx 1\,\text{MHz} \cdot \left(\frac{v_d}{1\,\text{m/s}} \right)$$

$$(2.16)$$

The detection band must accommodate for these frequency changes and a frequency plan must guarantee that the beat notes are always within a suitable range. While the frequency at the transponder is, for example, locked to a fixed value the beat note on the master spacecraft will be at a frequency given as two times the induced Doppler-shift plus the offset-lock frequency. Even though the detection treats all frequencies positive, the sign of each beat note needs to be known to later on compute the correct phases. The measured Doppler-shifts contain the desired information of the relative spacecraft movement, therefore this technique is also referred to as *Doppler-ranging*.

2.4.1. Laser frequency lock

The performance of the offset frequency phase lock is in general uncritical for the overall noise floor. Its residual phase error is also tracked and can therefore be used to correct the signal in post processing. The offset lock is usually implemented by the phasemeter. The measured beat note frequency signal is compared to a set-point and then fed via an additional controller back to the frequency actuators of the laser source [40, 41], as indicated in Figure 2.4. The open loop gain of this feedback control is denoted as $G(s)$. More details on the implementation of such a system are given in Chapter 9.

2.4.2. Laser frequency Noise

As already shown in Section 2.2, the instability of the laser frequencies directly couples into measurements with unequal light travel times. Using transfer functions one can deduce the effective, frequency dependent, coupling of frequency noise into the measurement for the master transponder scheme. The master usually has a pre-stabilised laser with a frequency noise term $\widetilde{f}_{\text{stab}}$, while the transponder has a free running laser with initial noise $\widetilde{f}_{\text{free}}$. The effective frequency noise measured at the transponder consists of the noise from the master, delayed by the light travel time τ_d, and the noise of the transponder. Both of them are attenuated by the offset phase-lock.

$$\widetilde{f}_t = \frac{\widetilde{f}_{\text{stab}}\, e^{-s\tau_d} - \widetilde{f}_{\text{free}}}{1 + G(s)} \tag{2.17}$$

The effective noise at the master therefore consists of three terms, the initial noise $\widetilde{f}_{\text{stab}}$ and the two terms from the transponder delayed by τ_d.

$$\widetilde{f}_m = \widetilde{f}_{\text{stab}} \left(\frac{G(s)e^{-2s\tau_d}}{1 + G(s)} - 1 \right) + \widetilde{f}_{\text{free}} \frac{e^{-s\tau_d}}{1 + G(s)} \tag{2.18}$$

If one assumes an offset frequency lock with a high gain in the measurement band (at low frequencies) the equations can be approximated. $\widetilde{f}_t$ becomes negligibly

small, while $\widetilde{f}_m$ simplifies to

$$\widetilde{f}_m = \widetilde{f}_{\text{stab}} \left(e^{-2s\tau_d} - 1 \right). \tag{2.19}$$

This shows that the effective frequency noise coupled into the measurement is completely dominated by the stability of the master laser and the distance between the satellites. If the light travel time is short in comparison to the measurement band frequencies Equation 2.19 can be further approximated. Considering that a frequency noise spectrum $\widetilde{f}(s)$ can be converted into a phase noise spectrum by integration ($\widetilde{\varphi}(s) = \frac{\widetilde{f}(s)}{f}$) and that the phase is directly proportional to length $\widetilde{l}$ one can deduce the effective length measurement noise $\widetilde{l}$. The phase measured here corresponds to two times the satellite distance, since it measures the induced Doppler-shift twice.

$$\begin{aligned}
\widetilde{l} &= \frac{\lambda_0}{4\pi} \frac{\widetilde{f}_{\text{stab}}}{f} \left(e^{-2s\tau_d} - 1 \right) \\
&= \frac{\lambda_0}{4\pi} \frac{\widetilde{f}_{\text{stab}}}{f} (4\pi f \tau_d) \\
&= \lambda_0 \tau_d \widetilde{f}_{\text{stab}} = \frac{l_d}{f_0} \widetilde{f}_{\text{stab}}
\end{aligned} \tag{2.20}$$

This is a typical formula for the coupling of laser frequency noise, it is applicable for laboratory experiments and even for GFO. One should note that the coupling of frequency noise due to delay of $2\tau_d$ reduces to the simpler coupling for the final length measurement, since the length is measured twice. Only for LISA the light travel time is so large that the above approximations are not suitable any more to deduce the frequency noise inside the measurement band (up to 1 Hz), since the light travel times are larger than 1 s.

As will be shown in Chapter 3 and Chapter 4 the design of the readout system depends not only on the noise inside the measurement band but on the whole spectrum. Therefore the original Equation 2.18 and Equation 2.17 have to be evaluated using an expected frequency noise spectrum. Figure 2.5 shows the free running frequency noise assumed in this thesis $\widetilde{f}_{\text{free}}$ in comparison to measurements of available lasers.

The strong influence of laser frequency noise on the achievable performance requires the master laser to be sufficiently pre-stabilised, both in LISA and GFO. For the design and tests in this thesis an effective frequency noise at the master $\widetilde{f}_m$ was chosen as shown in Figure 2.5. It is comparable to real laser sources at high frequencies and it has a very high frequency noise at low frequencies in comparison to proposed stabilisation requirements for LISA and GFO. This conservative shape was chosen to ensure that all later described experiments are compatible with a large variety of laser sources.
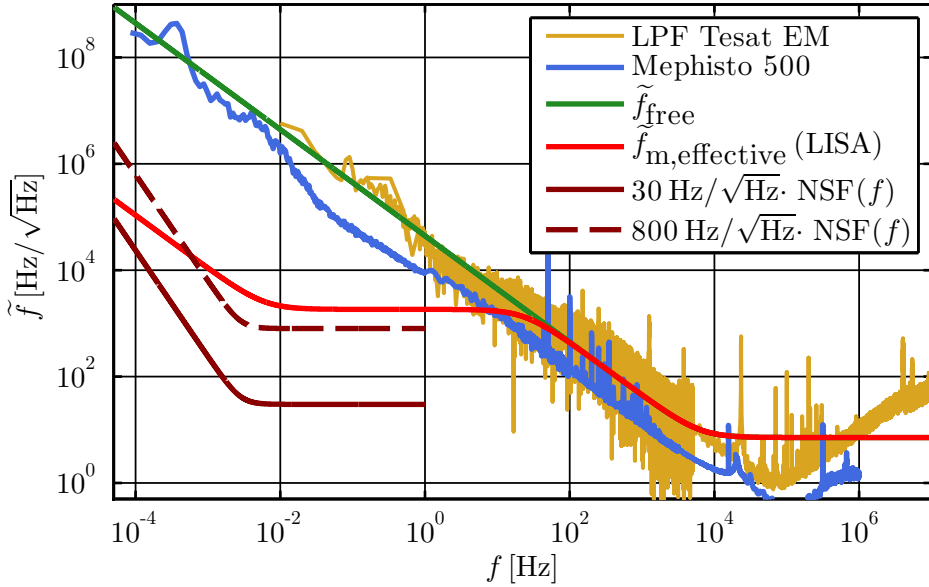
Figure 2.5.: The plot shows the free running frequency noise measured with an Innolight Mephisto 500 laser [24] used in laboratory experiments in comparison to the engineering model of the LISA pathfinder laser from Tesat [42]. Based on these two a free running noise was defined for the test in this thesis. The high frequency behaviour of both measurements is dominated by read-out effects and the assumed free running noise curve was based on separate measurements at high frequencies done by Kullmann [43]. The effective laser noise measured on the master satellite and used for performance simulations is also shown in comparison to two considered levels of pre-stabilisation.

2.4.3. Shot noise and signal to noise density

A signal beam received from a distant spacecraft with a power of P_s is in general interfered with a local oscillator beam with a power of P_{LO}. Both powers are measured in front of the beam splitter, which we here assume to split 50/50. As described earlier the received power is much smaller than the local oscillator, and in this low-power case we can rewrite Equation 2.7 as

$$
\begin{aligned}
i_s(t) =& R_{PD} \left(\frac{P_{LO}}{2} + \sqrt{\eta_e P_s P_{LO}} \cos\left[\varphi_r(t)\right] \right) \\
=& i^{DC} + i_s^{AC} \cdot \cos\left[\varphi_r(t)\right].
\end{aligned}
\tag{2.21}
$$

The DC current is fully dominated by the local oscillator, since the received signal power is negligible in comparison (e.g. 100 pW vs. a few mW). Due to the quantum nature of light the photons of the LO have a statistical distribution,

causing a random signal called shot noise [4]. This noise has a normal distribution and a white spectrum, i.e power equally distributed over all frequencies. The shot noise induced current noise density is

$$\widetilde{i}(f) = \sqrt{2q_e R_{\mathrm{PD}}\frac{P_{\mathrm{LO}}}{2}} = \sqrt{2q_e i_s^{\mathrm{DC}}}, \tag{2.22}$$

where q_e is the elementary charge ($q_e \approx 1.602 \times 10^{-19}\,\mathrm{C}$). The presence of $\widetilde{i}(f)$ introduces a fundamental limit to the phase measurement performance given by the ratio between the noise and the signal amplitude times $\sqrt{2}$. This additional factor can be understood to be caused by the folding of white noise from positive and negative frequencies, adding two uncorrelated noises, during the mixing process in the phasemeter [44].

$$\widetilde{\varphi}_{\mathrm{shot}}(f) = \sqrt{2}\frac{\widetilde{i}(f)}{i_s^{\mathrm{AC}}}\mathrm{rad} = \frac{\widetilde{i}(f)}{i_{s,\mathrm{RMS}}^{\mathrm{AC}}}\mathrm{rad} = \sqrt{\frac{2q_e}{\eta_e R_{\mathrm{PD}} P_s}}\mathrm{rad}. \tag{2.23}$$

The factor $\sqrt{2}$ can also be accounted for by using the root-mean-square (RMS) amplitude of the input signal ($i_{s,\mathrm{RMS}}^{\mathrm{AC}} = i_s^{\mathrm{AC}}/\sqrt{2}$). The phase measurement noise limit can also be described by the inverse ratio, which is equal to the carrier-to-noise density C/N_0 for shot noise with units of dBHz.

$$C/N_0 = 20\log_{10}\left(\frac{i_{s,\mathrm{RMS}}^{\mathrm{AC}}}{\widetilde{i}(f)}\right) = 10\log_{10}\left(\frac{\eta_e R_{\mathrm{PD}} P_s}{2\sqrt{2}q_e}\right) = 20\log_{10}\left(\frac{1}{\widetilde{\varphi}_{\mathrm{shot}}}\right) \tag{2.24}$$

Even though the power of the local oscillator defines the absolute level of shot noise, its ratio to the signal, and therefore the phase noise performance is only limited by the received signal power [44]. One should note that this is only true in the low-power case, where one beam is much weaker than the other.

In addition to shot noise two other additive noise sources are important, laser amplitude noise (described in the following section) and electronic noise (see Chapter 3). For all these in principle the same coupling mechanism into the phase measurement applies, they are undistinguishable for the phasemeter system. Therefore it is often useful to add these noises and to thereby define a single carrier-to-noise density to describe the input signal. LISA is designed so that its C/N_0 is limited by shot noise, making it the dominating effect by design.

2.4.4. Laser amplitude noise

So far the optical power of each beam was assumed to be constant. Real laser sources produce a non negligible amount of technical amplitude noise in addition to shot noise. That noise is often characterised as relative intensity noise (RIN), a spectral density of the power fluctuations normalised to the average power

[45]. The coupling of this amplitude noise into the phase measurement can be understood by extending Equation 2.21 with the RIN of the local oscillator [46].

$$
\begin{aligned}
i_s(t) =& R_{\mathrm{PD}} \left(\frac{P_{\mathrm{LO}} + \widetilde{P}_{\mathrm{LO}}}{2} + \sqrt{\eta_e P_s P_{\mathrm{LO}}} \cos\left[\varphi_r(t)\right] \right) \\
=& i^{\mathrm{DC}} + \widetilde{i}^{\mathrm{\,RIN}} + i_s^{\mathrm{AC}} \cdot \cos\left[\varphi_r(t)\right].
\end{aligned}
\tag{2.25}
$$

The additional current noise $\widetilde{i}^{\,\mathrm{RIN}}$ can now also limit the carrier-to-noise density of the signal. In contrast to shot noise the induced phase noise increases with the average power of the local oscillator P_{LO}, since the amplitude noise scales linearly, while the signal only scales with the square root. Amplitude noise of available laser sources also has a large frequency dependency and it usually increases at lower frequencies. To allow for a shot noise limited measurement the RIN of the LO therefore has to be characterised to determine a lower limit for the beat note frequency, and the LO power has to be adjusted to reduce its overall influence.

Equation 2.25 did not include the power fluctuations that modulate the signal amplitude. This was done because their coupling into phase is much weaker and not relevant here. The part of the noise that is mixed during the phase detection to MHz frequencies does not directly affect the measurement, since it is shifted to frequencies far away from the beat note. Fluctuations of the signal amplitude at low frequencies only influence the phase measurement indirectly and their effect is small [43].

Another major difference to shot noise is that RIN is correlated for all channels using the same LO, this means it is, for example, the same on all segments on a quadrant photo diode or on both outputs of an interference beam splitter. This can be used to implement a technique called balanced detection. Here the signals from both interference beam splitter ports are subtracted from each other. Thereby the signal amplitude doubles (due to the opposite sign), the RIN almost fully cancels, and the shot noise increases, but only by $\sqrt{2}$, effectively increasing the shot noise limited C/N_0 by the same amount. This technique has not yet been investigated for reducing RIN in MHz heterodyne interferometry, however, it has already been used to reduce stray light effects in pre-experiments investigating the back-link fibre for LISA [47, 48].

2.4.5. Differential wavefront sensing

In addition to measuring the distance between satellites, the interferometer can also measure the pointing towards each other, or to be more precise, the pointing between the interfered beams. This measurement principle is called differential wavefront sensing (DWS) [49, 50, 51] and is applied in many other interferometers, for example, also for measuring the tilts of the test masses in
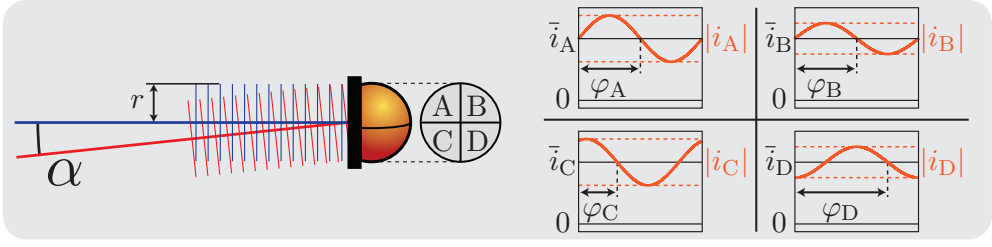
Figure 2.6.: Differential wavefront sensing on a quadrant photo diode. Each segment generates a DC and an AC current. The AC signal amplitude and phase depend on the optical power, the beam overlap, the beam alignment, the beam parameters and the position of the PD along the beam axis. The left side shows a a case where two beams are only tilted in one direction of the QPD, generating a phase difference between the segments at the top (A,B) and at the bottom (C,D). The right side shows an example of the photo currents in a more realistic interferometer, with different phases and amplitudes on all four segments (the effects are exaggerated).

LISA pathfinder [25]. To utilise DWS for both angles (tip and tilt) a photo receiver with at least four separate segments is necessary. As shown in Figure 2.6 each segment receives different phase front alignments if the two interfering beams are tilted relative to each other. By using a photo diode with four segments, a quadrant photo diode (QPD), the two alignment angles can be fully resolved.

To predict the real DWS signal an analysis of the two interfering beam shapes and the photo diode geometry is necessary. In general the detectable phase difference in one dimension depends on the relative beam angle α, the laser wavelength λ_0 and a complex geometry function $g_{\mathrm{DWS}}(\alpha)$.

$$\varphi_\alpha = \frac{g_{\mathrm{DWS}}(\alpha)}{\lambda_0}\alpha \qquad (2.26)$$

For the case of two flat top beams with a radius r and small tilts one can approximate the geometric function [39]

$$\varphi_\alpha^{\text{flat top}} = \frac{16r}{3\lambda_0}\alpha. \qquad (2.27)$$

The small laser wavelength amplifies the physical tilt angle, allowing it to be measured with very high precision. Using a QPD the signal above can be calculated from the phases on the individual segments.

$$\varphi_\alpha = \frac{\varphi_{\mathrm{A}} + \varphi_{\mathrm{B}}}{2} - \frac{\varphi_{\mathrm{C}} + \varphi_{\mathrm{D}}}{2} = \frac{\varphi_{\mathrm{A}} - \varphi_{\mathrm{C}}}{2} + \frac{\varphi_{\mathrm{B}} - \varphi_{\mathrm{D}}}{2} \qquad (2.28)$$

The signal for the perpendicular angle β is computed accordingly.

$$\varphi_\beta = \frac{\varphi_A + \varphi_C}{2} - \frac{\varphi_B + \varphi_D}{2} = \frac{\varphi_A - \varphi_B}{2} + \frac{\varphi_C - \varphi_D}{2} \qquad (2.29)$$

The original length signal, as measured on a single element diode, is recovered by summing the phase of the four quadrants. One should note that for most set-ups not only the phase depends on the beam tilts but also the individual signal amplitudes on the quadrants, as sketched on the right side of Figure 2.6. Depending on the tilting range and the beam shapes it might therefore in some cases be more useful to add the phases weighted by their corresponding amplitudes [25, 52, 53].

The use of DWS increases the number of readout channels, requiring four phasemeter inputs for each QPD. In contrast to the length measurements the phase between the segments needs to be measured absolutely, i.e. without unknown constant offsets, which has to be addressed by the readout system (see Section 4.3.2). Laser frequency noise and Doppler-shifts do not couple into the DWS measurement, since they are common mode between all segments. Shot noise on the other hand does limit the maximum DWS performance. DWS is not sensitive to roll, rotation around the beam axis, as long as the beams are circular and therefore rotation invariant, and not circularly polarised.

2.4.6. Laser link acquisition

Before an inter-satellite interferometer can start to operate, it first needs to establish the laser link. This requires the pointing between the beams to allow for a minimum contrast. Also the difference between the laser frequencies of the master and transponder, including the Doppler-shifts, need to be within the observable beat note range.

If such a state is achieved the phasemeter can lock its PLLs to this frequency and all other subsequent control loops can be closed. The full acquisition procedure for a mission is a complex problem and not discussed here [54]. Assuming a suitable beat note is present the phasemeter itself can acquire lock by detecting the signals with a Fourier transform and initialising its PLLs accordingly.

2.5. Laser Ranging Interferometer for GRACE Follow-On

Using the basic elements of inter-satellite interferometry, the architecture of the Laser Ranging Instrument (LRI) for GRACE Follow-On is outlined here. Figure 2.7 shows a sketch of the LRI. Since the GFO satellites are essentially a copy of the original GRACE design, the line of sight between the two centres of mass is blocked by the K-band ranging instrument. The LRI therefore utilises
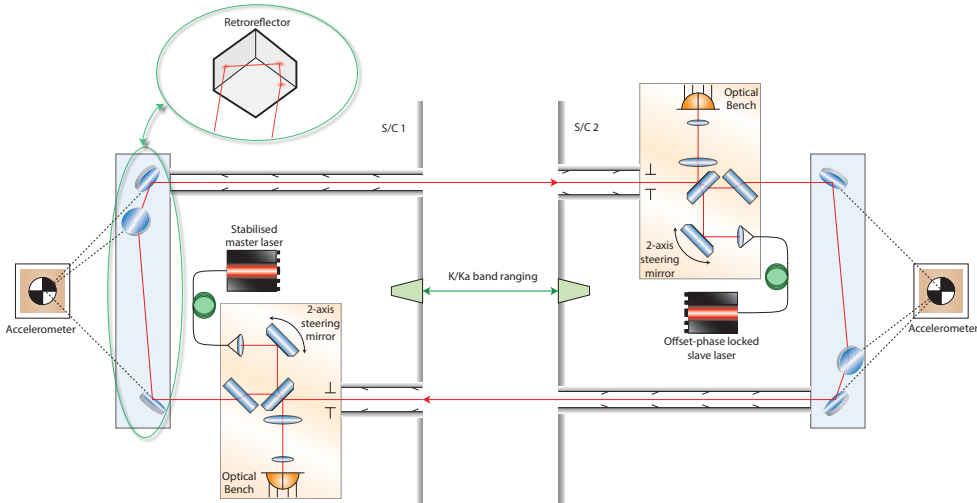
Figure 2.7.: Layout of the laser ranging interferometer for GRACE Follow-On from [39]. The LRI consists of two virtual corner cubes, two optical benches (OB) two lasers (both of which can be stabilised to a reference cavity, not shown) and two phasemeters (not shown). Two steering mirrors are used to correct for satellite pointing in an active control loop.

triple-mirror assemblies (TMAs) to redirect the laser beams in a racetrack like configuration. The TMA consist of three mutually perpendicular mirrors and acts as a retro reflector, which, to first order, has no coupling of rotations around its virtual mirror intersection point into the optical path length of the reflected beams. By placing this intersection point into the centre of mass of each satellite the LRI is able to measure only the desired distance variations. In contrast to the earlier described scheme, the LRI uses only a single beam splitter, one on each spacecraft, for the signal interference and the transmission of the local laser beam to the remote satellite.

Laser frequency noise is one of the limiting noise sources of the LRI, which only consists of a single arm. Therefore a laser pre-stabilisation system is included that reduces the frequency noise below $30\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ [55, 56]. Both lasers can either be stabilised or locked to the incoming light, allowing the master and the transponder operation to be switch and thereby increasing the system redundancy.

The relative pointing stability expected between the two satellites in GFO is not sufficient to sustain interferometric contrast. Hence an active pointing correction is implemented by steering a mirror in such a way, that the DWS signal, proportional to the pointing between the local and the remote beam, stays sufficiently small. This is done by implementing a feedback control loop, which is explained in detail in Chapter 7. Any length variations induced by
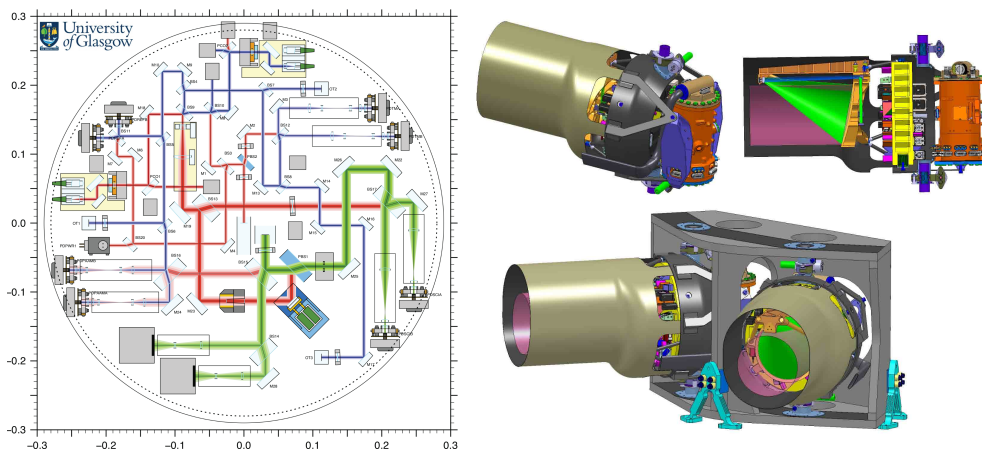
Figure 2.8.: The left side shows a layout of an optical bench design for LISA [57]. Shown are the received beam (green), the local beam (red) and the reference beam (blue). The right side shows a design for the optical assembly of NGO/eLISA [11]. Each consists of a telescope structure, an optical bench and a gravitational reference sensor (GRS) that contains the test mass.

the steering mirror or the fibres transmitting the light to the optical bench do not couple into the length measurement, since they are not part of the racetrack. An imaging system consisting of two lenses resizes the beams to fit fully into the photo diode radius and to remove beam walk on the photo diode. They thereby also influence the geometric coupling factor for the DWS. A more detailed description of the full interferometry concept is available from B. Sheard et al. [39]

2.6. LISA interferometry

For space based gravitational wave detection an interferometric distance measurement between free floating test masses over a long baseline is necessary. For LISA this is realised by dividing the distance measurement of one arm in three parts, one from each test mass to an optical bench on each satellite and one between the satellites. The latter interferometer connects the two optical benches and uses a telescope to increase the effectively received light power (see the right side of Figure 2.8).

On the optical bench (shown on the left side of Figure 2.8) all interferometers are connected and the optical signals are converted into electronic ones which are then read out via the phasemeter. Each bench has three beams at different frequencies, a local one, one form the remote spacecraft and a reference beam. The local and the reference beam are necessary to perform the local interferom-

etry and they need to be stabilised to each other. The different interferometers are listed in the following:

- Science IFO: Beat between the local and the remote laser. Measures the distance variation and pointing to the other satellite and thus contains the gravitational wave signal. It operates with low received power levels, variable beam pointing, and a high amount of initial laser frequency noise. It also needs to accommodate the auxiliary functions listed later and is, therefore, the most difficult benchmark.

- Test mass IFO: Beat between the local and the reference laser. Measures the movement and attitude of the test mass relative to the OB. Uses convenient power levels and measures almost no laser frequency noise.

- Reference IFO: Measures the relative noise between the local laser and the reference, supplied either from the other OB on the same SC or from a separate laser. In a SC connected to two interferometer arms the reference laser is usually locked to the local one to allow the common mode rejection of laser frequency noise. It does not need to measure any DWS signals.

- Auxilliary IFOs: The original LISA design also included interferometers to measure the influence of a point-ahead angle mechanism (PAAM) and the stability of the telescope via optical truss [10]. For NGO and eLISA neither of these additional IFOs are currently planned to be implemented.

LISA will, in contrast to GFO, have a sufficient attitude control of the satellite and of the test mass to maintain the interferometric signals without a separate steering mechanism. This is achieved by the drag-free and attitude control system (DFACS), which uses the DWS signals in a feedback loop to control the satellites thrusters and the capacitive actuators for the test masses.

The overall performance requirements of LISA consist of various noise allocations. As a rule of thumb, each subsystem of the metrology should contribute less than $1\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ multiplied by a noise shape function (NSF), defined in the measurement band between $0.1\,\mathrm{mHz}$ and $1\,\mathrm{Hz}$.

$$\mathrm{NSF}(f) = \sqrt{1 + \frac{2.8 \cdot 10^{-3}}{f/(1\,\mathrm{Hz})}} \tag{2.30}$$

The phase readout of the interferometers therefore also requires a performance level better than $1\,\mathrm{pm}/\sqrt{\mathrm{Hz}} \cdot \mathrm{NSF}$ which corresponds to a phase noise level of $\approx 6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}} \cdot \mathrm{NSF}$. One pm is roughly one ppm of the laser wavelength.

2.6.1. Laser frequency noise removal

Laser frequency noise in LISA is one of the major noise sources, since it is initially several orders of magnitude above the performance goal. Similar to

Michelson type ground base interferometers the frequency noise for equal arm length vanishes completely when the differential phase between the two arms is calculated. This is, however, only true if the two independent laser sources on the central spacecraft are locked to each other and their frequency noise is therefore the same. Due to the achievable orbits of the satellites the real length of two arms will however vary by up to 1%, (this corresponds to 10 000 km for the eLISA design). This weakens the common mode rejection drastically and makes additional methods necessary. By using a technique called time-delay interferometry (TDI) the original rejection can be regained by post processing. These algorithm require the measured phases and additionally the absolute time of flight τ_d between the satellites [58, 59, 60] as input.

PRN ranging

Since the absolute distance between the SC is not known with a high enough precision, an additional ranging system with a resolution of better than 1 m is required. The current base line for LISA therefore includes an additional phase modulation on the light that is sent to the remote SC. This modulation consists of a pseudo-random noise (PRN) code that can be used to measure the time of flight by locking to it a delay-locked loop [61, 62, 63, 44, 64, 65].

The modulation consists of a pattern of ones and zeros that change with the modulation frequency f_m, also called chip rate. The patterns have a minimal cross correlation, but since they are only pseudo-random they repeat after one code length. The PRN modulation for ranging in LISA only uses a small fraction of the overall light power and modulation frequencies in the order of 1 MHz. It thereby adds small but broad side lobes to the signal that are not interfering with the phase measurement. Figure 2.9 shows the influence of different PRN modulation on a carrier signal spectrum. By encoding additionally data into the PRN codes a communication between SC via the laser link is also implemented.

Alternative ranging methods

Two main alternatives to PRN ranging delays are also under research. One is TDI ranging [66], which uses the common mode rejection of an injected signal (or the frequency noise itself) to deduce the correct delays for the measured data. The second alternative is to lock the master laser not to a stabilised reference but to one of the interferometer arms, which is by design much more stable than any reference. This allows the reduction of the overall laser frequency noise below the required values [67]. This technique is referred to as arm-locking [68, 69] and some variants of it have been studied theoretically and experimentally [70, 71, 72, 73].
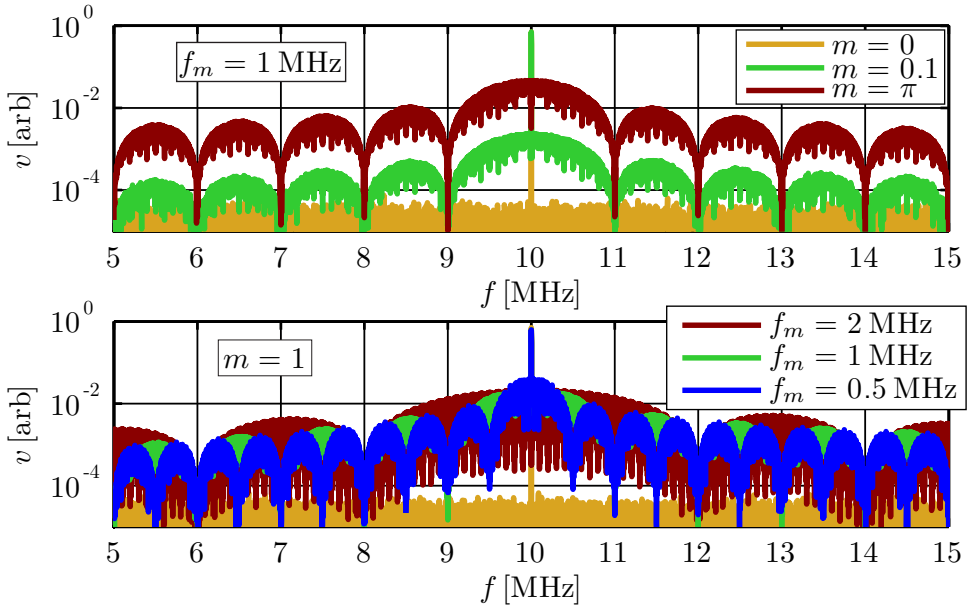
Figure 2.9.: Shown are simulated spectra of different PRN modulations of a 10 MHz beat note. The top plot shows the influence of modulation depths for $f_m = 1$ MHz. One should note that for a modulation depth of $m = \pi$ the carrier is fully destroyed. The bottom plot shows the influence of modulation frequency for $m = 1$.

2.6.2. Clock noise and frequency distribution

The signals on each LISA spacecraft are digitised with a clock at a frequency of, for example, 80 MHz. Any clock noise between the SCs will introduce phase noise (as described in Section 3.3.1) and currently none of the available clocks are stable enough to fulfil the LISA requirements. Hence, the introduced clock noise has to be removed and to be able to do that, it has to be measured. Though the absolute clock stability can not be determined the relative noise between the clocks can be measured by comparing them. For LISA this is done by adding phase modulation side bands onto each laser beam, with a different frequency for each SC. This will generate two additional beat notes in the science interferometers that can be tracked independently. The measured phase noise of these clock noise side bands can then be used in post correction [74, 75]. Since the detection precision of the science IFO is limited by the received power the clock side bands have to be rather small, which also decreases their C/N_0. To be able to measure the clock noise with sufficient precision, multiplied versions of the original clocks are used as modulation side bands, with frequencies in the GHz region. For the clock correction to be successful the GHz tones have to be

sufficiently phase stable to the sampling clock.

Since the sampling of signals itself inevitably introduces another clock noise, the ADC sampling jitter, a third tone is required to correct for this, the pilot tone [43]. Using a jitter correction tone it turns out that the phase stability between the pilot tone and the GHz side bands is the relevant factor. The stability of the sampling clock is only of minor importance, since it is corrected by the pilot tone. The architecture that generates and distributes these different tones with the required stability is called the frequency distribution subsystem (FDS).

Chapter 3

Phase Readout: The analogue part

The readout of inter-satellite interferometers requires studying the effects occurring in the analogue domain. The design of these components plays a crucial role for precision interferometry. This chapter describes the basic building blocks in the analogue signal chain, also called analogue front-end (AFE), and introduces the necessary analysis and design elements to achieve the performance levels required in LISA. The effects discussed here also apply for GRACE Follow-On, though a specific analysis might have a different focus due to the different signal dynamics and a less stringent required phase measurement performance.

The initial conversion of the photo diode signal into a voltage is shortly described in the first section. Then a discussion of the active and passive components used to prepare the signals for digitisation is given and their influence on phase noise is investigated. This is directly followed by a description of the effects occurring during the digitisation process in the analogue-to-digital converter (ADC). Finally, the unwanted interference of analogue signals, the so-called crosstalk, is discussed and based on this analysis the influence on the readout design is derived.

3.1. Converting optical to electronic signals

The optical signals generated by interferometers are converted into voltages by a photo diode (PD) and a trans-impedance amplifier (TIA). Together these elements are also called a photo receiver (PR). For inter-satellite interferometers with small initial signals both functions are usually integrated into a single device and placed directly onto the optical benches. Thereby the generated signals are immediately amplified and can then easily be transmitted to the phase measurement system using standard radio frequency (RF) cables and
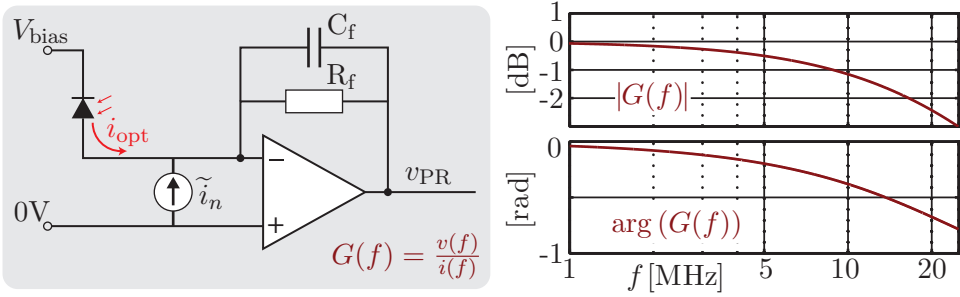
Figure 3.1.: The left side shows a simplified model of a photo receiver that converts the PD current i_opt into a voltage v_PR and adds a general input current noise $\widetilde{i}_\mathrm{n}$ of the TIA. The right side shows a transfer function of the circuit, the amplitude in the top plot and the phase in the bottom one. In this example it is dominated by a simple low pass filter defined by the feedback resistor R_f and capacitance C_f.

electronics. In the case of large optical signals with lower frequencies, like in LISA Pathfinder, the TIA and the photo diode can also be split.

3.1.1. Photo diodes

The PDs investigated and considered for space-based laser interferometry are either based on silicon or on indium gallium arsenide (InGaAs). The choice of material influences the most important parameters of the PD, the capacitance, the responsivity and the size of the active area. The capacitance limits the achievable bandwidth as well as the noise floor of the TIA. For quadrant photo diodes more parameters are relevant, for instance the crosstalk between the individual segments and the size of the slit, the non responsive area between the segments. The size of the active area is an important parameter for the imaging systems, since the beams should not be clipped on the PD. A non continuous efficiency over the active area can strongly influence the differential wavefront sensing (DWS) response. A comparison of different PD candidates is found in [48]. Due to their higher efficiency and bandwidth at 1064 nm, InGaAs photo diodes are preferred for weak MHz signals at that wavelength.

3.1.2. Trans-impedance amplifier

The conversion of the photo current into a voltage in the TIA is one of the most critical points in the measurement chain for low signal powers. Several noise sources are added at this stage and each one has to be below the shot noise limit. Detailed models of these noise sources are available [76, 77, 78]. The conventional approach for these models is to describe each noise source by a

corresponding equivalent input current noise $\widetilde{i}_n$ of the TIA, as shown in Figure 3.1. This notation allows one to easily compare their influences with shot noise (see Equation 2.22) and the signal. Depending on the circuit noise and power levels the AC and DC signals are either amplified by a single TIA or by two, normally separated by low and high pass filters.

The bandwidth or, to be more precise, the complex trans-impedance of the TIA $G(f)$, defines the dependency of amplitude and phase on the signal frequency. An example of such a transfer function with a -3 dB corner frequency at 25 MHz, the upper limits of the beat note band for LISA, is shown in Figure 3.1. For reasons described in the next section a high bandwidth and therefore a low phase drop is desirable.

Depending on the signal levels and the system design a PR might include additional amplification stages as well as other components mentioned in the following.

3.2. Amplifier, buffer, adder

An idealised analogue measurement chain transfers the sufficiently amplified, single ended PR signals with impedance matched cables to the phasemeter. There the pilot tone is added and the sum is converted into a differential signal and fed into the ADC.

With the exception of signal amplification all of the above mentioned steps can either be implemented using active or passive components. In general the active components introduce an additive noise that has to be considered, while the additive noise from passive components (for example, Johnson noise from resistors) is usually negligible after the TIA. Each component or set of components has a frequency dependent amplitude and phase behaviour and can be modelled by a transfer function similar to the one shown in Figure 3.1.

3.2.1. Phase noise in the analogue front-end

The complex, frequency dependent phase of each circuit is defined by the values of its sub components, like the feedback resistance in Figure 3.1, including their parasitics. These values depend on temperature and consequently temperature changes can easily couple into phase measurements. This is one of the major noise sources observed in the experiments conducted in this thesis and it is described here in more detail with a simple example.

The amplifier circuit from Figure 3.1 is analysed for its phase behaviour in case of a changing feedback resistor R_f. The temperature dependency of the capacitor and the amplifier itself are assumed to be negligible. A change of the resistor value will move the corner frequency of the transfer function $G(f)$. Figure 3.2 demonstrates this for a relative change of the feedback resistance by $\pm 2\%$. A value change of 2% leads to a phase error of 10 mrad at 25 MHz.
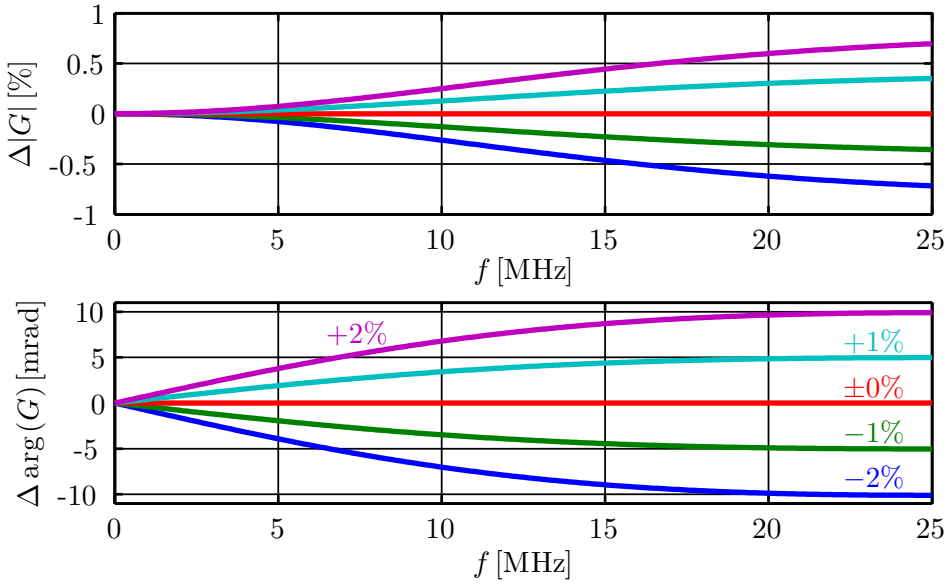
Figure 3.2.: Shown is the change in amplitude and phase over frequency for a simple analogue first order filter while changing the feedback resistor value from $+2\%$ to -2%.

To achieve a performance in the order of μrad this means that the resistor value should at maximum change by 2 ppm. The rather stable Mini-MELF [79] resistors, often used in this thesis, have in their standard package a temperature coefficient of ± 50ppm/K. Assuming such resistors are used the outcome of this exercise is that a temperature stability of 10^{-2}K/$\sqrt{\text{Hz}}$ is necessary to achieve the desired performance. At low frequencies this temperature stability is not easy to achieve, especially if active components are in close proximity. Therefore a comprehensive system design either utilises more stable components or it needs to increase the temperature stability.

Another, more complex option is to use the temperature dependence of one component to compensate for another. For this to be successful the complex temperature dependence of each component needs to be well known. This approach was not used during this thesis, but an example of such a design is included in the frequency distribution system shown in Chapter 9. One should note that these calculations were simplified by the fact that the signal only had a single, fixed frequency instead of varying between 2 and 25 MHz.

Some effects influencing the phase can not be described by linear analysis, like the up conversion of low frequency noise to RF phase noise or flicker noise [80], which can often limit the performance at low frequencies. Components showing such phase noise have so far simply been discarded in the development, since

methods to reduce these noises would require further detailed studies.

3.2.2. Reflections and impedance matching

Since the analogue signals are in the MHz regime and they have to be transmitted over a few metres of cable, their wave properties become relevant. This issue is normally addressed by using standard impedance matching techniques [81]. The impedance used during this thesis is $50\,\Omega$. The precision of impedance matching using standard components is in the order of a few percent. The impedance at RF frequencies includes various parasitic effects and is a complex function of frequency. Reflections due to small impedance mismatches can cause noise which depends on the level of the mismatch and on the dynamics of the input signal. In a simple picture the signal arriving at the phasemeter consists not only of the original one but also contains a small fraction of itself, delayed by twice the length of the impedance mismatched line. The coupling of this small additional signal into the phase measurement is expected to be similar to crosstalk, which is described in Section 3.4. Experimental evidence for relevant noise influences due to impedance effects is shown in Chapter 8 and Chapter 9.

3.2.3. Anti-aliasing

To prepare the digitisation an anti-aliasing filter (AAF) is necessary, especially for signals with low signal to noise ratio [82]. During the digitisation process the signal components above the Nyquist frequency are folded into the resolvable frequency band. For a signal containing a single beat note and a white noise floor this will decrease the signal-to-noise ratio (SNR) and thereby reduce the phase noise measurement performance, as shown in Figure 3.3. Depending on the desired performance levels the required suppression of the AAF becomes quite demanding. This is worsened by the fact that analogue filters with high suppressions cause in general a strong phase drop, which makes them very sensitive to temperature fluctuations, as described earlier. The design of an AAF is somewhat simplified by taking the inherent transfer functions of the PR and the other analogue components into account. For LISA the AAF filter can be placed directly into the PR on the optical bench, which is the most temperature stable environment.

3.3. Analogue to digital conversion

The digitisation process of the analogue signals is implemented in ADCs. Their most crucial parameters for phase readout performance are the sampling frequency, the number of generated bits, the additive input noise and the sampling jitter. To avoid effects of limited bandwidth in the ADC itself normally they are
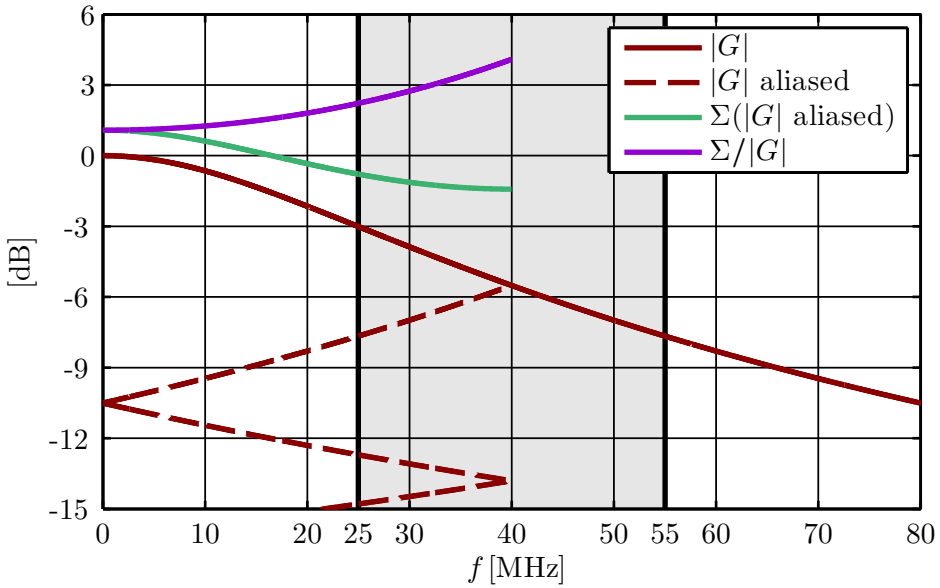
Figure 3.3.: The effects on the noise floor due to aliasing are demonstrated here in an example case. The simple transfer function used earlier is shown extending to frequencies above the beat note band limit at 25 MHz. The signal is digitised with an 80 MHz clock, therefore all components above 40 MHz are folded down. For a white noise background the additional noise in the signal band is directly given by the magnitude of the transfer function (the individual contributions from each folding are shown as dashed red lines). These contributions can then be added quadratically (since the noise is uncorrelated) to determine the effective resulting noise floor after digitisation (green line). The effect on the phase noise performance is then evaluated by computing the ratio of the signal amplitude and the noise floor (purple line). With a simple first order low pass filter the effective noise limit is increased by 1-2 dB in this example. One should note that a larger ratio between the sampling frequency and the highest signal frequency reduces the required AAF suppression drastically.

chosen such that their analogue input bandwidth is much larger than the signal bandwidth. The high bandwidth needs of telecommunication devices is the reason that various types of ADCs with these specification are available. Some examples of suitable components are presented later on in the experimental chapters.
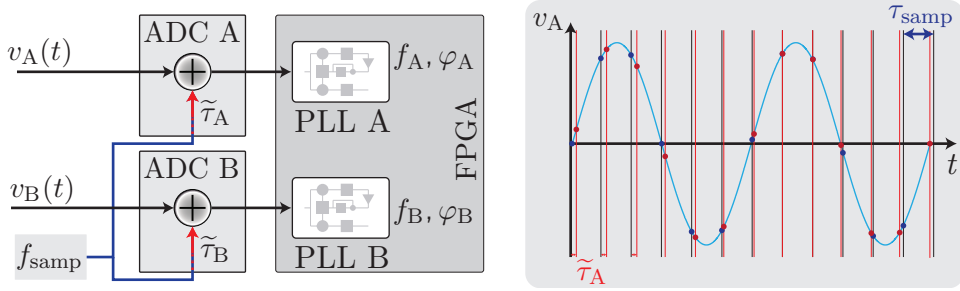
Figure 3.4.: Even though two ADCs (or two ADC channels) use the exact same sampling clock each digitisation has a low frequency time variation, the sampling jitter, as shown on the left side. The + symbols here are not normal additions but they symbolise additional time delays caused by the sampling jitter $\widetilde{\tau}$. The right side shows the sampled voltage of a sinusoidal signal for a perfect sampling (blue dots) and under the influence of sampling jitter (red dots).

3.3.1. Sampling Jitter

The measured phase of a beat note depends on the sampling frequency of the system. Any unwanted variations of this frequency and its corresponding sampling time therefore cause an unwanted phase noise, as described in Figure 3.4. Studies of phase measurement systems for LISA have revealed that available ADCs have an inherent variation of the sampling time at low frequencies, the so-called sampling jitter, that spoils the μrad performance for signals in the MHz regime [36, 43, 83].

Since phase and frequency have a simple relation (see Equation 2.1) the sampling jitter $\widetilde{\tau}_A$ induced phase noise $\widetilde{\varphi}_{\tau_A}$ for an input signal at frequency f_A can easily be calculated:

$$\widetilde{\varphi}_{\tau_A} = 2\pi f_A \cdot \widetilde{\tau}_A. \tag{3.1}$$

This equation does not only lead to the understanding of this coupling but also to its correction. Since the sampling jitter couples into all sinusoids in the input signal one can add a reference tone, the pilot tone, and use its measured phase variations to correct the optical beat notes. Assuming a pilot tone f_P is present in the input signal and its phase φ_P is read out, the phase correction term is

$$\varphi_{A,cor} = -\varphi_P \frac{f_A}{f_P}. \tag{3.2}$$

The implementation of such a correction scheme is described in Chapter 8. The jitter between channels is corrected if a common pilot tone is used. But each pilot tone also carries an overall jitter from its imperfect source (since there are no perfect clocks). This is the exact reason why there is a need for a clock noise

correction in LISA (see Section 2.6.2), otherwise the jitter between the pilot tones on each satellite would spoil the system performance.

In the case of an input signal with low SNR the pilot tone phase measurement will (depending on its amplitude) also be limited by the white noise floor. Thereby the signal correction not only decreases the sampling jitter, but it also increases the noise floor of the corrected signal. This effect can be controlled by choosing higher pilot tone amplitudes (increasing its SNR) and by using tones at higher frequencies, as indicated by Equation 3.2. Note that clock and pilot tone jitter are not only introduced by their sources and in the ADCs, but also in their distribution. These effects are similar to the ones described for the phase noise of the analogue components in the signal chain, but they need to be evaluated for different and normally constant frequencies [74, 75].

3.3.2. Signal levels and digitisation noise

One important function of the AFE is the correct amplification of the signals in preparation for their digitisation. Next to standard limits on the signal levels, like the maximum ratings of the ADC, some additional constraints have to be met.

For calculating the maximum signal level one has to take into account the full content of the input signal. For the LISA science IFO this includes one main beat note, two side band beat notes, a pilot tone, the pseudo-random noise modulation and additive noise (mainly shot noise), which adds with a normal distribution. To keep the linear models of the phase readout intact the digitisation should never limit the maximum signal amplitude, which is also known as clipping. Though the effects of this have not been studied in detail a single sine wave that is clipped will, for an increasing amplitude, slowly turn into a rectangular signal, introducing various additional frequency components and a phase noise equivalent to a 1-bit quantisation (as described below). In this thesis input levels of the dominating signals, the combined main and pilot tone, were chosen such that they fill, at maximum, around $1/4$ of the available input range to leave enough room for other signal components.

The lower limit of the input signal amplitude is given by the noise floor of the ADC, which consists of a voltage noise floor and the digitisation noise. The latter depends on the number of bits N and the sampling frequency f_s, as described in the next chapter. For calculating the relative phase noise induced by ADC digitisation ($\widetilde{\varphi}_{\mathrm{ADC}}$) the ratio of the input signal amplitude V_{in} (the peak voltage) to the ADC maximum input range V_{p} (maximum peak voltage for an AC coupled signal) needs to be included to determine the required number of bits [84].

$$\widetilde{\varphi}_{\mathrm{ADC}} = \sqrt{2}\frac{V_{\mathrm{p}}}{V_{\mathrm{in}}}\frac{2^{-\mathrm{N}}}{\sqrt{6 \cdot f_s}}. \tag{3.3}$$

For a phase noise performance of $6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ for LISA, a sampling frequency
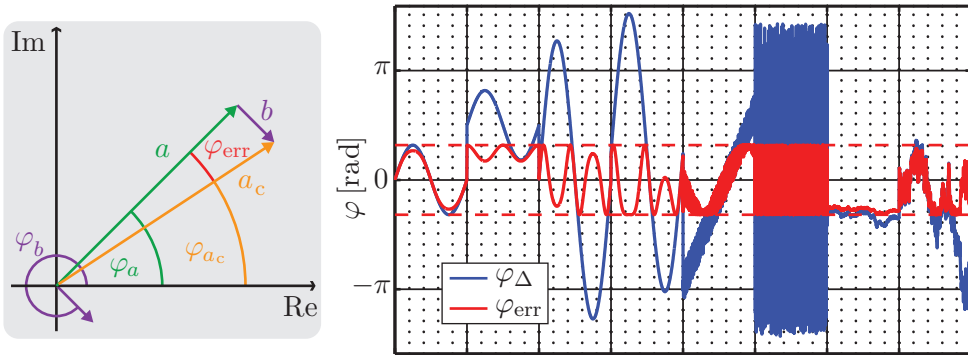
Figure 3.5.: The influence of crosstalk is understood by applying simple vector algebra, as shown on the left. Often a smaller parasitic phasor is moving relative to the actual signal, causing a phase error by changing the phase of their sum. Therefore this type of coupling is also sometimes referred to as small vector coupling, generating small vector noise. The induced noise by uncorrected crosstalk is shown for a few examples in the simulated time series on the right. It includes a coherent φ_δ with and without a phase offset, a white noise and a $1/f$ noise. Each case is shown with deviations smaller and larger than π.

of $80\,\mathrm{MHz}$ and an input ratio of $V_\mathrm{p}/V_\mathrm{in} = 4$ a minimum of 6 bits is required. The available ADCs often have more than 10 bits and therefore they have extra margin to account for varying input signal levels. The voltage noise at the ADC input is normally on a similar level as the digitisation noise or slightly above. In general it is either specified for an ADC or has to be determined by measurements.

3.4. Crosstalk

Interference between analogue signals before their digitisation can compromise the measurement performance. Reducing this crosstalk between independent readout channels in a phasemeter is another critical aspect of the analogue design, especially for experiments like LISA, where several interferometers are read out with a single device [83].

The influence of crosstalk is calculated here and used to compute requirements. For a signal $a(t)$, a parasitic influence by a signal $b(t)$ and a crosstalk factor of c, the contaminated signal a_c is

$$a_c(t) = a(t) + c \cdot b(t) = a \sin\left(\omega_0 t + \varphi_a\right) + c \cdot b \sin\left(\omega_0 t + \varphi_b\right). \qquad (3.4)$$

Since the IQ demodulation determines the amplitude and phase of the input

signal in the complex plane, the resulting measured phase of the contaminated signal φ_{a_c} can be calculated with simple trigonometry (as shown in Figure 3.5). For a phase difference between the original and parasitic signal of $\varphi_\Delta = \varphi_a - \varphi_b$ the measured phase is

$$\varphi_{a_c} = \varphi_a + c\frac{b}{a}\sin(\varphi_\Delta). \qquad (3.5)$$

If the phase of both signals is well aligned ($\varphi_\Delta \ll 1$) the phase error due to crosstalk can be approximated by

$$\varphi_{\text{err,cross}} = \varphi_{a_c} - \varphi_a \approx c\frac{b}{a}\varphi_\Delta. \qquad (3.6)$$

For a phasemeter with several independent beat notes the influence on the final performance is not easily evaluated. Any crosstalk mechanism will likely transfer small signals from one channel to another and vice versa, meaning that both channels will have an additional phase error that couples in a non-linear fashion. Measuring this effect is not simple, since the often performed zero or null measurements (see Section 5.2) are not directly sensitive to it.

Critical elements for crosstalk coupling

Crosstalk can be introduced at various points in the measurement chain. Depending on the types of signals that influence each other the effect on the overall system performance varies considerably. Two extreme cases are discussed here that relate to two critical examples in inter-satellite interferometry.

Quadrant photo diode (QPD)

In the first case there is a cross-coupling of two or more signals which are rather stable to each other and also have a well measured absolute phase relation and amplitude. The corresponding example of this case is the crosstalk between segments on a QPD used for DWS. Even though the expected DWS dynamics vary greatly, depending on the mission, the absolute phase relation is always well measured and therefore the effect is somewhat predictable. The full influence of crosstalk in QPDs on DWS is still under research [78]. While the crosstalk in QPDs is often stated as a single number in the order of a few percent (see [48] for some examples), the full description of crosstalk between QPD segments requires to determine spatial interference and responsivity patterns, as well as frequency dependent effects.

The pilot tone distribution in LISA

The second case is more critical, it is the crosstalk between readout channels of independent interferometers without any existing or measured absolute phase

relation. Due to the unknown phase relations and the potentially extreme phase dynamics between the signals, this type of crosstalk can limit the achievable system performance. The LISA phasemeter is the relevant example for this type of crosstalk. It reads out various channels from different interferometers that are physically close to each other and share a common pilot tone. Coupling due to electro magnetic interference or shared power supplies and ground planes can potentially be reduced by circuit layout, shielding and simply increasing physical distances. The pilot tone distribution required for LISA is the most demanding connection between the channels. The easiest possibility of reducing crosstalk is to use different heterodyne frequencies for each IFO, but the LISA architecture, the limited resources, and the highly increased complexity make this approach unfeasible. The alternative is to ensure a sufficient suppression of crosstalk by adding various components like filters, isolators and buffers, though these, of course, have to be compliant with the phase noise performance.

Assuming all signals have roughly the same amplitude ($a \approx b$), Equation 3.5 shows that a coupling factor of $c \lessapprox 10^{-6}$ is required to achieve μrad performance. This is based on the worst case, a phase shift between both signals from $\varphi_\Delta = \pi/2$ to $-\pi/2$. The required suppression of $120\,\mathrm{dB}$ is larger than the one quoted for most multi-channel ADCs, making them an inappropriate choice for the readout of independent IFOs. Consequently, separate ADC chips should be used for such signals.

Chapter 4

Phase readout: The digital part

Core of the readout of inter-satellite interferometers is a phase-locked loop (PLL, see Section 2.3), which performs the phase and frequency tracking. These PLLs are implemented in field-programmable gate arrays (FPGAs) that allow large amounts of parallel signal processing. Since the algorithms are fully implemented in the digital domain they are also called all-digital phase-locked loops (ADPLL). Their implementation and influence is critical for the measurement performance and stability and therefore the dedicated analysis developed during this thesis is presented here.

First an overview about the relevant fundamentals of digital signal processing (DSP) is given. Thereby a notation is introduced, as well as the underlying effects in the digital algorithms. The second section of the chapter is focused on the core element of the phase readout, the ADPLL. A linear model of the phase tracking is presented, the influence of the input signal phase and additive noise is calculated and the detection of the signal amplitude is discussed. The third section describes how the phase information is extracted from the ADPLLs and how the signals are filtered and decimated to rates suitable for storage and communication to ground. PLLs are inherently non-linear and therefore additional effects occur in the presence of input signals with a large dynamic range like in LISA or GFO. The last section therefore describes these non-linearities and the resulting design aspects of the ADPLL.

4.1. Relevant fundamentals of digital signal processing

The signal processing investigated and applied in this thesis is implemented mainly in FPGAs, using fixed-point integer arithmetic and VHDL [85] coding,

but also in PCs and microprocessors using floating-point arithmetic and the programming languages C and C++. For the readout of inter-satellite interferometers the use of fast integer arithmetic is the crucial factor and it is therefore described in more detail in the following. Since the floating-point arithmetic used is less critical it is not focused on here and its limitations are only discussed when relevant.

4.1.1. Scaling

The ADPLL and the other algorithms running in FPGAs are implemented using fixed-point arithmetic, where each value is represented by X bits and one can use different scalings to map these to numbers that represent phase etc. used in the analysis of the algorithm. Here it is chosen to scale each integer by 2^{-X}, which leads to the following ranges for signed and unsigned numbers:

$$-0.5 \leq \left[\frac{-2^{X-1}}{2^X} \leq \text{ signed } \leq \frac{2^{X-1} - 1}{2^X} \right] < 0.5$$

$$0 \leq \left[\frac{0}{2^X} \leq \text{unsigned} \leq \frac{2^X - 1}{2^X} \right] < 1$$

(4.1)

Numbers with this scaling have no units, since they only represents values in the digital computation. Only appropriate further scaling maps them to real physical quantities with units.

4.1.2. Quantisation

The most distinguished property of DSP is quantisation, since it only operates with discrete numbers. For any number represented by X bits the smallest distinguishable value is given as 2^{-X}, also called the value of the least significant bit. The effects of quantisation have to be well understood and treated to make them predictable and compliant with the requirements of each subsystem. Since each signal can only represent a finite number of distinct values with a constant, non-zero separation between them, the process of quantising a signal introduces noise and wrong DC averages, as described in the following.

Truncation Noise

The truncation of a continuous signal to a digital number with X bits at a sampling rate f_{samp} can be modelled as an addition of uniformly distributed white noise with a linear power spectral density of

$$\widetilde{x}_{\text{trunc}} = \frac{q}{\sqrt{6 \cdot f_{\text{samp}}}} = \frac{2^{-X}}{\sqrt{6 \cdot f_{\text{samp}}}}.$$

$$[\widetilde{x}_{\text{trunc}}] = \frac{1}{\sqrt{\text{Hz}}}.$$

(4.2)
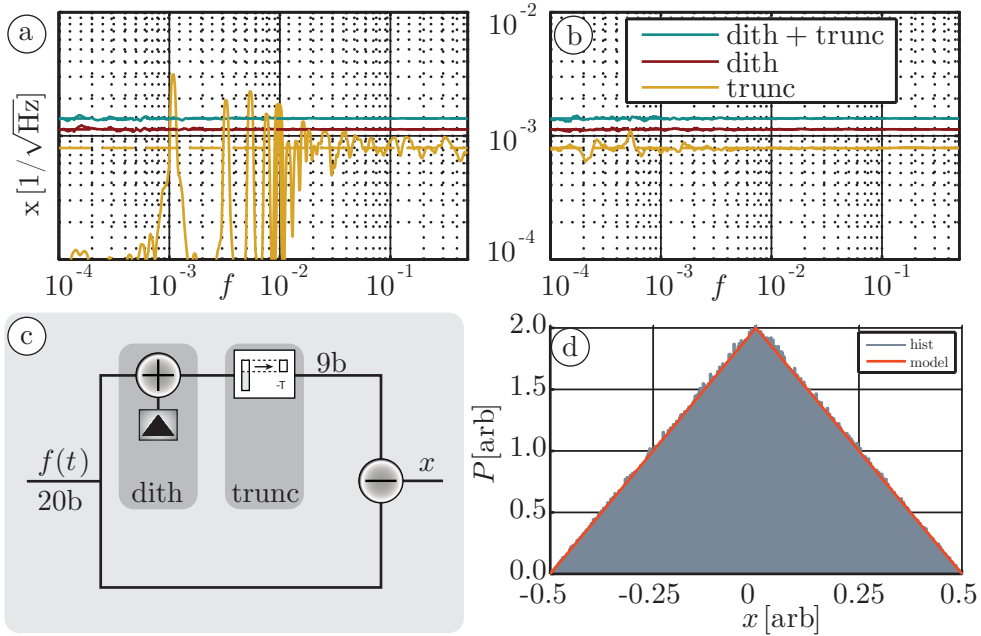
Figure 4.1.: (a) Power spectral density of a signal consisting of a single sine wave being truncated, dithered or both. (b) Power spectral density of a signal consisting of two non-harmonic sine waves being truncated, dithered or both. (c) Set-up of the truncation and dithering noise analysis. (d) Histogram of the triangular dither in comparison to the expected shape. The overall probability is normalised to 1.

This applies also for the re-truncation of digital signals. However the assumption of additive white noise is only valid for signals that move through a significant range of digital values without any coherent relationship to the sampling frequency, like e.g. two or more sine waves at non-harmonic frequencies [86]. Otherwise the quantised signal will show artefacts and peaks from the coherent interaction with the truncation process. This difference is shown in Figure 4.1, the yellow line in plot (a) shows the spectral density of the noise introduced by truncating a single sine wave, while (b) shows the same plot for a signal containing two sine waves at non-harmonic frequencies. The dashed yellow line shows the assumed additive white noise given by Equation 4.2. While (b) shows the predicted noise floor (a) shows various artefacts, many of which are additional tones at harmonic frequencies of the input signal.

Dither

To avoid such artefacts an intentional noise floor was added to any digital signals before truncation done in FPGAs in this thesis, so called **dither**, with triangular

dither being the preferred implementation [86]. For a truncation from K bits to $K - L = X$ bits the dither has a length of $L + 1$ bits.

A triangular dither generator was implemented in VHDL by subtracting the outputs of two independent linear feedback shift registers (LFSR) that generate pseudo-random noise with repetition lengths longer than $10000\,\mathrm{s}$. This ensured that no artefacts of the dither repetitions are visible in the desired signal spectrum range ($0.1\,\mathrm{mHz}$ to $1\,\mathrm{Hz}$) The histogram of 800k samples generated by the triangular dither generator is shown in Figure 4.1 (d).

The performance of the dither generator and its influence on the signals was separately investigated by simulating all three possible combinations of truncating and dithering a signal. This was done with the same signals previously introduced to demonstrate the truncation noise influence. The results of this are shown in 4.1 (a) and (b).

By only adding dither and not truncating the signals (red lines), a purely white noise floor of $\sqrt{2} \cdot \widetilde{x}_{\mathrm{trunc}}$ is measured. This is expected and understood to be caused purely by the addition of the two independent white noise sources used for the dither generation.

The addition of dither with a subsequent truncation showed that a purely white noise floor is achieved for both signals, without any additional artefacts. The level of the noise floor is slightly higher, with a value of

$$\widetilde{x}_{\mathrm{trunc+dith}} = \frac{2^{-N}\sqrt{3}}{\sqrt{6 \cdot f_{\mathrm{samp}}}} = \frac{q\sqrt{3}}{\sqrt{6 \cdot f_{\mathrm{samp}}}}.$$
$$[\widetilde{x}_{\mathrm{trunc}}] = \frac{1}{\sqrt{\mathrm{Hz}}}.$$

(4.3)

The increase by a factor of $\sqrt{3}$ can in general be tolerated, since the introduction of spurious signals that may cause other unwanted behaviour is now suppressed. The noise can of course always be arbitrarily reduced by using more bits and therefore more computational resources.

Rounding

Truncations can also introduce small signal offsets due to rounding errors [87, 88]. For truncations used here this was prevented by using offset-free rounding algorithms based on simple integer arithmetic. Specific rounding blocks were created in VHDL and included into all implementations. The blocks truncate symmetrically around zero, and they are linear, keeping the amplitudes of signals constant. To generate the correct offset for some truncation cases a dithered bit is used to determine the rounding direction. This is necessary to solve the ambiguity for some integer numbers that prevents a simple decision between rounding up or down and it thereby keeps the average offset error at zero.

4.1.3. Discrete time

Digital systems highly depend on their sampling time t_s, which is in general determined by a sampling clock at frequency f_s. Depending on the architecture used for DSP the computation time might have any ratio to the sampling time. Though for the FPGA based systems discussed in this thesis the computations are mostly done with the same clock as the signal sampling. Any exceptions from this will be discussed separately with the corresponding implementations.

Sampling

By sampling a signal with f_s it contains all its phase and frequency information if the sampling frequency is perfectly constant and the signal frequency f_{sig} satisfies the Nyquist-Shannon sampling theorem, where $f_{\text{sig}} < f_s/2 = f_{\text{Nyq}}$. The effects of sampling jitter on the phase of a incoming signal are described in Section 3.3.1.

Aliasing

As already described in Section 3.2.3 components at higher frequencies are folded/aliased back below f_{Nyq} [82]. An aliased tone will thereby not only be at a new frequency but it will can change the sign of its phase variations. Especially for pilot tones, which are often at high frequencies (see Equation 3.2) this has to be taken into account during jitter correction. Even though a pilot tone above Nyquist will be aliased, it will still carry the phase noise introduced by sampling jitter corresponding to its original frequency [43].

z-domain analysis

To describe the digital system in frequency domain the z-transformation transfer functions will be used. These transfer functions describe the frequency dependent behaviour of time discrete systems relative to the sampling frequency. They are described by defining a delay Operator z^{-1} that represents a single delay at the sampling frequency f_s:

$$z^{-1} = e^{-i2\pi \frac{f}{f_s}}. \tag{4.4}$$

4.2. All digital phase-locked loop

The core of the phase readout system is the ADPLL which was studied in great detail as part of this thesis. The basis for the understanding of the ADPLL is its linear model. Most of the contents of this and the following two sections has been published in Classical Quantum Gravity [89].
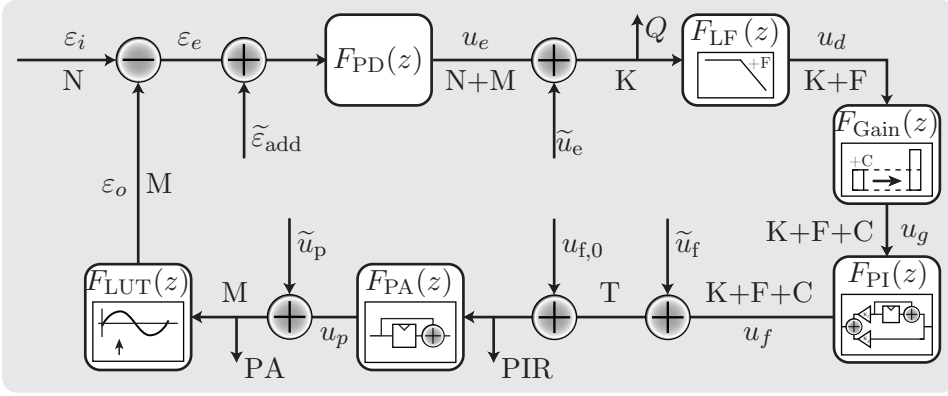
Figure 4.2.: Block diagram of the linearised PLL model. Included are bit length indicators (N,M,K,F,C,T), possible truncation noise additions ($\widetilde{u}_e$, $\widetilde{u}_f$, $\widetilde{u}_p$), input additive noise $\widetilde{\varepsilon}_{\mathrm{add}}$ and markers for signal readout points (Q, PA, PIR). Not shown here are the amplitude readout and the additional computation delay transfer function. The controller and the low pass filter can be implemented according to signal and requirement specifications. The input and output of the look-up table (LUT) is kept at an equal bit length in the linear design of the model. (A larger output does not lead to phase noise improvement, since the phase information is already lost before the LUT)

4.2.1. Linear model

The specific linear model presented here is a modification of well known ADPLL models described by Gardner [37]. Figure 4.2 shows the block diagram of the model, which considers phase as the quantity that is sensed and actuated [90, 91]. The signals and blocks shown are described in the following.

Assuming an input signal with a peak amplitude V_{in} in Volts and a maximum peak-to-peak voltage range of the analog-to-digital converter of $V_{\mathrm{p\text{-}p}}$, the digitised input signal $i[n]$ [1] is

$$i[n] = \frac{V_{\mathrm{in}}}{V_{\mathrm{p\text{-}p}}} \cdot \sin\left(\omega_0 n + \varepsilon_i[n]\right) = A \cdot \sin\left(\omega_0 n + \varepsilon_i[n]\right),$$
$$[i[n]] = 1; \ (-0.5 \leq i[n] < 0.5)$$
$$[A] = 1; \ (0 \leq A < 0.5)$$

(4.5)

with $\varepsilon_i[n]$ as phase, the signal of interest, and ω_0 ($[\omega_0] = \mathrm{cycle}$) as value corresponding to the beatnote frequency f_0 ($[f_0] = \mathrm{Hz}$) in a digital system sampled by a sampling frequency f_s. This leads to time steps between two samples n and $n+1$ of $\tau_s = 1/f_s$. One should note here that ω_0 is only necessary for

---

[1] In the following a unit $[x]$ and range $(\min < x < \max)$ is specified for each quantity x.

the initial loop acquisition and not for the linear model, it is included here to keep the resemblance to the actual signals inside the PLL. Additional terms for additive noise and additional tones are not included here. The output $o[n]$ of the numerically controlled oscillator (NCO) is described as

$$o[n] = \frac{1}{2} \cdot \cos\left(\omega_0 n + \varepsilon_o[n]\right),$$
$$[o[n]] = 1; \quad (-0.5 \le o[n] < 0.5) \tag{4.6}$$

with $\varepsilon_o[n]$ as NCO output phase, the current PLL reference. The ideal error signal of the loop ($\varepsilon_e = \varepsilon_i[n] - \varepsilon_o[n]$) is not directly accessible by arithmetic operations, therefore it is approximated by multiplying both signals to compute an error signal $u_e[n]$,

$$u_e[n] = i[n] \cdot o[n]$$
$$u_e[n] = \frac{A}{4} \cdot \left[\sin\left(\varepsilon_e[n]\right) + \sin\left(2\omega_0 n + \varepsilon_i[n] + \varepsilon_o[n]\right)\right]. \tag{4.7}$$

At this point two linearisations are introduced to complete the linear model. First we assume a small phase error ($\varepsilon_e[n] \ll 1$), which implies the loop is locked with sufficient loop gain, and second, we assume a suppression of the second harmonic term by appropriate filtering (this also includes the suppression of additional tones). This simplifies Equation 4.7 to

$$u_e[n] \approx \frac{A}{4}\left(\varepsilon_i[n] - \varepsilon_o[n]\right) = \frac{A}{4}\left(\varepsilon_e[n]\right),$$
$$[u_e[n]] = \mathrm{rad}. \tag{4.8}$$

The phase detector can now be described with a linear transfer function including the signal amplitude as part of its gain,

$$F_{\mathrm{PD}}(z) = \frac{u_e(z)}{\varepsilon_i(z) - \varepsilon_o(z)} = \frac{u_e(z)}{\varepsilon_e(z)} = \frac{A}{4},$$
$$[F_{\mathrm{PD}}(z)] = 1. \tag{4.9}$$

Here z is the discrete-time operator defined in Equation 4.4. A generic low pass filter follows the phase detector and provides the suppression of higher harmonics. The design and implementation of this filter depends on the exact loop design and should be adapted accordingly. A more detailed discussion is shown in the non-linearity section.

$$F_{\mathrm{LF}}(z) = \frac{u_d(z)}{u_e(z)},$$
$$[F_{\mathrm{LF}}(z)] = 1. \tag{4.10}$$

The open-loop gain of the PLL is determined by a controller, for example, a simple proportional-integral controller. For our implementation the full loop

model shows that an overall gain reduction in the loop is necessary to achieve a stable condition. Therefore we include a constant gain reduction before the servo, to allow the system to operate at the correct range and to prevent any overflows in the digital accumulators, where the fixed-point arithmetic is performed. For convenience we use bit shifting, adding a number of C bits from the left to the signal leads to

$$F_{\text{Gain}}(z) = \frac{u_g(z)}{u_d(z)} = 2^{-C},$$

$$[F_{\text{Gain}}(z)] = 1.$$

(4.11)

In the servo amplifier the desired bandwidth and loop response can then be set, by tuning the κ_p and κ_i values:

$$F_{\text{PI}}(z) = \frac{u_f(z)}{u_g(z)} = \kappa_p + \kappa_i \frac{z^{-1}}{1 - z^{-1}},$$

$$[F_{\text{PI}}(z)] = \frac{\text{cycle}}{\text{s} * \text{rad}}.$$

(4.12)

The frequency signal u_f is now representing the frequency of the NCO and, assuming the PLL is locked, it can be used to determine the phase of the incoming signal. The register containing this value is also denoted as phase increment register (PIR). For lock acquisition, this value must be pre-set close to the incoming frequency.

By accumulating the PIR value in a register called phase accumulator (PA) the phase driving the NCO is generated:

$$F_{\text{PA}}(z) = \frac{u_p(z)}{u_f(z)} = \frac{z^{-1}}{1 - z^{-1}}$$

$$[F_{\text{PA}}(z)] = \text{s}.$$

(4.13)

This phase is then fed into a sine and cosine look-up table to generate the local oscillator. In the loop, this operation is described by the transfer function

$$F_{\text{LUT}} = \frac{\varepsilon_o(z)}{u_p(z)} = 2\pi$$

$$[F_{\text{LUT}}] = \frac{\text{rad}}{\text{cycle}}.$$

(4.14)

One additional element not yet included are the delays of the signal processing. These delays become important for high bandwidth and they can directly be computed from the number of registers used in the loop logic. For a total delay of D clock cycles they are included as z^{-D}. If parts of the loop are running at slower frequencies, the delays should be scaled according to the sampling rate of the signal.
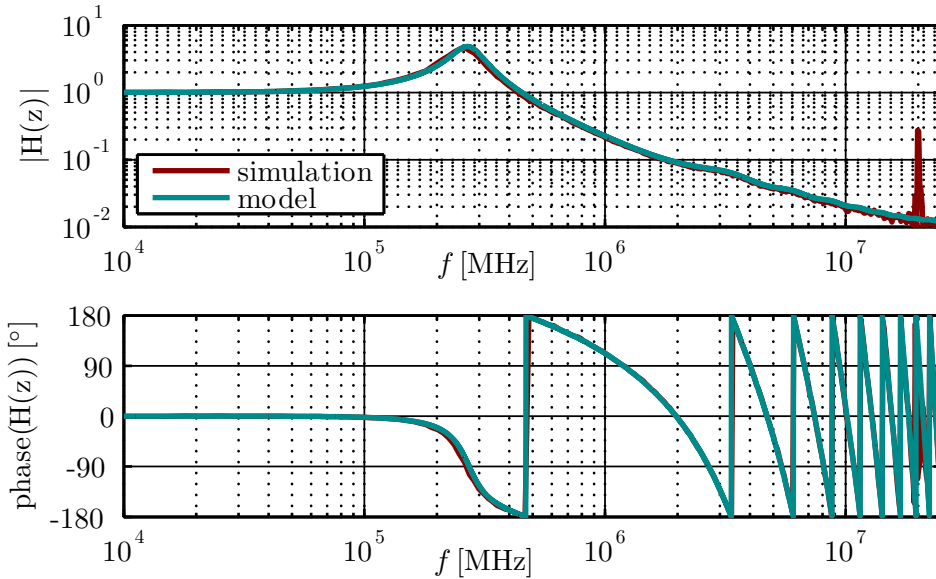
Figure 4.3.: Example of a closed-loop PLL transfer function $H(z)$ from a simulation and the linear model. Both curves are in very good agreement, save for the second harmonic, present at $20\,\mathrm{MHz}$.

Continuing, the above results in an open loop transfer function $G(z)$, which allows to to determine loop stability, noise suppression and suppression of higher harmonics:

$$G(z) = \frac{\varepsilon_o(z)}{\varepsilon_e(z)} = \frac{A\pi}{2} \cdot F_{\mathrm{LF}} \cdot 2^{-\mathrm{C}} \cdot \left(\kappa_p + \kappa_i \frac{z^{-1}}{1 - z^{-1}}\right) \cdot \left(\frac{z^{-1}}{1 - z^{-1}}\right) \cdot z^{-\mathrm{D}}. \quad (4.15)$$

The system transfer function $H(z)$ and the error function $E(z) = 1-H(z)$, which describes the untracked parts of the input signal and therefore the tracking error, are derived from the open-loop transfer function.

$$H(z) = \frac{G(z)}{1 + G(z)} = \frac{\varepsilon_o}{\varepsilon_i}, \quad (4.16)$$

$$E(z) = \frac{1}{1 + G(z)} = \frac{\varepsilon_e}{\varepsilon_i}. \quad (4.17)$$

This model was validated by performing a loop gain measurement of an early VHDL implementation of an ADPLL. This was done by injecting digital noise (using the dither generator) into the ADPLL and by measuring the in- and output of the adder used for the injection. The results in Figure 4.3 show a very good agreement between the model and the simulation, with the exception of the peak corresponding to the second harmonic frequency at $20\,\mathrm{MHz}$. Due to

the approximations done at the beginning of the linear model this loop analysis is only correct if the gain of the loop is sufficient to maintain the error point close to zero. This is discussed in more detail in Section 4.4.

4.2.2. Additive and phase noise

We now extend the input signal (Equation 4.5) by including additive noise $\widetilde{A}$ (including shot noise, electronic noise and relative intensity noise) and phase noise $\widetilde{\varepsilon}$. Since the PLL can not distinguish between phase noise and phase signal ε_s both terms can again be described by a single term ε_i.

$$i[n] = \widetilde{A}[n] + A \cdot \sin\left(\omega_0 n + \widetilde{\varepsilon}_i + \varepsilon_s[n]\right) = \widetilde{A}[n] + A \cdot \sin\left(\omega_0 n + \varepsilon_i[n]\right). \quad (4.18)$$

For phase noise the standard deviation of the residuals σ_{phase} can be computed by integrating the product of the noise with the loop error function $E(z)$:

$$\sigma_{\text{phase}}^2 = \int_0^\infty \widetilde{\varepsilon}_i^2(z) \times E(z)^2 df \qquad (4.19)$$

This is a measure of untracked residual phase error.

The standard deviation of the error generated by input additive noise σ_{add} is computed by integrating the product of the effective phase noise with the system transfer function $H(z)$, since this transfer function describes how noise added to the error signal $\varepsilon_e(z)$ is attenuated in a closed loop. Due to the mixing process the amplitude noise induced phase noise $\widetilde{\varepsilon}_{\text{amp}}$ is also increased by $\sqrt{2}$

$$\sigma_{\text{add}}^2 = \int_0^\infty \left(\frac{\sqrt{2}\widetilde{A}}{A}\right)^2 (z) \times H^2(z) df = \int_0^\infty \left(\widetilde{\varepsilon}_{\text{add}}\right)^2 (z) \times H^2(z) df. \qquad (4.20)$$

The different treatment of phase noise and additive noise, both of which are present in the input signal, can be understood if one considers phase as the quantity that propagates around the loop. Input phase noise directly represents that and an increased bandwidth allows tracking this phase more precisely. Additive noise as such does not represent a phase error. It only gets converted into phase noise by the action of the mixer (phase detector) which is *in the loop*, hence the different transfer function.

4.2.3. Amplitude detection

The detection of the amplitude A of the incoming signal is performed by multiplying $i[n]$ with an in-phase output of the NCO $I[n]$.

$$\begin{aligned} I[n] &= \frac{1}{2} \cdot \sin\left(\omega_0 n + \varepsilon_o[n]\right), \\ [I[n]] &= 1. \end{aligned} \qquad (4.21)$$

The multiplication of $i[n]$ and $I[n]$ gives

$$u_I[n] = \frac{A}{4} \cdot \big[\cos\left(\varepsilon_i[n] - \varepsilon_o[n]\right) - \cos\left(2\omega_0 n + \varepsilon_i[n] + \varepsilon_o[n]\right) \big]. \qquad (4.22)$$

Assuming a locked PLL ($\varepsilon_i - \varepsilon_o = \varepsilon_e << 1$) and a sufficient filtering of the second harmonic, this can be reduced to

$$u_I[n] = \frac{A}{4} \cos\left(\varepsilon_e[n]\right) \approx \frac{A}{4}. \qquad (4.23)$$

Therefore the readout of u_I (I) yields directly the amplitude $\frac{A}{4}$ of the tracked tone, while DC offsets and signals at sufficiently different frequencies average to zero. Knowledge of the signal amplitude is required to understand the loop bandwidth, to track changes in the interferometry, like in contrast and optical power, and to perform calculations involving the vector properties of the input signal, like, for example, stray light corrections [92].

4.2.4. Noise shaping inside the PLL

The linearised ADPLL model can also be used to understand the effect of in-loop truncation noise on the phase tracking performance, by applying standard control theory. Truncations inside the ADPLL are necessary to reduce the required number of bits and computational resources to a suitable amount. Without any truncations an ADPLL would require an infinite amount of bits.

Truncations at various points in the loop are possible and have been considered, as indicated in Figure 4.2. The effects of these truncations can, assuming dithering and rounding are applied, be computed by using the loop transfer functions, Equation 4.3 and the appropriate scaling.

As an example we evaluate here the influence of a truncation of the frequency value $u_f(z)$, it shows a strong noise shaping due to the loop transfer functions and also allows a strong reduction of bits.

Frequency truncation

A naive PLL implementation would use a high number of bits at $u_f(z)$ [35]. This is because truncations at this point in an open-loop system are especially critical, since they introduce a white frequency noise, which leads to a $1/f$ phase noise inside the PLL. The linear model shows, however, that this noise is suppressed directly by the loop error function $E(z)$. Since the PLL includes a f^2 suppression at low frequencies, due to the double integrator in the open loop gain (see Equation 4.15), the effective phase noise is easily reduced below $1\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ in the LISA signal bandwidth for only 12 bits at $80\,\mathrm{MHz}$.

$$\widetilde{\varepsilon}_{u_f}(z) = \sqrt{3}\, \frac{f_s}{f}\, \frac{2^{-T}}{\sqrt{6 \cdot f_s}} * E(z) \text{ rad}. \qquad (4.24)$$

Such a truncation was used during all ADPLL implementations used in this thesis, not limiting the performance, even though the absolute LSB frequency resolution for this example was only ≈ 20 kHz.

As already mentioned the bit reduction of this specific signal is especially useful, since the readout of the PLL frequency requires the highest dynamic range of all PLL signals. The noise shaping decreases the required number of initial bits to be decimated and therefore also reduces all bit lengths in further processing.

4.3. Readout

The ADPLL operates at MHz sampling frequencies and contains various signals. Even though one would ideally collect all of this information it is just not possible to store these amounts of data or to communicate them to ground. (A single ADC channel sampling with 14 bit at 80 MHz generates more than 1 Gbit of data per second). Other subsystems also require information from the phasemeter in real time, for example, the drag-free and attitude control system for LISA and LISA Pathfinder. Therefore the extraction of the phase from the ADPLL and the decimation to lower sampling rate are described in the following.

4.3.1. Frequency readout

The phase ε_i tracked by the ADPLL can be reconstructed by reading the frequency value u_f (PIR) or the phase value u_p (PA) of the PLL, which represent the frequency/phase of the incoming signal, respectively. For the PA:

$$
\begin{aligned}
u_p[n] &\approx (\omega_0 n + \varepsilon_i[n])/2\pi \quad \text{(for } \varepsilon_e[n] \ll 1\text{)}. \\
[u_p[n]] &= \text{cycle}; \quad (-\pi \text{ rad} \le (2\pi \times u_p[n]) < \pi \text{ rad}).
\end{aligned}
\tag{4.25}
$$

The frequency offset in a heterodyne system requires a constant phase ramp in the ADPLL in addition to the actual phase signal. Due to this large ramp a direct readout of u_p is not practical since this value will overflow very quickly. A decimation of such a sawtooth function is difficult, and the dynamic range for a non overflowing value of u_p is very large. The preferred possibility for the phase readout of a single loop is the frequency value u_f.

$$
\begin{aligned}
u_f[n] &\approx (\omega_0 + \frac{\delta \varepsilon_i}{\delta \tau_s})/2\pi \quad \text{(for } \varepsilon_e[n] \ll 1\text{)}. \\
[u_f[n]] &= \frac{\text{cycle}}{\text{s}}; \quad (0 \text{ Hz} \le (f_s \times u_f[n]) < f_s).
\end{aligned}
\tag{4.26}
$$

This value is not overflowing and allows for standard decimation and filtering algorithms to be implemented, though one has to keep in mind that this signal has a large dynamic range. Any requirements on decimation filters and

bit length have to take into account that the signal of interest (phase) is not directly processed, but its derivative, the frequency, which changes its spectral properties. The phase fluctuations can easily be reconstructed afterwards by integration.

4.3.2. Phase accumulator readout

If several channels track the same frequency and they only vary slightly in phase (Δu_p), the differences of the PLL phases can be read out directly by subtracting the PA values. The rapid ramp present in the individual loops is thereby completely subtracted and only the small signal of interest remains. This method was conceived as part of this thesis and it is an ideal candidate for the phase readout for differential wavefront sensing (DWS). An implementation of this technique is described in Chapter 7.

Even though the small differences in phase can also be reconstructed from the PIR readout, the PA readout is preferred. This is because the PIR values need to be tracked continuously to reconstruct the correct absolute phase values. This means that any glitches or cycle slips (see Section 4.4) will break the reconstruction. Even though the reconstruction can be restarted, a new initialisation of the PLLs would be required. In contrast to that the PA readout does not break by such an event, (assuming in both cases that the PLLs stay in lock) but it would automatically return to the correct value after the event passed, making it more reliable. The PA difference signals also have a much smaller dynamic range, allowing a reduction of the required bit lengths and computational efforts in further processing.

For a quadrant photo diode one therefore ideally reads out the sum of the four frequencies and the four PA differences required for DWS (see Equation 2.28). Thereby only one high dynamic range signal needs to be decimated together with four ones of small dynamic range. This simple scheme does however not include any consideration for redundancy as required for later space implementations.

4.3.3. Additional IQ readout

If the root mean square (RMS) of the phase error in the loop ε_e exceeds the acceptable noise level, because the loop reacts too slowly to track precisely the phase fluctuations of the incoming signal ($\mathrm{RMS}[\varepsilon_e(< 1 \text{ Hz})] > 6 \,\mu\mathrm{rad}$ for LISA), an additional corrective readout can be performed [36]. This might be necessary if the required PLL bandwidth needs to be rather low to achieve stable operations (see Section 4.4).

Since the untracked signal in a PLL is a vector and not a scalar, the readout of both quadrature components I and Q is required for the additional phase reconstruction. For a loop locked near but not exactly on zero phase difference,

they can be written as

$$Q = u_e[n] = \frac{A}{4} \sin\left(\varepsilon_e[n]\right)$$
$$I = u_I[n] = \frac{A}{4} \cos\left(\varepsilon_e[n]\right). \tag{4.27}$$

The residual phase ε_e can be reconstructed by computing

$$\varepsilon_e = \arctan\left(\frac{u_e[n]}{u_I[n]}\right) = \arctan\left(\frac{Q}{I}\right). \tag{4.28}$$

Equally the amplitude A of the vector in this situation can be computed as

$$\frac{A}{4} = \sqrt{u_e[n]^2 + u_I[n]^2} = \sqrt{Q^2 + I^2}. \tag{4.29}$$

The implementation of an Arctangent calculation in an FPGA is, however, not very efficient, therefore this calculation has to be done at a lower sampling rate and, preferably, in a different system, for example, a CPU. Which readout is required can be evaluated by comparing the PLL bandwidth with the dynamic range of the incoming signal. For the designs in this thesis a controller was used that has sufficient signal suppression at low frequencies to reach the required performance without additional IQ readout. Nevertheless it was still implemented for diagnostic purposes.

4.3.4. Decimation

The above mentioned signals need to be decimated to a desired sampling rate (typically of the order of a few Hz) for storage and further computation. The decimation can be implemented in one or in several steps and can make use of different computation methods, based on the hardware used. Here we only describe the initial decimation taking place inside the FPGAs, which is normally restricted to use integer parallel processing. A description of decimation using floating point processing is given for an example in Chapter 9.

Cascaded integrator-comb (CIC) filters [93] are chosen for the decimation inside the FPGAs. Their implementation is simple (they only require accumulators and differentiators), they are easily modelled and they provide notches of suppression exactly at the most critical frequencies, the ones that would be aliased to very low frequencies. Which order of filter is required can be computed for each signal by comparing the sum of all frequencies filtered and aliased into the signal band to the requirements. Since the suppression of CIC filters is increasing with frequency and the signal dynamic is, in the case of LISA like signals, decreasing with frequency, this calculation is completely dominated by the first notch of the filter. Such a calculation for a specific case is also shown in Chapter 9.

The use of CIC filters also allows the phasemeters to utilise an additional noise shaping technique [94]. This technique reduces the readout bit length of some signals by shaping the quantisation noise introduced within the filter. Figure 4.4
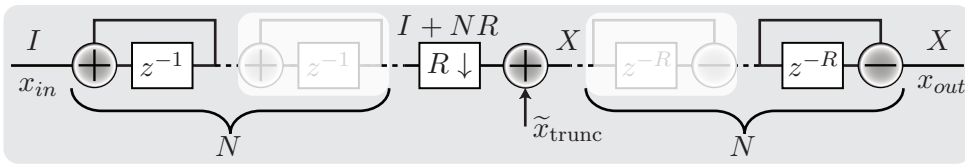
Figure 4.4.: Sketch of a 2nd order CIC decimation filter

shows the implementation of a CIC filter in a block diagram. The decimation is achieved by first integrating the signal in accumulators and then decimating the signal using combs, differentiators operating at the lower sampling rate. The order N of the filter is determined by the number of stages implemented. The integer rate change between the two sampling frequencies is R and for convenience this rate can be chosen as $R = 2^D$, to optimise the filter for binary computation. The transfer function of a CIC filter consists of two parts, the first includes the integrator stages.

$$F_{\text{CIC,int}} = \left(\frac{1}{1 - z^{-1}} \right)^N . \tag{4.30}$$

After this the transition between components operating at the fast rate and the slow rate is done. The comb stages are then described by another transfer function.

$$F_{\text{CIC,comb}} = \left(1 - z^{-R} \right)^N . \tag{4.31}$$

The individual functions are shown for an example in Figure 4.5. The product of the two stages yields the overall behaviour for a decimated signal.

$$F_{\text{CIC}} = F_{\text{CIC,int}} \cdot F_{\text{CIC,comb}} = \left(\frac{1 - z^{-R}}{1 - z^{-1}} \right)^N . \tag{4.32}$$

As described by Lehtinen et al. [94] any noise introduced between the individual stages by quantisation (Figure 4.5 shows an example) is attenuated by the combs, which decreases the noise below the corresponding truncation noise at that bit length. Additionally the combs do only transmit DC values that have been amplified by the integrator stages, small rounding error from a truncation inside the filters are therefore also suppressed, allowing simple re-truncations to be used without dithering and rounding.
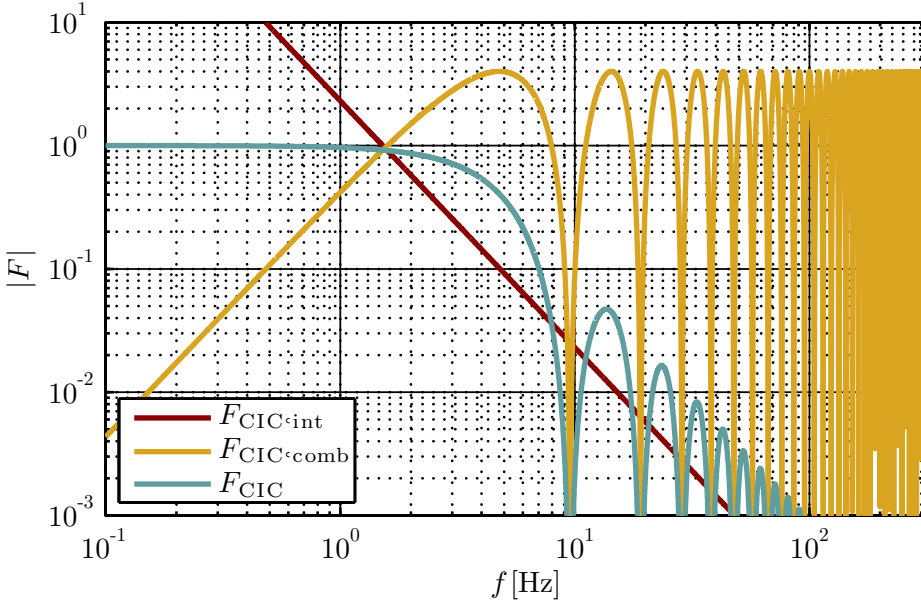
Figure 4.5.: Transfer function of a CIC decimation filter, also shown are the individual TF of the integrator and comb stages. The filter parameters are $f_s = 40\,\mathrm{MHz}$, $N = 2$, $R = 2^{22}$.

4.4. Non-linearity and cycle slips

The linear analysis of the ADPLL is valid for many applications. Phasemeters based on this analysis were already successfully tested and used in various laboratory experiments [58, 95, 43].

But a phasemeter in an inter-satellite interferometer has to operate under rather extreme conditions, a low signal-to-noise ratio of the input signal (due to the additive noise) and a large dynamic range of the signal phase, due to frequency noise and signal dynamics. Therefore the PLL could reach a state where it becomes non-linear and where cycle slips occur in the PLL tracking [96, 97, 98]. These slips are a jump of the phase tracking by one or more full cycles. The resulting phase noise from R slips during a measurement period is given as [99]

$$\phi_{slip}(z) = \frac{\sqrt{2}\sqrt{R}}{f} \ \mathrm{rad.} \tag{4.33}$$

This means that any slip spoils the system performance and introduces large artefacts in, for example, DWS signals. For LISA a cycle slip is therefore almost comparable to a loss of lock or another similar measurement disturbance. Earlier experimental investigations by Dick et. al. [99] and detailed modelling [96, 97] have shown that the relation between the bandwidth and the signal noise floor
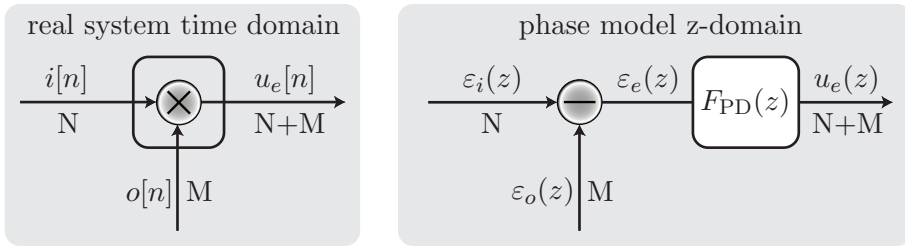
Figure 4.6.: Comparison of phase detector in the real system and the phase model.

is the critical factor for the probability of cycle slips.

Since the phasemeter should operate far outside any cycle slip region and therefore in the linear regime a model was compiled that determines a suitable loop bandwidth for a given set of signal parameters. This should minimise the probability of a cycle-slip and non-linear effects of the phase tracking.

The two most important reasons for non-linear behaviour are the sinusoidal response of the phase detector and the existence of second harmonics and other additional tones, like the side-band beat notes for inter spacecraft clock transfer or an ADC pilot tone.

4.4.1. Phasedetector

The non-linear output of the phase detector, omitting the second harmonic, is

$$u_e[n] = \frac{A}{4} \cdot \sin\left(\varepsilon_i[n] - \varepsilon_o[n]\right) = \frac{A}{4} \cdot \sin\left(\varepsilon_e[n]\right). \tag{4.34}$$

A direct comparison of the block diagram for the earlier used linearised model and the real implementation is shown in Figure 4.6.

For an error signal $\varepsilon_e[n] \ll \pi/2$ we can assume a quasi-linear behaviour of the phase detector. If the error signal exceeds $\pi/2$ the loop gain starts to reduce until it crosses zero and changes sign at $\varepsilon_e[n] = \pi$ (the turning point of the sine function). At any of these points the loop is potentially unstable and the error signal eventually slips by 2π or more, which results in a phase tracking error of the same amount.

4.4.2. Optimal bandwidth

Since the absolute error signal is directly related to the linearity of the ADPLL one can use the calculated standard deviations (see Section 4.2.2) to evaluate the size of the error signal. The internal quantisation noise influences can be added quadratically (assuming uncorrelated noise). The comparison of the resulting standard deviations can be compared for different bandwidth and phase margin (damping) configurations to find optimal parameters for an ADPLL. The

standard deviations are added quadratically to compute the resulting overall standard deviation σ_{sum}.

Figure 4.7 shows the modelled standard deviations for example parameters (discussed in Chapter 6) and their dependency on the loop gain. To verify the model the standard deviation of the PLL error signals were directly measured as part of this thesis. The procedure and the parameters used for this are described in Chapter 6. The measured values are shown as dots. The model shows excellent agreement between the predictions and the measured values, which verifies that the linear model is appropriate for this range of operation and to determine an optimal bandwidth. Operating the PLL at this point should minimise any non-linear phase artefacts and the probability of cycle slips. Though one can not deduce the exact probability for a given σ_{sum} yet, one can test the system for stability and performance using the so determined bandwidth, which is the approach used in this thesis.
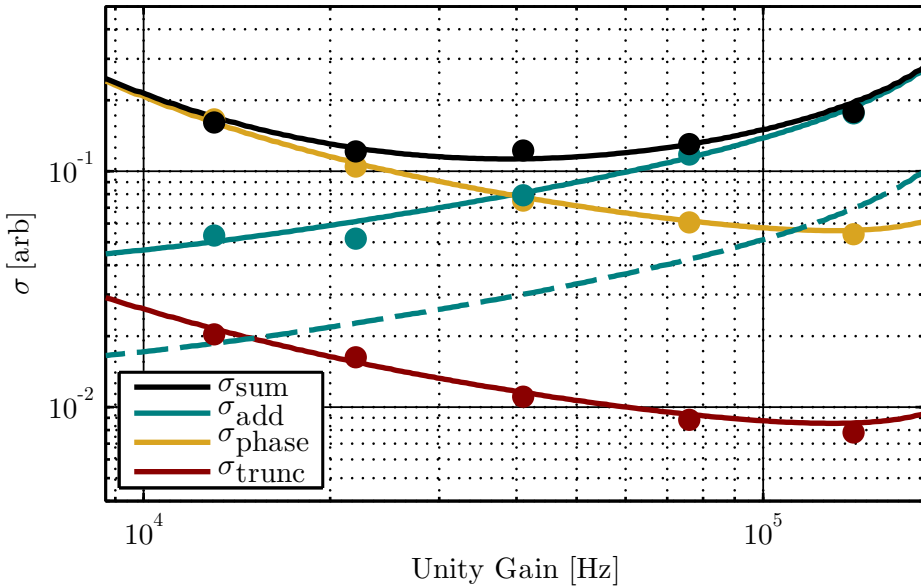
Figure 4.7.: Shown is the $1\,\sigma$ standard deviation of the error point from additive noise (blue), phase noise (yellow), truncation noise (red) and of their quadratic sum (black). The level of the additive noise is $72\,\mathrm{dBHz}$ (solid line) and $80\,\mathrm{dBHz}$ (dashed line) respectively. The phase noise is the laser frequency noise described in Chapter 2. The dots show the measured values of σ. The optimal bandwidth for this example is found as $\approx 40\,\mathrm{kHz}$.

4.4.3. Second harmonic

The second non-linear behaviour of the phase detector is the generation of a second harmonic of the input signal, as shown in Equation 4.7.

One can split the effects of the second harmonic into two parts. The first effect creates parasitic phase noise in the signal band, which is described in detail in the following. For convenience the second harmonic part of Equation 4.7 is rewritten in the continuous time domain

$$u_{e,2f}(t) = \frac{A}{4} \cdot \sin\left(2\omega_0 t + \varepsilon_i(t) + \varepsilon_o(t)\right). \tag{4.35}$$

This equation is simplified by assuming the PLL to be tightly locked ($\varepsilon_o \approx \varepsilon_i$) and by defining an effective phase value $\varepsilon_{\text{eff}}(t) = \omega_0 t + \varepsilon_o(t)$, with an effective frequency $\omega_{\text{eff}} = \frac{\delta\varepsilon_{\text{eff}}}{\delta t}$. The second harmonic propagates through the PLL in a time τ_g and creates an effective phase modulation. The output of the NCO at the time t can therefore be written as

$$\text{NCO}_{\text{out}}(t) = \frac{1}{2} \cos\left(\varepsilon_{\text{eff}}(t) + m \cdot \sin\left(2\varepsilon_{\text{eff}}(t - \tau_g)\right)\right). \tag{4.36}$$

Here m is a modulation index given by the attenuation of the second harmonic by the open loop transfer function ($m = |G(2\omega_{\text{eff}})|$), referred to the gain for the nominal low-frequency error signal, which in the signal range is ≈ 1. Using Bessel functions of the first kind one can expand this to

$$\begin{aligned}
\text{NCO}_{\text{out}}(t) =& \frac{1}{2} J_0(m) \cos(\varepsilon_{\text{eff}}(t)) + \frac{1}{2} J_1(m) \sin(\varepsilon_{\text{eff}}(t)) \sin(2\varepsilon_{\text{eff}}(t - \tau_g)) + \mathcal{O}(m^2) \\
=& o(t) + o_{2f}(t) + \mathcal{O}(m^2).
\end{aligned} \tag{4.37}$$

Assuming $m \ll 1$ one can approximate the first two Bessel function by $J_0(m) \approx 1$ and $J_1(m) \approx m/2$. This yields the original NCO output $o(t)$, the term from the second harmonic $o_{2f}(t)$ and higher terms $\mathcal{O}(m^2)$, which are discarded in the following. o_{2f} is then rewritten and the third harmonic term is immediately discarded.

$$\begin{aligned}
o_{2f}(t) =& \frac{1}{2}\frac{m}{4}\left(\cos(\varepsilon_{\text{eff}}(t - \tau_g)) - \cos(3\varepsilon_{\text{eff}}(t - \tau_g))\right) \\
o_{2f}(t) \approx& \frac{1}{2}\frac{m}{4} \cos(\varepsilon_{\text{eff}}(t - \tau_g))
\end{aligned} \tag{4.38}$$

The phase modulation side-band at $\omega - 2\omega = -\omega$ thus ends up at the same frequency as the nominal NCO output and results in a parasitic phase signal by the action of the mixer. The additional mixer output is

$$i(t) \times o_{2f}(t) \approx \frac{A}{4}\frac{m}{4} \sin\left(\varepsilon_{\text{eff}}(t) - \varepsilon_{\text{eff}}(t - \tau_g)\right). \tag{4.39}$$

The effective parasitic phase error $\varepsilon_{p,2f}$ is therefore (see Equation 4.8)

$$\varepsilon_{p,2f}(t) = \frac{m}{4}\sin\left(\varepsilon_{\text{eff}}(t) - \varepsilon_{\text{eff}}(t - \tau_g)\right). \tag{4.40}$$

Assuming a constant τ_g and an effective frequency ω_{eff} that varies on time-scales smaller than τ_g one can approximate this to

$$\varepsilon_{p,2f}(t) \approx \frac{|G(2\omega_{\text{eff}})|}{4}\sin\left(\omega_{\text{eff}}(t)\tau_g\right). \tag{4.41}$$

This parasitic noise couples very non-linear and depends highly on the suppression of the second harmonic m, the group delay in the PLL τ_g and the dynamics of the input signal $\omega_{\text{eff}}(t)$. The coupling is at its maximum in the linear range of the sine. Assuming this operating condition we can calculate the maximum parasitic phase error in dependency of the signal frequency noise spectrum.

$$\widetilde{\varepsilon}_{p,2f}(f) \approx \frac{|G(2\omega_{\text{eff}})|}{4}\widetilde{\omega}_{\text{eff}}(f)\tau_g. \tag{4.42}$$

Though this coupling is rather small it was observed during simulations with large laser frequency noise, as shown in Chapter 6.

The second effect caused by the second harmonic is an additional instantaneous root-mean square value of the error signal. Even though this does not cause a phase error at low frequencies, it does increase the probability to leave the linear range of the phase detector. The maximum additional error is $\varepsilon_{e,2f}(max) = |G(2\omega_{\text{eff}})|$.

Additional tones in the input spectrum only interfere with the measurement if they get very close to the tracked beat note (which is obvious) or if they are at its second harmonic frequency (due to the coupling just described). A frequency plan has to include these effects, as well as the aliasing of signals to other frequencies.

The above equations allow us to determine the necessary suppression by low pass filters in the ADPLL for a given system, by calculating the error signal residuals and by comparing the signal dynamics with the required phase performance. Since the choice of low pass filter is also limited by logic resources, a trade-off is necessary. Infinite impulse response (IIR) filters were found to be a good compromise between suppression and logic resources required. A second order IIR filter with a corner frequency of $300\,\text{kHz}$ is included in ADPLLs in this thesis for cases with large input channel dynamics. In critical cases, e.g. when the signal frequency can span a wide range, a more complex $2f$-filter could be used, for example one, that specifically filters the second harmonic frequency [100]. One should also consider that this analysis is only valid if the second harmonic is below the Nyquist frequency ($f_s/2$). If this is not the case the second harmonic will be aliased to another frequency and potentially not cause a parasitic phase error.
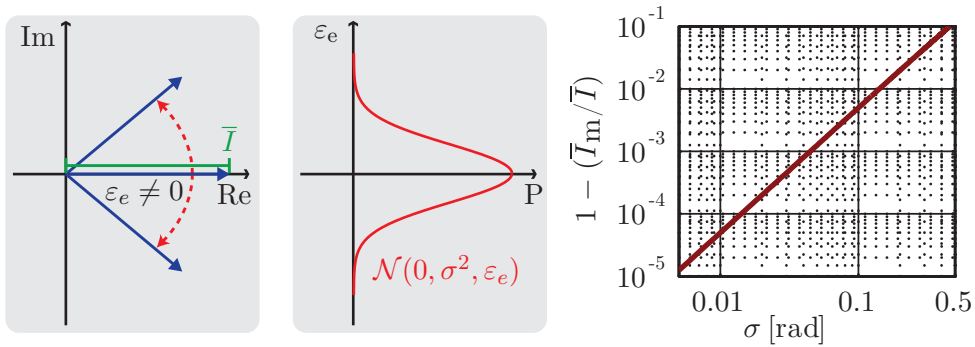
Figure 4.8.: The left side shows the deviation of the signal phasor if the PLL is locked. The decimated quantity is the projection of the phasor onto the real axis, reducing the measured length for a non-zero phase signal. The middle shows a sketch of the assumed normal distribution of the error signal. The right side shows the dependency of the signal reduction on the standard deviation of the error signal distribution (see Equation 4.44).

4.4.4. Amplitude readout non-linearity

The approximation for the amplitude readout (Equation 4.23) is only correct for small, average values of the phase error signal ε_e. In the case of larger RMS values the simple decimation of I will cause a deviation of the measured average value $\bar{I}_m$ to the real one ($\bar{I}$). The left sketch in Figure 4.8 illustrates this effect.

Analytically the measured amplitude can be calculated by integrating the original cosine dependency of the I value while weighing it with the probability distribution of ε_e.

$$\bar{I}_m = \int_{-\infty}^{\infty} \frac{A}{4} \cos\left(\varepsilon_e\right) \cdot \mathcal{N}(0, \sigma_{\text{sum}}^2, \varepsilon_e) \, d\varepsilon_e \tag{4.43}$$

Here the assumed probability is a normal distribution $\mathcal{N}$ with an average value of zero and a standard deviation of σ_{sum}. This equation can be solved analytically to determine the relative error of the decimated DC value for a given σ_{sum}:

$$\bar{I}_m/\bar{I} = e^{-\frac{\sigma_{\text{sum}}^2}{2}} \tag{4.44}$$

The resulting reduction of the measured amplitude is shown in the right plot in Figure 4.8. The correct amplitude value can only be recovered if the average phase noise is well known.

An independent readout of the RMS phase error might be implemented in the future, which would allow us to directly determine the ADPLL stability as well as the amplitude error in real time. It could also be useful for set-ups with non-stationary noise, where the noise floor and therefore the RMS error

signal changes during the measurement. Such a readout could, for example, be implemented by directly decimating Q^2 or by comparing the $\overline{I}_\mathrm{m}$ value with an additional readout of $I^2 + Q^2$, where the later one would potentially yield the correct signal amplitude.

Chapter 5

Testing phase readout performance

During the work for this thesis several phase readout systems were designed, implemented and tested. The most crucial aspect was to achieve the required phase measurement performance for each system. Each method available to test these phase noise levels has different sensitivities to non-linear effects and specific noise sources. The large dynamic range between, for example, signals initially measured in LISA and its final sensitivity (about nine orders of magnitude [95]) requires to perform tests that are able to identify such non-linearities. This chapter describes some basic testing methods, which were used and are considered most relevant, and their fundamental differences in the coupling or non-coupling of unwanted effects into performance measurements.

The first section concentrates on the measurement of a single, well defined signal. Even though this scheme was only rarely used in this thesis, studying it is useful to introduce some fundamental concepts. For phasemeter testing most often a test set-up was used where a signal is split and measured by different channels. This is described in the second section. Even though this set-up is very useful for investigating several noise sources, it has strong limitations to measure non-linear effects. Using three signals a test set-up can be constructed that is sensitive to almost all noise sources and non-linearities. This is described in the third section. The effective coupling of crosstalk in the three signal test is discussed in more detail in the last section.

5.1. Absolute signal measurement

For an ideal test a well known phase signal φ_t is transformed into a beat note with the desired properties and is then injected into a phase measurement chain (see Figure 5.1). Afterwards the measured phase φ_a is directly compared to the
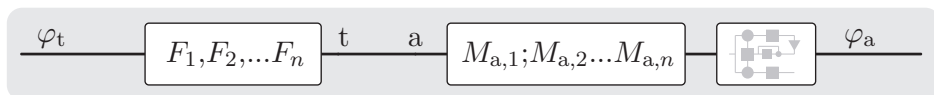
Figure 5.1.: Sketch of a generalised absolute signal measurement test set-up. A phase signal φ_t is converted into a beat note signal t by some elements F_x. This signal is then injected into the phase measurement chain 'a' consisting of elements $M_{a,x}$. After the phase φ_a is determined by a PLL it can be compared to the input signal.

input phase and any differences can be used to determine the noise floor of the system.

While one is only interested in the influence of the elements in the measurement chain $(M_{a,x})$ the signal generation elements (F_x) also contribute to the noise and one can not simply distinguish between their influences. The construction of a well defined input signal generator is possible (for example in digital, as shown in Section 6) but it has almost the same complexity as the phase readout system itself. However, as part of this thesis the design of an absolute analogue signal generator was investigated, which is described in Section 8.2. This scheme was in general not suitable for measuring analogue or optical signals, since this would have included various additional components that required a phase stability similar or better than the ones in the readout chain.

Even though no implementation of a signal generator with a performance suitable for μrad tests in the mHz regime is yet available it would be a valuable device. The reason for this is that this measurement principle is sensitive to all non-linear effects in the measurement chain, including crosstalk and non-linearities in the phase-locked loop (PLL). Each one of these effects generates an additional phase signal that could be detected immediately.

One aspect not captured either in Figure 5.1 or the above description is the influence of decimation filtering and aliasing. Both the input and the output phase signals would be defined at a certain rate and a comparison between them at a different rate would need to include sufficient filtering of at least the input signal. If the same filter is used for the input and the output signal the risk remains that aliased components are subtracted and thereby not detectable. Even though this is crucial, the filtering of data streams is a well known area of digital signal processing (DSP) and correct results can be achieved by applying standard techniques.

Performance requirements are often given for a single channel and are therefore directly applicable to $\varphi_t - \varphi_a$. No further scaling is required. Assuming a perfect signal generator the performance will directly reveal all additive noise and phase noise in the measurement chain.
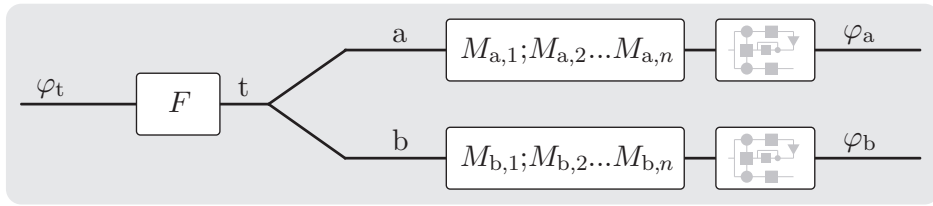
Figure 5.2.: Sketch of a generalised split-measurement test set-up. A signal t is split into two or more measurement channels with elements $M_{a,x}$ and $M_{b,x}$. The two measured phases φ_a and φ_b will only differ due to effects in the two measurement chains.

5.2. Split measurements

The test set-up that is used most often is referred to here as split measurement. It is often also called 'null' or 'zero' measurement, though these names are misleading. A phase signal that has the desired properties is converted into a beat note, split and then fed into at least two measurement chains, as shown in Figure 5.2. In this scheme neither φ_t nor t have to be known with the desired precision, since the comparison of φ_a and φ_b will allow one to detect any non-common noise sources in the two measurement chains.

The realisation of this measurement set-up is simple, since the phase performance of the signal generator is uncritical. This makes it very useful to measure noise that is uncorrelated between individual channels. Even though split measurements are very useful for hunting various noise sources they are also almost completely insensitive to non-linear effects. Crosstalk, for example, can, even if it couples differently between channels, not be detected, since all signals are almost exactly the same. The PLLs might even have cycle slips in their phase tracking, which would be completely hidden in the measured phase differences, if they occur in all channels at the same time.

In addition to this, the absolute phase performance of a device is also not easily testable. Common temperature fluctuations might, for example, induce a phase noise term in all channels. This would only be detectable if two devices are used for measuring a and b that are temperature-wise isolated from each other. The use of different signal frequencies is also not feasible, since this would require a perfect signal generator or some other very stable architecture to generate the two signals or to separate the split one into two frequencies.

Another property of split measurements is, that additive noise in the input signal t will not limit the phase performance, since it is common in all channels. This is somewhat counterintuitive, since the measured phase noise between two signals with low signal-to-noise ratio (SNR) wrongly appears to be much better than possible for those SNRs (see measurements in Section 9.5 for an example).

One variant of this measurement type is to use the two outputs of an interference beam splitter as inputs to the measurement chains. This is called a π measurement and its main difference to a split measurement is a phase shift of π between the signals a and b, hence its name. One can also split optical signals after they passed the photo diodes, such a measurement, however, is not sensitive to shot noise, because it will be common after the splitting, as described in the previous section. A π measurement is sensitive to shot noise and electronic noise in the photo diode, since it uses two, independent conversions of optical power into photo currents and further into voltages. Another variant is to split the optical signal with an additional 50/50 beam splitter after the interference beam splitter, thereby, two signals with a phase difference of $\varphi = 0$ are generated. Their difference includes shot noise, as well as electronic noise from the photo diodes.

For a split measurement the performance requirement for a single channel can be scaled if the limiting noise is equally distributed between both channels and uncorrelated. In that case the requirement has a value of $\sqrt{2}{\cdot}6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}{\cdot}\mathrm{NSF} \approx 8.5\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}} \cdot \mathrm{NSF}$. Throughout this thesis this value will be plotted as LISA requirement for split measurements. Since noise in each measurement could also be caused by a single channel, which would require to apply the more stringent requirement, the plotted phase performance curves are never scaled. This also easily allows comparing them to other requirements. If more than two channels are tested the difference between one channel and the scaled average over all channels (except for obviously distorted ones) is usually plotted.

5.3. Three signal testing

Due to the complexity of a the absolute signal measurement (especially potential optical implementations) and the limitations of the split measurement another basic architecture is the most promising candidate for performing comprehensive phase measurement performance tests. This architecture is in general called a three signal or non-linearity test (see Figure 5.3). As the name indicates it uses three signals. They are generated in such a way, that a linear combination of the three measured phase signals combines to zero if no noise or non-linearities are present.

$$0 = (\varphi_1 - \varphi_2) + (\varphi_2 - \varphi_3) - (\varphi_1 - \varphi_3)$$
$$0 \overset{?}{=} \underbrace{M_{\mathrm{a,x}}(a)}_{\varphi_{\mathrm{a}}} + \underbrace{M_{\mathrm{c,x}}(c)}_{\varphi_{\mathrm{c}}} - \underbrace{M_{\mathrm{b,x}}(b)}_{\varphi_{\mathrm{b}}} \tag{5.1}$$

This method combines many positive aspects of the other two. The absolute phase signal $(\varphi_1, \varphi_2, \varphi_3)$ can be constructed with the desired properties but the phases do not have to be known with the required performance, meaning no
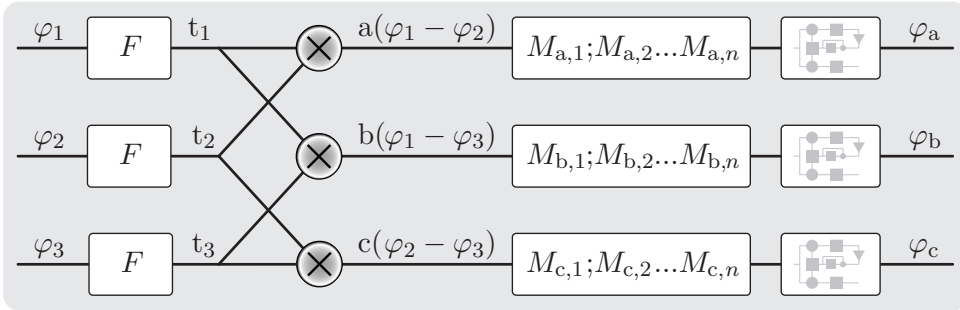
Figure 5.3.: Sketch of a generalised three signal measurement test set-up. The signals t_x are generated so, that the phase measurements can be combined to zero.

perfect signal generator is necessary. Since three signals are used they can be very different ($a \neq b \neq c$), allowing beat notes with completely different frequencies to be used, as well as additional tones, modulations and additive noise. This test is sensitive to all additive and phase noise in the three measurement channels.

The critical point in generating the three signals is to combine the initial tones (t_1, t_2, t_3) without introducing extra noise. Examples of possible implementations are shown in this thesis for three domains, digital (Chapter 6), analogue (Section 8.6) and optical (Section 8.7).

Even though this test makes it possible to investigate many non-linearities, effects of aliasing are not necessarily visible in the signal combination. This was found experimentally during this thesis and is understood to be caused by the fact, that even the down aliased signal components can combine to zero.

For the three signal test the requirement plotted for LISA throughout this thesis is $\sqrt{3} \cdot 6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}} \cdot \mathrm{NSF} \approx 10.4\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}} \cdot \mathrm{NSF}$, following the same approach as described at the end of the previous section. The measured phase combination signals are again not scaled, to allow an easy comparison to other requirements.

5.4. Crosstalk in the three signal test

The ability to use a three signal test to investigate crosstalk is discussed here. This is especially relevant since crosstalk measurements using standard equipment like oscilloscope will only partly be able to resolve suppressions in the order of $120\,\mathrm{dB}$ that are required for LISA (see Section 3.4). To test for this effect two of the three signal combinations need to be at the same beat note frequency. Depending on which beat notes from the combinations are chosen the sensitivity to crosstalk varies greatly.

Assuming signals a and b are at the same frequency their crosstalk induced phase noise is given as (see Equation 3.5)

$$\varphi a_{\mathrm{cross,b}} = C_{\mathrm{cross}} \frac{b^*}{a^*} \sin\left(\varphi_{\mathrm{a}} - \varphi_{\mathrm{b}}\right) \tag{5.2}$$

and

$$\varphi b_{\mathrm{cross,a}} = C_{\mathrm{cross}} \frac{a^*}{b^*} \sin\left(\varphi_{\mathrm{b}} - \varphi_{\mathrm{a}}\right) \tag{5.3}$$

with a^* and b^* as respective amplitudes and C_{cross} as crosstalk factor. One can now also write the second parasitic signal in dependency of the first one.

$$\varphi b_{\mathrm{cross,a}} = -\left(\frac{a^*}{b^*}\right)^2 \varphi a_{\mathrm{cross,b}} \tag{5.4}$$

The two general cases of coupling are the following.

Case 1: Full sensitivity If the signal combination shown in Figure 5.3 is chosen ($\varphi_{\mathrm{A+C-B}}$) the two parasitic signals will be subtracted and due to their opposite sign the results will directly reveal the crosstalk influence (neglecting all other effects for simplicity).

$$\varphi_{\mathrm{A+C\text{-}B}} = \varphi a_{\mathrm{cross,b}} + \left(\frac{a^*}{b^*}\right)^2 \varphi a_{\mathrm{cross,b}} \tag{5.5}$$

Case 2: Conditional sensitivity The three signals can also be combined in another combination, where a and b have the same frequency (A+B-C). For this set-up the measured phase performance $\varphi_{\mathrm{A+B-C}}$ will only be sensitive to crosstalk if the two signal amplitudes a^* and b^* are not equal.

$$\varphi_{\mathrm{A+B\text{-}C}} = \varphi a_{\mathrm{cross,b}} - \left(\frac{a^*}{b^*}\right)^2 \varphi a_{\mathrm{cross,b}} \tag{5.6}$$

If they are equal this combination is not at all sensitive to crosstalk.

One therefore has to be very careful to construct each measurement set-up in a way that fits the desired sensitivities. For optical interference signals the availability of a phase inverted output of each interference is very helpful, since it allows the signs of the signal combinations to be changed by simply switching between the two output ports.

Chapter 6

Phase readout simulations

The investigation of phase readout performance requires to test isolated parts of the measurement chain to identify noise sources and their coupling effects. Since the readout chain uses an all-digital phase-locked loop (ADPLL) to extract the phase from the beat notes, its performance has to be evaluated first. By using digitally generated signals for testing, all analogue and optical effects are fully excluded. A digital implementation also allows an easy generation of complex signals as input to the ADPLLs. This chapter describes the performance tests and simulations of the ADPLLs that were done during the course of the thesis. Large parts of the ADPLL model presented in Chapter 4 were based on and validated by these simulations.

The first section of the chapter describes shortly the framework that was used to implement the tests. It is very close to the one used for the later described hardware implementations and the benefits of this approach are also discussed. In the second section the generation of the simulated input signals is described, including the digital generation of additive noise and frequency noise as expected in LISA. This is followed by the results of digital three signal tests that were used to investigate the ADPLL performance under realistic conditions. Besides the measurement of non-linearities, the performance achieved with an optimal ADPLL design is also presented. The last section describes how the simulations were used to validate the error point distribution model presented in Chapter 4.

6.1. FPGA based performance simulations

The ADPLL simulation were implemented using VHDL code running either in a dedicated simulator software on a PC or directly on FPGAs, using the hardware

described in Appendix A.

The first approach allowed one to observe and influence all signals inside the simulation at the full sampling rate. The ability to acquire signals at this rate is very useful and was, for example, used to measure the ADPLL loop gain (see Figure 4.3). The speed and length of these simulations is limited by the available CPU power and storage. Measurement times longer than 1 s were not feasible. Therefore this approach was not suited to implement performance test that require measurement times equivalent to mHz frequencies.

By using FPGAs to host the simulations their full computation rates could be utilised. The numerical performance measurements were therefore done in quasi real time. Later on the use of additional debugging tools also allowed chunks of signals to be read directly from the FPGA implementations at the full sampling rate. The FPGAs used for the simulations were the same ones that were later used for the phasemeter hardware implementations. Therefore the same readout interfaces and software could be used. This also reduced the efforts of switching from a simulation to a measurement of real signals. The VHDL code containing the ADPLLs could also be ported without any changes and this reduced possible programming errors from adapting code.

In contrast to commercially available DSP software (like Matlab Simulink or Labview [101]) the direct use of VHDL code also made it possible to influence all computations and bit lengths directly. The VHDL coding in this thesis was mainly done in the IEEE 1076 standard from 1993 [102]. The syntax and the usage of more sophisticated programming elements (like records, functions and packages) have evolved during the thesis and no particular style was adopted.

6.2. Digital signal generation

Testing ADPLLs for linearity in a digital simulation requires to generate realistic signals. Various noise sources can be produced by filtering a white noise signal to shape it to the desired frequency dependency. Since physical noise sources like shot noise and laser frequency noise have a Gaussian probability distribution an almost Gaussian noise $\widetilde{n}$ was generated by adding 8 of the earlier described triangular dither signals [103]. Each of the contained PRN generators was initialised with an individual seed to suppress any cyclic signals. The measured distribution is shown on the left in Figure 6.1 together with a model of the assumed Gaussian noise distribution. The power spectral density of the simulated noise was also analysed and confirmed the desired white spectrum.

6.2.1. Simulating frequency noise

The generation for a single signal with the desired frequency noise is sketched in Figure 6.2. The spectral properties of the frequency fed into the NCO is defined
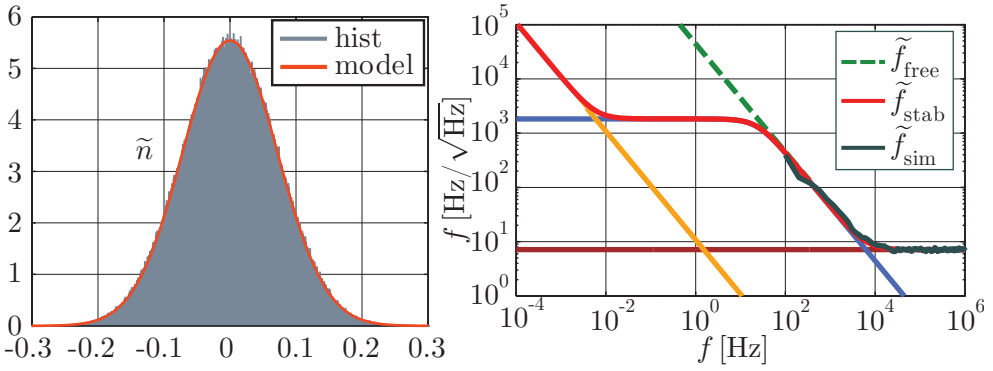
Figure 6.1.: (left) Normalised histogram of 800k samples of simulated Gaussian noise in comparison to a model assuming $\sigma = 0.072$. (right) Shown are the individual components and the compiled frequency noise shapes together with a measurement. The shapes are colour coded according to their generation, shown in Figure 6.2.

by three independent filters. The spectral composition of the assumed free running laser noise ($\widetilde{f}_{\text{free}}$) and effective laser noise at the master laser ($\widetilde{f}_{\text{m}}$) for LISA, as introduced in Chapter 2 is shown on the right in Figure 6.1. The specific levels shown here correspond to simulations using 80 MHz sampling frequency. These simulations were used to test the ADPLLs for the LISA phasemeter breadboard described in Chapter 9. The implementation of the gains was done using simple bit shifting, the same method used in the ADPLL. Therefore all gains are factors or inverse factors of two.

Each frequency noise consists of a white noise floor $\widetilde{f}_{\text{white}}$ that is simply generated by proportional scaling of $\widetilde{n}$.

$$\widetilde{f}_{\text{white}} = \underbrace{2^{-7}}_{\kappa_{\text{p}}} \frac{2^{-2} * 80\,\text{MHz}}{\sqrt{6 * 80\,\text{MHz}}} \tag{6.1}$$

The second component of both $\widetilde{f}_{\text{free}}$ and $\widetilde{f}_{\text{m}}$ is a 1/f noise. It is generated by simple integration of $\widetilde{n}$ with a gain that depends on the desired shape. The gains used were $K_{\text{int}} = 18$ for the free running noise and $K_{\text{int}} = 30$ for the effective noise at the master.

$$\widetilde{f}_{1/\text{f}} = \underbrace{2^{-K_{\text{int}}}}_{\kappa_{\text{int}}} \frac{z^{-1}}{(1 - z^{-1})} \frac{2^{-2} * 80\,\text{MHz}}{\sqrt{6 * 80\,\text{MHz}}} \tag{6.2}$$

For $\widetilde{f}_{\text{m}}$ a low pass filter is used to generate the flat response around 1 Hz and the drop to higher frequencies. A simple first order IIR filter was therefore implemented, again using only bit shifting. It is the same filter architecture that

97
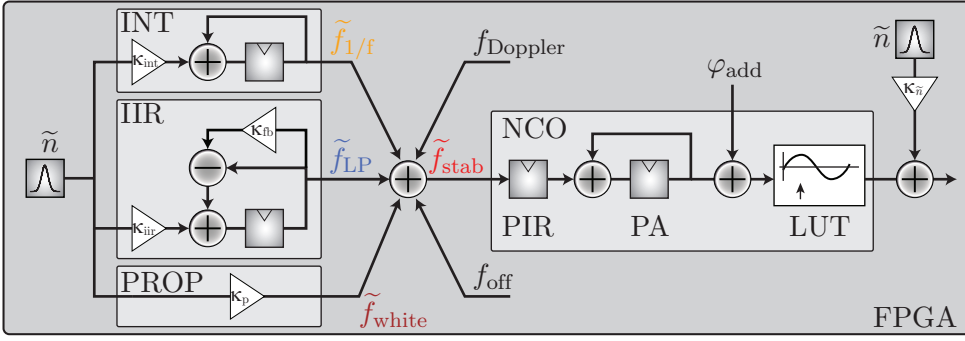
Figure 6.2.: Sketch of the signal generation inside an FPGA for digital performance measurements. The frequency noise is generated by three shaping filters, an integrator (INT), an infinite-impulse response filter (IIR) and a proportional element (PROP). In addition the frequency contains an offset f_{off} that defines the start frequency and potentially a ramp or sine wave simulating Doppler shifts $f_{Doppler}$. In the numerically controlled oscillator (NCO) the frequency value is integrated to a phase and here another phase signal φ_{add} might be added. The phase is then converted into a beat note with a look-up table (LUT). The desired additive noise is added in the last step.

is used in the ADPLL. This filter required an overall gain κ_{iir} and a feedback gain κ_{fb}.

$$\widetilde{f}_{LP} = \underbrace{2^{-18}}_{\kappa_{iir}} \frac{z^{-1}}{(1 - (1 - \underbrace{2^{-19}}_{\kappa_{fb}}) * z^{-1})} \frac{2^{-2} * 80\,\text{MHz}}{\sqrt{6 * 80\,\text{MHz}}} \tag{6.3}$$

The resulting frequency noise $\widetilde{f}_{m}$ was simulated for $10\,\text{ms}$ (800k samples) and the measured spectral density $\widetilde{f}_{sim}$ shows a good agreement with the predicted noise shape as shown in Figure 6.1.

6.2.2. Simulating additive noise

Additive noise was included by adding $\widetilde{n}$ to the output of the LUT (see Figure 6.2). The effective phase noise measurement limit $\widetilde{\varphi}_{\widetilde{n}}$ was determined by the additive noise gain $K_{\widetilde{n}}$. (The factor of -11 in the following equation is determined by bit shifting and the output amplitude of the LUT).

$$\widetilde{\varphi}_{\widetilde{n}} = \underbrace{2^{K_{\widetilde{n}}-11}}_{\kappa_{\widetilde{n}}} \frac{\sqrt{2} \cdot 2^{-2}\,\text{rad}}{\sqrt{6 * 80\,\text{MHz}}} \tag{6.4}$$

For $K_{\widetilde{n}} = 15$ this results in a noise floor of $\widetilde{\varphi}_{\widetilde{n}} \approx 260\,\mu\text{rad}/\sqrt{\text{Hz}}$, this corresponds to $C/N_0 \approx 72\,\text{dBHz}$ (see Equation 2.24) and a potential length measurement precision of $\widetilde{l}_{\widetilde{n}} \approx 41\,\text{pm}/\sqrt{\text{Hz}}$.
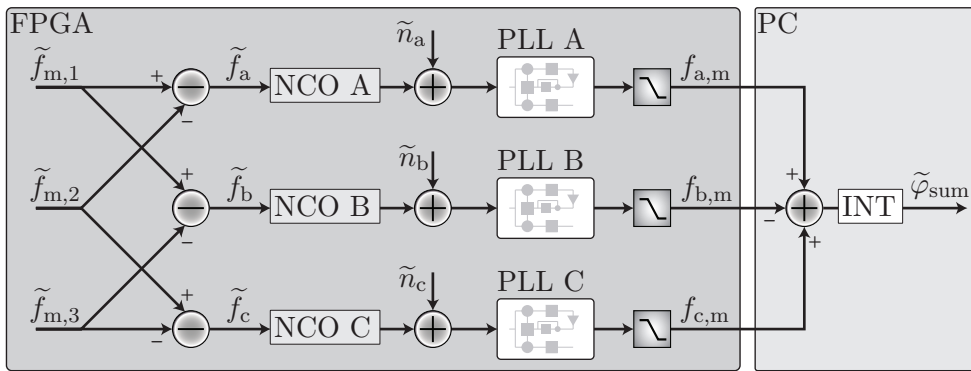
Figure 6.3.: Shown is the test set-up used for the digital three signal measurement. The three initial frequencies $\widetilde{f}_{m,x}$ are combined so that the measured phases/frequencies allow one to determine the linearity and noise floor of the ADPLLs.

6.3. Digital three signal test

Testing the ADPLLs is a perfect application of the three signal test configuration described in the previous chapter. This was first done by Shaddock et al. [36] to test the JPL version of the LISA phasemeter and an extended version of this approach was used here. Implementing this type of test using DSP allows for an optimal generation of the three signals by combining the frequency inputs of three NCOs instead of using a mixing process as required in the analogue or optical domain as shown in Figure 6.3.

6.3.1. Demonstration of LISA performance and stability

The ADPLL design for the LISA PMS was tested using the digital three signal test. The laser frequency noise described in Section 6.2.1 was used for the signal generation and an additive noise that corresponds to $C/N_0 \approx 72\,\mathrm{dBHz}$ was also added. First it was observed that using bandwidth values within a factor of two of the optimal value no cycle slips could be observed during several hours of simulation time. This demonstrates that the ADPLL design has a large robustness.

Figure 6.4 shows the two most significant performance tests. One measurement ($\varphi_{\mathrm{sum},\widetilde{n}}$) included additive noise and the result showed the expected white noise floor for three uncorrelated sources ($\sqrt{3} \cdot \widetilde{\varphi}_{\widetilde{n}}$). A second measurement without additive noise (φ_{sum}) revealed the underlying performance and linearity of the ADPLLs to be sufficiently below the requirement. The roll-up at low frequencies was attributed to the non-linearity caused by the second harmonic. This noise source can potentially spoil an absolute phase measurement (due to
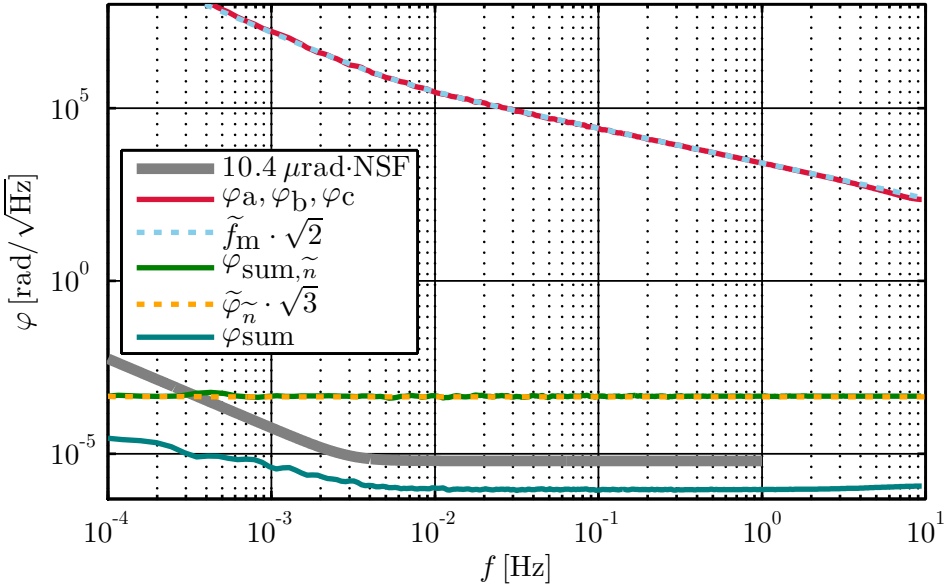
Figure 6.4.: Shown are two measurements of the digital three signal test. The beat note frequencies in both cases are $f_a = 7\,\text{MHz}$, $f_b = 25\,\text{MHz}$, $f_c = 18\,\text{MHz}$.

the coupling this depends on the stability of the laser frequency), but this type of readout is only required for DWS, where the non-linearity would be common between quadrants and therefore completely subtracted. Based on these simulation results it was concluded that the ADPLL is sufficient or at least close to a design required for a final implementation.

6.3.2. Non-linear influence of the second harmonic

Earlier simulations were done without an additional IIR filter inside the ADPLL, with larger ADPLL bandwidth and with even higher dynamic ranges of the input signals. These simulations, shown in Figure 6.5, have clear indications of the non-linearity induced by the coupling of the second harmonic described in Section 4.4.3. A quantitative analysis of this effect has however not yet been done, since it was not limiting the relevant measurements and since it requires to combine the measured input signal with a detailed transfer function model of the specific ADPLL implementation.
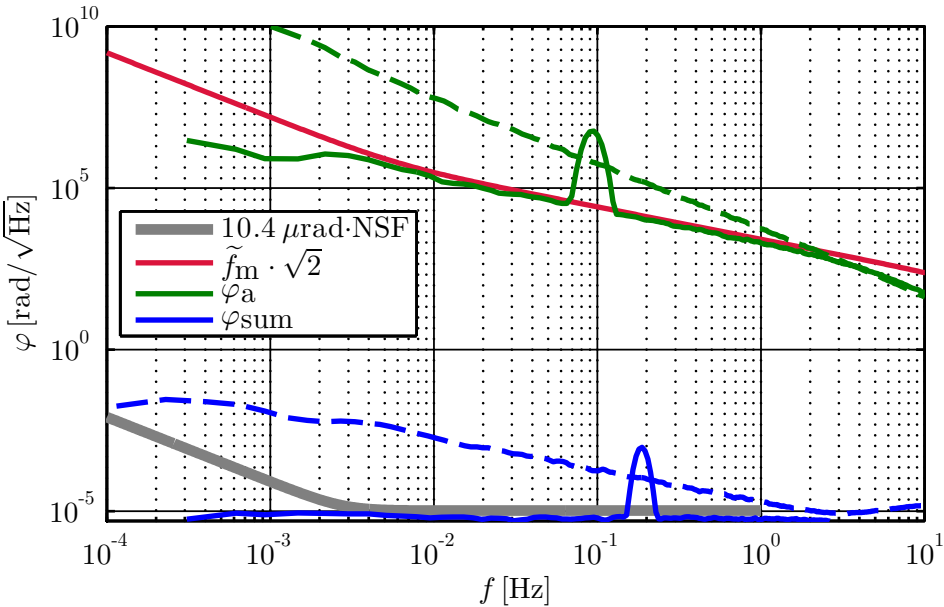
Figure 6.5.: Shown here are two examples of three signal simulations showing an excess phase noise. The ADPLLs used an intermediate design without additional low pass filters. The straight lines show an example with an additional tone at $\approx 0.1\,\mathrm{Hz}$ and an overall noise shape close to the one used in Figure 6.4. Two effects are observable, firstly, the white noise floor of the signal combination is slightly increased and, secondly, a peak appears in the combination at twice the modulation frequency. A phase noise closer to a free running laser (dashed lines) with a $1/f^2$ shape was added to the signal and the resulting combination is dominated by a $1/f$ noise, as predicted by Equation 4.42.

6.4. Testing the linear ADPLL model

A test bed in a single FPGA also allows absolute signal tests to be performed (see Chapter 5). This can be done by directly comparing the phase of the signal generating NCO φ_{NCO} with the phase in the tracking ADPLL φ_{PLL}, by comparing their respective phase accumulator values. Though this method has not been used during this thesis for performance measurements it is the most promising candidate for the further investigation of the non-linearities presented before. In addition it uses only a single NCO and a single ADPLL, which reduced the required resources, computation time and data storage space.

This method was used however to measure the noise and bandwidth dependencies of the error point distribution described in Chapter 4. This was done by computing the difference of φ_{NCO} and φ_{PLL} for one channel in the three
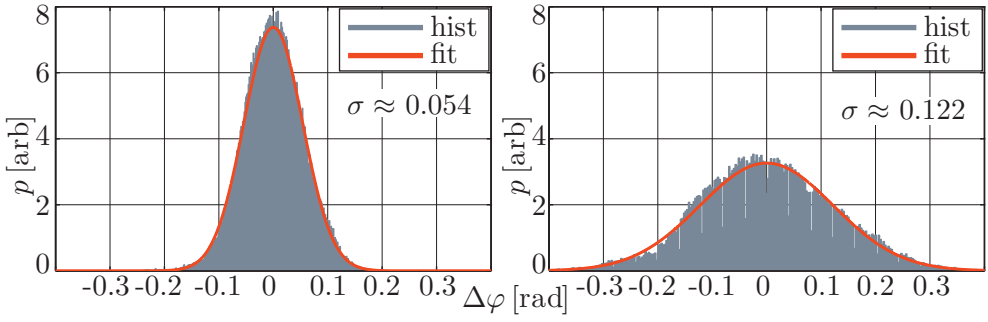
Figure 6.6.: Shown here are two histograms of measured phase differences together with their respective fit of a Gaussian distribution. The left plot shows the distribution for a bandwidth of $\approx 160\,\text{kHz}$ with only frequency noise present in the input signal. The right side shows the distribution for a bandwidth of $\approx 40\,\text{kHz}$ under the influence of frequency and additive noise. This corresponds to the optimal operation point for the assumed noise levels shown in Figure 4.7.

signal test set-up. The calculated differences of 131k samples were used to fit the standard deviation of a Gaussian distribution. The histograms and fits of two of these measurements are shown in Figure 6.6.

Chapter 7

Phasemeter for GRACE Follow-On: Laboratory experiments

During the preparation of this thesis the Albert-Einstein Institute took part in the planning, development, and testing of the Laser Ranging Instrument (LRI) [39], a technology demonstration add-on for GRACE Follow-On (GFO). Experiments to test and verify parts of the interferometry concept (described in Section 2.5) were conducted at the Albert Einstein Institute (AEI). As part of this thesis a phasemeter system was designed and tested that allowed us to perform these optical experiments with MHz heterodyne interferometry, as will be done during the actual mission.

The first section of this chapter describes the phasemeter system implementation for these experiments and its performance. The demonstration of differential wavefront sensing (DWS, see Section 2.4.5) using this phasemeter is shown in the second section. It is explained together with the optical set-up that was used to test a prototype for the LRI optical bench. The third section of the chapter describes the implementation of the steering mirror feedback control loop in this set-up using the phasemeter. The required additions to the digital signal processing (DSP) and the hardware are shown together with a discussion on the measured closed loop behaviour.

7.1. Phasemeter implementation

The LRI experiments required a minimum of five readout channels. Since no such phasemeter prototypes were available at the AEI a new one was constructed using an FPGA evaluation board (see Appendix A) together with a commercially available extension card FMC107 from 4DSP [104]. The card hosts four
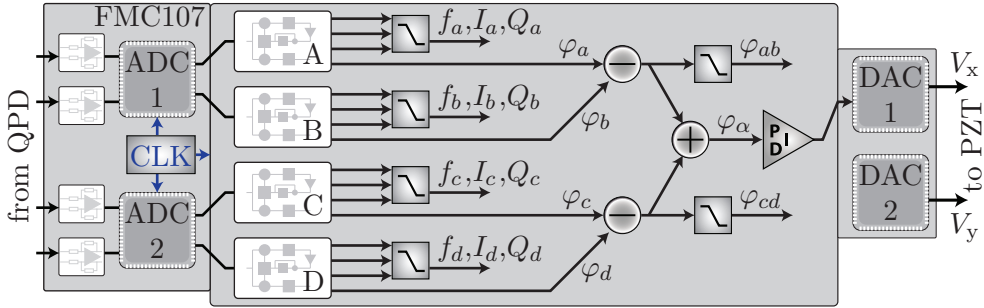
Figure 7.1.: Shown is the phasemeter architecture for the readout of four channels. After the input signals have been conditioned and digitised on the FMC107 they are fed into four ADPLLs. Each of them determines the signal frequency f_x, amplitude I_x and out-of-phase component Q_x (see Section 4.3 for a description of these signals). These signal were all directly decimated with a second order CIC filter to a rate of $\approx 2\,\mathrm{Hz}$ ($\approx 38\,\mathrm{Hz}$ in some experiments) and stored in a PC. If the four channels were used to read out the segments of a quadrant photo diode the phase values of two ADPLLs were each directly subtracted to form the phase signals required for DWS (see Section 2.4.5). The computed differences were once decimated and they were also directly added to generate the error signal for the steering mirror control loop. Shown are only the phase combinations for the horizontal tilt, the vertical tilt computes respectively.

two-channel ADCs and therefore allows one to read out a maximum of eight channels. The sampling frequency was adjusted using the available clock control on the FMC107 to 40 MHz and each channel was read out with a single dedicated ADPLL. Figure 7.1 shows a sketch of the phasemeter scheme for four of the eight implemented channels.

The use of a pilot tone was not necessary since the known ADC jitter influences were all far below the required performance levels, which are on the order of $\mathrm{nm}/\sqrt{\mathrm{Hz}}$, which corresponds to phase noise levels of $\mathrm{rad}/\sqrt{\mathrm{Hz}}$. The effects of crosstalk were considered not to be critical either at these performance levels, even though the FMC107 would potentially be sensitive to them due to its dense placement of RF components.

The performance of the phasemeter was tested for a single heterodyne frequency (6.25 MHz, which was convenient and within the range to be used in the actual satellite missions) The simulation of realistic signals was not part of the investigations. The achieved performance levels in a split measurement using analogue signals from a commercial signal generator is shown in Figure 7.2. The measured levels are far below the desired performance for the LRI experiments and even below the $\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ level required for LISA for large parts
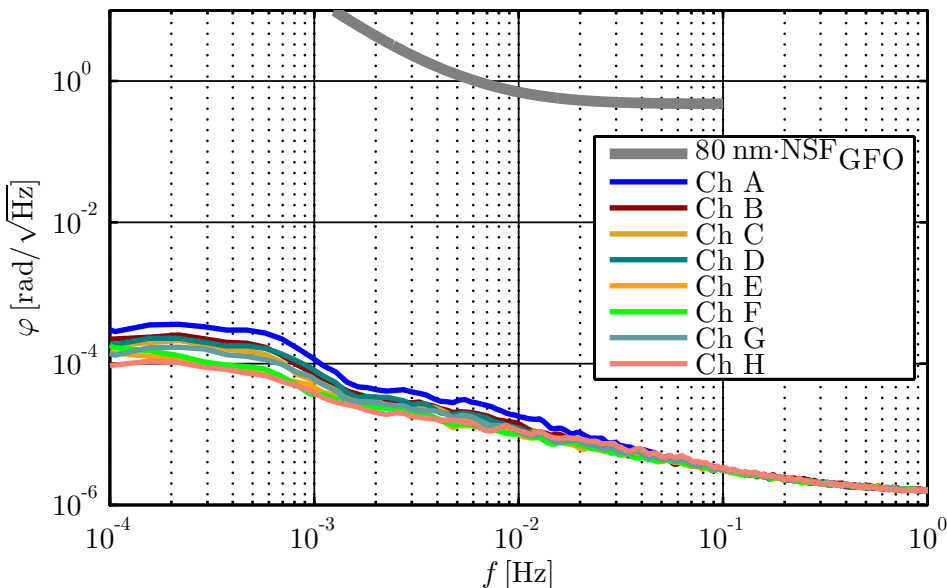
Figure 7.2.: Shown is the performance of a split measurement of the phaseme-
ter for GFO laboratory experiments. The signal was at 6.25 MHz and it
was split using an 8-way resistive power splitter. The performance is purely
limited by ADC sampling jitter, as later proven in Chapter 8.

of the spectrum. This was a promising result that showed on the one hand
that the phasemeter is sufficient for the optical test and on the other hand that
the FMC107 is potentially able to achieve LISA performance levels, as will be
shown in Chapter 8.

7.2. Differential wavefront sensing with MHz signals

To demonstrate and test the feasibility of DWS and the steering mirror con-
trol loop for GFO the phasemeter was integrated into an optical set-up with a
prototype of the LRI optical bench (OB). The set-up, which was designed and
built by D. Schütze [105] and G. Stede [106], is shown in Figure 7.3.

The prototype of the LRI OB was placed on a carbon fibre breadboard, which
was mounted on a hexapod system (M-824 6-axis hexapod from Physik Instru-
mente [107]). This hexapod was used to move the OB in all six degrees of
freedom, allowing us to simulate the pointing instabilities of the satellites ex-
pected in GFO. The optical layout was tuned to simulate the beam shapes that
are expected in the actual set-up on the OB. Therefore the beam representing
the light from the far spacecraft was brought onto the optical bench via a beam
expander and a fixed aperture. This generated a flat top beam that was inter-

Figure 7.3.: The top sketch shows a simplified layout of the set-up used to test differential wavefront sensing for GFO. Two lasers are used, one represents the local oscillator (LO) and one the received beam (RX). They are interfered once on the optical bench (OB) and once in the reference interferometer (REF). The two are locked to each other to an offset frequency of 6.25 MHz with an analogue offset phase lock (OPL). The received flat top beam is generated by a beam expander (BEX), increasing the beam radius to 40 mm, and an aperture (AP) with a slit diameter of 8 mm. It is interfered with the LO delivered with a commercial fibre injector, producing a collimated beam of approximately 5 mm diameter. These beam parameters were chosen to resemble the current baseline of the LRI [39]. The optical bench contains a piezo-type steering mirror (PZT), an imaging system (IMS) and a quadrant photodiode (QPD). The photographs on the bottom show parts of the actual set-up. Here the beams were delivered onto and from the OB using periscopes.

fered with the local Gaussian beam fed in via the steering mirror from an optical fibre. The interfered beams were brought onto a quadrant photoreceiver via a dedicated imaging system. This imaging system was used to image the aperture and the steering mirror planes onto the QPD, which minimises the beam

Figure 7.4.: Shown on the top are the commanded tip and tilt angles (v, w) of the hexapod relative to the centre of the aperture. The bottom shows the measured DWS phase signals, which are proportional to the physical angles between the two interfered beams α and β. For perfect alignment of the hexapod axis and the beam axis these angles are equivalent to v and w. The measured signals shows phase wrapping for signals larger than 2π. Note that in the real implementation the DWS signals could not be detected for misalignments larger than $\approx 200\,\mu$rad, because the interference contrast and the signal power would be too low [39]. The artefacts in the DWS signals are caused by unwanted rotations and motions of the hexapod during changes of v [106].

walk under tilts around these points. A reference interferometer (using simple gaussian beams) was implemented to complete the racetrack-like configuration.

After an initial alignment of the interferometers the ability to sense rotations of the optical bench or the steering mirror with the DWS signal was tested. Figure 7.4 shows such a measurement where the hexapod was commanded to perform tip and tilt movements around the centre of the aperture. These rotations were clearly visible in the DWS signal generated by the phasemeter. The DWS response is very linear in the observed range ($\pm 0.5\,$mrad), and the coupling factor $g_{\mathrm{DWS}}/\lambda_0$ (see Equation 2.26) was determined to be $\approx 9.4 \cdot 10^3\,$rad/rad for horizontal rotations (v) of the optical bench and $\approx 9.9 \cdot 10^3\,$rad/rad for vertical rotations (w). In a second step the coupling was again measured, but this time under rotations of the steering mirror. The determined coupling factors from

the mirror angles into the DWS are increased by a factor 2 for horizontal and by $\sqrt{2}$ for the vertical rotations, as expected from a beam deflected off a mirror at 45° incidence. All measured coupling factors were roughly a factor of two below the calculated values. The current assumption is that this is caused by wrong beam parameters or the quality of the wave front, which contained significant diffraction patterns from the aperture.

7.3. Active Steering mirror control

After testing the angular readout an experiment was conducted to prove the feasibility of the active steering mirror control using the DWS signals. Therefore additional servo controllers were programmed into the FPGA. They used the DWS angles as input to generate actuation signals, as shown for the horizontal angle in Figure 7.1. These actuation signals were then converted to analogue signals via digital-to-analogue converters (DACs). At the time no DACs were available that could easily be connected to the FPGA board. Therefore, two simple 1-bit DACs were implemented. Each of them used as input two digital FPGA pin signals. These pins were differentially switched with the actuator signals after a dithered truncation to 1-bit (this was done at the full sampling rate of 40 MHz). This implementation was very similar to a standard pulse-width modulation with a theoretical noise floor of less than $1\,\mathrm{mV}/\sqrt{\mathrm{Hz}}$ (This is based on the DAC output range of 10 V and the quantisation noise formula 4.3). The differential pins were amplified, buffered, and low-pass filtered with an operational amplifier (circuit shown in Figure B.1). The so generated voltages were then fed into a piezo driver, which rotated the steering mirror accordingly. The gains of the servo controllers were chosen based on a loop model taking into account the transfer functions of the piezo mirror and its driver, the DACs, the DSP, and the measured DWS couplings. A unity gain frequency of 30 Hz was calculated with a phase margin of approximately 70°. Adjustments of the gains and the actuation sign were possible during the experiments by sending simple commands to the phasemeter from a PC.

Even though no transfer function measurements were done, which could have been used to fully verify the applied model, the control loop could be operated successfully. This was determined by reaching a locking state where: The DWS signals were minimal, there was sufficient contrast present on all QPD segments, and the interferometric contrast of the reference interferometer was kept constant. The latter showed that the beam sent out from the OB was kept stable relative to its interfered beam. A performance of the DWS signals during a 12 h stability measurement with active steering mirror control is shown in Figure 7.5. The high noise induced by the 1-bit DACs was assumed to be the limiting factor for the in-loop phase performance. Its white noise floor is shaped by the closed loop error function, suppressing it at lower frequencies. Future imple-
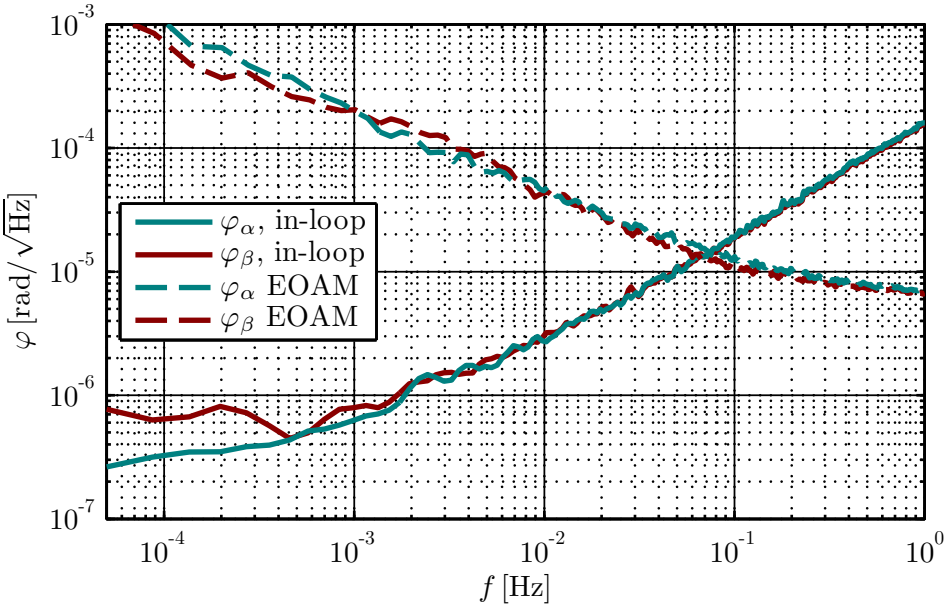
Figure 7.5.: Shown is the performance of the in-loop DWS phase signals with an active steering mirror feedback. The noise increase to higher frequencies is induced by the 1-bit DAC white noise, which is shaped by the closed loop control. A simple transfer function analysis also revealed that the noise shaping is independent of the DWS gains. Since both DAC channels have the same noise floor this explains the almost perfect overlap of the two in-loop measurements for most of the frequency band. A separate measurement was conducted to determine the differential phase performance of the phasemeter combined with the QPD. An electro-optic amplitude modulator (EOAM) was used to modulate the LO with a 6.25 MHz signal. By blocking the RX beam it was therefore possible to measure the DWS signals with an amplitude modulated beam that was shining on the QPD. The large wavelength of this modulation causes the phases of the MHz signals on each segment to be almost perfectly equal. Therefore this was a very useful reference measurement, which was also used to determined DC phase errors between channels. The DWS performance achieved during these measurements is shown for both angles as dashed lines. This noise level is representative of the sensor noise of the DWS loop (excluding other optical effects caused by clipping or higher order mode content). The combined in-loop and sensor noise measurements give a rough estimate of the actual DWS phase performance, which is better than $1\,\mathrm{mrad}/\sqrt{\mathrm{Hz}}$ between $0.1\,\mathrm{mHz}$ and $1\,\mathrm{Hz}$.

mentations using a dedicated DAC card (see Appendix A) will further improve these results. Assuming the earlier determined coupling factors the physical

angular stability achieved for α and β can be calculated from the actual DWS phase performance. This performance is estimated from the in-loop measurements and a separate characterisation of the sensor noise (see Figure 7.5) to be better than $1\,\mathrm{mrad}/\sqrt{\mathrm{Hz}}$ below $1\,\mathrm{Hz}$. This leads to a performance of the angular stability of better than $\approx 0.1\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ in the same frequency range.

Chapter 8

Experimental pre-investigations for a LISA phasemeter

Individual aspects of the phase measurement chain were investigated experimentally during this thesis. Some of these experiments were directly related to the LISA phasemeter prototype, others investigated alternative architectures of analogue components. The results presented here also include the initial performance levels achieved in analogue and optical three signal tests with phasemeter prototypes. Even though these particular investigations were not done in the presented order, they form the logical bridge between the rather uncritical phasemeter implementation described in the previous chapter and the full LISA phasemeter system discussed in the subsequent one.

As described in Section 3.3.1, a pilot tone correction is necessary to achieve μrad performance levels. The principle of such a correction is demonstrated in the first section together with the investigation of an ADC candidate for the LISA phasemeter. The use of a pilot tone is further discussed together with a proposed design for a dedicated signal generator that could be used to perform absolute signal tests at LISA performance levels. The results achieved with a prototype of such a generator using commercial hardware are also shown. The third section of the chapter describes a short investigation of an analogue gain controller. Afterwards some alternative analogue front-end configurations for the pilot tone correction are presented together with the achieved performance levels. Based on these results the implementation and testing of a 16 channel phasemeter with LISA performance is presented. This phasemeter was originally planned for testing a LISA optical bench prototype and during its testing a new noise source was discovered. The last two sections of the chapter describe some initial non-linearity test done with phasemeter prototypes. The first test generated analogue beat notes by mixing of GHz signals. The second test used a

specially designed hexagonal interferometer [108] to generate three optical beats and it is the planned set-up to perform tests of the full LISA phase measurement chain in the future.

8.1. ADC Evaluation

During the initial planning for the LISA phasemeter an ADC had to be chosen based on various parameters like interface type, bandwidth, number of channels, power consumption and most importantly phase measurement and sampling jitter performance.

While performance tests were already done successfully with single channel ADCs [43], the multichannel variants showed promising features like a lower power consumption, a simpler PCB layout, less synchronisation effort and potentially a lower sampling jitter between the channels on a single chip. Four channel ADCs were especially interesting, since they perfectly match the readout of QPDs used in most interferometers (IFOs), including the critical science IFO in LISA. The potential disadvantage of an increased crosstalk between channels was not found to be critical. The multichannel ADC described in the next paragraph is specified with a crosstalk suppression of 100 dB at 10 MHz, which is much higher than the suppression between the segments of a QPD. Joshi et al. [78] report a suppression value of 45 dB for a QPD that was specifically manufactured to achieve low crosstalk.

A study conducted by the Danish industry parter Axcon Aps. showed that the ADS6445 [109], a four channel ADC, is a good candidate. It is available in a space qualified version, it has a large analogue input bandwidth (500 MHz), which potentially makes it insensitive to thermal fluctuations, and it has a low power consumption, roughly half of what four single channel ADCs that have been investigated earlier [43] would require. Since no data on the low frequency phase noise performance or sampling jitter was available, an experimental evaluation was conducted. The goal of this was to measure the differential phase noise between two ADC channels and to prove that a pilot tone correction would bring it below the requirement ($6 \, \mu\mathrm{rad}/\sqrt{\mathrm{Hz}} \cdot \mathrm{NSF}$, see Section 2.6) with sufficient margin.

A very similar hardware architecture as described in Section 7.1 was therefore implemented. Instead of an FMC107 an evaluation board of the ADS6445 [110] was used together with an ML605 FPGA board (see Appendix A). Two of the ADC channels were connected to four ADPLLs, two for each channel. One PLL was used to track one tone, for example the beat note, while the other one tracked an additional tone, the pilot (for this test both tones only differ in frequency, making this distinction arbitrary). Since only the performance of the ADC was to be measured, the two tones were first combined by a standard power combiner, then split via a resistive power splitter and finally fed into two
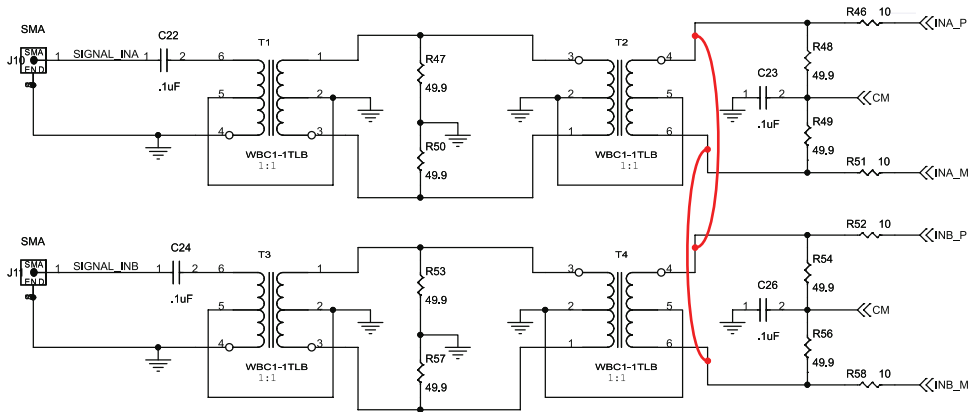
Figure 8.1.: Schematics of the ADS6445 evaluation board analogue front-end
[110]. The two transformer stages build a standard ADC driving circuit that
provides impedance matching and the transformation from single ended to
differential signals for the ADC inputs. For LISA performance levels these
circuits are not optimal since the transformers have an inherent strong phase
sensitivity to temperature fluctuations. For the ADS6445 evaluation the dif-
ferential signals driving the ADC were therefore cut short (indicated by red
lines).

ADC input circuits. This allowed us to completely neglect any uncommon phase
noise induced by separate pilot tone adders, which is one of the most critical
components in the phasemeters described later.

By tracking two tones at different frequencies the influence of sampling jitter
was directly observable by performing a correction of one tone using the other
(see Equation 3.2). For a purely sampling jitter dominated measurement the
corrected phase noise would only be limited by the digitisation noise of the ADC,
truncations inside the ADPLL and phase noise in the analogue components of
the front-end. The first measurements showed almost no improvement due to
jitter correction and they were dominated by an excess noise rising to lower fre-
quencies. From earlier experience this was attributed to the temperature noise
induced phase variations in the transformer based analogue front-end driving
the ADC (see Figure 8.1).

Instead of implementing a sufficiently stable AFE an alternative method was
used to directly measure the ADC noise influence. A direct short cut between
the two differential signals driving the ADC was implemented using small wires
soldered onto the evaluation board, as sketched in Figure 8.1. This shorted out
the differential phase noise of the transformer stages and it kept the impedance
matching of the circuit intact. By this modification differential phase noise
measurements directly probed the ADC noise influence.
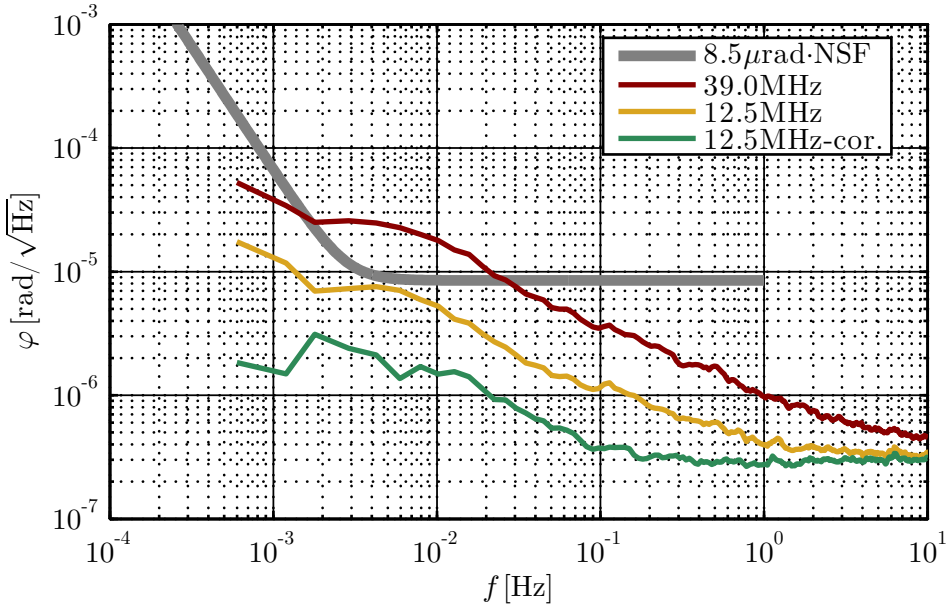
Figure 8.2.: Performance measurement of the ADS6445 evaluation using a short-circuit at two ADC inputs. The differential phase noise measured between the two channels is dominated by ADC jitter noise, causing the noise to scale linearly with frequency. The uncorrected noise of the 12.5 MHz signals is already below the requirement.

For the final measurement a 12.5 MHz and a 35 MHz tone were fed into the ADC. The achieved differential phase noise spectral density is shown in Figure 8.2, with and without correction of the 12.5 MHz tone. The observed sampling jitter is very low, which is understood to be caused by the fact that both ADC channels are implemented on a single chip, causing several noise influences to be common mode. This level of jitter was judged to be very promising, since it might even allow one to use uncorrected phase data for DWS, which has a relaxed requirement in contrast to the length measurement. One should mention here, that this measurement was done in an office, a very unstable thermal environment, which makes the achieved performance even more impressive.

Based on these measurements, and the aspects described at the beginning of the chapter, the ADS6445 was chosen to be used in the elegant breadboard (EBB) model of LISA phasemeter (see Chapter 9).

8.2. Signal Simulator

The technology development activity to build an EBB model also included the design of an appropriate test environment for the LISA phasemeter. Therefore,

Figure 8.3.: The principle scheme to perform an active jitter correction is shown here. In addition to the main tones (A,M and B,M) each DAC channel also includes a pilot tone generated by separate NCOs. Their frequencies are outside of the measurement bandwidth and they are different for each channel. In principle this allows for the timing jitter of each DAC channel to be tracked individually. The possible addition of a pilot tone from the PM under test can be used to correct for the ADC sampling jitter, either by an additional phase correction term for each NCO or by an active feedback to a phase shifter in the ADC clock supply, similar to the implementation shown in Chapter 12.

a dedicated signal simulator had to be constructed to generate beat notes as expected in LISA, including realistic levels of additive and phase noise, as well as the clock tone transfer side bands and pseudo-random noise modulations.

In parallel to the development of a dedicated digital signal simulator (DSS) that can be used to perform split measurements (see Section 9.5.1) an alternative architecture was investigated to potentially perform absolute signal measurements as described in Chapter 5.

A simple implementation of such a *noise free* signal generator is hindered by the inherent sampling jitter of the DACs, mirroring the effect in the ADCs used for the phase measurement. This fact led to the idea of using a self correcting architecture for the signal simulator that would have a sufficiently reduced phase jitter between channels and even relative to a pilot tone provided by the phasemeter under test.

8.2.1. Self correction of sampling jitter

The conceptual design of a built-in sampling jitter correction is shown in Figure 8.3. Instead of only generating simulated beat notes (or main tones) each

output channel also includes an additional pilot tone at a higher frequency, which is different for each channel to make it possible to track each corresponding sampling jitter separately. The outputs are split into the actual output signals v_a, v_b and signals used for DAC jitter correction. These signals are added together to generate a comb of reference tones. These tones are then digitised again by a single ADC adding its sampling jitter to each phase noise. By tracking these back looped signals with PLLs one can now deduce, after scaling with the appropriate frequencies, the sampling jitter of each DAC relative to the others. This information can then also be used to implement an active feedback to the NCOs generating the main tones, actively correcting their low frequency phase noise induced by sampling jitter.

The final step not shown in Figure 8.3 for using such an implementation for an absolute signal test is, to also compensate for the timing jitter added by the ADC, which is imprinted on the DAC outputs via the feedback loop. This requires the independent measurement of the ADC jitter relative to the absolute phase of the phasemeter, which is defined by its pilot tone. Adding this tone and tracking it by an additional PLL potentially allows for this jitter to be corrected as well, which would make an absolute signal test possible.

The implementation of such a correction scheme makes sense when all phase measurements are dominated only by sampling jitter. Also the pilot tones have to be filtered out before the main tones reach a phasemeter, otherwise they could spoil the measurement performance on that side of the chain. Splitting the pilot tones from the simulated signals is done by using higher frequencies, which requires fast sampling DACs.

The calculation of the phase correction values might also require to perform floating point computations. But since sampling jitter is only dominant at low frequencies, these computations can easily be done in a PC instead of the FPGA.

The biggest difficulty of such an implementation is the phase noise of the analogue components, which masks the sampling jitter influences and can change the phase relation between the main tones and the pilot tones, as described in the next section. The input range of the ADC also has to be large enough, since it needs to detect four or five signals with sufficient performance.

8.2.2. Hardware implementation and performance limits

Since no experience with the DAC sampling jitter correction was available, it was investigated experimentally using commercial components. An ADC extension card (FMC122 from 4DSP [111]), hosting a fast sampling converter (2 channels up to 1.25 GHz), was combined with a DAC card (FMC204 from 4DSP [112]), hosting two DACs (2 channels up to 1 GHz, 4 channels total), on the ML605 platform (see Appendix A). The output of the four DAC channels were split by resistive splitters, added and fed back into one of the ADC channels by means of a specifically designed circuit (see Figure B.3).
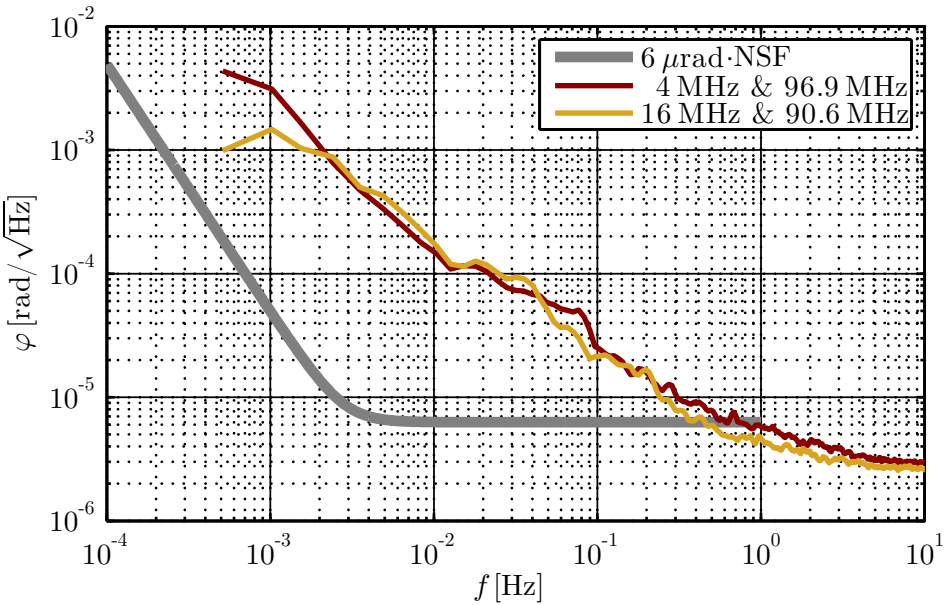
Figure 8.4.: Spectral density of the noise differences corrected for timing jitter for two signals generated and sampled by a single ADC and DAC. Shown is also the same measurement from a second DAC, using different set of frequencies.

A 400 MHz sampling clock was provided by an FPGA internal PLL. It was fed into the ADC card clock distribution and from there distributed to the ADC itself, which includes a fixed clock division by a factor of two. It was also divided by two and sent to the DAC card. Thereby all sampling and digital signal processing was running at 200 MHz, directly generating signals with a maximum frequency of 100 MHz.

Before phase performance tests were conducted parts of the LISA signal generation were implemented, including additional side bands and PRN modulations. This architecture was later ported and extended by Ioury Bykov to the DSS and is described in Figure 9.16.

After the DSP system and the interfaces to the converters were working, the phase noise performance of the system was investigated. Each channel produced a main and a pilot tone at a different frequency, generating a comb of eight sine waves that were re-digitised by the ADC. The phase of each of these tones was tracked by a dedicated ADPLL, as shown in Figure 8.3. No active feedback was implemented at this stage. The most crucial test for determining if the system is dominated by jitter, was to correct the main tones with their respective pilot tones, since they would have been influenced by the exact same sampling jitter. The results of such a measurement after applying the corrections

are shown in Figure 8.4 for two channels. The performance showed that none of the main tones could be corrected below $6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ of phase noise with their respective pilot tones, indicating that the tones are fully dominated by phase noise, probably due to thermal effects. The comparison of the two tones compared in Figure 8.4 also indicates that this phase noise is rather independent of the main signal frequency.

Due to this high initially measured phase noise floor the alternative signal simulator was not investigated further. The analogue components of the ADC and the DAC were not well known and replacing or optimising them would have been an extensive effort, since they are complex, ready made devices with dense designs.

A useful implementation of this self correction technique would likely require to design dedicated hardware and to perform a more comprehensive noise analysis, very similar to what is necessary for the LISA phasemeter. Even though one might consider this effort to be unnecessary the implementation of such a signal generator would be an elaborate experiment that allows one to demonstrate the robustness and understanding of the techniques used to build phase measurement systems. A self corrected signal simulator would only be used on ground. Hence, it must not comply with limitations of power consumption, size or weight that apply to a phasemeter system for space. Consequently, such a signal generator could utilise more extreme measures for thermal control, like large thermal masses or active stabilisations with large power consumptions.

8.3. Variable gain amplifier

Phasemeter systems in realistic IFOs will have to handle amplitude variations of the beat notes, which can potentially bring the ADPLL bandwidth away from the optimal operation point or increase the ADC truncation noise. The amplitude variations depend on the full optical design and are often directly influenced by the alignment between the interfering beams.

To compensate for changing input amplitudes one can, for example, adjust the gain of the PLLs controllers, keeping the bandwidth constant (the I-value measured by the PLLs allows the system to easily sense such amplitude variations). Another approach, that also keeps the amplitude of the digitised signal constant, is to introduce an analogue variable gain amplifier (VGA), that can act as an actuator in a potential amplitude stabilisation.

A simple version of such an amplifier was built and tested during this thesis. Based on a concept by G. Heinzel a non-inverting amplifier circuit was implemented using a PIN diode [113] as gain resistor and a fast current-feedback op-amp [114], as shown on the left in Figure 8.5. A DC current applied to the diode changes its resistance and thereby the amplification of the op-amp circuit. Hence, the signal amplitude can be actuated by controlling the DC current.
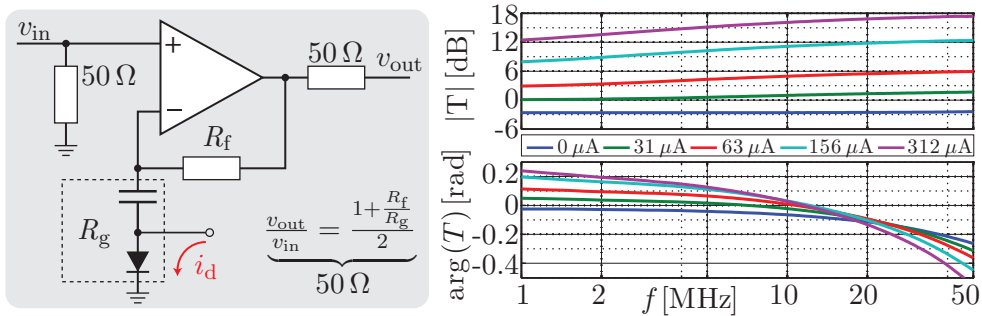
Figure 8.5.: The left side shows a simplified circuit diagram of the VGA prototype using a PIN diode. The feedback gain of the amplifier is determined by the resistance R_{g}, which is tunable via the DC current through the diode i_{d}. The full circuit is shown in Figure B.2. The right side shows the measured transfer function of the VGA for five different diode DC currents.

For different DC currents the transfer function of the circuit was measured. The right side of Figure 8.5 shows that the amplification does indeed increase with the applied current. But it does not scale equally in the observed frequency range, leading to a high pass like behaviour for the high amplifications.

One also observes that the signal phase is strongly dependent on the DC current, which might spoil the phase measurement performance. Assuming a linear coupling of $0.1\,\mathrm{rad}/100\,\mu\mathrm{A}$, based on the measurements, the applied DC current would require a relative stability of $10\,\mathrm{ppm}$, corresponding to $\approx 1\,\mathrm{nA}/\sqrt{\mathrm{Hz}}$, to achieve μrad performance. This added complexity and stability requirement are not easily met, making this type of VGA a questionable choice for LISA like performance levels that would require a much more detailed investigation. The range of signal amplitudes expected in LISA has not been studied in detail and, thus, it is not yet known if a VGA will actually be required. The phasemeter systems developed in this thesis are, therefore, designed and tested with signal amplitude variations that do not change the phasemeter gains or the ADC truncation noise significantly. Commercially available, controllable amplifiers and attenuators have not yet been investigated.

8.4. Analogue front-end investigations

The good phase measurement performance achieved with FCM107 ADC card (see Chapter 7) without the use of a pilot tone correction made it a promising candidate to achieve the μrad performance required for LISA. Since, at the time, no multichannel phasemeter with such performance was available, it could be used to perform various experiments and the investigation allowed us to empirically test alternative front-end concepts. The FMC107 card used for
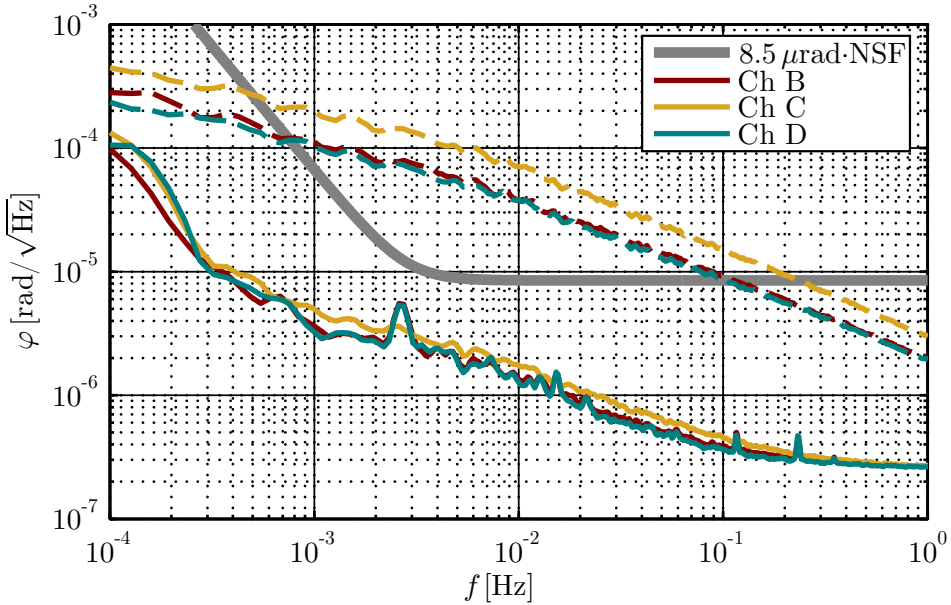
Figure 8.6.: Performance of active pilot tone adder and power-splitter distribution. The signal frequency was at 11 MHz, the pilot tone at 33 MHz. Channel A was defect. The dashed lines show the performance without jitter correction.

the LRI experiments used an AC coupled design which was confirmed by the manufacturer to be based on transformers. Since these were known to cause excess phase noise [43] a DC coupled variant of the FMC107 was used in the following experiments. It uses amplifiers to drive the ADCs and is therefore potentially much more phase stable.

8.4.1. Active pilot tone addition

Together with C. Diekmann [48] an additional analogue front-end (AFE) circuit was designed that added a pilot tone, distributed with a commercial power splitter, to four input signals by using the operational amplifier LMH6624 [115]. The schematics of this circuit is shown in Appendix B.4. By implementing additional ADPLLs into four channels of the phasemeter, described in Chapter 7, both tones could be tracked. The phase of the common pilot tone was used to correct the signal phases. The signal was split by means of a resistive power splitters, connected with short cables to the additional AFE, which was connected to the FMC107 inputs.

By packaging this front-end in bubble wrap foil the influence of external thermal noise was reduced. The FMC107 was actively cooled by an uncontrolled fan, which efficiently coupled its temperature to the rather unstable environ-
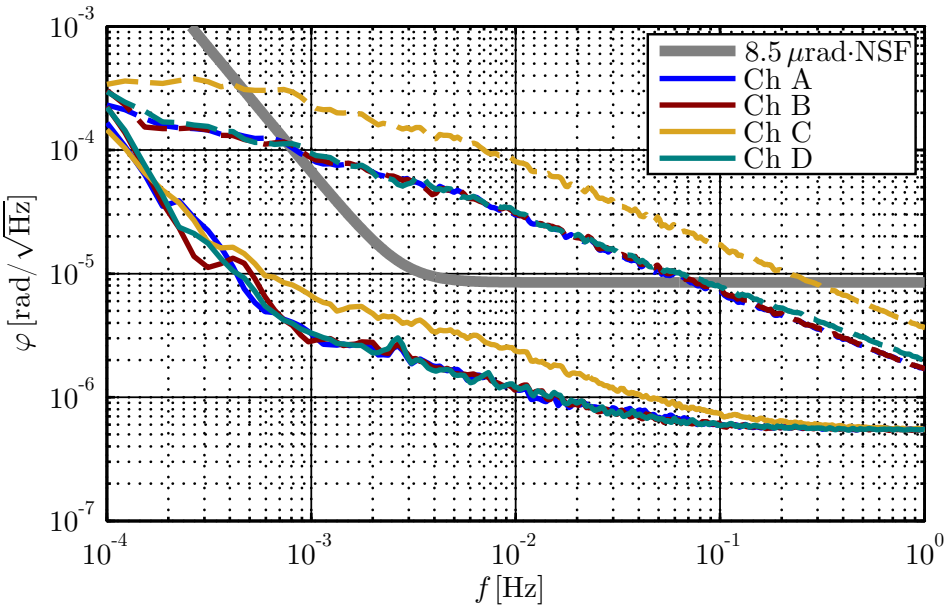
Figure 8.7.: Performance of active pilot tone adder and resistive distribution. The signal frequency was at 11 MHz, the pilot tone at 33 MHz. Excess noise in C is not understood. The dashed lines show the performance without jitter correction.

ment in the laboratory. The achieved phase measurement performance with this device is shown in Figure 8.6. The dashed lines correspond to the uncorrected signals while the straight lines use the pilot tone correction. The LISA requirements were fully met at the signal frequency without any complex temperature management or changes in the AFE. This demonstrated two things, first the addition of a pilot tone with an additional amplifier before an ADC driver is a suitable alternative, as already indicated by results from Burnett et al. [83], second the ADC driving circuit on the FMC107 is fully capable of achieving LISA performance levels even with very low thermal stability and a dense design. No temperature measurements were done for neither the AFE nor the FMC107, this is a disadvantage for their comparison to other front-end results. Still, the simple isolation of the AFE was enough to achieve the shown performance in the first measurement.

The use of a power splitter for the pilot tone at 33 MHz was also shown to be feasible. To investigate its influence another AFE board was developed that used a resistive pilot tone distribution together with the same active adder concept as before. The circuit is shown in Appendix B.5. The same measurement was performed with this AFE, again using only bubble wrap foil for thermal isolation. The results of this are shown in Figure 8.7.

This measurement showed almost the exact same performance (with the exception of some excess noise in channel A) with only a slightly increased white noise floor at higher frequencies. This might have been caused by lower signal amplitudes and is not considered to be caused by the resistive splitting. Comparing the results one finds that both methods of pilot tone distribution are feasible for a frequency of 33 MHz using only very simple means of temperature isolation.

8.4.2. Passive pilot tone addition

The performance measurements from the preceding chapter showed that it might be possible to use power splitters not only for the pilot tone distribution, but also for their addition to the input signals. This would make it possible to implement this function without using active components, which can save power and simplify thermal management. This was especially interesting for the design of a 16 channel phasemeter described in the next section.

The first test of such a scheme was done using only one type of power splitter (PS), the ADP-2-1W [116] that is available in a convenient surface-mounted device (SMD) package. The therefore built AFE (shown in Appendix B.6) uses three PS to split a pilot tone to four channels and it uses another four PS to combine these tones with the signals. This AFE was again tested by isolating it with bubble wrap and performing a split measurement, the results of which are shown in Figure 8.8.

The performance of the corrected signal was again below the requirement, though a noise increase was observed at higher frequencies in comparison to the active pilot tone addition. The low frequency performance however improved, pointing towards a lower influence of temperature fluctuations. Since this AFE was fully passive the power consumption was zero and therefore approximately 0.25 W below the one of the active circuit of the preceding section. This corresponds to a potential power reduction of 1 W for an implementation with 16 channels.

The coupling of temperature fluctuations also depends on the package of components, with some being inherently better isolated than others. Therefore another power splitter was investigated, the PSC-2-1 [117]. It has very similar parameters to the ADP-2-1, but has a solid metal enclosure instead of a plastic housing. Using this power splitter another passive AFE was constructed and its performance was again measured with a split measurement using only a passive isolation of the AFE. The circuit is shown in Appendix B.7.

Figure 8.9 shows the performance obtained with PSC-2-1 passive pilot tone adder and distribution is almost indistinguishable from the one obtained using the ADP-2-1. This indicates that both components, even though they have a completely different packaging, either cause the same low frequency noise, or that the residual noise at low frequencies is caused by the circuit on the FMC107
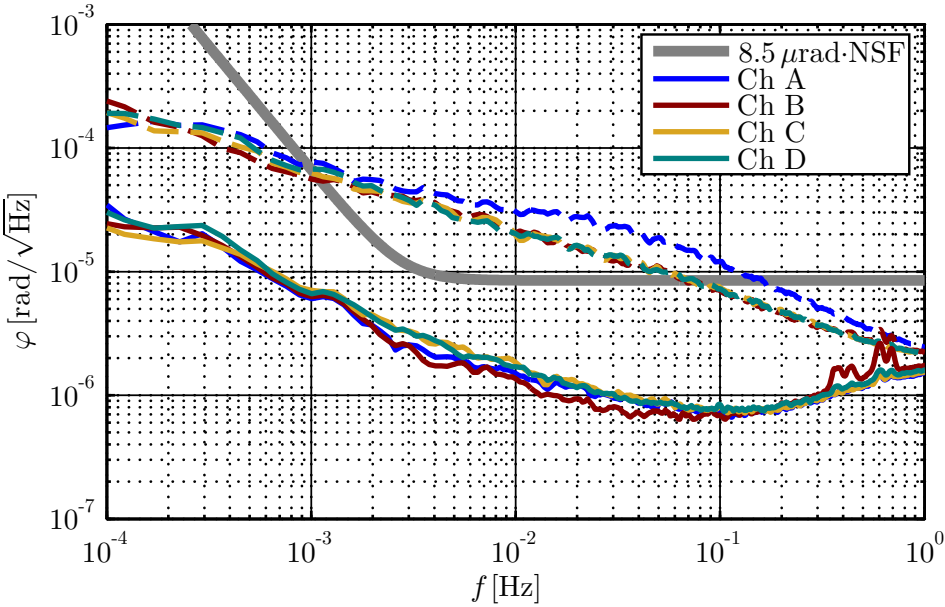
Figure 8.8.: Performance of passive pilot tone adder and distribution with ADP-2-1W. The signal frequency was at 15 MHz, the pilot tone at 35 MHz. The dashed lines show the performance without jitter correction.

itself.

The crosstalk coupling from a signal channel into the pilot tone for the passive pilot tone addition it is determined by the isolation between the signal ports of the power combiners. For the two components used here this value is roughly 30 dB in the corresponding frequency band. A signal coupling into the pilot tone must, in the here designed circuits, also pass at least one additional power splitter to reach another signal channel (see Appendix B.7). Therefore, the effective crosstalk suppression in this type of pilot tone distribution is on the order 60 dB.

8.5. Phasemeter for testing the LISA optical bench

During this thesis the space interferometry group at the AEI was involved in a technology development activity, funded by ESA, to develop and test a prototype of the LISA optical bench [57, 118, 48]. The planned test procedure included the read out of four QPDs of heterodyne interferometers with MHz signals. This required a phase measurement system with 16 input channels, able to achieve LISA performance. This phasemeter did not require to be functional over the full frequency range, which simplified its design and testing.

Since no such multi-channel phasemeter was available at the time, an alterna-
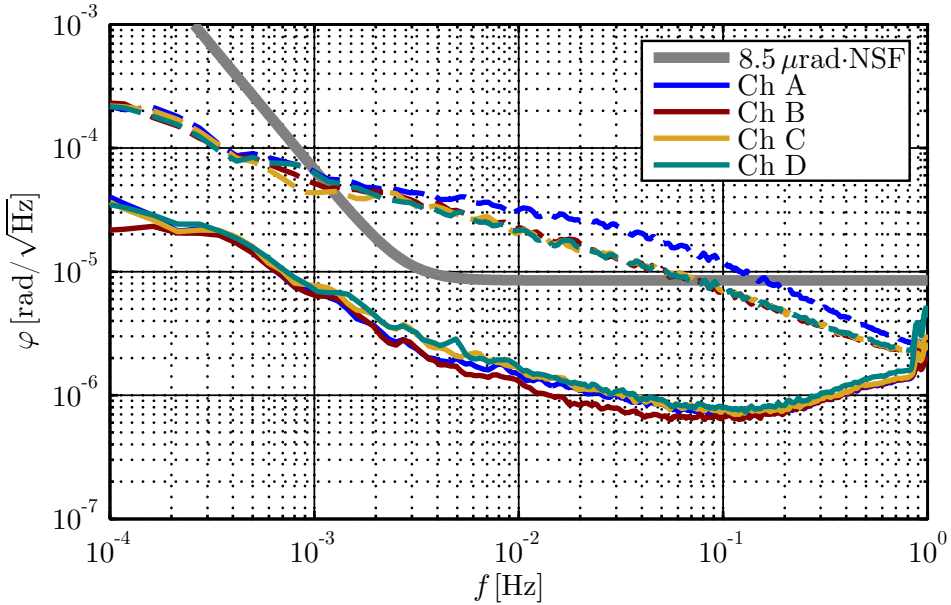
Figure 8.9.: Performance of passive pilot tone adder and distribution with PSC-2-1. The signal frequency was at 15 MHz, the pilot tone at 35 MHz. The dashed lines show the performance without jitter correction.

tive read out scheme using a down mixing process and a LISA pathfinder style phasemeter was tested by A. Schreiber et al. [119]. Even though this scheme was able to achieve the required performance levels, this type of readout would have increased the phase noise limit induced by additive noise sources by a factor of $\sqrt{2}$, due to additional folding of this noise onto the carrier by the additional down mixing. Therefore a LISA-style phasemeter system was constructed using commercial components and an AFE for the pilot tone addition, based on the investigations in the previous section.

This phasemeter used the ML605 FPGA platform (see Appendix A,) together with the ADC card FMC116 by 4DSP [120]. The card carries eight two channel ADCs, leading to a total of 16 channels. The manufacturer quotes a value of 60 dB of crosstalk suppression between the channels. The effective parasitic phase noise induced by crosstalk depends on its suppression, as well as on the phase stability between the signals, as described in Section 3.4. The OB tests would use signals that are sufficiently phase stable to each other to reduce the crosstalk influence below the required levels, making the FMC116 a suitable choice.

The AFE for this phasemeter was mounted, together with the FPGA and ADC card, into a common housing, as shown in the top pictures in Figure 8.10. It consists of three circuit boards, two of which contain twice the passive pilot
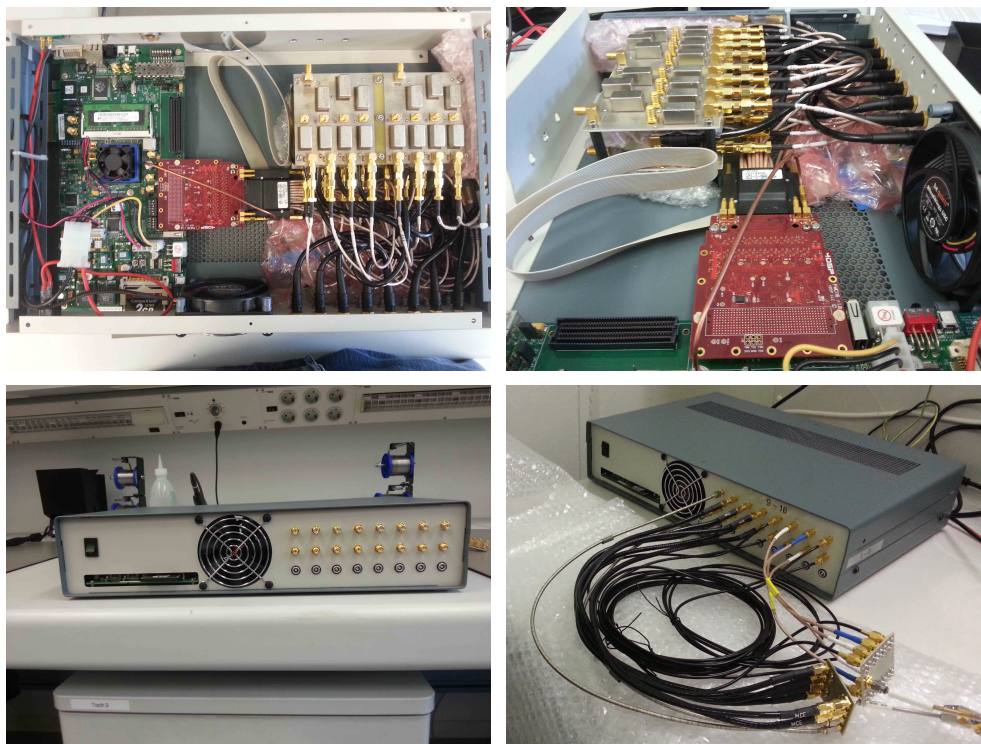
Figure 8.10.: Photographs of the 16 channel phasemeter for testing the LISA optical bench during assembly and testing.

tone adder circuit described in the previous section (schematics in Figure B.7). These two boards were mounted on top of each other and connected directly to a front panel and the FMC116 via equally long cables. The pilot tone itself was distributed to these four AFEs by three power splitters on an additional circuit board, which was mounted vertically at the back. After construction the AFE was thermally isolated by bubble wrap foil, as indicated in Figure 8.10.

The clock distribution on the FMC116 was configured such that all ADCs sampled at 80 MHz. This also allowed us to use the same DSP algorithms and ADPLLs as developed for the LISA phasemeter. For the readout of each channel one signal PLL and one pilot tone PLL were implemented, leading to a total of 32 ADPLLs. An additional direct readout of the PA differences was also implemented for each four channels to allow the easy generation of DWS data for the potentially connected four QPDs.

Influence of cables on phase performance

The performance of this phasemeter was tested by performing a split measurement with all channels. First measurements using equal cables after the power
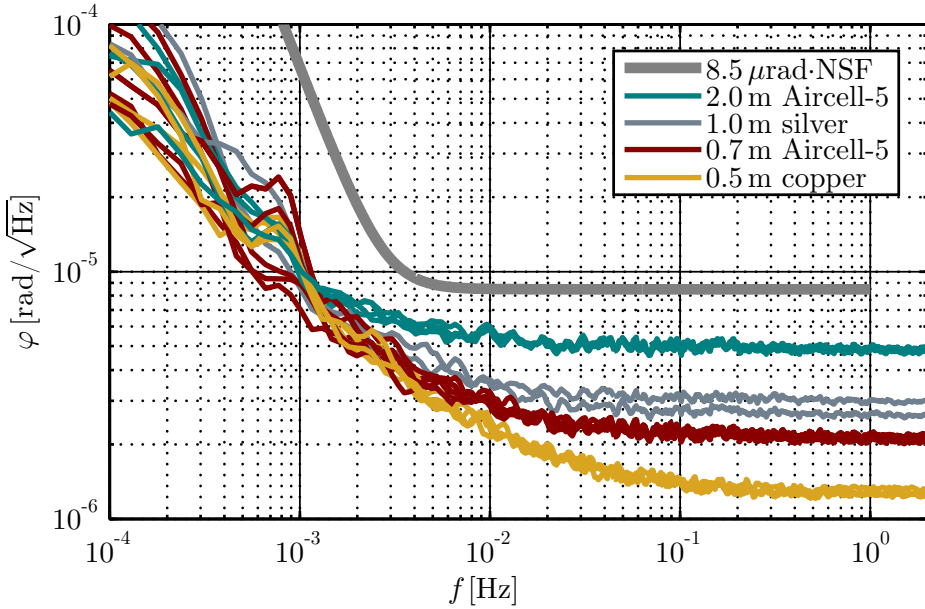
Figure 8.11.: The phase fluctuations of 16 channels relative to their average value are shown. The measurements were done using a 35.3 MHz pilot tone fed directly into the PT distribution, and a 13 MHz signal was injected into a 2-way and then two 8-way resistive splitters to inject it into all channels. The cables and the splitter were thermally isolated with bubble wrap foil. The set-up and the cables used are shown in the lower right corner of Figure 8.10. The initially measured signal phases (not shown) increase slightly faster than $1/f$ to lower frequencies and have a value of $5\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$ at 1 Hz, leading to a dynamic range of 9 orders of magnitude at 1 mHz between the measured signals and the shown phase performance.

splitters showed a similar noise performance in all channels below the requirement. But further measurements revealed a change in the noise floor if different cable types and lengths were used, as shown in Figure 8.11.

This effect indicated a strong influence of impedance matching and reflections on the measurement performance, as discussed in Section 3.2.2. The exact coupling of this noise source is not yet understood, however a very similar effect was also encountered during the testing of the LISA PMS in Chapter 9. Even though the effect does increase the noise, it is does not violate the performance requirement and it could potentially be suppressed by using the same cables for all channels.
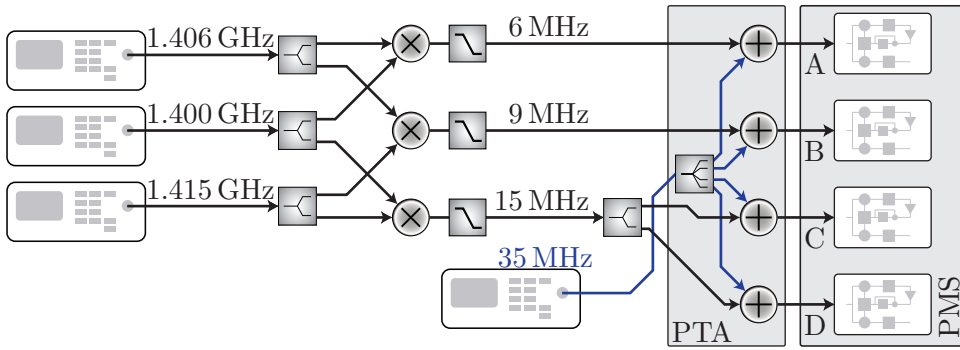
Figure 8.12.: A three signal test with analogue signals was implemented by mixing three GHz tones to generate beat notes in the MHz regime, which are then measured using a pilot tone correction as described in Section 8.4.1.

8.6. Analogue three signal test

The availability of an analogue front-end and an ADPLL that had been tested for linearity digitally, allowed performing three signal tests (see Chapter 5) using analogue signals. For this purpose an ML605 FPGA board together with an FMC107 ADC card and the AFE described in Section 8.4.1 were used as phasemeter system. The three analogue signals were generated by mixing three GHz signals from commercial signal generators as depicted in Figure 8.12.

The results of this measurement are shown in Figure 8.13. The sampling jitter corrected three signal combination (A+B−C cor.) lies above the requirement for almost the full spectrum and has a 1/f like noise shape. To ensure that this results is not limited by the phase noise of the AFE or jitter, the signal at the highest frequency was also split and used for a simultaneous split measurement. This measurement (C−D in Figure 8.13) shows the expected noise performance. The jitter corrected signal (C−D cor.) also demonstrated simultaneously the feasibility of the used pilot tone correction, indicating that the AFE was not limiting the performance. The fact that all three signals are strongly separated in frequency also indicates that crosstalk is not relevant. In the end the GHz analogue mixers were identified to cause the excess noise. Measurements of their phase stability done by Schreiber et al. [119] have shown that mixers produce a low frequency phase noise that strongly increases with the output frequency. These measurements were done at output frequencies between 160 Hz and 16 kHz. The observed effect might be the same one that limits the performance for MHz signals.

A reduction of the mixer noise was not investigated further, since an optical three signal test was already in preparation. Small amplitude modulations of the GHz signals revealed to cause a significant phase noise, making an amplitude stabilisation a potential candidate for improving the mixer noise performance.
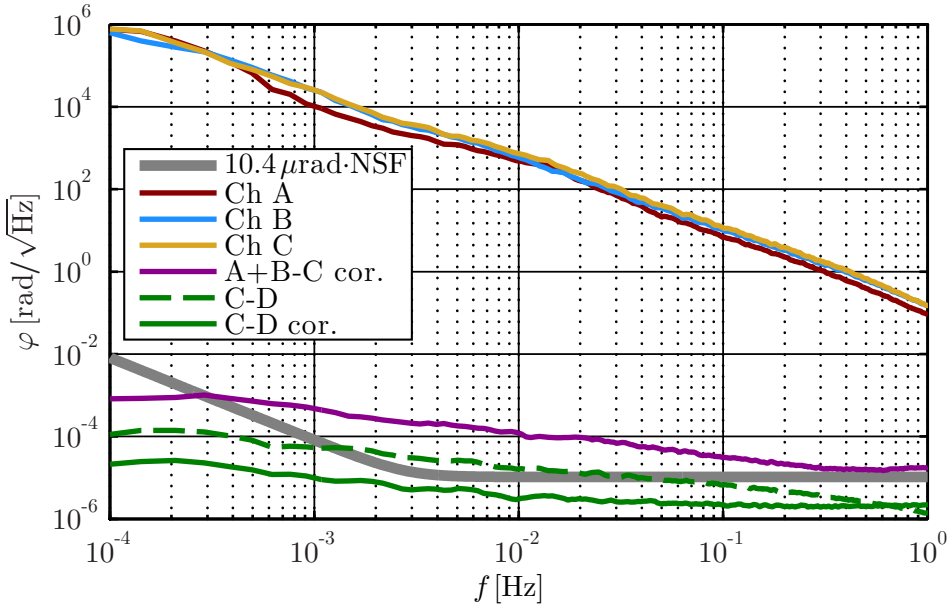
Figure 8.13.: Performance of the analogue three signal test. The pilot tone adder B.5 and a pilot tone of 35 MHz were used. Channel A was at 6 MHz, B at 9 MHz and channel C and D at 15 MHz. The three signal combination A+B-C is shown including a pilot tone correction. A parallel split measurement is shown with (C-D cor.) and without jitter correction (C-D).

The analogue three signal test would make it possible to utilise a simple set-up for non-linearity investigations without the effects caused in photo diodes, therefore it might be revisited again at a later point, during a more detailed noise hunting.

8.7. Optical three signal test

The most direct test of the phase measurement chain can be done with an optical three signal test (see Chapter 5). Using optical signals for generating the signal combinations puts all analogue and digital components, as well as the photo receivers into the sensitive path. Therefore such a measurement allows one to test the noise levels of all components and their linearity over the full signal spectrum.

The construction of such an interferometer is not trivial, since it should not introduce any additional noise by changing optical path lengths or by stray light beams, which act like an optical *crosstalk* [47]. Such an interferometer was designed by G. Heinzel and implemented by M. Dehne [108]. The layout of this hexagonal interferometer was optimised with the optical simulation software
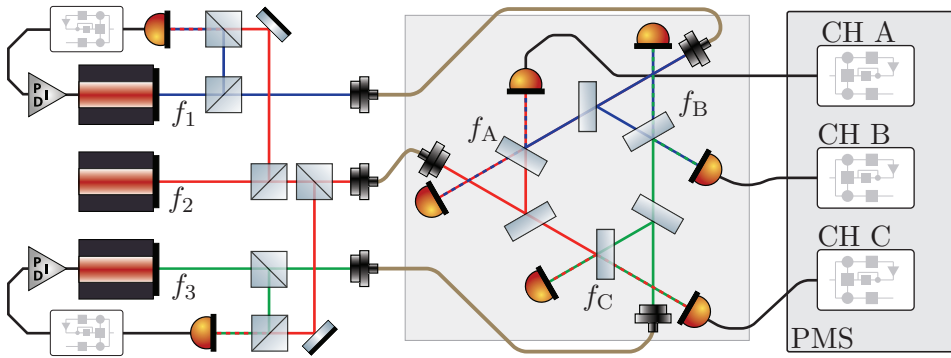
Figure 8.14.: A three signal test with analogue signals was implemented by mixing three GHz tones to generate beat notes in the MHz regime, which are then measured using a pilot tone correction as described in Section 8.4.1.

Ifocad [29] to minimise stray light influences, and it was built using the silicate bonding technique, which made it possible to achieve the high thermal stability required for testing the LISA phase measurement chain.

After the initial construction of this interferometer its working principle was tested by performing a three signal test using a phasemeter prototype, the same one as used for the analogue three signal measurement. This first test was done together with N. Brause and M. Dehne, using the set-up shown in Figure 8.14. Three Mephisto 500 NPRO lasers [121] were used as light sources. One of them acted as master and the other two were weakly locked to it via digital offset-frequency locks, based on an implementation by N. Brause, investigated during his Diploma Thesis [122]. The three weakly locked beams were then fed into single mode optical fibers and brought onto the interferometer via commercial fibre injectors [108]. Each beam was then split on the hexagon and interfered with the two other beams respectively, generating the optical beat notes required for a three signal test. Each beat note was available also with an opposite sign at the respective second beam splitter output. Three high bandwidth (> 50 MHz) photo receivers were used to convert the signals into analogue beat notes which were then fed into the phasemeter via equally long cables. A balanced detection using six photo receivers, two for each interference, was not possible at the time due to the lack of channels on the AFE board. The whole set-up was implemented in air, in a laboratory on an optical table, which is only a marginally stable thermal environment.

The performance results achieved during these measurements are shown in Figure 8.15. The individual signals show indirectly the suppression of the offset-frequency locks, which reduces their relative noise from the one of free running laser to a level similar to the one expected in LISA. The signal combination shows again a strong reduction, though this time a shoulder shaped noise is lim-
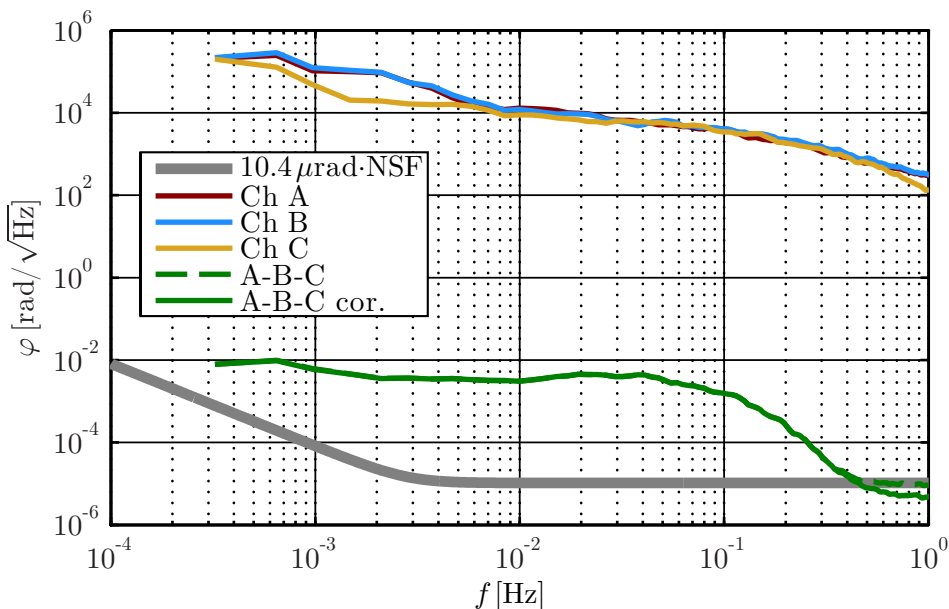
Figure 8.15.: Performance of the hexagon optical three signal test. The pilot tone adder B.5 and a pilot tone of 16 MHz were used. Channel A was at 6.25 MHz, and B and C at 3.125 MHz. The three signal zero combination is shown with (A-B-C) and without (A-B-C cor.) pilot tone correction.

iting the performance. The standard LISA requirement was only met around 1 Hz, where the pilot tone correction was able to further increase the performance. The cause for the high excess noise is currently not known. The fact that two beat notes used the same frequency makes it plausible that crosstalk might have been the limiting effect. Similar noise shapes have been measured under the influence of stray light [47], which also couples as small vector noise (see Section 3.4). The pilot tone adder configuration used for the measurements had a crosstalk suppression in the order of 60 dB (see schematics in Figure B.5). This levels fits roughly to the height of the shoulder of 10^{-3} rad/$\sqrt{\text{Hz}}$ expected from a small vector influence. These initial measurements were strongly aggravated by several cycle-slips occurring randomly during the measurements. The cause of these slips is currently unknown. Since the well tested ADPLLs described in Chapter 4 were used, they are rather unlikely to be the cause of this problem. Potentially the offset-frequency locks might have introduced these effects and, at the time, no in-loop data of these locks was available to investigate their performance and stability.

As an additional noise source angular jitter of the fibre injectors was identified [108]. This jitter couples into the length measurements because the optical components include a wedge angle to reduce stray light. This causes a change

of the effective path length through the components if the beams tilt, and due to the different refractive index of the components and the free beam medium (air or vacuum) this directly causes phase noise that is not common between the interferometers. To address this, the hexagonal interferometer will be modified for future experiments. It will use monolithic fibre injectors [123], as used in LISA pathfinder and LISA, which are expected to cause far less beam jitter. Since the experiment was not performed in vacuum, various thermal effects and thermally induced changes of the refractive index in air also contributed to the measurement noise.

Even though the measured noise is almost three orders of magnitude above the requirement this test was able to demonstrate the basic feasibility of the measurement concept. Various steps for improvement are well know from earlier experiments aiming at μrad performance levels [108, 47, 48] and they will make it possible to improve the measured phase noise performance in the future. This test also showed, for the first time, the feasibility of a pilot tone correction, though only in a small range, for an optical three signal test.

Chapter 9

Elegant breadboard model of the LISA phasemeter

A large part of the work performed for this thesis was related to our participation in a technology development activity initialised by the European Space Agency (ESA), called 'LISA Metrology System' (ESA Invitation to Tender: AO/1-6238/10/NL/HB). The goal of this activity was to design, build and test an elegant breadboard model (EBB) of the LISA phase measurement and the frequency distribution system. This activity was performed by a consortium composed of the National Space Institute of the Danish Technical University (DTU Space), the Danish industry partner Axcon ApS in Copenhagen, and the Albert-Einstein Institute in Hannover, Germany. This chapter gives an overview about the EBB and it summarises the most important results of the development activity.

The specific goals of the project are described in the first section. They are also put into perspective to the existing overall LISA technology development and future investigations. The second section of this chapter introduces the architecture of the EBB and explains critical design choices. This includes an overview of the hardware itself and is followed by a description of critical elements of the infrastructure. The implemented algorithms and parts of the programming of the FPGAs and the CPU are presented in the third section. The EBB was tested for key functionalities to validate the system architecture and the results of these test are presented in section four. The last section focuses on the testing of the most critical and most demanding function of the EBB, the phase measurement. Using a dedicated signal generator, LISA like signals were used to test the performance of the EBB in an extensive test campaign, of which the most important results are presented.

9.1. Goals

Measuring the phase of the inter-satellite and the additional interferometers continuously, with the full performance of $6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ and over the whole frequency range is the main function of the EBB, and demonstrating this was the first and foremost goal of its development. To perform a representative test the generation of signals containing the full LISA characteristics was required, which led to the development of a dedicated digital signal simulator (DSS). After proving the phase measurement capabilities the EBB could then be used to test various other aspects of the LISA metrology in future set-ups, including tests of photodiodes, DWS, clock noise transfer and many more.

The phase stability of the GHz modulation side bands in relation to the phasemeter sampling clock and, more importantly, to the pilot tone signal is crucial for the LISA metrology concept as described in Section 2.6. Without it, noise and drifts between the clocks in the three satellites would spoil the measurement performance. The generation and distribution of the relevant tones is therefore named as a separate item in the LISA technology development, the frequency distribution subsystem (FDS). The aim of the EBB development was to provide the most critical parts of the FDS, excluding optical and uncritical external analogue components, and to integrate them into the EBB. Especially the goal was to provide a pilot tone to all ADC channels, with the required phase stability to the GHz modulation side bands. By demonstrating this, the EBB could be used in experiments with more than one breadboard, using clock noise transfer to achieve the desired noise levels. This would allow us to perform more integrated tests of this critical aspect of the LISA metrology.

Absolute ranging between the LISA satellites is necessary to cancel the laser frequency via the TDI algorithms. Therefore it was required of the EBB to include the ability to read out PRN ranging modulations with the desired performance and to potentially generate them. Also the ability to encode data into the modulations and to decode data from the measurements was required, not only to demonstrate these functions, but also to allow testing them in future, more realistic experiments.

As described in Chapter 2, the phasemeter in an inter-satellite interferometer does not only read out interferometric beat notes, but it also controls them, or some of them, by actuating on the lasers. This is, for example, done with an offset-frequency phase lock. Therefore, the EBB was required to be able to perform such locks with sufficient performance and bandwidth. By adding the function to actuate two lasers and to use two beat notes the EBB would also be able to test more sophisticated control schemes developed for LISA, like dual arm-locking [72] or alternative laser stabilisation techniques [124, 125].

While the phasemeter can be seen as the final item in the LISA phase measurement chain, it is only one component in the drag free attitude control scheme (DFACS), that keeps the various degrees of freedom of the satellites and the
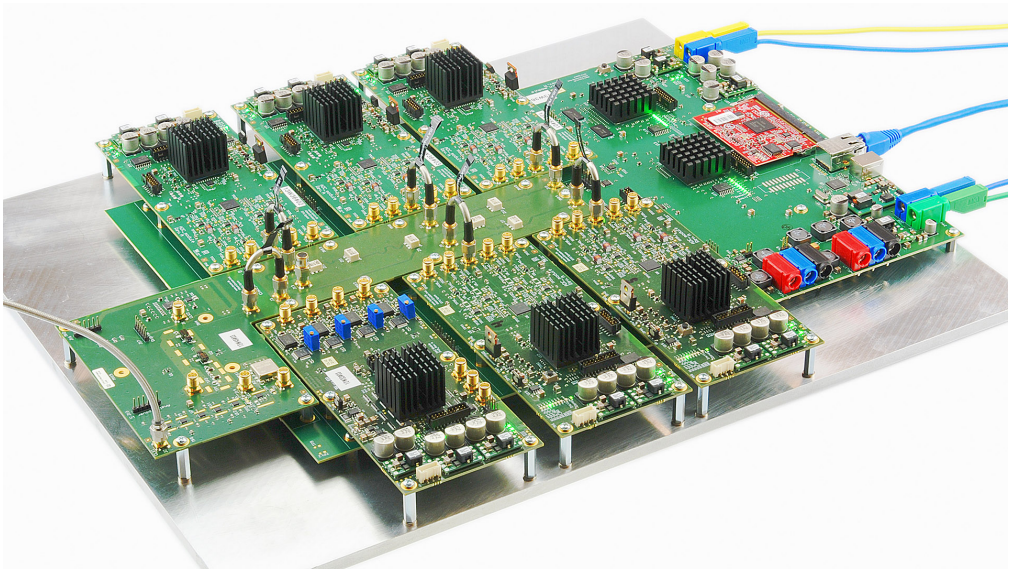
Figure 9.1.: Photograph of the assembled PMS EBB. Shown is a main module with one clock, one DAC, and five ADC modules.

test masses controlled using closed-loop algorithms. The EBB therefore aimed to include and demonstrate all functions that would be required for it to act as a real time read out. This ensures that the EBB could be used in future experiments that simulate part of the DFACS control in LISA or other future missions.

The last main function, required of the EBB, was the ability to acquire lock of a beat note and to potentially include algorithms to acquire lock between free running lasers. The EBB should also be able to perform this function, and potential variants of it, during the more complex acquisition between satellites, as necessary in LISA and GFO. Again, these capabilities enable the EBB to be used in all future experiments investigating this procedure.

Next to the described primary functions the EBB was also aimed to allow simple debugging and experimenting by providing additional sensors, resources and interfaces, as well as flexibility in the used algorithms and critical analogue components.

9.2. System architecture

The driving aspects of the EBB design were the required functionalities and the performance goals. But also considerations for cost, power consumption, flexibility and special aspects of known critical elements (like the analogue front end) were taken into account. Figure 9.1 shows a photograph of the EBB hard-

ware. While for critical components space qualifiable parts were investigated and partly implemented, standard commercial parts were used for the bulk of the system. A system sampling clock of 80 MHz was chosen, which is higher than in previous prototypes and simplified some aspects of the design, like the anti-aliasing filters (see Section 3.2.3). FPGAs were selected from the Spartan 6 series from Xilinx, since they provided sufficient logical resources, including large margins for future developments and experiments, together with reasonable power consumption and price. The hardware was split into four individual modules, each containing different parts of the functionality. One full EBB consists of one main module, one clock module, one DAC module and five ADC modules. Hereinafter, each module is presented together with its key components and functions. Afterwards two critical aspects of the infrastructure are described in more detail, the signal and data flow between the components and the synchronisation subsystem that is required to ensure a common system timing.

9.2.1. ADC module

The analogue signals from the interferometers are digitised and processed on the ADC module. The main components on this module are an analogue front-end circuit, which conditions the input signals and adds the pilot tone, an ADC, that performs the digitisation and an FPGA, that houses the readout algorithms and communicates with the other subsystems. The layout of the ADC module is shown in Figure 9.2. The module can either be connected to the main board or it can be used in stand-alone operation. The latter variant simplifies testing, since a single ADC module can be used as an independent phasemeter prototype and does not contain the various additional features of the full system which could potentially cause unwanted disturbances.

An ADC module is plugged via a multi pin connector directly into one of the six slots of a main module. This connector provides the system clock and the supply voltages to the module. It also connects several interfaces, most importantly the main, high bandwidth link between the local FPGA and the FPGAs on the main module. One ADC module was designed to measure four input channels. Therefore it contains four equal analogue front-ends and an ADC with four input channels. The ADC is the ADS6445 from Texas Instruments [109] and its phase measurement performance was tested beforehand, as described in Section 8.1. The main function of the FPGA on this card is to perform the phase and ranging readout algorithms, to decimate this data to a suitable rate, and to transmit it to the main module.

In stand-alone operation the FPGA can directly be controlled and read out via a USB or generic serial interface, added to the module by using free FPGA pins connected to an expansion header. Another connector provides power and an additional crystal oscillator has to be soldered onto the module to provide a
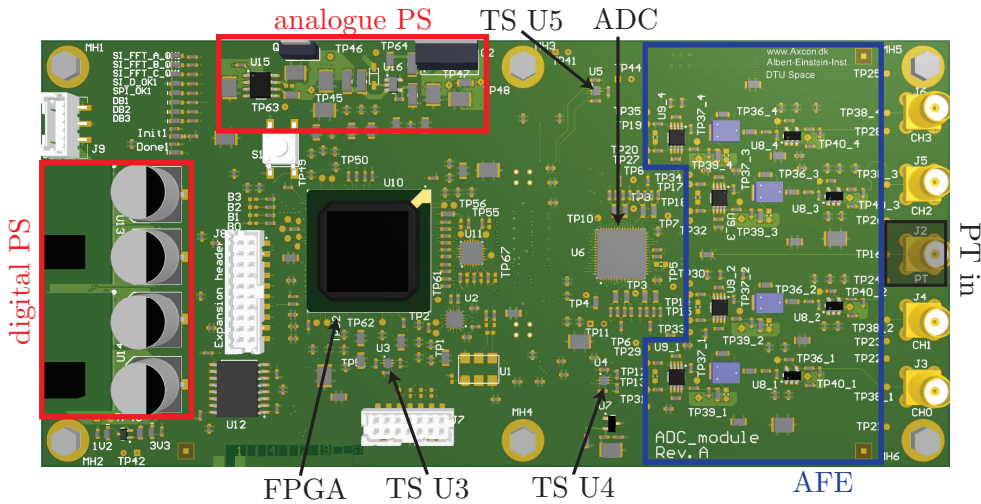
Figure 9.2.: Layout of the ADC module. It contains four signal inputs, one pilot tone input (PT in), four analogue front-end circuits, one ADC, one FPGA, three temperature sensors (TS), an analogue and a digital power supply (PS) as well as several other peripheral components and connectors. The connector to the main module is located on the bottom.

system clock.

Analogue front end

The signals are fed into the ADC module via four SMA connectors. A fifth connector is used to fed in the pilot tone signal directly from the clock module. The AFE circuit for one channel (shown in Figure B.10) includes three main components: a gain and buffer stage for signal conditioning, an anti-aliasing filter, and a differential amplifier, that is used to simultaneously add the pilot tone and to generate a differential signal required for driving the ADC. This architecture was based on the pre investigations using phasemeter prototypes developed at the AEI [43].

The gain stage was implemented using a current-feedback op-amp with a high gain-bandwidth product (LMH6702 [114]). The aim of this was to generate a flat amplitude and phase response for various signal gains, to reduce temperature induced noise. No fixed requirements on the gain were available and, since they depend highly on the signal power and the photo receiver (PR) gain, the implementation allowed for flexibility by simply changing the feedback resistors and therefore the amplification.

The anti-aliasing filter (AAF) is a simple 2nd order low pass and was included for testing purposes. A detailed study of the required filter suppression and

phase linearity, as discussed in Section 3.2.3, was not conducted, since they depend strongly on the PR transfer function and other not specified parameters. A full measurement chain design might move the AAF, as well as all required gain stages, fully into the PR, depending on the available thermal stabilities and power budgets.

The differential amplifier, used to drive the ADC and to add the pilot tone, has a large bandwidth ($> 300\,\mathrm{MHz}$), to reduce thermally induced phase noise. The pilot tone was fed into the second input port of the amplifier and it was simply split into the four channels without additional impedance matching (except for a $50\,\Omega$ matching to the SMA input). This is considered feasible since the dimensions on AFE were small compared to the signal wavelengths.

Since all of these analogue components were known to be critical for the measurement performance, a possibility for exchanging not only the front-end components, but also the layout was included into the design. Several test-points were added, that allow one to solder a hook-up board onto the original AFE. Such a board directly connects to the signal inputs and the ADC inputs, and can therefore replace the original AFE circuit, thereby most of the ADC module can be re-used. More details of the AFE and modifications are discussed in Section 9.5, together with the phase performance measurements.

9.2.2. DAC Module

The EBB generates up to three digital and four analogue signals with the DAC module. Its main components are an FPGA, four DACs, each with an analogue offset correction and amplifier circuit, and three buffered digital outputs. The FPGA infrastructure of this module is almost an exact copy of the one of the ADC module. This allowed us to use some of the same programming code and it made it possible to easily exchange the two module types. Therefore a main board can potentially be used with none, one or even more than one DAC module. The layout is shown in Figure 9.3. The DAC module can also be used in stand-alone operation.

The FPGA on this module has two main functions. The first function is to host offset-frequency phase locks [40, 126] or similar algorithms. These algorithms use the frequency information of the beat notes from the ADC modules, communicated via the main board, to generate actuator signals. These signals are then converted into analogue ones by the DACs and their offset and amplitude is further conditioned by analogue back-end circuits, before they are made available at separate SMA connectors. Each output was designed to have a range of $\pm 10\,\mathrm{V}$, which is convenient for laboratory experiments. While the DACs are sampling with $80\,\mathrm{MHz}$, the analogue outputs were low pass filtered with a corner frequency of approximately $10\,\mathrm{MHz}$. This is a more than sufficient bandwidth for laser control and allows the DACs to also produce RF signals for other purposes. The outputs also include a resonant filter at $80\,\mathrm{MHz}$ to reduce
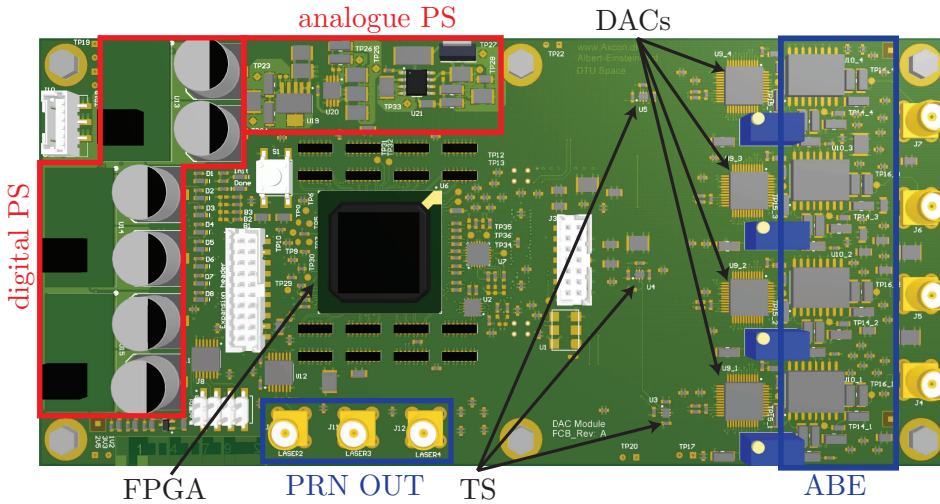
Figure 9.3.: Layout of the DAC module. It contains three 1-bit digital outputs, four DACs, four analogue back-end (ABE) circuits, one FPGA, three temperature sensors (TS), an analogue and a digital power supply (PS), as well as several other peripheral components and connectors. The connector to the main module is located on the bottom.

any residuals of the clock signal in the outputs. The actuator signals can also be made available via a serial interface, if a laser with digital interfaces has to be controlled.

The second function of the FPGA is to generate PRN signals for the phase modulation required to implement ranging and data communication (see Section 2.6.1). For three channels, data received from the main board is therefore encoded with individual PRN signals and fed via simple digital buffers to SMA connectors.

9.2.3. Clock Module

The clock module generates the system clock and it provides and distributes the pilot tone to the EBB. It is plugged into a dedicated connector on the main module which supplies the supply voltages and takes in the clock signal from the clock module. Figure 9.4 shows the layout of the two clock modules designed for the EBB.

Both variants contain a pilot tone distribution section, where the pilot tone is split via power splitters and made available at six SMA connectors with equal delays to ensure a common overall phase. This type of distribution was chosen to ensure a high phase stability and cleanliness of the pilot tone, both of which might have been spoiled by using a distribution via the main board. The
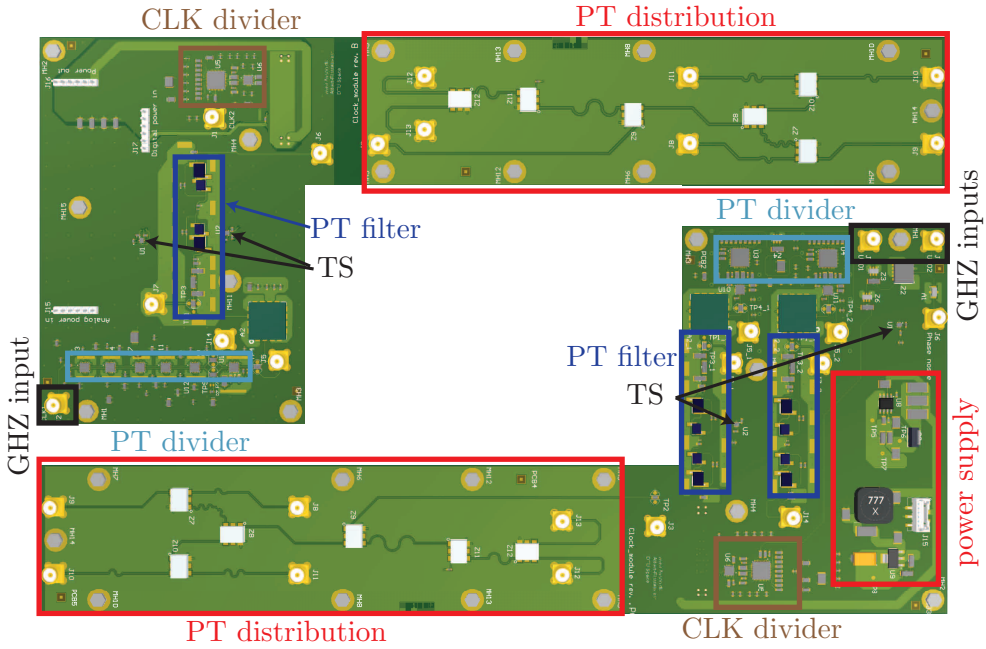
Figure 9.4.: Layout of the second generation (top) and first generation (bottom) clock module. The elongated part of the modules contains the pilot tone distribution with delay matched signal lines and six power splitters. Except for an additional connector the design of this part is equivalent for both modules. The other part of the modules contains the frequency generation. It consists of the clock divider and of the pilot tone divider and filter (the first generation board contains two pilot tone chains for testing). The power supply was removed from the second generation board and has to be added externally, since it was found to cause large temperature gradients over the pilot tone signal chain.

distribution part of the clock module is in the middle of the ADC modules (see Figure 9.1), this shortens the signal path and thereby decreases the influence of strong heat dissipation by the FPGAs and power supplies, which are situated on the outside of the EBB.

The second part of the modules is the tone generation. On both boards a 2.4 GHz signal, which is one of the GHz clock noise side bands, is divided twice, one time by 30 to generate the 80 MHz clock tone and one time by 32 to generate the 75 MHz pilot tone. While the clock is directly fed to the main module, the pilot tone is further filtered and amplified. Since the initial pilot tone is a rectangular signal (due to the digital dividers) a temperature compensated passive filter is used to reduce its content of higher harmonics.

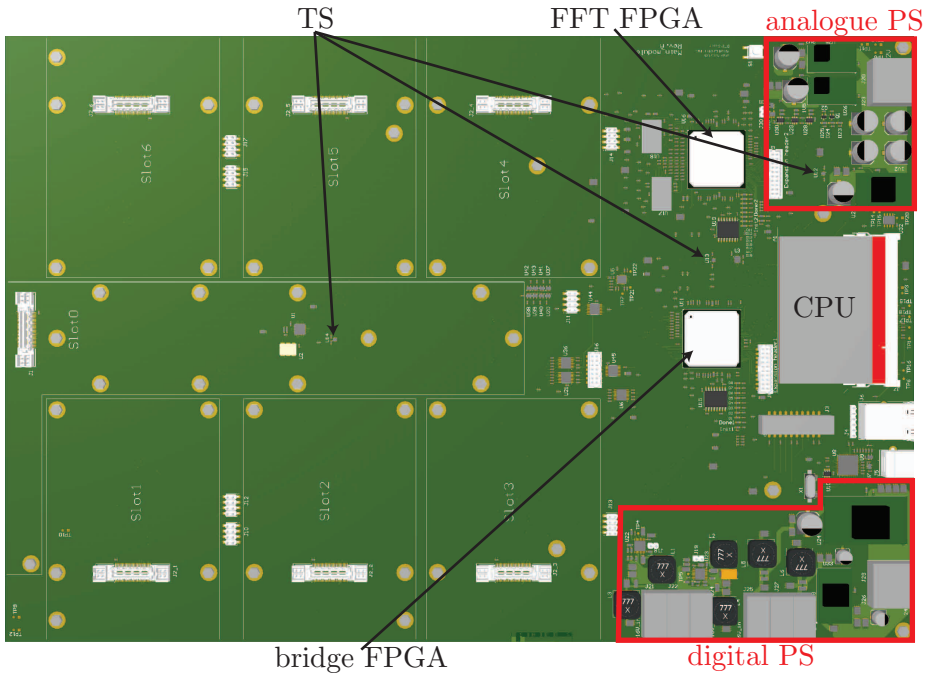The original clock module contained two such pilot tone divider and filter

Figure 9.5.: Layout of the EBB main module. The left side contains seven connector slots for sub modules. Slot 0 is reserved for the clock module, slots 1 to 6 can be equipped with ADC and DAC modules. The right side of the EBB contains the FFT FPGA, the bridge FPGA, the CPU, the power supplies and various connectors and additional peripherals.

chains, to be able to easily compare them. Measurements using this module revealed however, that an excess noise was present in the chain and therefore a second, revised board was developed. It uses a different divider architecture, carries less components and uses an external power supply. These changes should bring the performance of the pilot tone chain below the requirement. The clock module also has an external supply voltage connector to allow for stand-alone testing.

9.2.4. Main Module

The main module hosts all other sub modules, provides power, distributes the system clock, it hosts an FPGA dedicated to performing fast Fourier transformations (FFT FPGA), it hosts the CPU module that controls the system, and it hosts an FPGA that establishes the interfaces between the FPGAs and the CPU (bridge FPGA). The main module can also use a crystal oscillator as source for the system clock. This allows using the system without a clock module or if, for example, no 2.4 GHz source is available. A 12 V supply voltage is passed to

all sub modules and used to generate all voltages for digital circuits. A 24 V supply voltage is used to generate the local supply voltages of ± 6 V and ± 16 V for analogue components, which are also supplied to all sub modules.

Bridge FPGA

The bridge FPGA works as the central switch for most communications between the FPGAs and the CPU. It multiplexes the various data streams that are send back and forth in the system and frames them into the required formats. The data streams are discussed in Section 9.2.5. It also handles the FPGA synchronisation, which is described in more detail in Section 9.2.6. For simplification this FPGA was chosen to have the same type as the FPGAs on the ADC and DAC modules.

Central processing unit

The EBB required a central processing unit with floating point capabilities and high-level programming options. The LPC3250 OEM Board from Embedded Artists was chosen for this purpose. It houses an NXP ARM926EJ-S LPC3250 processor and includes various peripherals. It is simply mounted onto the main board with an expansion connector, which provides power and the required connections. The CPU functions can be divided into two main tasks. The first one is to handle the measurement data. The incoming data packages have to be formatted, filtered and then packed and transmitted via Ethernet to a connected PC. Later on the data might also be send to a DFACS system, via a not yet defined interface, or to another phasemeter, via sending it through an optical link using the PRN modulations. The filtering of the incoming data is the most demanding task in terms of required processing power. The chosen CPU therefore includes a dedicated vector floating point co-processor, which gives the system a computational overhead that can be utilised during later developments. The second task of the CPU is to control the system and to perform housekeeping functions. The availability of a general-purpose programming environment allows one to implement complex algorithms that can be used to operate the EBB autonomously. This is necessary for using the EBB in experiments that simulate flight-like situations, where the communication to the satellites would be non-permanent and very limited in bandwidth. For example, the autonomous acquisition of the ADPLLs had to be implemented as part of the EBB development. An important additional feature of the CPU is the readout of temperature sensors spread across all modules. This data is also packed and send to the PC. The CPU runs the RTEMS (Real-Time Executive for Multiprocessor Systems) operating system.

Fast Fourier transform (FFT) FPGA

To allow for maximum acquisition performance a dedicated FPGA was implemented to perform fast Fourier transformations. This FPGA is connected directly with high bandwidth interfaces to some of the ADC modules. The full ADC data sampled at 80 MHz is fed through these interfaces and made available on the FFT FPGA. There it is fed into FFT blocks that perform full duty-cycle FFTs on four input data streams, for example, on four segments of a QPD. To achieve this amount of computational power the FFT FPGA is a larger Spartan 6 type, with roughly twice the logic resources than the other FPGAs in the system. The FPGA also hosts peak detection algorithms which operate on the FFT outputs and produces the data send to the CPU.

Additionally the FFT FPGA connects to random access memory (RAM) with a capacity of 256 megabyte. This allows a burst of data to be stored at the full sampling rate, which can then be read out via the CPU. The ADC modules providing this data can also be modified to transmit, for example, signals from inside the ADPLL. These types of burst data can be used to perform debugging and signal analysis over a wide range of frequencies, something which is not possible otherwise, due to bandwidth limitations.

9.2.5. Data flow

The modular design approach of the EBB required to implement interfaces that have sufficient bandwidth for the different functionalities. The various interfaces can be divided into two types, streamed data, which is send at regular intervals and contains well defined data frames, and command data, which consists of the register read and write packets that are controlled by the CPU. An overview of these interfaces is shown in Figure 9.6. The register interface uses the bridge FPGA as a switch and connects from there to all other FPGAs. There are three types of streamed data interfaces implemented into the EBB. The first one is the science data, it contains all the various DSP readout signals and additional meta data, like the counter values from each FPGA. This data is send to the bridge, where it is re-framed and then send as one large package to the CPU. The rate of the science data is rather slow ($80\,\mathrm{MHz}/2^{17} \approx 610\,\mathrm{Hz}$), but it is sufficient for providing signals to a DFACS system and it has to be even further decimated for the data storage. More details about the science data and its decimation are described in Section 9.3.3. The second streamed type of data is for the laser lock implementation. For these locks the EBB acts as part of a fast control loop, with a requirement of an absolute delay of less then $\approx 3\,\mu\mathrm{s}$. To implement these locks the frequency/phase information of the beat notes, measured in the ADC FPGAs, has to be send to the DAC FPGAs without causing excess delays. A rate of 10 MHz was chosen for this interface, which kept the delays caused by the interface sufficiently low. Only two of the laser lock signals sent from the five
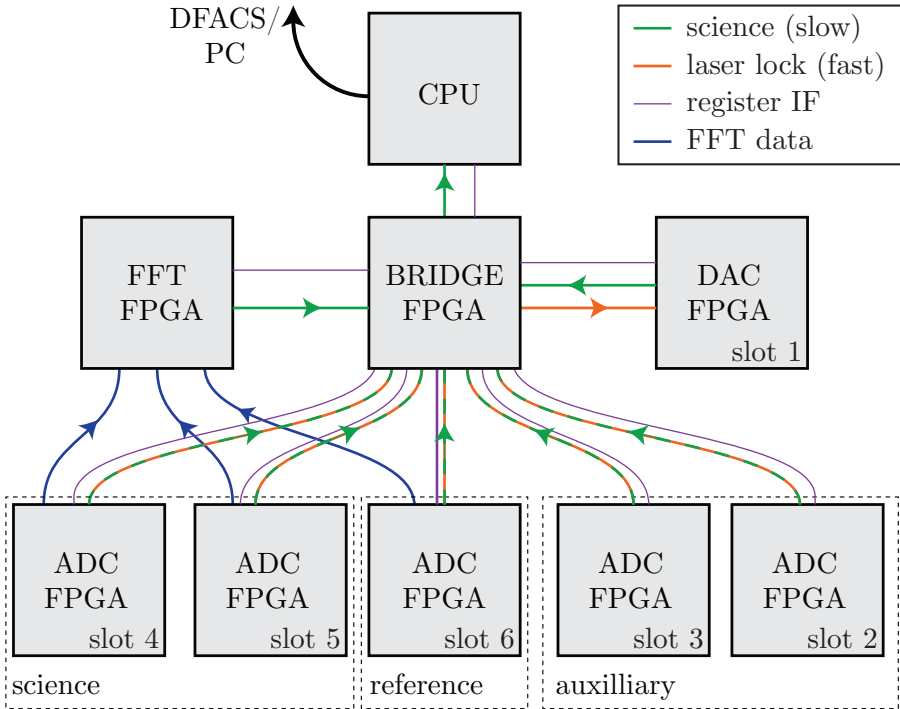
Figure 9.6.: Sketch of the data flow between the FPGAs and the CPU on the EBB. The streamed data interfaces (science, laser lock, FFT data) are all one directional and they send pre defined data frames at fixed intervals. The register interface is bi-directional and is used to send and read data. It is not synchronised, since it is controlled by the CPU. The data, encoded with PRN codes on the DAC card for inter-satellite communication, is communicated also via the register interface, since it only has a low bandwidth requirement.

ADC cards can be send further to the DAC FPGA, therefore a multiplexer in the bridge chooses which ones are required. The physical interface between the bridge and the other FPGAs is based on a single serial link with a bandwidth of several Gbits/s.

The last type of streamed data interface is the delivery of all the ADC sampling data from three ADC modules to the FFT FPGA. The large required bandwidth (4 ADCs $\cdot$ 14bit $\cdot$ 80 MHz $\approx$ 4.5 Gbit/s) was achieved by adding three physical serial interfaces to two of the ADC slots. A single additional link was implemented for a third ADC slot, with the reasoning, that a reduced amount of data is sufficient for acquiring lock with the local reference interferometers in LISA , since it is mechanically fixed and local it has a constant high contrast with low SNR signals.
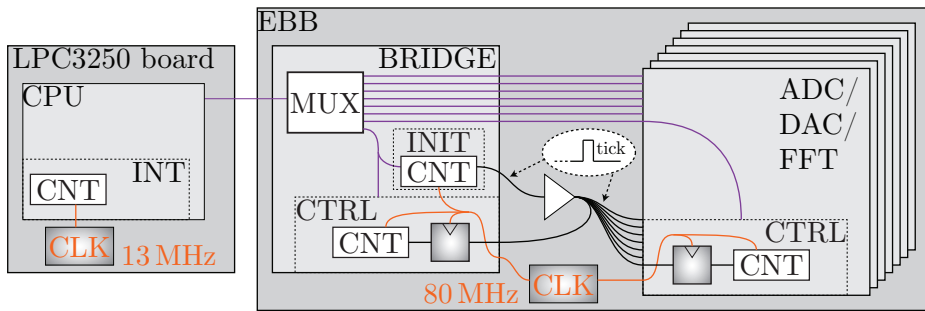
Figure 9.7.: Shown is a sketch of the EBB synchronisation scheme. A counter (CNT) in the bridge FPGA, with a dedicated state machine (INIT), is triggered by a CPU command to generate ticks at regular intervals. The ticks are distributed via a buffer and delay matched signal tracks to all FPGAs. There the tick is sampled, with the same system clock, and used to trigger the pre armed counters in the respective FPGA control blocks (CTRL). All data packages send to the CPU contain this counter value from the time of transmission. The value can be compared to a software package counter in the CPU to detect missing packages. The CPU also contains an internal hardware counter (INT), driven by an independent clock, which is also read out for debugging purposes.

9.2.6. Synchronisation

A dedicated synchronisation scheme was implemented to generate a system wide absolute time. This is necessary to ensure that all sampling and decimation, as well as the PRN generation, start predictably at the same edge of the system clock. The scheme is based on a synchronisation signal, the so called tick, that is generated by the bridge FPGA and fed to all FPGAs (also back to the bridge) via delay matched lines (see Figure 9.7). Each FPGA contains a state machine that controls a counter that is triggered by the received tick. After a successful synchronisation each FPGA has an exact copy of this counter, that can be used for all local timings. This counter is, for example, used to generate the decimation clocks and the PRN codes. The state machine controlling the counter is armed (it waits for the tick), by a simple command from the CPU. The tick generation produces a tick at regular intervals, but it can be switched on and off via a CPU command. Therefore triggered state machines can also be used to time other algorithms or changes in the FPGAs, allowing one to exert full control of timing in the system, if necessary.

If the synchronisation works, only determined, constant time offsets remain, which have to be integrated into the signal post-processing. One should note, that the full timing in the EBB has three components, the 80 MHz clock, which defines the timing intervals, the tick synchronisation, that sets the absolute

system time i.e the counters, and the pilot tone, which corrects the timing intervals of each ADC channel relative to the actual system clock which, in this case, is the 2.4 GHz side band tone.

9.3. Algorithms & programming

Six devices on the EBB required individual programming, the four types of FPGAs were all programmed using VHDL, the CPU and the PC interface programs were written in C. The hardware interfaces on all devices were implemented by Axcon, together with additional abstraction layers, most of the infrastructure for the register interface and the operating system and drivers for the CPU. The digital signal processing, control, framing and data streaming was then implemented on top by AEI and as part of this thesis. In the following relevant aspects of the programming are described in more detail.

9.3.1. Control

The EBB is initialised via a PC program. The program establishes a TCP/IP link via Ethernet to the CPU and provides it with various parameters, depending on the chosen mode of operation. The CPU then performs the FPGA counter synchronisation and initialises the DSP blocks in the system by writing various parameter into the corresponding register, like the gains and initial frequencies for the ADPLLs. After this, the CPU activates the streamed data interface on the bridge and it begins the data readout loop. During this loop it waits for packages to arrive, decodes and filters them, and sends the resulting packages to the PC, where they are stored.

The system is controlled via a register based interface handled by the CPU. Each register in the EBB has a specific address and it is written via a simple CPU command. Registers can also be read out by the CPU if they are configured for that in the VHDL code (readable registers required significantly more resources and therefore they had to be instantiated specifically). For debugging purposes a notification for each register write and read is also send to the PC. It contains the data and address of the register, as well as the CPU internal counter value for timing. This enables us to reconstruct every behaviour of the CPU, including reaction times, which is especially important if the EBB works partly autonomously.

9.3.2. Science readout

The readout algorithms implemented on the FPGA and each ADC module are described here. The here presented specific scheme is designed for detecting the four beat notes generated on a QPD for the LISA inter-satellite link. This scheme was implemented during the development, since it is the most complex
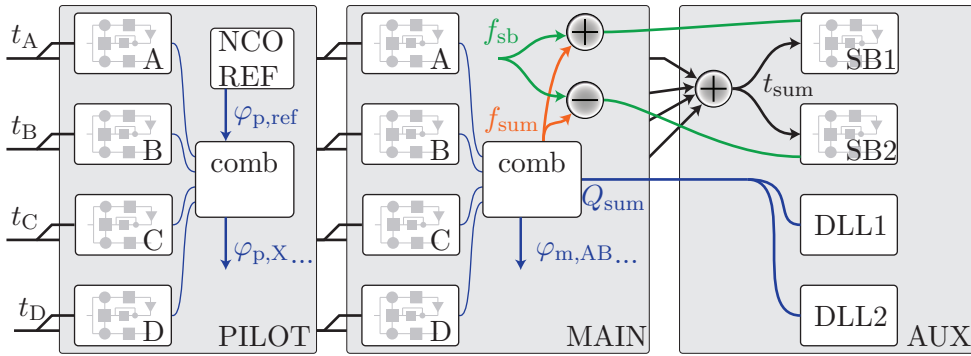
Figure 9.8.: Shown is an overview of the DSP implemented in an ADC FPGA for the science readout. The four digitised signals from the ADC (t_A,...t_D) are fed into four ADPLLs for the pilot tone readout and four ADPLLs for the main beat note readout. The added time series signal (t_{sum}) is used as input to the ADPLLs tracking the upper and lower clock noise side bands. The added mixer outputs of the main PLLs (Q_{sum}) is used as input for two delay locked loops (DLL), tracking the local and remote PRN ranging codes and performing the data readout.

and critical one for the missions. The implementation of the algorithm might be modified and simplified for the readout of, for example, the auxiliary interferometers. The scheme is sketched in Figure 9.8.

Pilot tone PLL

The pilot tone is measured for each channel by a dedicated ADPLL. The current implementation used almost the same design as developed for the interferometric beat notes. A more specific design might utilise a very low bandwidth, since the absolute pilot tone phase changes are very small. The frequency of such a dedicated pilot tone PLL might even be fixed to the specific frequency by additional logic. The values that are read out from these ADPLLs are the frequency values, used to reconstruct the pilot tone phase changes, and the in-phase (I) value, to monitor the signal strength. A readout of the out-of-phase (Q) value is not necessary due to the high bandwidth compared to the absolute signal. The readout of the frequency has a large dynamic range which is not required for the pilot tone. A simplified readout mechanism was therefore included additionally. Since the phase of each PLL contains a ramp at the signal frequency, a constant ramp from an internal NCO (synchronised to the system counter) is subtracted from each of them. The so generated phase values are directly proportional to the sampling jitter in each channel, and they use less then half the amount of bits and therefore bandwidth as the frequency values.

Science tone PLL

The interferometric beat notes are read out in the so called main ADPLLs. Their internal design is based on the simulations and pre-experiments presented throughout this thesis. Read out from these PLLs are the individual frequencies, I and Q values as well as the direct phase differences for DWS. In addition the sum of the four frequency values is computed and read out. This value is also decimated to a rate of 10 MHz and used for the laser locks.

Side band PLL

The added ADC signals are fed into two separate ADPLLs for tracking the clock noise transfer side bands. Since the side bands only use a small amount of the signal power their effective SNR is even lower than the one of the main beat notes. At the same time the side band PLLs have to track the same laser frequency noise dominated signal as the main beat notes, in addition to the clock noise and the Doppler shifts. The side bands are 2.4 GHz away from the carrier and their effective Doppler shifts are thus also slightly different. Since the largest part of the phase signal is already tracked by the main beat notes this information can be used to reduce the required bandwidth of the side band PLLs. Therefore the summed frequency of the four main PLLs is used, after adding/subtracting the corresponding offset, as offset frequencies in the SB PLLs (see Figure 9.8). Thereby the frequency noise and Doppler shift information is transferred in a feed forward scheme and the SB PLLs only need to track the clock noise and the slight Doppler variations. Individual frequencies and the I and Q values are read out from these PLLs.

Delay-locked loop

The readout of the inter-satellite ranging is performed in two delay-locked loops (DLL), which are tracking the local and remote pseudo-random noise (PRN) codes. They use the combined sum of the main PLL mixer outputs (Q_{sum}) as inputs. These DLLs also read out the raw data from the PRN stream. For the DLLs a VHDL version of the algorithms developed at the AEI by J.J. Esteban [44] and implemented by N. Brause is used. The readout signal of the DLLs are the correlation values, the measured delay and the raw data. The raw data normally contains additional bits for data correction algorithms like Reed-Solomon encoding [127], to account for bit errors. The decoding of received data is planned to be implemented into to CPU.

9.3.3. Decimation chain for science data

For the science data the decimation is implemented in three steps, shown in Figure 9.9. The initial data inside the FPGAs, available at the sampling speed
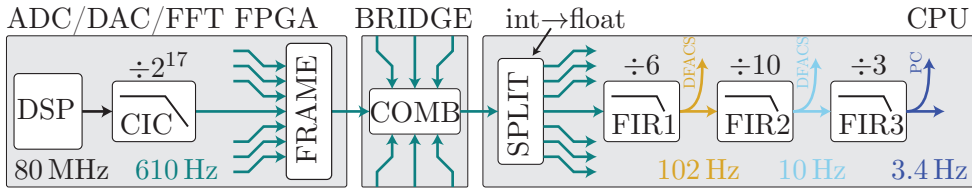
Figure 9.9.: Shown is the decimation chain of science data in the EBB. Each individual signal is decimated in the FPGAs and packed into a frame. These frames are transmitted to the bridge, where they are combined and send to the CPU. The large frame received by the CPU is split into the original signals which are then mapped onto floating point numbers and further decimated. Signals like counters and received DLL data are not filtered and are available in the CPU as integer values.

of 80 MHz, is down sampled by a CIC filter (see Section 4.3.4) to a much lower rate of ≈ 610 Hz. After being transmitted to the CPU, the data is converted from integer to floating point numbers and then decimated in three consecutive steps by FIR filters. The final rate of 3.4 Hz is then used to store the data on a PC or to transmit it via the laser link. The intermediate higher sampling rates of ≈ 102 Hz and ≈ 10 Hz are required for feeding data to a DFACS system. The consecutive filtering also simplifies the design of the three FIR filters.

The most important signal, in terms of filtering, is the frequency information of the main PLLs, since it contains the satellite distance information. It also has the largest dynamic range, making its decimation the most demanding. The other signals can potentially use much simpler filters that, for example, cause less signal delays, as required for DWS. Due to time constraints only the strong frequency decimation filters were implemented and they were used for all signals. The frequency signals are fully dominated by laser noise, therefore the filter suppression can be computed by comparing the expected noise with the filter transfer functions. Figure 9.10 shows the effective laser frequency noise model and its attenuation via the implemented down sampling filters. One should note that decimated variable is a frequency and not the phase, which changes the spectral properties, as well as the shape of phase requirement. Future implementations might also include compensation filters to gain the signal flatness that is required for DWS, post processing and TDI. However, these filters might be computationally very expensive due to the strong constraints on the signal flatness required for the TDI algorithms that suppress noise by more than nine order of magnitude. Therefore the filters might just be better situated in the on-ground post processing.
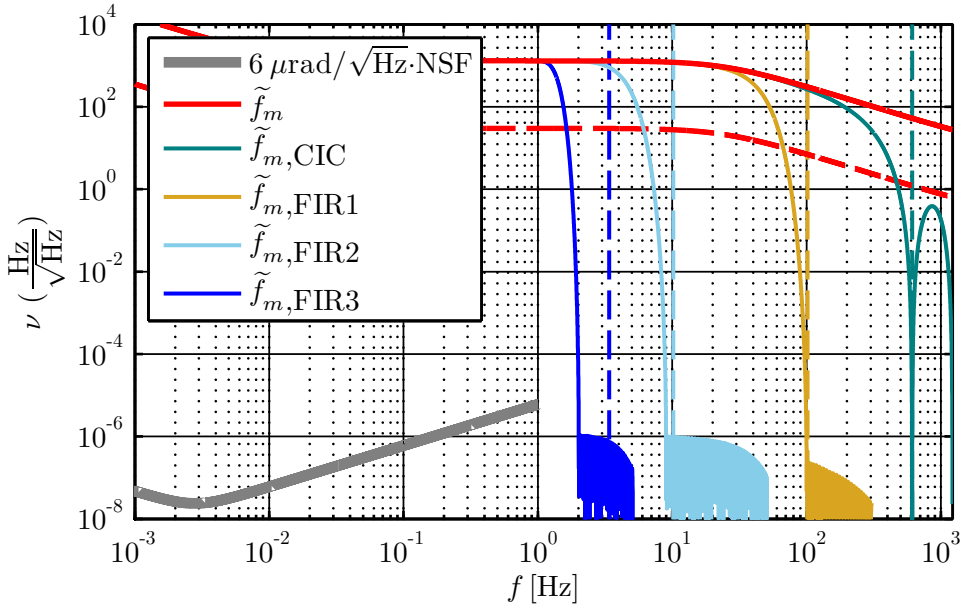
Figure 9.10.: Shown is the effective laser frequency noise model used in the signal simulation and its attenuation in the different decimation steps. The dashed vertical lines mark the sampling frequencies, which are also the lowest frequencies that will be aliased into the signal band ($> 1\,\mathrm{Hz}$) during each decimation. The first stage is a 3rd order CIC filter which has notches of suppression at all critical frequencies (only the first two are shown). The second, third and fourth stages are 54, 104 and 77 tap FIR filters respectively (see Appendix C). For the earlier assumed laser noise model (red line) the suppression of the stages is not sufficient. Using a more realistic frequency noise, closer to the current LISA base line design ($30\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ in the measurement band, dashed red line), relaxes the required suppression and makes the implemented filters compliant with the performance. One should note that the later described measurements use the larger frequency noise, but that no aliasing effects are visible, since the aliased components are equal in all channels (see Chapter 5).

9.3.4. Laser control

So far, the simple laser frequency offset phase lock has been implemented, but none of the other, more complex algorithms. The used algorithms are based on analogue implementations of such a lock that have been used and studied before [40, 41].

The active laser control, performed by the phasemeter, uses the frequency information of the beat notes, determined in the ADC FPGAs. From these frequencies the desired offset is subtracted and the residual value is integrated

to generate a phase error signal. This signal is then fed through a simple PI controller to generate a fast actuation signal that is send to one of the DACs. The fast actuator signals is also fed into a another PI controller that generates a second slow actuation signal. Both of these signals can then be used to drive the frequency of an NRPO laser by actuating on the laser crystal with a piezo (fast) and with temperature (slow). More details of the implementation are shown with the tests in Section 9.4.2.

9.4. Implementation and functionality tests

The previously described interfaces and algorithms were implemented and verified by simple functionality tests throughout the development. As part of the technology development it was also required to run some more extended test of the critical functionalities, to evaluate the feasibility and performance of the EBB design. Some of these test results are presented in the following.

9.4.1. Test of PRN generation

As described in Section 2.6.1, the LISA phasemeter has to generate PRN codes that are phase modulated onto the light to allow the implementation of a ranging and data communication system. A specific set of six suitable codes, each with a length of 1024 bits and a minimal cross correlation, was designed by J. J. Esteban and G. Heinzel [44] and they were implemented into the EBB. Since a full test of the data encoding and decoding was not part of the project the PRN generation was tested by itself. The codes are generated with a modulation frequency of 1.25 MHz and they are either put out directly or they are multiplexed with a clock at 2.5 MHz. The later scheme is used to implement Manchester encoding, a technique that changes the spectral shape of the PRN modulation and thereby reduces the high pass filter effect of the ADPLL on the modulation side bands during the read out.

The six codes were generated by register controlled state machines and small memory blocks in the DAC FPGA, they were fed through digital buffers and time series were taken with an oscilloscope. Two of these measurements are shown in Figure 9.11. Both of them are shown twice, once with an additional delay of the code repetition length. The perfect identity of the delays codes shows that the codes are indeed pseudo-random and that their repetition rate is correct.

To validate that the six codes have the correct properties a cross correlation diagram of the measured time series was computed, which is shown in Figure 9.12. As desired, the correlation peaks only on the diagonal elements and for zero delay. All other correlations are minimal, proving that the codes and their implementation are as desired.
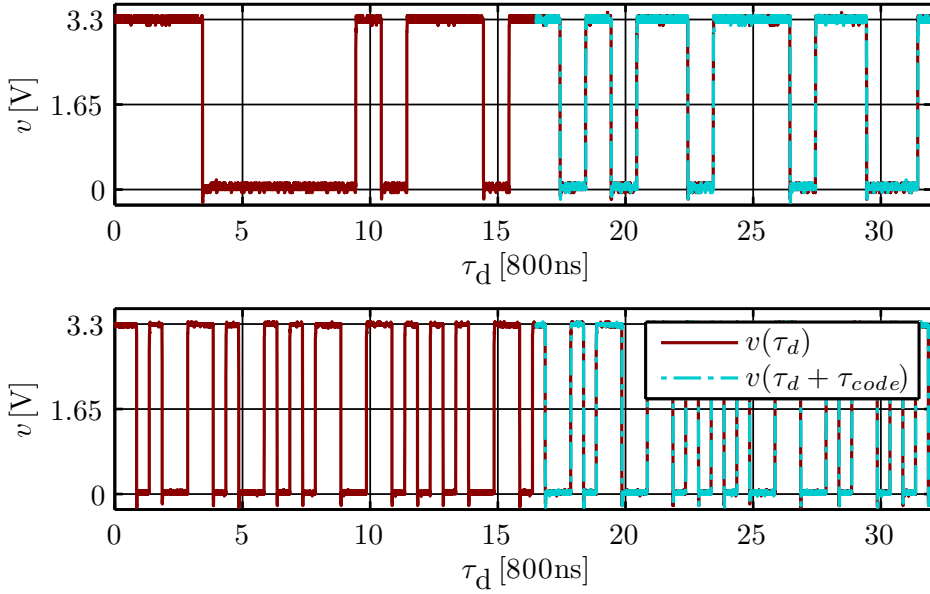
Figure 9.11.: Time series of PRN code 1 generated from DAC card. The upper graph shows the signal with a direct PRN output, while the lower graph shows it with an additional Manchester encoding.

9.4.2. Test of offset-frequency phase lock

A test of the laser lock capabilities of the EBB was done by implementing a frequency offset phase lock with signal generator. Thereby the lock could be tested in a well defined environment, without the complexity of an optical experiment. The used set-up is shown in Figure 9.13. A frequency modulation input of the signal generator (SMB100A by Rhode & Schwarz [128]) was used as actuator to influence the frequency of the generated beat note that was fed into an ADC module on the EBB.

A model of the closed-loop behaviour was compiled to determine the gains in the PI controller to achieve a stable control loop with a bandwidth of more than 10 kHz. The initial model included the transfer function of the ADPLLs (from the input to the frequency output), the decimation filter in the ADC module, the interpolation filter in the DAC module, the phase accumulator (PA), the PI controller, the DAC and the gain of the frequency modulation. Initial locking attempts showed that dynamic range of the phase error signal needed to be increased far above the single cycle limit used in the ADPLLs. This was done by simply extending the bit lengths in the registers of the phase accumulator (PA) in the DAC FPGA. Even though the phase lock reduces the error signal to values smaller than 2π when it is locked, it requires a larger range during the acquisition to avoid register overflows. After fixing this issue, reproducible
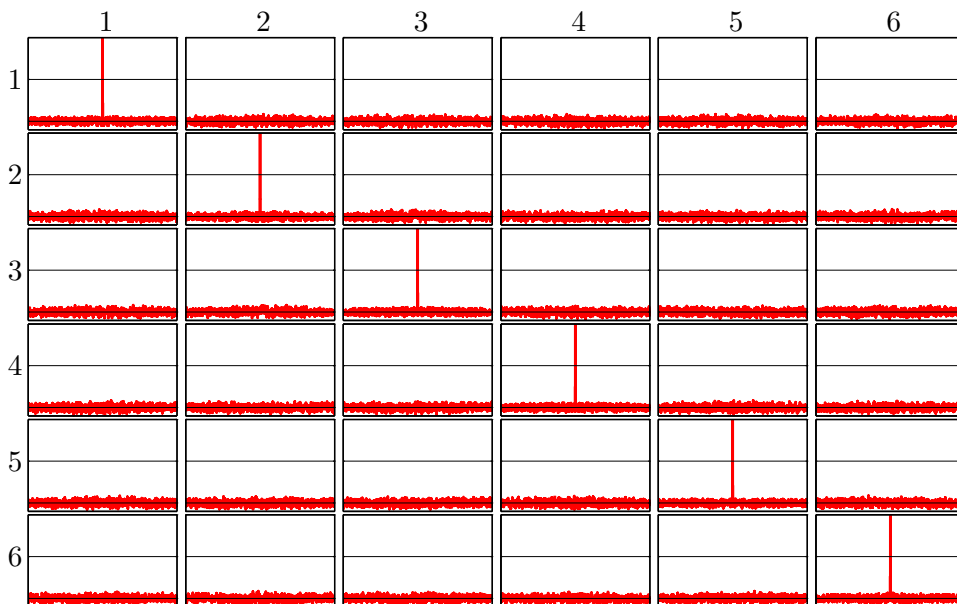
Figure 9.12.: All possible cross correlations between the six PRN codes generated by the EBB. The correlation is plotted between -0.1 and 1.0 and the delays are plotted from -511 to 512.

locks to various frequencies were achieved by first locking the ADPLLs and then activating the algorithms in the DAC FPGA a few seconds later.

An additional analogue adder, introduced into the loop (see Figure 9.13), was used to measure the open loop gain of the system. Figure 9.14 shows the measured open loop transfer function in comparison to an adjusted model. Additional delays of $\approx 6.1\,\mu s$ had to be introduced to explain the phase behaviour of the transfer function. A measurement of the signal generator reaction time showed that it contributed more than $2\,\mu s$ and an additional first-in-first-out memory, that is not part of the standard EBB interface structure, contributed $1.7\,\mu s$. This reduced the delays, caused the basic EBB infrastructure, to less than $2.4\,\mu s$. This value is suitable for the laser locks and it proved the feasibility of the implementation, without further adapting the interfaces.

The prompt activation of the laser lock showed an unexpected behaviour during some of the experiments. The four ADPLLs, that were used as frequency sensors, started to lock to frequencies a few Hz to kHz around the set point. Their combined frequency value was, however, close to the desired value, generating a quasi-locked state with a highly increased overall in-loop phase noise. This meta stable lock, apparently possible due to the use of more than one independent sensor, might be excluded in the future by slowly increasing the gain of the laser lock. This would reduce any start-up transients that might kick the
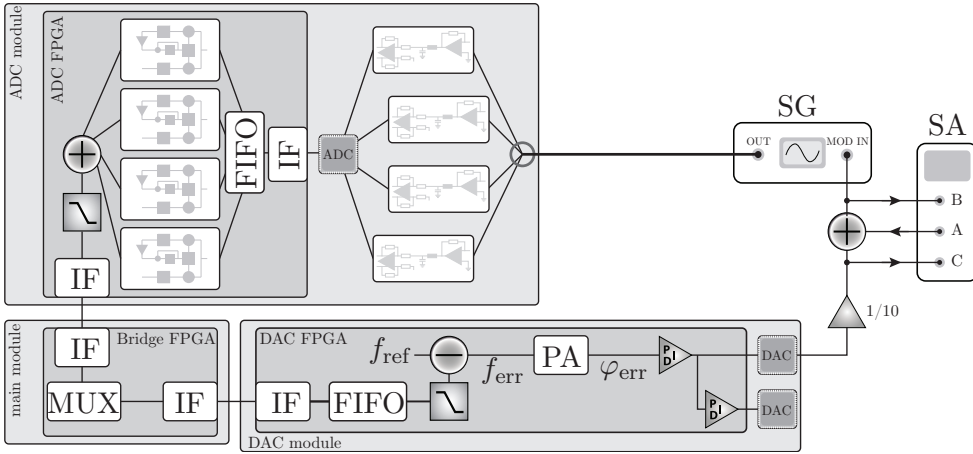
Figure 9.13.: Shown is a schematic overview of the test set-up used to investigate the frequency offset phase-lock capabilities of the EBB. The control loop includes a signal generator (SG), an ADC module, a main module, and a DAC module. Additionally an adder and a spectrum analyser (SA) are included to measure the closed-loop transfer function. The large signals from the DAC are also attenuated by an additional one to ten voltage divider to protect the SG modulation input.

ADPLLs away from the carrier. This has however not yet been investigated in detail and it remains a potential challenge for the use of the combined frequency values for the laser lock.

Finally, for three frequencies the locks were run for two hours to test the long term stability and to determine the in-loop performance. Figure 9.15 shows the phase spectra of the combined frequency values (the error signals) for locks to three different frequencies. The shown performance is below even the science phase measurement requirement, but since it is only an in-loop measurement is does not show the sensor noise of, for example, the analogue front-end. The test did not include a pilot tone correction, which already shows that the performance is not representative of the real phase tracking, since that would be spoiled by sampling jitter. However it was shown that the locks are stable and their absolute phase performance is not critical (since any phase deviations will be well measured). The addition of a pilot tone correction for this lock could be implemented in the future by adding correction terms to the phase error via simple CPU commands. This should be possible, since the sampling jitter noise is slow, but it would require to perform out of loop tests to validate that no additional noise is introduced into the error signal.
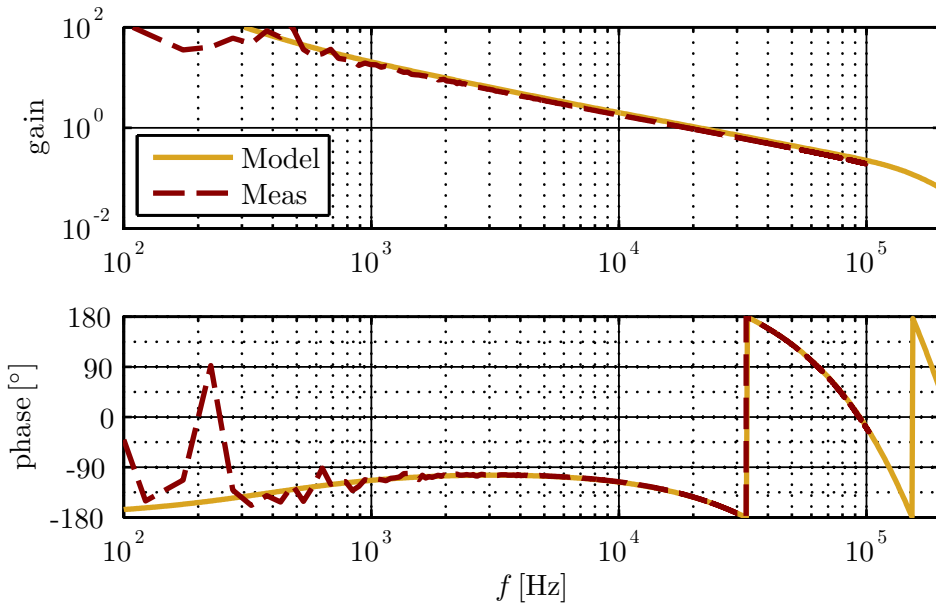
Figure 9.14.: Shown is a comparison of the modelled and measured transfer function of offset frequency phase lock of a signal generator to the EBB. Unity gain is a $\approx 20\,\mathrm{kHz}$ with a phase margin of more than $30\,°$.

9.5. Phase measurement performance

The most critical aspect of the EBB is the phase measurement performance. An extensive measurement campaign was done to reach the desired noise levels with realistic signals. During this campaign several modifications of the set-up were implemented to reach the full performance. In the following the results of these test are presented. Beforehand the source of the test signals, the digital signal simulator, and the generated signal properties are described.

9.5.1. Digital Signal Simulator

For testing the EBB it was decided to perform split measurements with LISA-like electronic signals. For this purpose the dedicated digital signal simulator (DSS), constructed by I. Bykov was used. The DSS generates the full signal spectrum in an FPGA and a microprocessor and converts it into analogue signals using digital-to-analogue converters sampling at 80 MHz. The DSP core used for the signal generation is an extension of the core originally written for the self-correcting DSS described in Section 8.2. For generating the additive noise and the laser frequency noise the VHDL implementations described in Chapter 6 were adapted to the DSS FPGA hardware. An overview of the DSS signal generation is shown in Figure 9.16, together with some example spectra of the
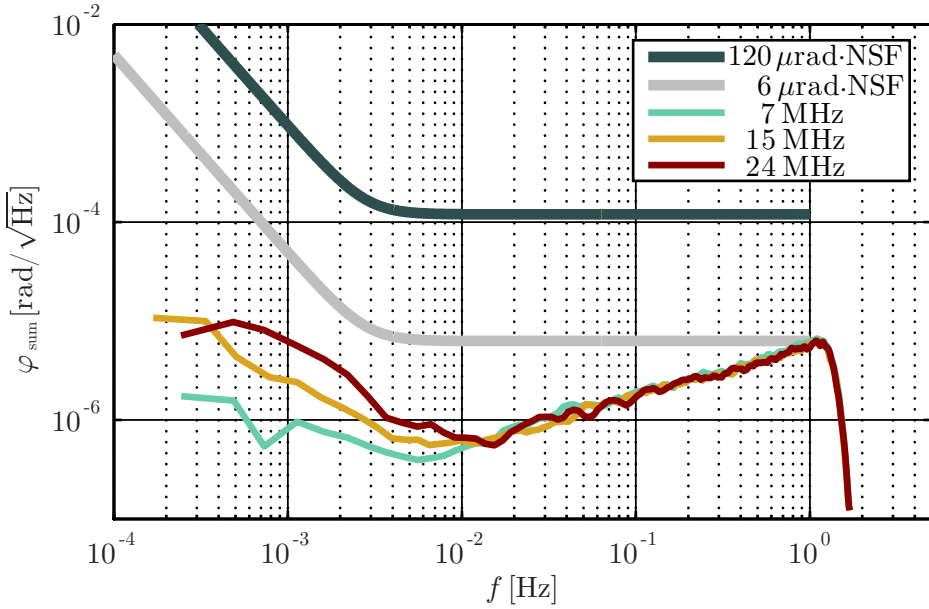
155

Figure 9.15.: In-loop phase performance of the offset frequency phase lock for three frequencies in comparison to the requirement given by ESA ($120\,\mu$rad) and the phase measurement requirement ($6\,\mu$rad). The noise floor shows an increase to higher frequencies that is common for all three frequencies. This might be a noise source in the actuation chain that is decreased to lower frequencies by the open loop gain. A noise increase to lower frequencies shows a strong dependency on the signal frequency and its cause is currently not understood.

DSS output that were taken in preparation of the measurement campaign.

All measurements using the DSS shown in this thesis include the following signal contributions.

- Laser frequency noise was chosen to be the same as used in the phase readout simulations (see Chapter 6). It corresponds to the effective noise at the master laser for an extreme weakly pre-stabilised laser. Therefore it exceeds levels expected in a real implementation, especially at low frequencies.

- Frequency drifts of approximately $4\,$Hz/s, to account for Doppler shifts change rates expected in LISA.

- Two clock tone side bands were placed $1\,$MHz away from the carrier with an amplitude of -13 dB relative to the carrier. This corresponds to roughly $10\,\%$ of optical power being used for the side bands.
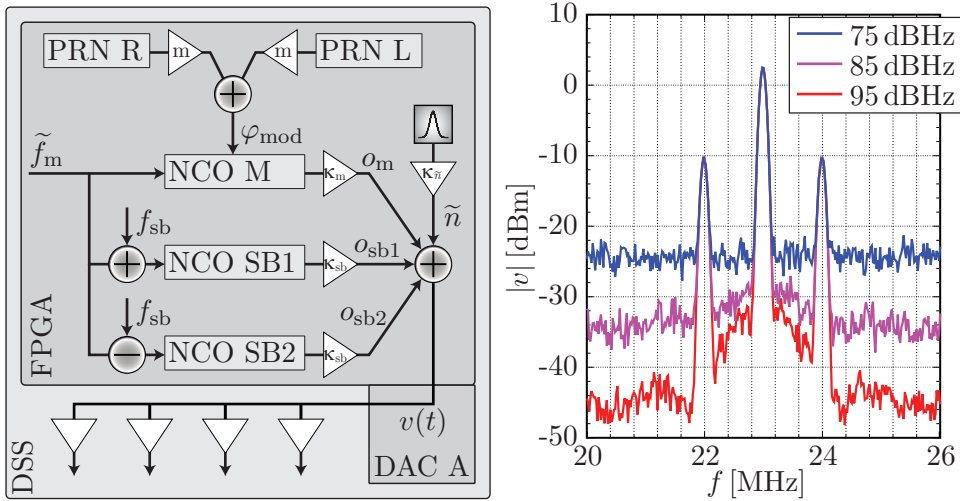
Figure 9.16.: The left side shows the working principle of one channel of the DSS. Inside the FPGA three NCOs are used to generate the three beat notes and two PRN generators add an additional phase modulation to the main beat to generate the ranging modulation. All signals are then scaled and added together with an additional white noise. This signal is fed to the DAC and the analogue output is buffered by four RF signal buffers, which distribute the signal to the EBB and simultaneously to a spectrum analyser or oscilloscope. The right side shows the output of the DSS for a 23 MHz signal with PRN modulation, side bands and three different levels of additive noise. The measurement was done by J. Kullmann and I. Bykov for the EBB test campaign.

- Two PRN modulations, as described in Section 9.4.1, were added, with an amplitude equivalent to about 1 % of the total signal power.

9.5.2. Initial performance in air

Before the performance of the full system was measured, an ADC stand alone module was tested by J. Kullmann. This measurements showed that a better performance could be achieved by using a different amplifier for driving the ADCs. The identified amplifier (LMH6551 [129]) was, therefore, also used for the full system measurements. The pre measurements also showed that the phase performance was strongly limited by temperature fluctuations. Therefore the anti-aliasing filter in the AFE was short circuited to reduce the number of components that produce temperature dependent phase fluctuations.

The initial measurement set-up used a pure sine wave, produced by a commercial signal generator (SG) (Tektronix AFG3102 [130]), to drive the eight
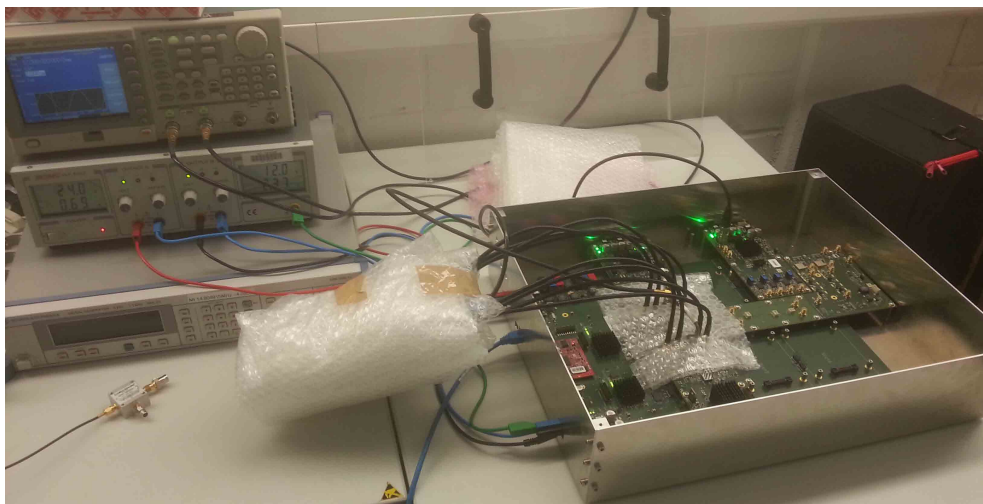
Figure 9.17.: Photographs of the initial EBB testing without thermal isolation of the system. The signal splitter, distributing the signals, is packed in three layer of bubble wrap foil to ensure a high thermal stability.

inputs of two ADC cards mounted into slots 3 and 4 of a main board. The set-up is shown in Figure 9.17. The signal was split by means of an $50\,\Omega$ impedance matched resistive 8-way power splitter, that was built using temperature stable Mini-Melf resistors [79]. The pilot tone used during the measurement campaign was also generated by a signal generator and it was split into the two pilot tone inputs by means of a commercial 2-way resistive power splitter. The signal generator was locked to the 80 MHz clock of the EBB with a 10 MHz reference tone, generated by one of the PRN outputs of a mounted DAC card. This ensured that the pilot tone was be stable with respect to the system clock.

An example of the achieved phase measurement performance is shown in the top of Figure 9.18. Low frequency excess noise spoils the performance below 0.1 Hz. The temperature fluctuations on the two ADC cards are shown in Figure 9.18 and they far exceed the temperature fluctuations that were required in the stand-alone tests ($< 0.1\mathrm{K}/\sqrt{\mathrm{Hz}}$) to achieve full performance. Even though full performance was not yet reached at this point, the measurements showed that there was probably no excess noise between the two ADC modules, an important result for the feasibility of the modular approach of the EBB. The results shown here and in the following are for frequencies around 25 MHz, the highest beat note frequency assumed during the project. The temperature induced phase noise scales with the frequency (as explained in Section 3.2.1), and therefore only the most temperature critical measurements are shown.
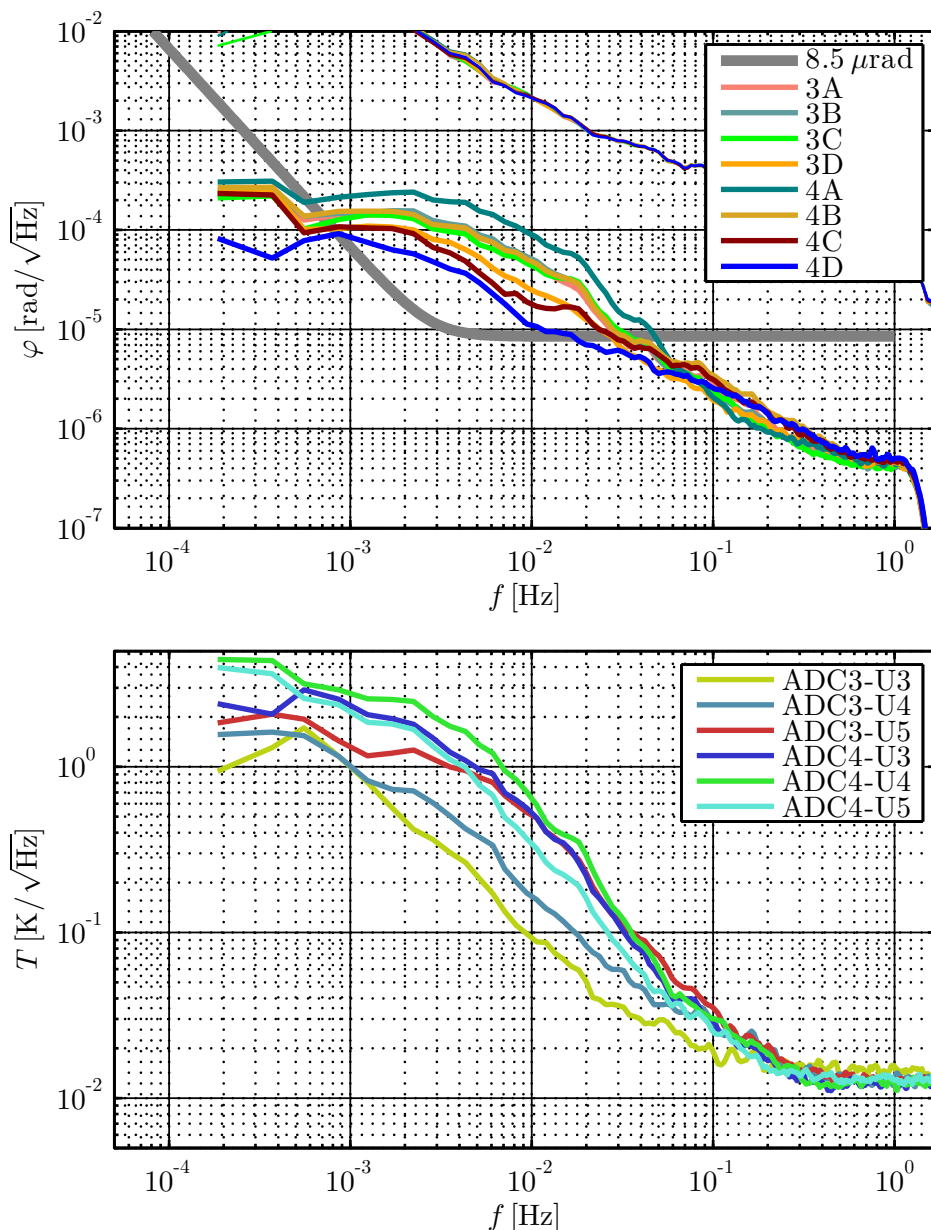
Figure 9.18.: Jitter corrected phase and temperature performance of tracking a 24.9 MHz signal from a SG. The EBB was not thermally isolated and the signal was distributed via an 8-way power splitter and measured using the full AFE. The temperature stabilities were measured with the sensors on the ADC modules.
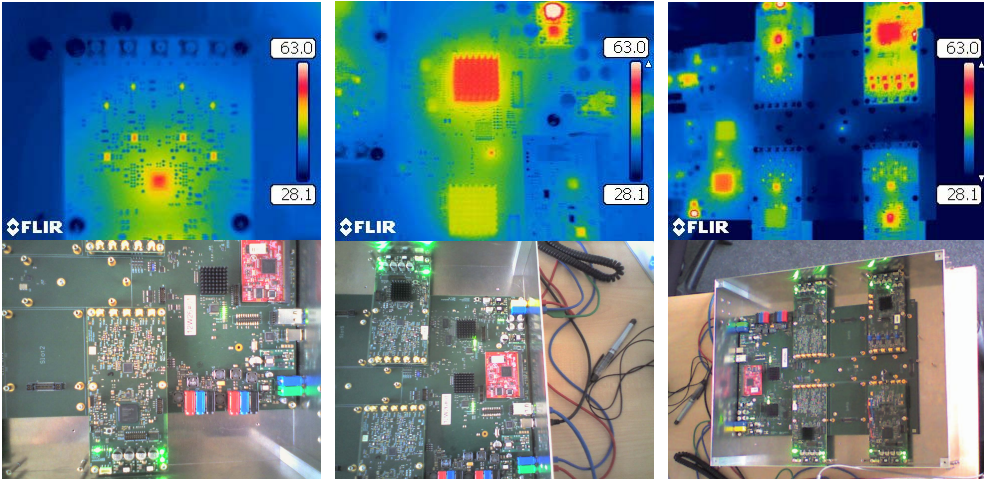
Figure 9.19.: Thermal images of parts of the EBB 20 minutes after start up. The left pictures show the analogue front end, the ADC and the amplifiers are clearly visible. A strong temperate gradient extends over the critical analogue front end. The picture in the middle shows the FFT and bridge FPGA, as well as the CPU (idle at the time) and the digital power supplies in the upper right corner. One can see that the power supply is very hot, since it provides most of the current consumed by the EBB. Even though both FPGAs have a passive cooler it is easy to identify the FFT FPGA, which, due to its larger computational resources, produces significantly more heat. The right side shows the full EBB with three ADC modules and a DAC module. The DAC module produces clearly most of the heat and potentially also large parts of the temperature fluctuations. Therefore all measurements during the campaing were done with a DAC module present, to ensure a realistic environment.

9.5.3. Active temperature stabilisation

After some, only partially successful tests, using means of passively isolating the analogue front-ends, it became clear that an active temperature stabilisation was necessary. Isolating the full system, as done in the stand-alone test, was considered not feasible, due to the large heat production of the EBB. The temperature of some parts of the EBB already exceeded $60\,^\circ$C after 20 minutes of operation in air, as shown in Figure 9.19. Using a standard cooling technique like fans was not considered either, since even though it would have cooled the system, it would have coupled its temperature stability to the surrounding environment, which, even for most laboratories, is not stable enough.

Hence an active temperature stabilisation of the full EBB was implemented. An aluminium casing was manufactured to create an enclosed thermal environ-
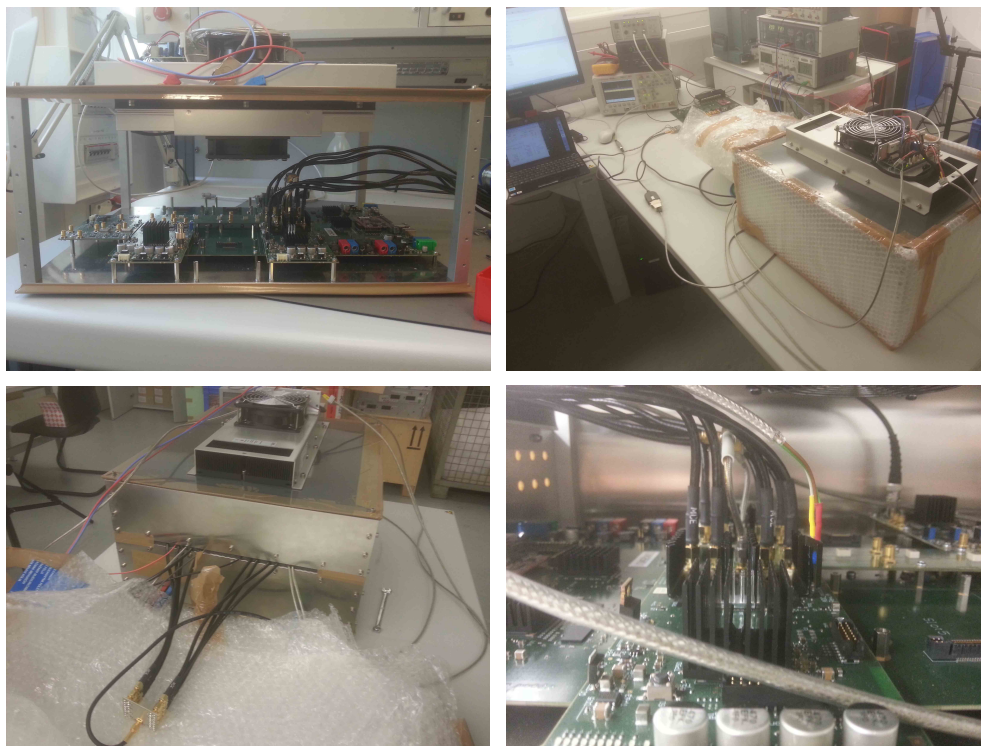
Figure 9.20.: Photographs of the EBB active thermal control during construction and testing. The top left picture shows the thermoelectric assembly mounted into the top plate of the enclosure. The bottom and top plate are double walled, each with a layer of aluminium on the inside, plastic on the outside and 5 mm of air in between. The top right picture shows the experimental set-up during the measurement campaign. The EBB is additionally isolated by bubble wrap foil, which was later removed and did not effect the thermal stability. The bottom left picture shows the 8-way resistive signal splitter, as well as the air tight cable port of the EBB enclosure. The bottom right picture shows one of the ADC modules inside the enclosure with passive heat exchangers mounted on the FPGA, the ADC and the AFE.

ment. The temperature inside this enclosure was then stabilised with a thermoelectric assembly (TA) (AA-200-24-22 from Laird Technologies [131]) mounted into the cover plate. A photograph of the case during assembly is shown in the upper left corner of Figure 9.20. The TA consists of two heat exchangers, one on the inside and one on the outside, which are connected via Peltier elements. Additional fans on both sides ensure a strong coupling of the heat exchanger temperatures to the respective air temperatures. For the inside of the EBB enclosure this also ensure a constant airflow over the EBB circuits. The

temperature gradient between both heat exchangers is actively controlled by a temperature controller (PR-59 from Laird Technologies [132]), which modulates the amplitude and the polarity of the current through the Peltier elements. The temperature is sensed via a negative temperature coefficient thermistor (NTC) that is read out by the PR-59. The NTC was placed at the heat exchanger on the inside of the enclosure, directly into the air stream, to ensure a strong coupling and a short reaction time. Other positions of the NTC have also been investigated (close to the analogue front-end, in the air stream away from the heat exchanger) but they were found to be less optimal, mainly due to an increased delay in the temperature sensing. The temperature control of the PR-59 was optimised by using standard PID controller design techniques to reach the maximum stable bandwidth. A maximum of 9 A of supply current was used and both fans were set to run at full constant speed. The enclosure of the EBB has inside dimensions of 50 cm x 40 cm x 20 cm, leading to an overall volume of 40 l. Assuming a standard airflow for a 12 cm fan of 150 l/h the air inside the enclosure is circulated roughly once per second. However this simple calculation does not include the direction of the air flow and the turbulences inside the box.

Pilot tone input

Using the active temperature control, first measurements of the performance were done by feeding a combined signal and pilot tone into both ADC cards via the pilot tone input. Thereby the temperature induced phase noise of only the differential amplifiers, and the sampling jitter could be measured. The phase noise performance of this measurement for 24.9 MHz is shown in the top of Figure 9.21 and the corresponding temperature performance is shown in the bottom. It is clearly visible, that the performance is below the requirements and that the temperature fluctuations on the ADC card have been reduced drastically. One should note that this measurement also included means of passively directing the airflow inside the EBB enclosure over the analogue front end. This was done by building a chimney like structure inside the box out of card board that ensured that the air, sucked into the fan from below, would be in good contact to the AFE. The heat exchange between the air and the circuits was further enhanced by placing heat exchangers on top of the AFE that are in contact with the amplifiers via thermal paste. The heat exchangers can be seen in the lower right photograph in Figure 9.20.

Full front-end

After confirming the performance of the differential amplifiers, first measurements were done using the same means of temperature stabilisation, the full AFE, and signals from the DSS. A signal without additional amplitude noise, but including all other components, was first amplified by a simple operational
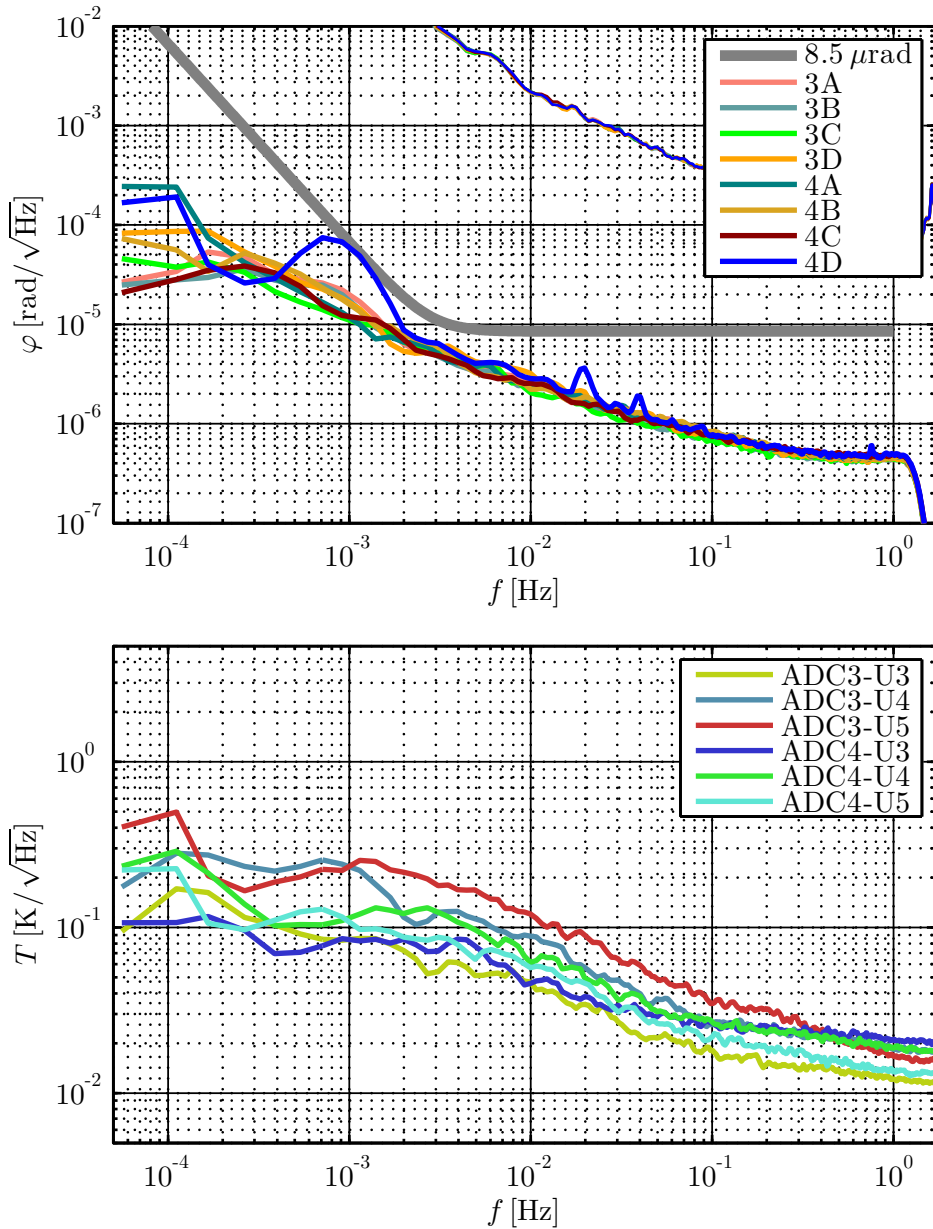
Figure 9.21.: Jitter corrected phase and temperature performance of tracking a 24.9 MHz signal from a SG. The signal was distributed together with the pilot tone via a 2-way resistive splitter and measured using only the differential amplifiers. The EBB was actively temperature controlled and the air flow was passively directed over the AFE.

Figure 9.22.: Jitter corrected phase and temperature performance of tracking a 25 MHz signal from the DSS without additive noise. The signal was distributed via an 8-way resistive splitter. The EBB was actively temperature controlled and the air flow was passively directed over the AFE.

amplifier circuit, and then fed into all eight channels by means of an 8-way resistive splitter. The pilot tone was again generated by a signal generator and
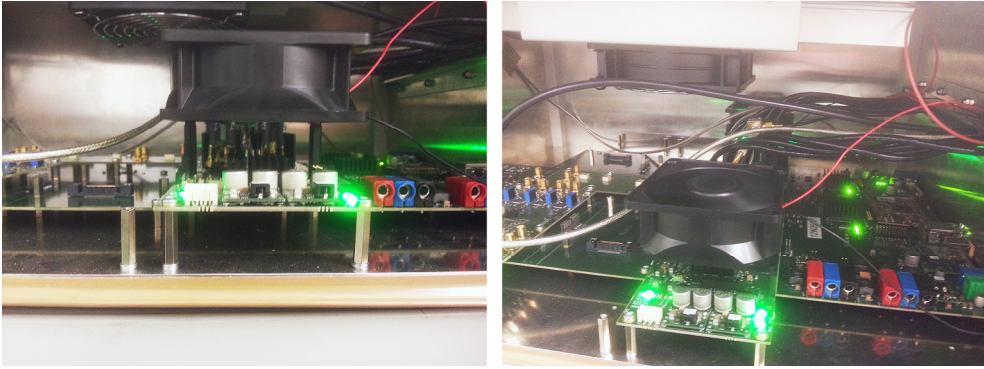
Figure 9.23.: Shown are photographs of additional fans placed on top the ADC modules to increase the thermal stability of the AFE, by increasing the thermal interaction between the AFE and the air flow. The fans have a diameter of 8 cm and they were mounted using two distance bolts screwed into the central mounting holes of the modules.

distributed with a 2-way resistive splitter. The phase and temperature measurement results are shown in Figure 9.22. Even though the measured thermal stability was the same, an excess noise in the phase is now present around 10 mHz. This indicates that the gain amplifier stage is much more temperature sensitive than the differential amplifier stage. Additionally, one can see an increased white noise floor for some of the channels (channel D on the module in slot 4 is the most prominent example). This noise was not expected and is discussed further in the next section.

9.5.4. Performance with full signals

To achieve the desired performance, a further reduction of the temperature fluctuations in the AFE was required. Since further passive redirection of the air flow was not feasible an alternative was implemented. Fans were mounted onto the ADC modules, that provided a constant air flow over the AFE, as shown in Figure 9.23. The resulting phase measurement performance (shown in Figure 9.24) did finally reach the requirement for all but one channel. The measured temperature fluctuations did not improve significantly, showing that the sensors and their positioning are not suitable to determine the performance directly. The physical distance between the sensors and the gain stages (the most temperature sensitive components) hinders a more detailed understanding of the temperature to phase coupling, as already indicated by the strong temperature gradients shown in Figure 9.19.

Measurement 9.24 also includes additive noise corresponding to the low SNR case. As desired the additive noise is almost completely subtracted. But an
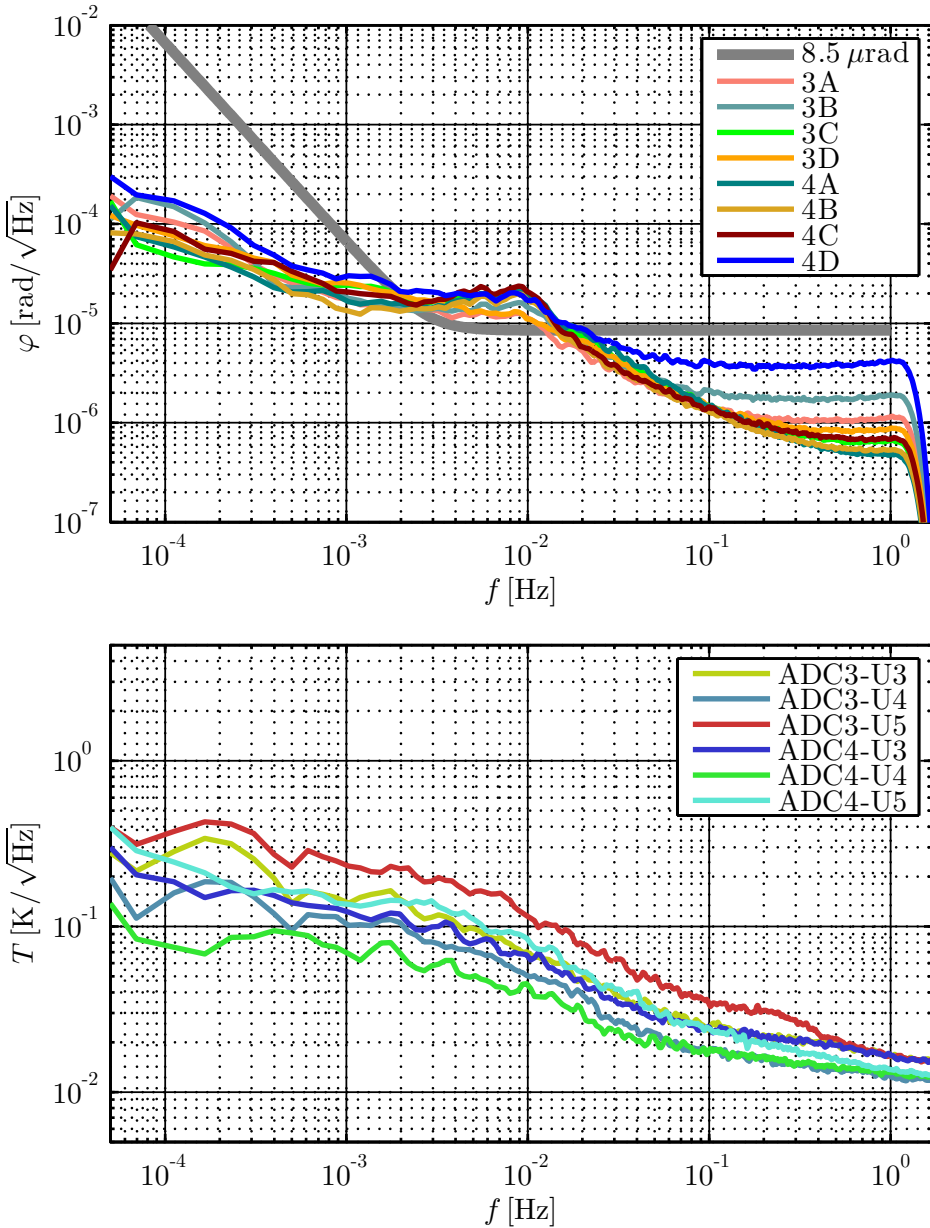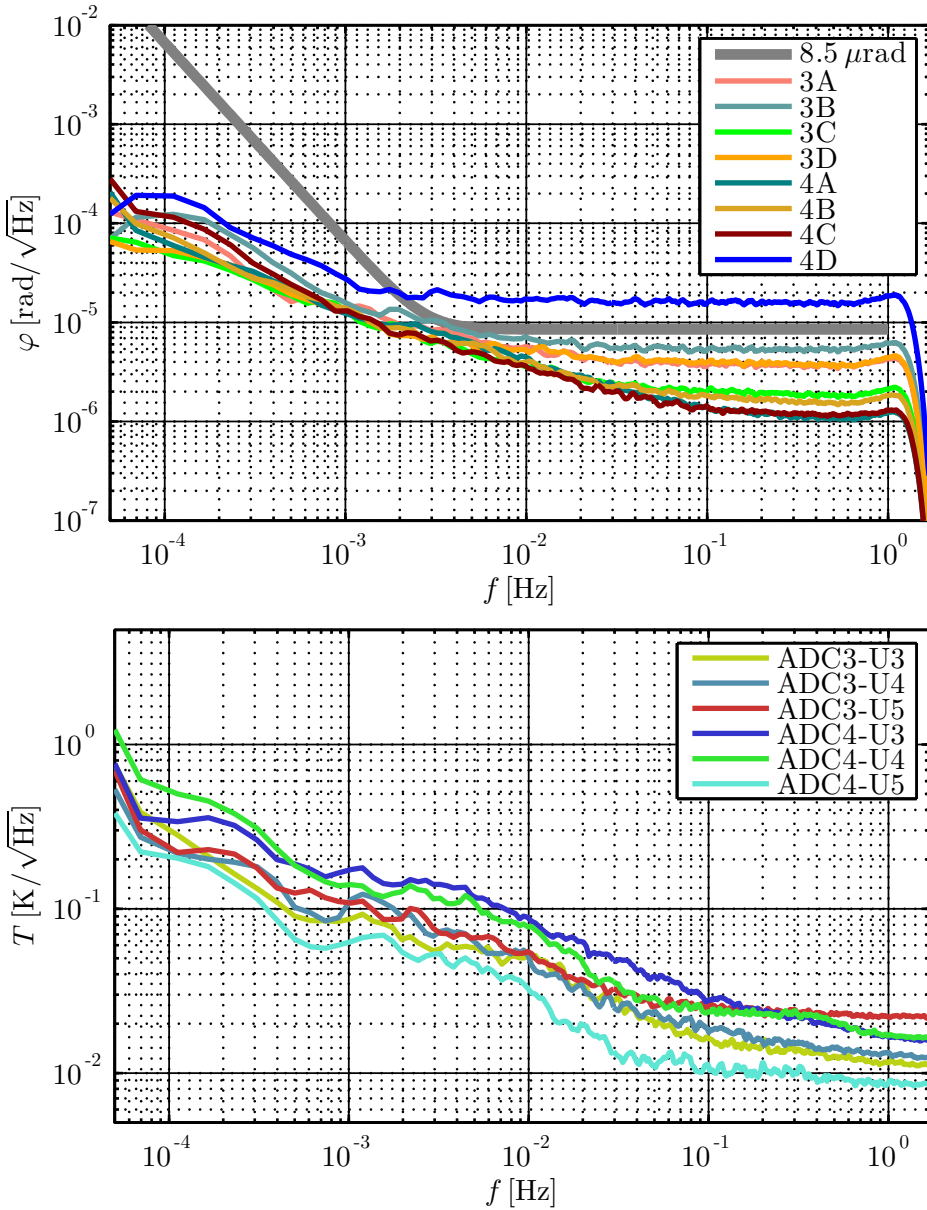
Figure 9.24.: Jitter corrected phase and temperature performance of tracking a 25 MHz signal from the DSS with additive noise of 75 dBHz. The signal was distributed via an 8-way resistive splitter and measured using the full AFE. The EBB was temperature controlled and the air flow was actively directed over the AFE with additional fans.

additional white noise floor (already observed in Measurement 9.22) is present. This noise was not visible in the measurements using a signal generator as signal source (see Figure 9.18 and 9.21) and was therefore understood to be caused by some interaction of the signal source and the EBB. Impedance was one parameter that could nominally change by using another signal source (though both the DSS and the SG are matched to $50\,\Omega$). Therefore the impedance matching of the 8-way resistive splitter was investigated. It was found that the resistor values ($38.88\,\Omega$ in a balanced 8-way splitter) were only slightly off by around $0.3\,\%$. In lack of another explanation at the time, the matching was improved to better than $0.1\,\%$ by simply exchanging the resistors. The impedance at MHz frequencies was expected to be dominated by parasitic capacitances and other effects. Measurement of the RF impedance showed also no significant difference between different splitter channels, probably because the measurement precision was to low. The subsequent measurements using the better matched power splitter showed indeed a reduction of the white noise floor, for some channels even by factors of two. The measurement in Figure 9.23 already includes this optimised impedance. No further analysis of this effect was done at the time, but the excess noise, especially in channel 4D, was present during the whole measurement campaign. This includes measurements at various frequencies and with low and high SNR (see Appendix D). The white noise was found to scale with frequency, which would fit to an impedance matching problem, and with the level of additive white noise, which does not fit to an impedance matching cause. Since the effect was not visible in the measurements using a locked signal generator (Figure 9.18 and Figure 9.21), they have to be caused either by an unknown interaction between the EBB and the signal source, or by the signal dynamics. This effect might also be the same one that was observed with the 16 channel phasemeter for the LISA optical bench, described in Section 8.5. Further measurement using different cable lengths, signal dynamics, and impedance matchings will allow us to investigate this noise coupling in the future. Even though this excess noise was present, the campaign was able to show the desired performance levels for most channels for various frequencies and for the low and high SNR signals. Figure 9.24 shows the measured performance for 25 MHz at high SNR in comparison to the input signals and to the actual phase measurement limit due to the additive white noise. The campaign was thereby able to demonstrate a performance of more than 12 orders of magnitude at 1 mHz with analogue signals.

9.5.5. Measurement with clock module

After the campaign the phase performance was tested using the main board together with a first generation clock module. The module was driven with a 2.4 GHz signal, which was converted to the system clock and a pilot tone The pilot tone was further distributed via the dedicated pilot tone distribution and
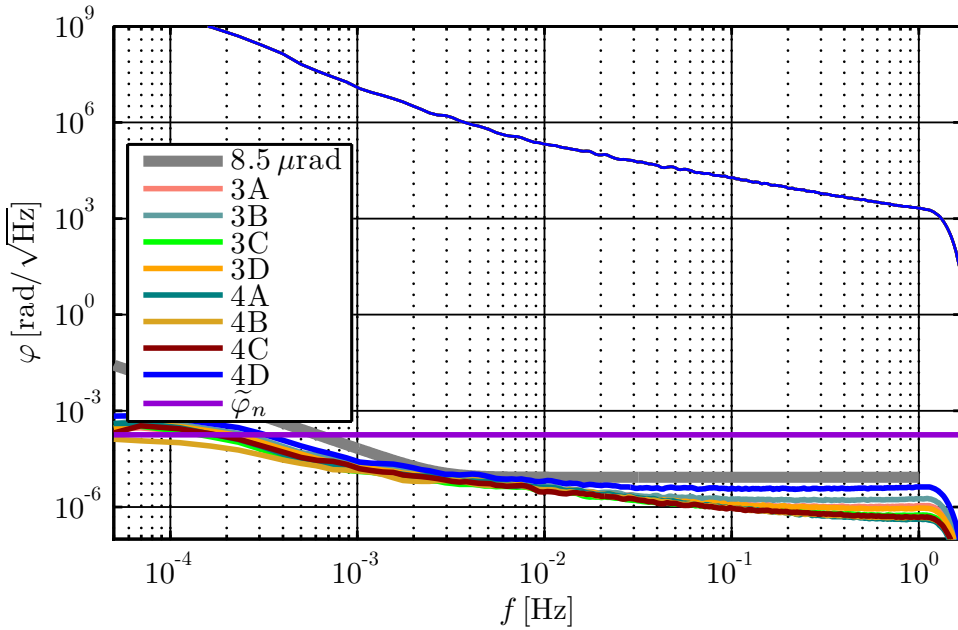
Figure 9.25.: Jitter corrected phase performance of tracking a 25 MHz signal from the DSS with additive noise of 95 dBHz. The signal was distributed via an 8-way resistive splitter and measured using the full AFE. The EBB was temperature controlled and the air flow was actively directed over the AFE with additional fans. For comparison the initial signal dominated by simulated laser frequency noise (thin blue line) is also shown.

fed from the clock module to the ADC cards via short, 'u'-shaped cables. Additionally, two high pass filters were added on both modules to reduce crosstalk. The filters allowed the pilot tone to pass and each one of them attenuated signals below 25 MHz, travelling in any direction, by more than 30 dB. Thereby these filters generated a suppression of crosstalk via the pilot tone distribution of more than 120 dB. Only one measurement with this set-up was performed so far, using the most critical signal parameters (25 MHz and low SNR) and an active temperature stabilisation with additional ADC fans. The performance was found to be equivalent to the one shown in Figure 9.24, proving that the pilot tone distribution, as well as the clock generation perform as desired.

9.5.6. Analogue front-end discussion

Measurements using only one input signal into an ADC module showed that all four ADPLLs locked to the signal. This was caused by a large crosstalk on the module. The measured amplitude was roughly 1/10 of the one measured in the actual channel, corresponding to an effective crosstalk suppression of only
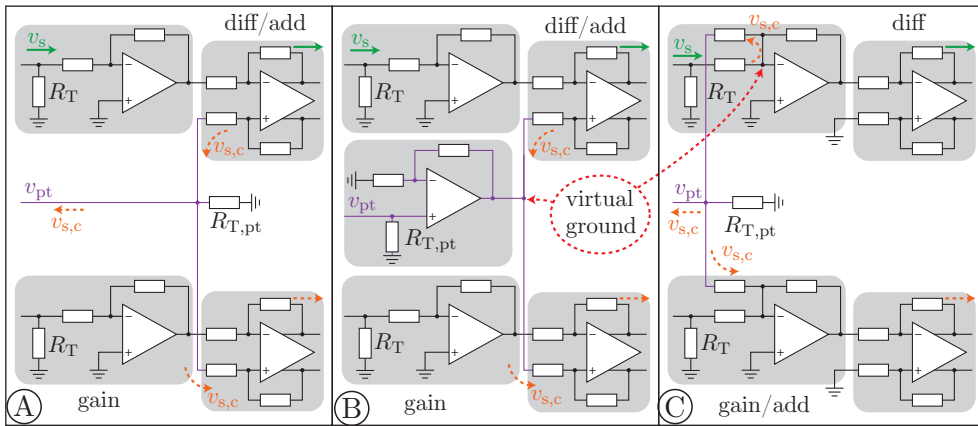
Figure 9.26.: Shown are three possible alternative schemes for the analogue front end of the EBB. For simplicity only two channels are shown instead of four. (A) Initial design implemented on the EBB using the differential amplifiers to add the pilot tone. (B) A design that uses an additional buffer in the pilot tone input to produce a virtual ground for the local distribution. (C) A third option that uses the gain stages for adding the pilot tone and thereby utilises the virtual ground of the single-ended amplifiers to reduce crosstalk.

20 dB.

This large value was found to be inherent to the used design of the analogue front-end. Figure 9.26 shows the initial design (A) together with two possible alternatives. The initial design distributes the pilot tone by a simple resistive network and the tone is added by the differential amplifier. Thereby the science signals are not suppressed at the pilot tone input of the amplifier and they are also partly added to the other differential amplifier inputs. One alternative (B) is to add an additional pilot tone buffer on each ADC AFE to generate a virtual ground, this will suppress the signals looped trough the differential amplifier and reduce the crosstalk between channels by more than 60 dB in the measurement band. The second alternative (C) is to use the gain stages for adding the pilot tone. Here simpler adder circuits can be implemented. They have a virtual ground that also suppresses the crosstalk much more efficiently than variant A, also reaching suppressions of around 60 dB between the channels on an ADC module.

Circuit simulations using LTSpice [133] have shown that the crosstalk suppression between channels on the same board is in general larger by 10 dB to 20 dB for variant B in contrast to C. Variant B also includes a strong suppression of crosstalk back to the clock module, with values around 120 dB, making the additional high pass filters, that are currently used, unnecessary. The draw back

of option B, however, is that an additional amplifier is necessary, which among other things increases the power consumption. The pilot tone buffer might also increase the sensitivity to temperature fluctuations, a problem that can in general be reduced by a more detailed thermal design of the overall system.

Chapter 10

Outlook

This chapter gives an outlook of the investigations necessary for inter-satellite interferometry that go beyond the work done in this thesis. Open issues are addressed and ideas and concepts for future experiments and developments are described.

10.1. Modelling of analogue components

The presented investigations of phase performance and crosstalk of analogue components were mainly based on empirical investigations. This approach was sufficient to construct phase measurement systems with the desired performance levels for laboratory use. But future implementations might benefit from a more systematic approach, which would allow one to better optimise the phasemeter on a system level. For ensuring crosstalk suppressions around 120 dB, simulations of electromagnetic interference effects might become necessary. A potential way to improve the models for temperature induced phase noise is to characterise isolated circuit parts using the now available phasemeters. By feeding the determined coupling factors back into the models, one could try to predict the stabilities of more complex circuits and thereby create a library of components with well-known behaviour. Such a tool-set could then, in the next step, be used together with thermal management models to perform trade off studies, leading to more sophisticated system designs. The development of photo receivers would also benefit from and contribute to such models, since they are an integral part of the analogue phase measurement chain.

10.2. Testing the full LISA metrology chain

The availability of phase measurement systems that achieve μrad performance levels in simple split measurements is another step towards more integrated tests of the LISA metrology. Performing an optical three signal test with such a phasemeter, as described in Section 8.7, will allow us to investigate the non-linearity of the full phase measurement chain and it can be used to measure the influence of crosstalk under realistic conditions. Similar experiments have already been done to demonstrate parts of the metrology concepts [95, 58].

One promising test of the entire LISA metrology can be done by performing an optical three signal test with three independent phasemeter systems, as shown in Figure 10.1. Each phasemeter adds a GHz clock side band and a ranging / communication signal to one of the laser beams sent onto the interferometer. The final measurement performance of this set-up can only be achieved by correcting for the three independent clock noises with the side band phase information. Thereby the full frequency distribution and phase measurement is tested in a single experiment. The ranging signal can be used to determine the constant time delays in the system. Such a test could be used to demonstrate the dynamic range not only of the phase measurement, but of all critical building blocks of TDI (clock synchronisation, delay measurements, time interpolation with ns precision and linear signal combinations). TDI itself however could not directly be tested with the shown set-up, since no optical delays between the interferometer arms are included. The readout of the experiment could also be performed by reading only data from one phasemeter, that collects the readout data from the other two using the PRN modulation based communication. If such an experiment is successfully implemented and performs well, it can also be used as a test bed for components. A laser candidate, for example, could simply replace one or all of the previously used lasers and any changes in the overall system performance could directly be observed. To perform such tests the phasemeter systems would require to operate almost completely autonomous. Therefore the algorithm for acquisition, error handling and other unexpected events have to be further developed as well.

10.3. Future investigations

Some others aspects of inter-satellite interferometry that might be covered by future investigations are summarised here.

Realistic signal phase dynamics The signal dynamics play an important role for various design aspects of the phase measurement chain. Implementations aimed for specific missions will require a detailed upper limit of this to design the readout system accordingly. The effective laser frequency noise assumed
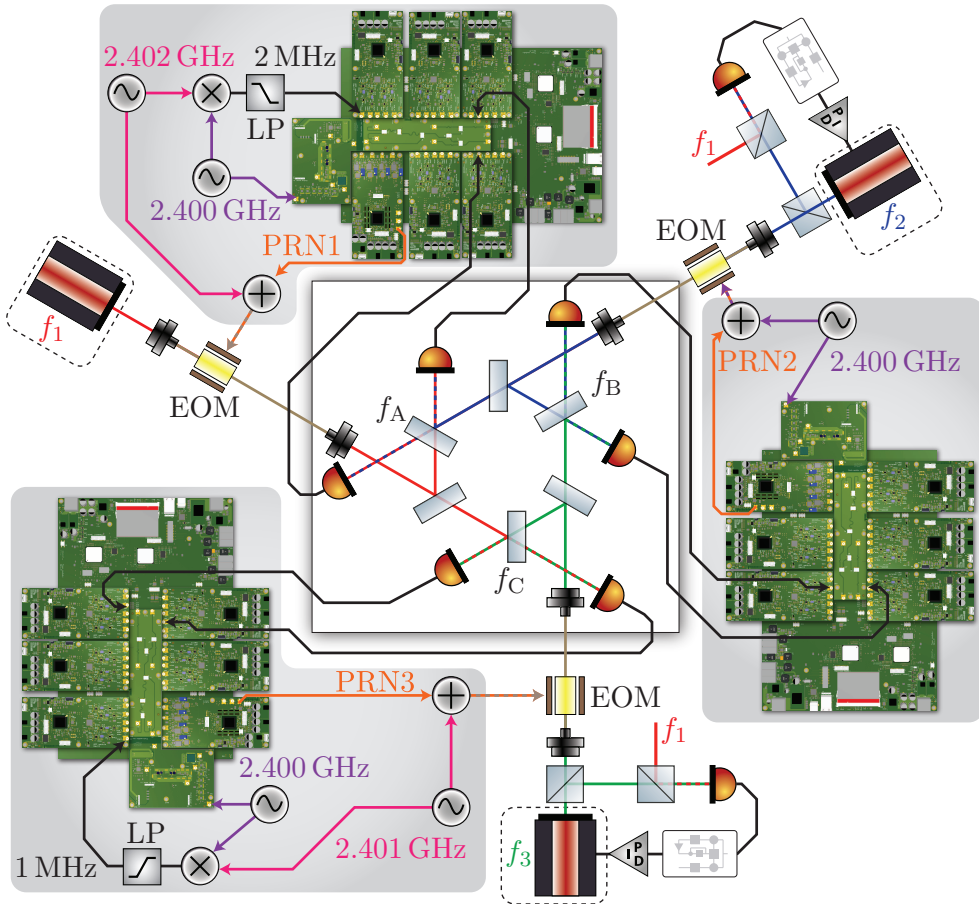
Figure 10.1.: Shown is a possible set-up for a three signal test using three independent EBBs. Each EBB would require a 2.4 GHz clock signal. Two would additionally require GHz signals at, for example, 2.401 GHz and 2.402 GHz to generate side band beat notes in all three interferometers. The deviations between the 2.4 GHz signals and the additional GHz tones would be determined by mixing both signals and by tracking the resulting beat note with an additional channel on the EBB. Two laser frequency locks could be used to tune the effective laser frequency noise by either increasing the loop bandwidth or by introducing additional noise.

for LISA during this thesis was, for example, very high, especially in the lower frequency range ($f < 100\,\text{Hz}$). Since it is likely that LISA will use a pre-stabilisation with a much better stability [55, 56], it would make sense to use the corresponding noise shape for the detailed system design. This would reduce the dynamic range and thereby simplify the ADPLL design, it would lower the requirements for the decimation filters and it might also reduce unwanted effects in the analogue front-end.

Amplitude dynamics The experimental investigations in this thesis did not yet include realistic signal amplitudes and their fluctuations. Future experiments will have to be done in the presence of such effects to determine the influence on the phasemeter. Strong amplitude changes in the presence of large additive noise might, for example, require to implement an automatic gain control in the ADPLL to keep its bandwidth in the optimal range.

Signal dynamics in future geodesy missions Currently investigated concepts for future geodesy missions aim to use orbits with increased relative velocities and therefore higher Doppler shifts [134]. These strong frequency fluctuations, which happen close to the measurement band, will potentially cause additional non-linearities in the phase measurement. Therefore it is worthwhile to investigate the potential limits of phase measurement systems regarding these high Doppler-shifts and strong frequency drifts.

Architecture for local interferometry The LISA phasemeter uses MHz heterodyne interferometry for the inter-satellite, as well as for the auxiliary interferometry. This means the same readout architecture has to be implemented for all interferometers, even though its basic design (low crosstalk, pilot tone correction, frequency tracking) is mainly driven by the inter-satellite interferometry. Especially the crosstalk requirements of $120\,\text{dB}$ are so high because the phasemeter has to read out independent interferometers using the same RF beat note frequency. Future investigations could therefore also include a trade off study of using different interferometry techniques for the auxiliary readouts. Some of these possible alternatives have been investigated in the second part of this thesis.

Part II.

Alternative techniques for intra-satellite interferometry

Chapter 11

Real-time readout for deep phase modulation interferometry

This chapter describes the implementation of a real-time phase readout system for a new interferometer scheme. The scheme uses strong phase modulations, in one arm of a homodyne-style interferometer, to generate a comb of heterodyne beat notes. The signals are then processed, using complex digital signal processing algorithms, to determine the interferometric phase. The scheme requires less complexity in the optical hardware than classic heterodyne interferometers, and it is, therefore, a potential alternative for future implementations. This readout, called deep phase modulation (DPM) interferometry, was first proposed and successfully tested at the Albert Einstein Institute [135, 136, 137]. The work presented here was carried out together with T. S. Schwarze, whose Master's Thesis [138] was supervised as part of this thesis.

The first section of this chapter describes the working principle of DPM interferometry. The signal generation, as well as the readout scheme are presented. The second section gives an overview about the implementation of the real-time phase readout, using a dedicated DPM phasemeter. The hardware and digital signal processing (DSP) are discussed together with advantages of the real-time readout system. The achieved phase measurement performance using the DPM phasemeter in digital, analogue and first optical tests is presented in the third section. Finally, future developments and possible alternatives to the here investigated algorithms are discussed.

11.1. Deep phase modulation

Deep phase modulation uses one laser with a single frequency that is split into two beams, as shown in Figure 11.1. One of the beams is fed trough a phase modulator, which introduces a sinusoidal phase modulation $m(t)$, with a modulation depth m, a frequency ω_m, and a phase of the modulation signal ψ. The resulting photo current i_{DPM} can be determined by rewriting Equation 2.8, including the phase modulation and the interferometric phase φ, the signal of interest.

$$
\begin{aligned}
i_{\mathrm{DPM}}(t) &= i_{\mathrm{DC}} \left(1 + c \cos\left[m(t) + \varphi(t)\right]\right), \\
i_{\mathrm{DPM}}(t) &= i_{\mathrm{DC}} \left(1 + c \cos\left[m \cdot \cos\left(\omega_m t + \psi\right) + \varphi(t)\right]\right).
\end{aligned}
\tag{11.1}
$$

As indicated by its name, DPM utilises modulation depths that are greater than π. Thereby, the DC component of the signal is largely reduced and it is spread to various harmonics of the modulation frequency. Mathematically this is understood by rewriting Equation 11.1 using the Bessel functions J_n of the first kind,

$$
\begin{aligned}
i_{\mathrm{DPM}}(t) = &\overbrace{i_{\mathrm{DC}} \left(1 + c \cos\left[\varphi\right] J_0(m)\right)}^{i'_{\mathrm{DC}}} \\
&+ \underbrace{i_{\mathrm{DC}} \left(c \sum_{n=1}^{\infty} 2 \cos\left[\varphi + \frac{\pi}{2}n\right] J_n(m) \cos\left[n(\omega_m t + \psi)\right]\right)}_{i_{\mathrm{AC}}(t)}.
\end{aligned}
\tag{11.2}
$$

In the following only the AC components of the signal are relevant. If one analyses the signal in Fourier space, it becomes clear that it can be fully described by the amplitudes i_n of the harmonics $n\omega_m$, and the modulation phase ψ.

$$
\begin{aligned}
i_{\mathrm{AC}}(t) &= \sum_{n=1}^{\infty} i_n \cos\left[n(\omega_m t + \psi)\right], \\
i_n &= \underbrace{2 i_{\mathrm{DC}}\, c}_{k} \cdot \cos\left[\varphi + \frac{\pi}{2}n\right] J_n(m).
\end{aligned}
\tag{11.3}
$$

Figure 11.1 shows the time series and the corresponding spectra for an example case.

The readout scheme makes use of the fact that all relevant information is contained in the harmonics, by performing IQ demodulations, also called single-bin Fourier transformations, at the corresponding frequencies (see Section 2.3). Thereby, one is able to recover a complex amplitude $\tilde{i}_n$ for all demodulated harmonics, consisting of the corresponding amplitude i_n and phase $n\psi$.

In the next step a sophisticated algorithm is used to determine the four unknown parameters (φ, m, ψ, k). This is done by minimising the sum of squares
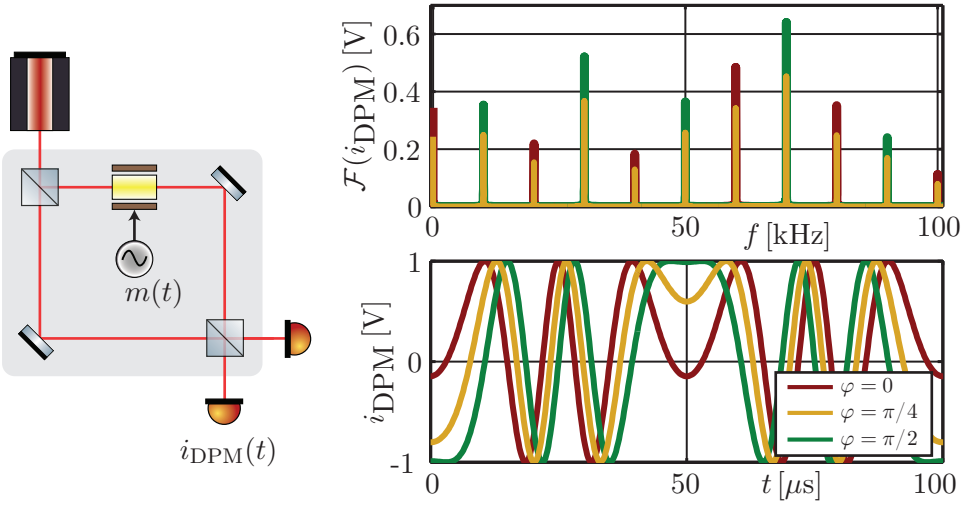
Figure 11.1.: The left side shows the basic interferometer set-up for DPM. A strong, sinusoidal phase modulation is introduced in one interferometer arm to achieve the desired spread of the signal over various frequencies. The right side shown the Fourier spectrum (top) and the time series (bottom) for three example cases with $\omega_m = 10\,\text{kHz}$, $m = 8$ and $\psi = 0$. For $\varphi = 0$ all power is concentrated at DC and in the even harmonics of ω_m. For $\varphi = \pi/2$ only the uneven harmonics contain the signal power, as predicted by Equation 11.3.

(χ^2) between complex amplitudes computed from an analytical model (i_n^c), and the measured ones.

$$\chi^2 = \sum_{n=1}^{N} \left(i_n^c - \tilde{i}_n \right)^2 . \tag{11.4}$$

The critical parameters for this approach are the number of bins N, used to perform the fit, and the average modulation depth, which in general determines how far the signal is spread to higher frequencies. The ratio between the number of bins and the modulation index determines how much power of the signal is actually used for the fit. A high modulation index and a low N will cause a significant loss in the usable signal amplitude. The values of m and N determine how far the equation system (see Equation 11.4) is over-determined. In general higher numbers lead to better fitting results, but physical constraints (processing capabilities and achievable modulations) also set an upper limit.

The previous experiments performed at the AEI used a data acquisition system to capture time series data [135]. This data was divided into segments, which were processed using a Discrete Fourier Transform to determine the complex amplitudes. A Levenberg-Marquardt fit algorithm, implemented by G. Heinzel, was then applied to perform the minimisation for each time segment.
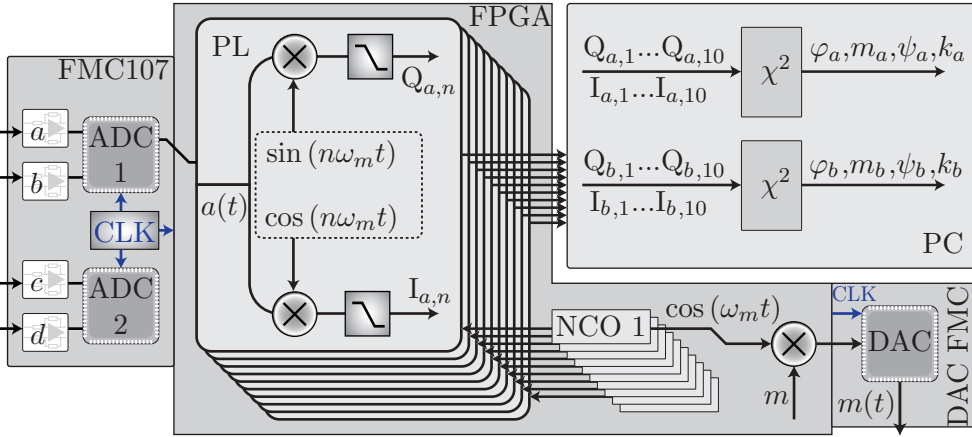
179

Figure 11.2.: Shown is a sketch of the phasemeter for deep phase modulation. Ten numerically controlled oscillators, operating at frequencies $n\omega_m$ (with $n = 1..10$), are implemented to generate local tones at each harmonic. The fundamental tone is also scaled with a modulation depth m and then converted into the analogue modulation signal $m(t)$ (see Equation 11.1). The I and Q values for each frequency component of the input signal are determined in dedicated demodulation pipelines (PLs). Here the digitised signal is first mixed with sine and cosine from the respective NCO and then decimated to an intermediate data rate. The resulting values are transferred to a PC. Here, the fit algorithm is implemented in a c-program, that performs the χ^2 minimisation for each individual readout channel.

The resulting phases could then be used to reconstruct the interferometer phase behaviour and performance. The investigations at the time have shown that modulation index values between 6 and 10 show promising results, if 10 or more bins are chosen to perform the fits.

11.2. Phasemeter system

A dedicated DPM phasemeter was implemented by using an FPGA evaluation board (see Appendix A), together with an eight channel ADC extension card (FMC107 from 4DSP [104]) and a DAC extension card (see Appendix A). The parallel processing capabilities of the FGPA were used to implement IQ demodulation pipelines (see Figure 11.2) that determine the complex amplitudes of the signal harmonics in real-time. The decimated signals are then fed into the fit-algorithm, which was integrated into the FPGA readout software running on a connected Linux PC. The implementation used $N = 10$ bins with configurable modulation depth m and frequency ω_m. The data sampling and demodulation

is performed at 40 MHz, the decimation rate, and, therefore, the repetition rate of the fit, was chosen to be $40\,\text{MHz}/2^{12} \approx 9.7\,\text{kHz}$ for initial simulations and $40\,\text{MHz}/2^{15} \approx 1.2\,\text{kHz}$ for performing low frequency measurements. To reduce the data volume to manageable size the 1.2 kHz fit results were also further decimated to $\approx 5\,\text{Hz}$ via an finite impulse response (FIR) filter, using a triangular window.

The available DAC outputs (up to 8 with the current implementation, including the modulation signal output), can be used to control various components of an interferometer, making it possible to actively reduce various noise influences. Prominent examples of this are laser amplitude and frequency stabilisations, as well as active path length stabilisations. The real-time readout enables us to use the determined interferometer signals (φ_x, k_x) to implement such active feedback by sending the respective sensor or actuator signals back to the FPGA. There they can be used to control the DAC outputs and, therefore, the connected actuators. The main application of this feature is to actively control the modulation depth m. Initial optical experiments using DPM have shown significant deviations of m during measurements, probably induced by hysteresis effects in the used phase modulator. Simulations show that the readout noise floor of DPM depends highly on the absolute value of m [135], therefore, its active stabilisation would potentially further improve the phase measurement performance.

11.3. Experimental results

To evaluate the phase performance of the DPM phasemeter a step-wise noise investigation was conducted. In an approach similar to the development presented in Part I, the system was tested using first digital signals and then analogue and optical ones. This allows us to easily distinguish noise sources, by step-wise introducing the influence of the DSP, the ADCs (including the analogue front-end) and the photo diodes / the optical phase modulator. The following gives an overview of the development steps and shows the performance levels achieved during the measurements.

11.3.1. Digital signals

The digital signal processing, which would later be implemented on the FPGA, was first simulated on a PC using floating-point capabilities. Thereby truncation noise influences could initially be avoided. Afterwards, the DSP core was converted into a VHDL based, fixed-point implementation. To test the performance of this system absolute phase measurements (see Chapter 5) were conducted. For this purpose a DPM signal simulator was implemented into the FPGA, together with the DSP required for the readout (see Figure 11.3). The
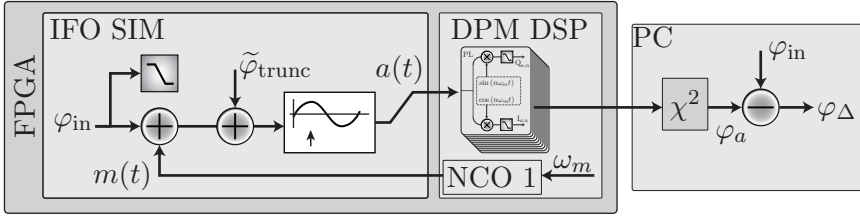
Figure 11.3.: Sketch of the digital absolute phase measurement set-up used to test the DSP of the DPM phasemeter. The simulated input signal is generated in an interferometer simulator (IFO SIM). Its phase includes the modulation signal $m(t)$, generated in the DSP readout core, and a configurable absolute phase signal φ_{in}.

phase signal fed into the simulator φ_{in} was decimated, together with the complex amplitudes, and was then directly compared with the phase φ_a, determined from the fit. The resulting phase difference φ_Δ was used to investigate noise sources and non-linearities.

For small signal dynamics ($\varphi_{\text{in}} \ll 2\pi$) the set-up was quickly able to achieve µrad performance levels. However, several optimisations were required to rech such levels also under the presence of larger input signals, since they caused significant non-linearities in the system. Aliasing during the decimation filtering in the FPGA was found to be critical. This issue was addressed by choosing modulation frequencies that have harmonics at exact notches of the implemented CIC filters (see Section 4.3.4). Additionally, the order of the filters was increased from $N = 1$ to $N = 2$. This was done to further increase the aliasing suppression for signals with large phase deviations, where the harmonics are spread away from the frequencies of the filter notches. One should note that this change, even though easily implemented, drastically increases the required logic resources. Each readout channel requires ten pipelines and therefore twenty filters, which now needed twice as much space.

The second most critical influence were harmonic artefacts, also called spurs, caused by the numerically controlled oscillators (NCOs). These artefacts are mixed, together with the actual tone, with the input signal and, thereby, they couple power of *wrong* harmonics into the detected amplitudes. This in turn reduces the ability of the fit algorithm to find an appropriate model of the modulation parameters, leading to non-negligible errors in the computed phase values. The application of simple phase dithering in the NCOs (see Section 4.1.2) was by itself not sufficient to decrease the magnitude of the spurs to desired levels. Reasonable results were achieved by choosing modulation frequencies that are not integer division factors of the sampling rate. Such signals have a much longer time before they perfectly repeat themselves, hence, they produce less parasitic spurs. The final performance achieved with the purely digital set-
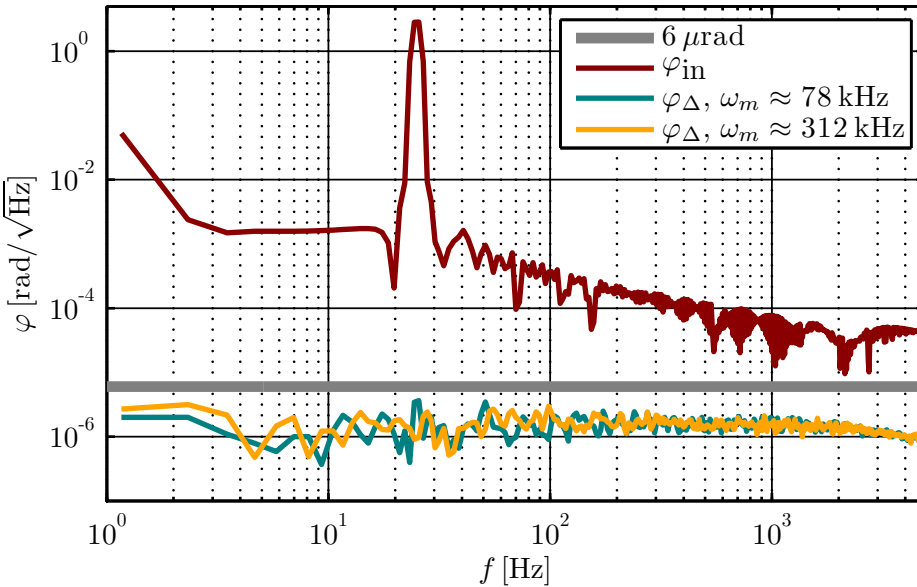
Figure 11.4.: Shown is the phase performance achieved during digital absolute signal measurements described in Figure 11.3, for two modulation frequencies. The input φ_{in} contained a strong signal at 25 Hz, thereby non-linearities were investigated. The white noise floor level is caused by the truncations in the NCOs. Later measurements with analogue signals (Figure 11.6) used NCOs with more bits and correspondingly showed a reduction of the noise level.

up is shown in Figure 11.4. For the implementations in the following sections the bit length of the sine look-up tables (LUTs) was also increased from 12 to 14 bit, to further decrease the white noise floor limiting the performance.

The white noise floor is higher than if the same underlying set-up (same hardware and 12 bit LUTs) is used as a phase-locked loop phasemeter (see Figure 8.6). A simple explanation for this is, that a phase-locked loop based system only uses a single NCO and a single mixer to demodulate the full signal power, thus, it is only once influenced by the phase noise induced by truncations before the LUT. The DPM phasemeter uses 10 NCOs and 20 mixers to demodulate the signal, coupling in the same amount of noise multiple times. The exact propagation of such a white phase noise through the fit is not yet understood, preventing a quantitative analysis.

The modulation depth used during this and the following experiments is $m = 9.3$. For this value the χ^2 value of the fits became minimal, leading to the potentially best fit results. This value also corresponds to one of the optimal points calculated during the initial investigations of DPM [135] for using 10 complex amplitudes.
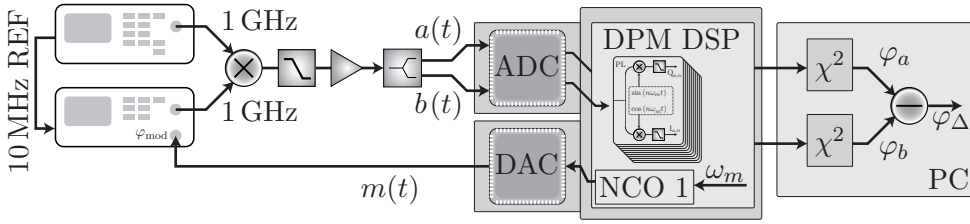
Figure 11.5.: Sketch of the split measurement set-up used for determining the performance of the DPM phasemeter with analogue signals. Two signal generators, locked to each other via a 10 MHz reference, are used to generate the modulated signal. Both produce a signal at 1 GHz, one of which is phase modulated by the DAC of the phasemeter. By mixing the two tones the desired phase modulated DC signal is generated. The tones around 2 GHz are suppressed with a low pass filter. The signal is amplified, split via a resistive splitter, and injected into two channels of the phasemeter.

11.3.2. Analogue signals

The next step in the performance testing was to use analogue input signals, including the analogue front-end, as well as the ADCs into the measurement chain. Since no dedicated signal generator was available another test set-up was implemented. Split measurements were carried out using the scheme described in Figure 11.5. Accordingly, the following measurements are mainly sensitive to additive noise, phase noise, and sampling jitter influences. But non-linearities are potentially not visible, as discussed in Chapter 5.

During the analogue and the later shown optical measurements the phase data was stored using the lower rate of ≈ 5 Hz. The low pass behaviour of a required signal amplifier made it necessary to reduce the modulation frequency to $\omega_m \approx 20$ kHz. The results obtained at low frequencies are shown in Figure 11.6. The achieved noise level is very promising, even the requirement for LISA is met at all frequencies. A noise increase to low frequencies is also observed. Based on the experience with the LISA-style phasemeter this noise can either be attributed to ADC sampling jitter or to thermally induced phase noise. Even though the modulation frequency is rather low, its 10th harmonic, the highest used signal component, is at ≈ 200 kHz, where sampling jitter might be relevant.

11.3.3. Optical signals

Finally, first tests using optical signals have been conducted. A simple experiment was constructed, that allowed one to perform π measurements, with signals from both outputs of an interference beam splitter (see Figure 11.7). The achieved performance for phase measurements between the two photo diodes showed a shoulder shaped excess noise, with a level of ≈ 1 mrad/$\sqrt{\text{Hz}}$ at fre-
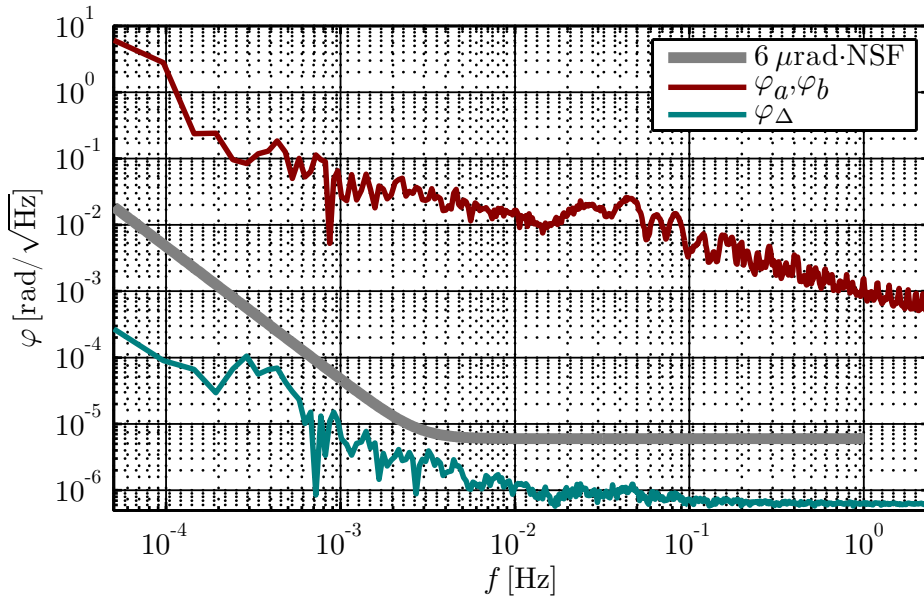
Figure 11.6.: Shown is the phase performance achieved during the split measurements using analogue input signals, as described in Figure 11.5.

quencies between $0.1\,\mathrm{mHz}$ and $0.1\,\mathrm{Hz}$ (not shown). Similar noise levels have been observed in the earlier DPM experiments [135]. There it was found that the low pass like transfer function of the photo diodes had caused the complex amplitudes to be distorted. A post correction of the complex amplitudes was performed by measuring the transfer function of the diode. Correction factors were computed and multiplied with the amplitudes before they were fed into the fit. Such a correction has so far not been implemented for the real time phasemeter described here, therefore, the underlying performance could not yet be fully determined.

The signal was also split behind the photo diode, which also revealed a performance worse than in the analogue measurements shown in the previous section. By exchanging the signal splitters and the used cables the noise performance of this split measurement could be brought to the levels shown in Figure 11.8. The thereby chosen components were optimised to have maximally flat transfer functions. However, the current understanding indicates that the already distorted signals are much more sensitive to small differences in the channels. The current understanding of this is, that a low pass filtered DPM signal not only has a reduced signal power, but that it also has a shape that is not predictable with the analytic model. For this reason, a fit operating on such a signal and its amplitudes will always give an incorrect answer, and in extreme cases it might not even converge.
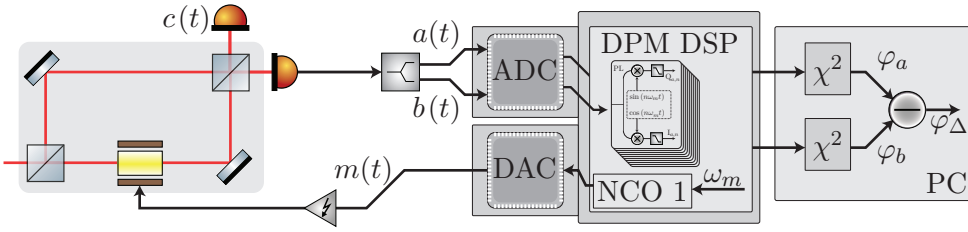
Figure 11.7.: Sketch of the split measurement set-up used for determining the performance of the DPM phasemeter with optical signals. The actual set-up was partly implemented in fibre, most notably a fibre wound around a ring piezo was used as phase modulator [136] to achieve the high modulation depth. An additional high voltage amplifier was used to drive the piezo with the modulation signal from the DAC. An active stabilisation of the modulation depth was not yet tested.

11.4. Future developments

The obvious next step of the development is to perform optical measurements using photo diodes with higher bandwidth or implementing a transfer function correction. Depending on the achievable noise floor a closed-loop control of the modulation depth might be implemented. Since π measurements are using the same signal (except for an inverted sign) their coupling to effects of noisy modulation phase and depth might be common. A measurement using more than one interferometer, with unequal beam path, will therefore be necessary to exclude such couplings and to prove the final system performance.

Even though the optical set-up required for DPM is extremely simple, the here implemented algorithms are rather excessive in terms of required processing capabilities. A single DPM channel (10 pipelines), uses ten times the FPGA resources required for a LISA pathfinder like readout and three times what is required for a PLL phasemeter (including a pilot tone PLL). Additionally, DPM requires a CPU to perform a fit running at kHz frequencies, in comparison to a simple arctangent computation required for the pathfinder-style readout, and no additional computations required for the PLL-style phasemeter. Hence it is desirable to reduce to amount of processing required for DPM in the future by using different, not yet developed, algorithms. These could make use of the fact that two of the signal parameters (m, ψ) are rather stable. If m is sufficiently stabilised by a feedback loop from the phasemeter, it can be assumed constant for all channels. Depending on the interferometer scheme this only partly true for ψ. It might have different delays in each interferometer between the time it was generated by the DAC and digitised by an ADC, leading to different phase values. The noise of ψ can however be assumed to be very low. In a suitable implementation its influence will be dominated by low frequency drifts,
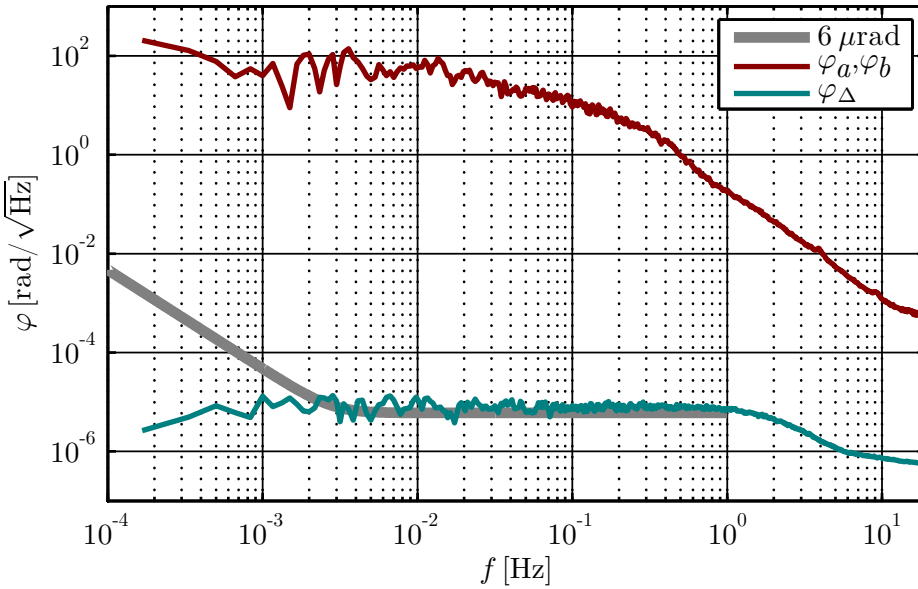
Figure 11.8.: Shown is the phase performance achieved during the split measurements using optical input signals, as described in Figure 11.7. The plotted graphs show the performance for signals that are split **after** the photo diode. The shoulder-like noise floor indicates a small-vector like coupling, as described for crosstalk in Section 3.4.

caused by the DAC and ADC jitter, as well as analogue phase noise. With these assumption it might be possible to implement a Kalman filter into the FPGA to track the signal phase, that uses less logical resources than the pipeline-fit design. A single set of pipelines could still be used, together with the fit, to determine values for the filter initialisation. However, no such algorithm has yet been realised and it therefore remains a conceptual idea.

Chapter 12

Digitally enhanced interferometry

Interferometers that operate over more than one optical fringe normally require a minimum of one readout chain (interference, photo diode, phasemeter channel) for each length measurement. However, a technique developed by D. A. Shaddock [139, 140, 141, 142, 143] is able to multiplex interferometric signals into a single measurement chain. Therefore, it can be used to strongly reduce the optical complexity of interferometers. This technique requires an extended digital signal processing and, hence, it is called digitally enhanced interferometry or simply Digital Interferometry (DI). This chapter describes the experimental investigation of DI that were conducted at the Albert-Einstein Institute in cooperation with the Australian National University (ANU). Goal of this endeavour was to demonstrate the multiplexing of signals with small relative optical delays ($< 1\,\text{m}$) and to determine possible performance limitations. Thereby, the feasibility of DI for performing local interferometry in future space missions was evaluated. The work presented here was carried out together with S. Köhlenbeck and K.-S. Isleif, whose Diploma and Master's Theses [144, 145] were supervised as part of this thesis, and with A. Sutton from the ANU.

The working principle of DI in a heterodyne interferometer is described in the first section. The role of pseudo-random noise (PRN) and its modulation rate on the multiplexing capabilities are explained together with the expected signals. The second section describes the implementation of a phasemeter system for the experiments. The generation of PRN codes with rates of up to $1.25\,\text{GHz}$ is explained together with the digital signal processing required for decoding and phase readout. Section three presents the optical set-up that was used to perform the experiments. The achieved multiplexing capabilities in the table-top experiment are presented. Noise performance measurements and their results are discussed in section four, along with an overview of experiments conducted
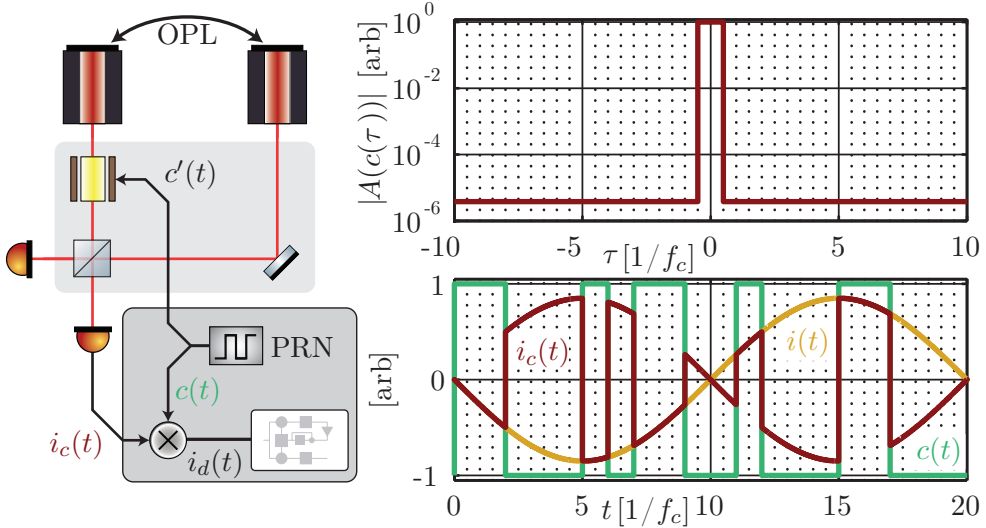
Figure 12.1.: The left side shows the modulation and demodulation scheme utilised in digital interferometry with two offset phase-locked (OPL) lasers. The bottom right side shows a time series of the original heterodyne signal $i(t)$, the PRN modulation $c(t)$, as well as the resulting signal $i_c(t)$. The heterodyne beat note can only be recovered if $i_c(t)$ is demodulated with a correctly delayed version of $c(t - \tau)$. The idealised autocorrelation of $c(t)$ in dependency of τ is shown for an example in the upper right plot using the following assumptions: a PRN code sequence with a length of 18 bits, a chip rate f_c equal to the sampling rate, and infinite bandwidth.

for noise hunting. In section five a new noise coupling effect is described that limited the phase performance. First experimental results for a reduction of this effect via an additional stabilisation loop are presented. The last section gives a short overview about a new variant of DI using a homodyne-style interferometer scheme. This technique was implemented and tested by A. Sutton at the ANU and it is described here, since it has some potential advantages over the heterodyne scheme.

12.1. Digitally enhanced heterodyne interferometry

A heterodyne interferometer generates a coherent photo diode signal with a carrier at the heterodyne frequency ω_0. In one interferometer arm a phase modulation is now introduced (as shown in Figure 12.1), described by a function $c'(t)$. The resulting photo diode signal (the DC components are neglected here) is

$$i_c(t) = i_{\mathrm{AC}} \cdot \cos\big(\omega_0 t + \varphi(t) + c'(t)\big). \tag{12.1}$$

The function $c'(t)$ is generated from a PRN code with a repetition time τ_r and a chip rate f_c. The PRN is amplified to generate phase values of $c'(t)$ that are either 0 or π (in contrast to 0 and 1). Thereby, the power at the signal carrier is fully suppressed in the photo diode signal (as shown in Figure 2.9). This type of modulation is also called binary phase-shift keying (BPSK). Since a phase modulation by π effectively inverts the sign of the AC signal, one can also describe the modulation by a function $c(t)$. This function has the same time dependence as $c'(t)$, but its values are 1 and -1 respectively, $c(t) = 2c'(t) - 1$. Using $c(t)$ the photo diode signal is given as

$$i_c(t) = i_{\mathrm{AC}} \cdot \cos\left(\omega_0 t + \varphi(t)\right) \cdot c(t). \tag{12.2}$$

An example of such a modulation is shown on the right side of Figure 12.1. One can now use a local copy of $c(t)$, artificially delayed by a time τ, to decode the signal before it is fed into the phase readout algorithm:

$$i_d(t) = i_{\mathrm{AC}} \cdot \cos\left(\omega_0 t + \varphi(t)\right) \cdot c(t)\, c(t - \tau). \tag{12.3}$$

If the signal is averaged over the repetition time τ_r, one can assume that the product of the code functions only changes the amplitude of the decoded signal $i_d(t)$. The resulting amplitude is then only dependent on the autocorrelation of $c(t)$:

$$A_c(\tau) = \int_{-\infty}^{\infty} c(t)\, c(t - \tau)\, dt. \tag{12.4}$$

For $\tau = 0$ the original beat note is recovered. Since $c(t)$ is a PRN the autocorrelation is minimal if $\tau \neq 0$, except for the situation where the delay equals a multiple of the repetition time τ_r. In the case of small correlation the signal will not be reconstructed. This means, that if two or more signals with sufficiently different delays of the PRN modulation are present, one can choose which beat note to reconstruct by simply tuning the demodulation delay. The unwanted beat note will be further decoded, and its residual amplitude is determined by the residual autocorrelation. Note that the assumption of a pure amplitude influence of $c(t) \cdot c(t - \tau)$ on the signal is only valid to first order, as described in Section 12.5.

Linear feedback shift registers (LFSRs) are a simple implementation of maximum length PRN sequences [146]. For an LFSR with N bits such a sequence has a repetition length of $2^N - 1$ and its residual autocorrelation, away from $\tau = 0$, has a value of $-1/(2^N - 1)$. Therefore, LFSRs are the tool of choice to produce the PRN used in DI. An example of the autocorrelation for $N = 18$ is shown in the top left plot of Figure 12.1. In a continuous system the correlation around $\tau = 0$ has triangular shape, given by the convolution of the two functions consisting of rectangular parts. However, in a sampled system this correlation is quantised, giving it a more step-function like appearance for chip rates closer to the sampling rate of the system. If both rates are equal, the correlation

is simply a rectangular function around zero with a width given by the chip delay $\tau_c = 1/f_c$. Any real system will include bandwidth limiting components and these will deform this ideal correlation. This is discussed in more detail in Section 12.3. Since τ_c corresponds to the width of the correlation it defines the minimum delay between signals that can be multiplexed. This means that an increase in chip frequency f_c can be used to detect signals with shorter light travel time or arm-length between them.

12.2. Phasemeter design

A phasemeter for the experiments was implemented using the FPGA infrastructure described in Appendix A. An ADC card (FMC112 from 4DSP [111]) was added to the system. It contains a fast ADC (EV8AQ160 from e2v [147]) to sample input signals with 1.25 GHz. The generation of such fast PRN signals was done by utilising a gigabit transceiver that is available in the FPGA (GTX from Xilinx [148]). These transceivers are complex serialisers that are normally used for ultra-fast serial interfaces. They can generate bits with rates of more than 5 gigabit per second. The interconnection between the components and the clock scheme is shown in Figure 12.2. The high modulation and sampling rates were used to allow multiplexing of signals with small optical delays between them, as required for table-top experiments. Earlier implementations of DI used sampling rates in the order of 250 MHz [143], which required larger optical set-ups or fibre based experiments.

Since the digital signal processing inside the FPGA can not be done at the full 1.25 GHz it is run at one eighth of this rate, 156.25 MHz. Therefore LFSR blocks were programmed that shift eight bits at a time. These blocks can also be controlled with a delay value, which triggers a state machine to produce the corresponding delayed PRN code. This is done by either shifting seven or nine bits in each clock cycle, until the desired delay is reached. One of these blocks, with a constant zero input delay, is used to generate the code that is serialised in the GTX and then sent out to the experiment. A LFSR with a length of $N = 18$ bits was used during most of the experiments. If the PRN code is generated at 1.25 GHz, the PRN has a repetition rate of $1/\tau_r \approx 4.7$ kHz.

The readout is divided into eight channels. Each one first decodes the ADC signal with a correspondingly delayed PRN, extracting the desired beat note. Then the signals are fed into a phase-locked loop (PLL), which tracks the beat note and determines its amplitude and phase. The PLLs use a variant of the design that was developed for inter-satellite interferometry (see Chapter 4). For correlation measurements the PLLs could also be used as simple heterodyne amplitude detectors. This was done by disconnecting their internal feedback, and by computing the amplitude from the I and Q values. Only seven of the channels actually decoded the signal with a local PRN. Channel eight simply
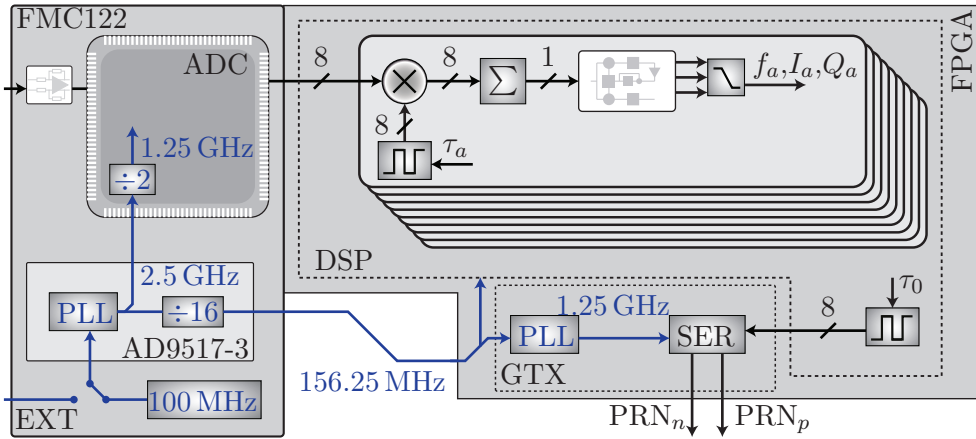
Figure 12.2.: Shown is a sketch of the phasemeter implemented for the digitally enhanced heterodyne interferometry experiments. The DSP block contains an additional LFSR, that is used to generate the initial code that is send to the experiment. The GTX component produces an analogue differential signal from these eight bits at 1.25 GHz in a serialiser (SER). The system either uses an external clock (EXT) or a 100 MHz crystal oscillator as initial clock. This clock is further distributed on the FMC122 via a dedicated clock chip (AD9517-3 from Analogue Devices [149]). The clock used for the digital signal processing and for the PRN generation is fed to the FPGA via an external cable. It is connected to a clock input of the GTX on the FPGA evaluation board.

measured the residual beat note amplitude of the interference. Since it depends on the modulation depth this was an indirect measurement of it, allowing it to be tuned during the experiment.

12.3. Optical set-up

A table-top experiment was constructed to test the multiplexing capabilities at 1.25 GHz and to determine the phase performance of the system. To generate the phase modulation with the required modulation index a fibre-coupled electro optic modulator with a bandwidth of 7 GHz was used (NIR-MPX-LN-05 from Photline [150]), which was driven by a dedicated BPSK amplifier (NIR-MPX-LN-05 from Photline [150]). The amplification and, therefore, the modulation index could be tuned with a simple potentiometer. The signal detection was done with a free beam photo receiver with a bandwidth of 12 GHz (1567-A from New Focus [151]). A heterodyne interferometer was constructed as depicted in Figure 12.3. The interference signals detected on the photo diode were generated
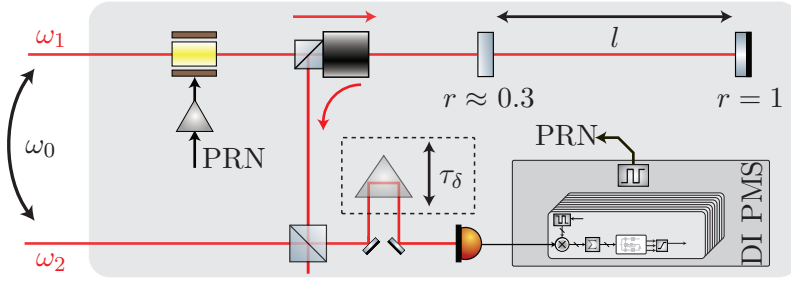
Figure 12.3.: Shown is a simplified sketch of the experimental set-up for digitally enhanced heterodyne interferometry. Two laser beams with an offset frequency ω_0 were used to establish the heterodyne scheme. A Faraday rotator is used, together with a polarising beam splitter, to send the reflections from two mirrors to the interference beam splitter. The reflectivity of the mirrors were chosen such that several reflections with sufficiently equal power were generated. An additional optical delay line was implemented with a translatable corner cube, enabling us to change the overall delay of all signals. The set-up is shown in more detail in Appendix E.

by reflections from a two mirror set-up, and the corresponding signal delays depend on the distance l between them.

The system was characterised for its multiplexing capabilities by determining the RMS amplitude of the decoded signal for a wide range of delays. Since the system sampled with 1.25 GHz, it should ideally be able to distinguish signals with relative delays larger than $\tau_c = 1/f_c = 0.8\,\text{ns}$. This corresponds to a length of $\approx 24\,\text{cm}$. The set-up detects reflected signals that travel back and forth to each mirror, therefore the minimal usable separation is 12 cm. However, the measured amplitudes were found to be strongly deformed, as shown in Figure 12.4. This effect is understood to be caused by the finite bandwidth of the system, which is probably limited by the analogue front-end circuit on the ADC card. The corresponding transfer function is not well specified by the manufacturer, therefore, no quantitative analysis of this effect was done. In a first model the influence of a filter on a step function, corresponding to the correlation at 1.25 GHz chip rate, was calculated. This model indicates that the transfer function has a bandpass like behaviour consisting of a first order high pass at a few MHz and a first order low pass at around 700 MHz. Due to this effect the mirror separation for the later performance measurements was chosen as three times the minimum separation, significantly reducing any crosstalk between reflections due to residual correlation.

Delay scans over a wider range have also revealed that the correlation has smaller maxima at delays before the actual signals (bottom plot in Figure 12.4). This effect is currently not understood. Bandwidth limitations of the system
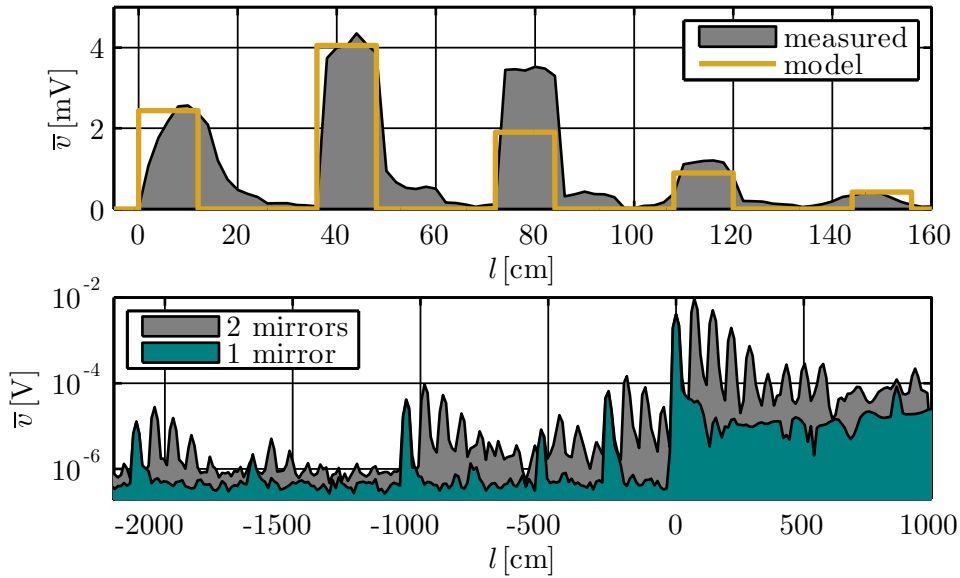
Figure 12.4.: The top shows the measured amplitude, and, therefore, the correlation for a delay scan with both mirrors present. The chip frequency was 1.25 GHz. The sub sampling delay resolution was achieved by tuning the overall delay τ_δ with the optical delay line shown in Figure 12.3. An ideal model is fit to the data, it assumes infinite bandwidth of the PRN generation and signal detection, as well as signal amplitudes that only depend on the power levels expected from the mirror reflectivities. The bottom plot shows a delay scan (without tuning τ_δ) over a much wider range for both mirrors and the second mirror blocked. In these measurements the chip frequency was 625 MHz.

are only expected to cause deformations of the delays in one direction. Signal reflections occurring on the analogue components should also only delay the signal and not advance it. Residual correlations from an earlier round trip are not expected to be the cause of this phenomenon, since the corresponding round trip length for an LFSR with $N = 18$ is ≈ 3 Mm. If the light path was blocked no correlations remained, this showed that the additional peaks are not caused by purely electronic crosstalk in the phasemeter.

12.4. Phase noise performance measurements

To determine the performance of the system the three largest reflections from the two-mirror set-up were measured. Their propagation is shown in the top of Figure 12.5. The phase change corresponding to the relative mirror motion φ_l
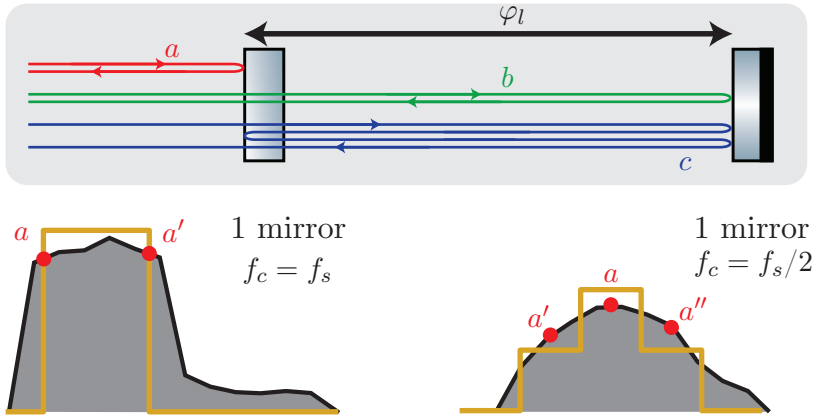
Figure 12.5.: The top shows the propagation for the first three reflections from the two-mirror set-up. Signal a probes only the first and b only the second mirror position. Signal c circulates the set-up twice and probes the first mirror position as well as twice the mirror separation.

can be computed from signal a and b, as well as from signal a and c:

$$
\begin{aligned}
\varphi_l &= \varphi_b - \varphi_a = \varphi_{ab} \\
&= (\varphi_c - \varphi_a)/2 = \varphi_{ac}/2.
\end{aligned}
\tag{12.5}
$$

All phase signals used here correspond to the one-way phase, which is directly proportional to the mirror distance changes. The two measurements of φ_l were used to form a zero combination,

$$
\varphi_{ab} - \varphi_{ac}/2 = \varphi_{l-l} = 0.
\tag{12.6}
$$

This measurement is not sensitive to the actual mirror motion, therefore it probed the system performance at low frequencies without a thermally stable measurement set-up. The best achieved performance in such a measurement is shown in Figure 12.6. To achieve this performance level some aspects of the system had to be tuned. For certain combinations of PRN repetition length and heterodyne frequency artefacts spoiled the measurement performance. These artefacts were identified to be caused by unwanted mixing products of the code and the sinusoid, which were aliased into the measurement band.

An excess noise at low frequencies spoiled the performance of the zero combination. Since phase differences induced by optical effects between the reflections could not be excluded, another type of zero combination was used. The broadening of the correlation peaks due to limited signal bandwidth allowed us to measure a single reflection twice, as shown in the bottom of Figure 12.5. The two corresponding phase values φ_a and $\varphi_{a'}$ were generated by the exact same
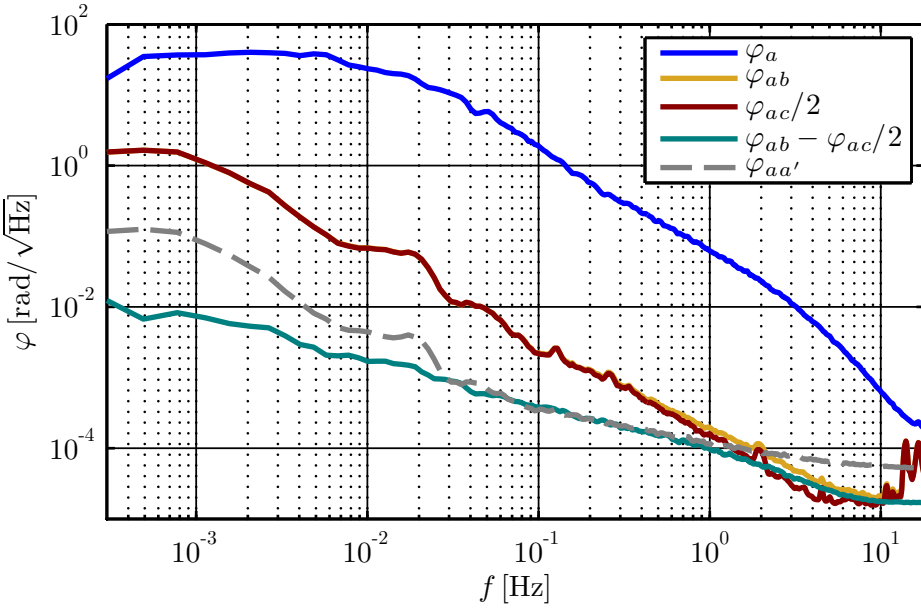
Figure 12.6.: Performance with 36 cm mirror separation and 1.25 GHz modulation frequency. Shown is one of the initial signals (φ_a), the mirror distance combinations φ_{ab} and $\varphi_{ac}/2$, as well as the two-mirror $\varphi_{ab} - \varphi_{ac}/2$ and one-mirror $\varphi_{aa'}$ zero combinations.

beam. Therefore no geometrical or optical effects should limit their zero combination,

$$\varphi_a - \varphi_{a'} = \varphi_{aa'} = 0. \tag{12.7}$$

By using the first mirror, and blocking the second one, it was also ensured that no crosstalk due to residual correlations from other reflection influence this measurement. However, the achieved performance levels were limited to similar values as the earlier zero combination, as shown in Figure 12.6. This indicated that the performance in both cases was limited by the same effect.

Several experiments were conducted to investigate influences from known noise sources to identify the excess noise at low frequencies. The two critical ones are the following:

- **Laser frequency noise** Both presented zero combinations should not be sensitive to frequency noise. For $\varphi_{aa'}$ it cancels completely and for φ_{l-l} the coupling is negligible. Correspondingly the use of an iodine stabilised laser, which has a much lower laser frequency noise at low frequencies, did not improve the performance in the zero combinations.

- **Laser amplitude noise** Low frequency amplitude fluctuations of one of the laser beams could potentially have caused the excess noise. By actively
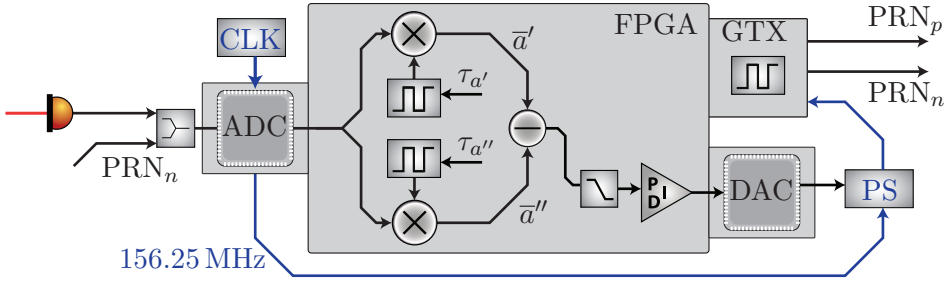
Figure 12.7.: Shown is a sketch of the delay-locked loop, implemented additionally into the phasemeter to reduce clock jitter between the ADC and the GTX. Part of the PRN code, available as differential signals with a positive (PRN_p) and a negative (PRN_n) signal line, was directly fed back into the ADC and decoded with two different delays. Linear operating points on the correlation ($\varphi_{a'}$ and $\varphi_{a''}$) were made available by reducing the chip rate to 625 MHz (see bottom left sketch in Figure 12.5).

modulating the amplitude of one laser the coupling was investigated. Significant coherent excitations of the amplitude did however not influence the phase measurement performance in either of the zero combinations. Therefore, amplitude fluctuations were discarded as cause of the excess noise.

Since all signals are passing through the same phase measurement chain, many noise sources are common and should therefore not limit the performance [139]. Phase noise in the analogue front-end due to temperature drifts could therefore also be excluded.

12.5. Clock noise coupling

Sampling jitter was initially assumed to not be limiting for heterodyne DI, since all signals would be influenced by the exact same jitter and it would, therefore, be common between all channels. To test this, an additional sampling jitter was actively introduced into the system by using a signal generator that phase modulated the 156 MHz clock that is sent from the ADC card to the GTX (see Figure 12.2). These measurements showed a strong correlation between the induced time jitter and the phase error in the zero combinations. The amplitudes of the decoded signals were also found to be modulated by the time jitter. The current understanding of this coupling is, that not only the amplitude of the decoded signals depends on the delay, but that also the phase is depending on it. Any clock jitter between the PRN generation and the ADC sampling would, therefore, couple in phase and amplitude fluctuations, as observed in the mea-
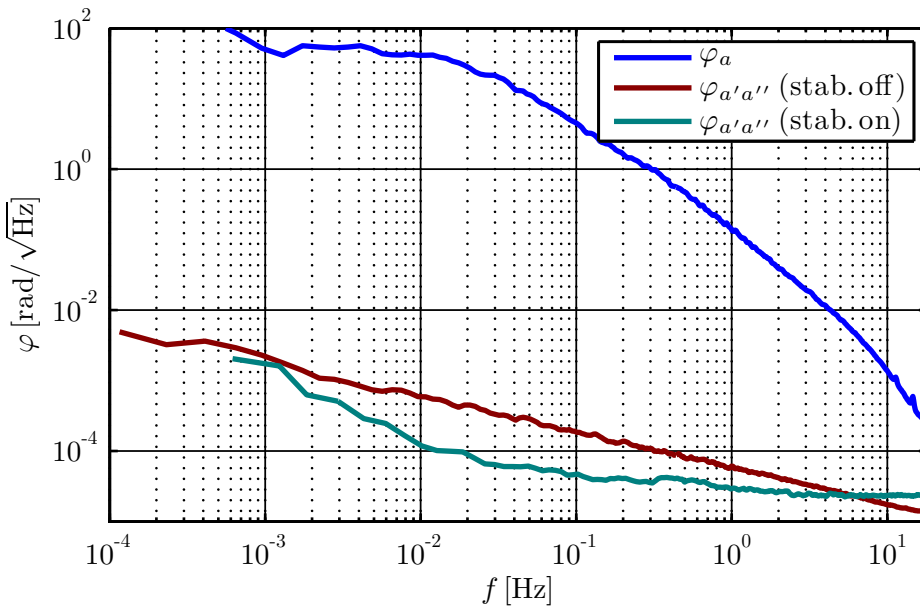
Figure 12.8.: Performance of the one-mirror zero combination measurement $\varphi_{a'a''}$ with and without active clock stabilisation. The chip rate was 625 MHz.

surements. This effect depends strongly on the deformation of the correlation due to limited bandwidth, hence, it might not have been visible in earlier experiments using slower chip rates and lower bandwidth signals. It is assumed that for an infinite bandwidth this effect would not occur, since the reconstructed sine waves after decoding would be perfect sinusoids. First simulations were also able to reproduce this coupling for a finite signal bandwidth.

To validate that this new noise coupling is indeed the limiting effect an active stabilisation scheme was implemented. A delay-locked loop (DLL) was constructed, that used the correlation of an analogue version of the PRN code to determine clock jitter between the ADC and the GTX (see Figure 12.7). The so generated error was fed into a controller and the resulting actuator signal was used to control the phase of an analogue phase shifter (see Figure B.9), which was included in the clock distribution from the ADC to the GTX. The unity gain frequency of the loop was measured to be around 40 Hz.

The stabilisation via the DLL reduced the excess noise in the single mirror zero combination measurement, as shown in Figure 12.8. The noise level at 10 mHz was reduced by almost one order of magnitude. This proved that the sampling jitter was indeed the limiting effect. The achieved results with the stabilisation were not easily reproducible. It is assumed that the analogue phase shifter was not well suited to drive the signal into the GTX. Its internal clock structure is rather complex and not designed for such an actuation scheme.

These results are preliminary, since only few experiments were conducted with the DLL stabilisation. So far no performance improvement could be achieved for the two-mirror measurements. But it is yet unknown, if this was caused by a problem with the stabilisation or another effect.

For the two-mirror zero combination the best achieved performance in the presented experiments has a white noise floor of $\approx 3\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ at high frequencies and a noise increase to low frequencies, with a level of $\approx 300\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ at 10 mHz. The one-mirror zero combination using the DLL stabilisation was able to improve this performance at 10 mHz to $\approx 20\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$. The multiplexing distance of 36 cm (3x 12 cm) is, to the authors knowledge, the smallest one achieved yet.

12.6. Digitally enhanced homodyne interferometry

An alternative method of implementing multiplexed interferometry has recently been investigated by A. Sutton [152, 153]. This scheme is called digitally enhanced homodyne interferometry. It utilises homodyne style interferometer set-ups, in which the sensing light beam is modulated. In contrast to the heterodyne scheme the modulated code has two components, which effectively combines to a single modulation function $c'(t)$ with values 0, $\pi/2$, π and $3/2\pi$. This type of encoding is also called quadrature phase-shift keying (QPSK). In the set-up developed by A. Sutton a dual parallel Mach-Zehnder modulator is utilised to generate the correct phase modulation. The homodyne technique has two distinct advantages over the heterodyne scheme.

- **Self homodyning** The QPSK scheme is able to produce interference signals that can be read out without an additional local oscillator. The interference of reflections with sufficiently different delays can be decoded with this technique, as demonstrated by A. Sutton. This simplifies the construction of interferometers significantly and can be used to develop new optical set-ups.

- **Supression of residual correlations** The homodyne scheme has an inherently higher suppression of undesired signals due to its QPSK modulation. For a code length of L the residual signal influence (which causes a white noise floor) is on the order of $1/L$ for the homodyne scheme [152], in contrast to an order of $1/\sqrt{L}$ for the heterodyne scheme [143].

Performance measurements with this scheme have revealed that it is able to achieve pm level performance at frequencies above 1 Hz. However, the achievable phase performance below 1 Hz has not yet been evaluated and it remains to be determined, if a similar clock noise coupling as in the heterodyne scheme is present.

Chapter 13

Comparison of techniques for local interferometry

Different techniques are available to perform the local position sensing in future satellite missions. Heterodyne interferometers with kHz signals, as used in LISA Pathfinder [33, 154, 155], or with MHz signals, as used in LISA and GRACE Follow-On, are the obvious choice for future missions. Both of these technologies have been used to demonstrate pm performance levels at frequencies below 1 Hz and space-qualified components are either already available or they are currently developed. Neither deep phase modulation, nor a variant of digitally enhanced interferometry have yet been able to demonstrate the same performance levels at low frequencies. However, based on ongoing developments, some of which have been presented in this thesis, it is reasonable to believe that the new techniques might be able to close the performance gap in the future. This chapter motivates such future developments by presenting the potential advantages of the alternative techniques for future implementations.

The first section compares the non-multiplexed interferometer schemes, classic heterodyne interferometry and deep phase modulation. All of them use a very similar design for the optical sensing part of the interferometer, therefore their main differences are in the light preparation and phase readout. The scaling of interferometer set-ups is critical for future gravity missions that aim to read out the position of more than one test mass in multiple degrees of freedom. Therefore, the readout of one test mass in two directions is used as an example in the second section to illustrate the potential simplifications possible with multiplexed interferometer schemes.
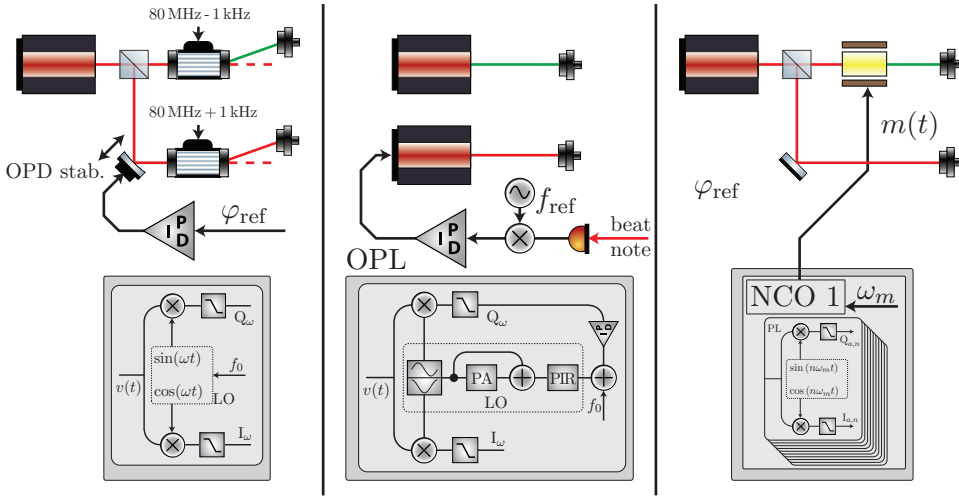
Figure 13.1.: The light preparation and the phase readout, for one channel, is illustrated for three non-multiplexed interferometer schemes. The left shows the LISA Pathfinder-style modulation used for kHz heterodyne interferometry. The middle shows the simple LISA-style scheme for MHz beat notes. The right shows the deep phase modulation set-up. Typically, the laser (or the master laser, in case of the scheme in the middle) would also be stabilised to reduce frequency noise (not shown). All shown components can also be implemented in fibre, which might reduce space.

13.1. Non-multiplexed interferometer schemes

Interferometers can conceptually be split into a stable, or phase sensitive, part and an unstable, or phase insensitive, part. In general the path length changes in the unstable part are common between all interferometers and, therefore, they do not couple into the measurement. This is achieved by including an additional "reference" interferometer in the stable part. In LISA Pathfinder, for example, these two parts are also physically split into a modulation unit, which is phase insensitive, and the optical bench, where the actual measurements are done [33, 154, 155]. Both heterodyne interferometer concepts (with kHz and with MHz signals), as well as deep phase modulation, require a very similar stable part of the interferometer (see the left sketch in Figure 13.2). This is mainly, because each sensing beam needs a separate local oscillator and dedicated photo diodes for its readout. Each stable bench also requires a reference interferometer, to subtract any phase noise induced in the light preparation or the optical fibres, which are assumed here to be used for connecting both parts of the interferometers. The unstable interferometer parts and the readout are sketched in Figure 13.1 for the three schemes.

The LISA Pathfinder-style scheme uses two acusto-optic modulators (AOMs),

operating at slightly different frequencies, to generate the kHz beat notes. Offset frequency locks between two lasers can not be used to generate kHz beat notes, since $\approx 10\,\mathrm{kHz}$ loop bandwidth is needed which in turn needs $\lesssim 100\,\mathrm{kHz}$ frequency offset for phase difference detection. A simple IQ demodulation (see Section 2.3) is used, together with an arctangent calculation implemented in a CPU (not shown), to determine the phase of the signals. This is possible for nearly stationary path lengths to be sensed, e.g. test masses with respect to a surrounding optical bench. The MHz signals used for driving the AOMs (here shown are 80 MHz, as used in LISA Pathfinder), must also be phase locked to the same clock that is used for the demodulation/sampling in the phasemeter. Additionally, a feedback loop, the so called optical path length difference (OPD) stabilisation [52], is required to achieve pm performance levels.

A LISA-like scheme with MHz beat notes consists simply of two lasers. An offset phase-lock (OPL) keeps their relative frequency stable. An error signal for this lock can, for example, be taken from the reference interferometer. The phasemeter can simply lock to the input signal with its phase-locked loops, as described in detail in Part I of this thesis. It requires a pilot tone correction to achieve pm performance levels at mHz frequencies. If one of the two lasers is only used for local interferometry, the heterodyne frequency can be set to a convenient value of a few MHz, above the relative intensity noise of the lasers and below frequencies that have strong temperature induced phase drifts in the measurement chain due to finite bandwidth (see Chapter 3). One can also generate MHz beat notes with AOMs. But to reach frequencies around or below 10 MHz at least two AOMs are required, since they only operate at frequencies above $\approx 30\,\mathrm{MHz}$. Due to the added complexity such a scheme is not considered here.

Deep phase modulation uses one laser and a phase modulator (for example an electro-optic modulator or PZT) that is driven from the phasemeter, as described in Chapter 11. The phasemeter is, however, currently the most expensive one in terms of required processing capabilities in comparison to the other schemes. At this point, the power consumption for the readout of each scheme can not be easily compared, but, deep phase modulation requires the fewest active components and will therefore probably be the most efficient scheme in terms of power consumption of analogue components. It has a lower optical complexity than the Pathfinder-style scheme and only requires one laser, making it a promising candidate for future missions.

One additional aspect that is crucial for these interferometer schemes is their sensitivity to stray light. Any residual photons that travel an unwanted path in the interferometer can end up in one of the interference signals and, thereby, cause small vector errors (see Section 3.4). This is, for example, the dominating noise effect for the back link fibre used in LISA for the reference interferometers [47]. One way to circumvent stray light coupling is to use one or more additional signal beams in different parts of the interferometer with non-equal frequencies
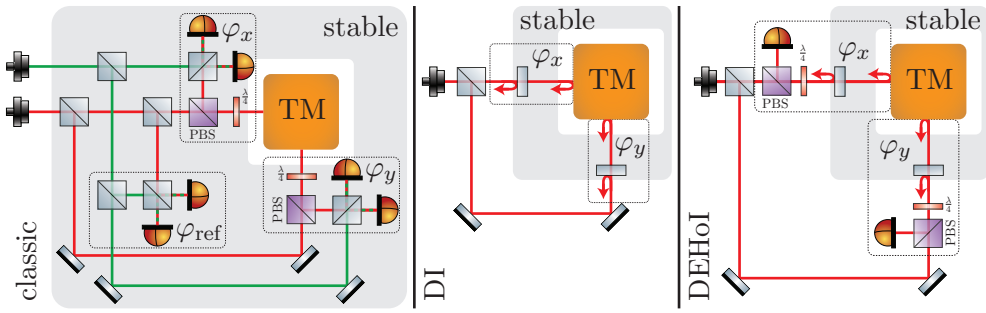
Figure 13.2.: Shown are the layouts for reading out two longitudinal degrees of freedom of a test mass. The left set-up uses non-multiplexed, classic interferometer schemes, described in the previous section. The set-up in the middle can be implemented with both types of digital interferometry (DI). The layout on the right uses self-homodyning and can therefore only be implemented using DEHoI.

or modulations [48]. Thereby, stray light from one beam does not influence the interference with another one, since its parasitic signal will be at a different frequency. In the case of MHz heterodyne interferometry this would require an additional laser for each beam. The kHz scheme would need additional AOMs and the correspondingly phase-locked modulation signals. Deep phase modulation requires only additional phase modulators and control channels, which makes it the probably most cost-effective scheme in terms of optical complexity, if many interferometers with large stray light couplings have to be read out.

13.2. New interferometer layouts with multiplexed interferometry

Multiplexed interferometer schemes using digital interferometry (DI) make it possible to reduce the number of phase stable components in an interferometer significantly. Figure 13.2 shows two layouts that can be implemented with DI in comparison to a layout using non-multiplexed interferometry. The components and the path lengths on the gray area have to be phase stable to achieve the desired performance levels. It is obvious that the DI schemes require much less optical components, and especially far less ones that are phase stable. Only the reference mirrors are critical. Note, that the distance between the mirror and the test mass is limited to values larger or equal to the chip rate. For the implementation described in Chapter 12, which uses 1.25 GHz, this would be 12 cm, though only if the bandwidth of the system is sufficiently high. The increased availability of data processing hardware and modulators with bandwidths above 10 GHz, developed for communication systems, will, however, make it possible

to achieve much higher chip rates and therefore smaller minimum distances in the future.

The scheme shown in the middle of Figure 13.2 performs the signal detection with the light that is sent back through the optical fibre. The modulation and detection is not shown. Only longitudinal signals can be measured with this set-up, since wavefront informations are destroyed in the fibre. This scheme can be operated either with digitally enhanced heterodyne interferometry (DEHeI) or with digitally enhanced homodyne interferometry (DEHoI).

The set-up on the right in Figure 13.2 is a scheme that is only possible with DEHoI. The signals reflected from a reference mirror and from the test mass are both sent onto a local photo diode and there they can be read out using the self-homodyning of DEHoI [65]. If a quadrant photo diode is used, differential wavefront sensing can also be performed. This combines the minimal amount of stable components for DI with the use of only a single fibre and angular readout capability. Note, however, that the photo diodes would require a sufficient bandwidth. They would therefore probably be quite small, which would make it necessary to implement additional imaging optics.

Due to the special properties of DI stray light effects can be suppressed very effectively if the delays between the parasitic beams and the actual signals are long enough in comparison to the chip rate. Therefore, DI is also a good candidate for future set-ups that have an inherent stray-light coupling, like the LISA back-link fibre [47].

Chapter 14

Summary

The first part of this thesis described the readout of inter-satellite heterodyne laser interferometers with digital phase measurement systems. A comprehensive analysis of critical effects and components in the analogue and digital parts of the phase measurement chain was presented. The central element of this was a linear model of the core phase readout algorithm, the all-digital phase-locked loop (ADPLL). The model was also extended to describe non-linear effects occurring in LISA or other missions due to large signal dynamics and low signal-to-noise ratios.

Digital signal simulations were used to test the ADPLL model by performing non-linearity tests with realistic, LISA-like input signals. The performance and stability of the ADPLL was verified together with several other aspects of the model, including non-linear noise couplings.

A phasemeter for GRACE Follow-On breadboarding experiments was used to demonstrate differential wavefront sensing with MHz beat notes. This demonstration also included an active feedback control, which made it possible to stabilise the angle between two laser beams to better than 100 nrad over 12 hours. This demonstrated the feasibility of one of the core functionalities of the Laser Ranging Interferometer for GRACE Follow-On.

Several phasemeter prototypes with different analogue front-ends were constructed that allowed measurement of $\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ performance levels at low frequencies.

Three signal tests were conducted with analogue and optical signals to investigate the non-linearity of the full measurement chain. These test have already been able to demonstrate dynamic range performance levels of more than six orders of magnitude at mHz frequencies.

The highlight of the first part of this thesis is the elegant breadboard model

of the LISA phasemeter, which has been successfully developed and tested. It includes all functionalities required for LISA and, with the use of an additional active thermal stabilisation, it was able to demonstrate the full phase measurement performance with realistic analogue signals. This demonstrated the feasibility of one of the core elements of LISA. The availability of such a system will allow for more integrated tests of the full LISA metrology in the future.

The second part of the thesis presented developments for new, alternative interferometry techniques. A real-time phase readout system for deep phase modulation interferometry has been constructed and successfully tested. Performance levels of $\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ have been achieved in analogue tests, and first optical tests have revealed $m\mathrm{rad}/\sqrt{\mathrm{Hz}}$ levels below $1\,\mathrm{Hz}$. The methods for improving these performance levels are well known and it is expected that deep phase modulation will be able to reach LISA like performance levels in the near future, using this new type of phasemeter.

A high-speed digitally enhanced heterodyne interferometry set-up and phasemeter with chip and sampling rates of up to $1.25\,\mathrm{GHz}$ has been constructed. It was used to demonstrate multiplexing for the shortest relative light travel distances yet and it achieved performance levels of $3\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ at $10\,\mathrm{Hz}$. An excess noise was found at low frequencies, which was identified to be a result of timing jitter between the pseudo-random noise generation and the signal sampling. First experiments with an active stabilisation were able to reduce its influence by one order of magnitude at $10\,\mathrm{mHz}$, opening the way for future implementations with increased low frequency performance and even higher sampling rates.

Finally, the advantages of multiplexed interferometry and deep phase modulation for future local satellite interferometry systems were presented. The significant simplifications of the optical set-ups were found to be a compelling case for further pursuing these new techniques.

This thesis has, among other things, shown that digital signal processing is not only a very efficient passive readout tool for interferometry, but that it can also be actively used to change the way interferometers operate. The increasing availability of computational resources and the advantages in processing algorithms will make it even more useful in the future and open up new, as yet unknown ways of performing precision metrology.

Appendix

Appendix A

ML605 FPGA infrastructure

Many phasemeter systems developed in this thesis used a common hardware infrastructure, which is described in this chapter. The hardware developed for the elegant breadboard of the LISA phasemeter is described in Chapter 9.

The core of the FGPA infrastructure is a Xilinx ML605 Virtex-6 evaluation board [156] shown on the left in Figure A.1. The ML605 carries a powerful FPGA that connects to an Ethernet interface and various other peripherals. The most important feature is the possibility to connect up to two FPGA mezzaine cards (FMC) that can carry various ADCs, DACs and other components. This allowed us to use the ML605 and its interfaces in the various projects presented in this thesis by connecting the required peripherals and adjusting the digital signal processing. The basic infrastructure implemented for the ML605 is shown on in Figure A.1.

A.1. FPGA internal organisation

The synthesizeable logic inside the FPGA was divided into blocks corresponding to their fundamental functions. This structure was also reflected in the VHDL code, by using separate files and VHDL entities for each block.

- UDP interface (UDP IF): This element reads UDP packets send to the IP address of the FPGA and to write back to the transmitter. The code is based on an UDP/IP stack [157] available on `opencores.org`. It was slightly modified to fit the desired functionality. Two different UDP ports are available to distinguish between register (REG) and stream (STR) communication.
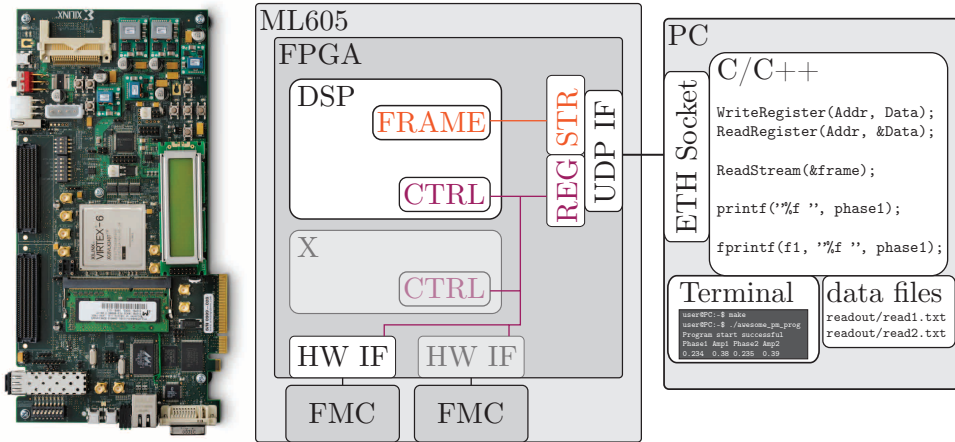
Figure A.1.: The left side shows a picture of the ML650 evaluation board. The right side shows a sketch of the infrastructure implemented into the ML605.

- Register interface: A rather standard implementation of communication between the PC and elements in the FPGA. A value can be written into a register by first sending its address and then the data over a 32 bit wide bus. Well defined address spaces allow one to reach each register individually from the PC. Registers can also be read out by receiving a read command from the PC. An automatic reply is generated and the desired value is sent back.

- Streaming interface: The large amount of data generated by the various DSP elements is not suitable to be sent via the register interface. Therefore a large packet of data is collected into a frame, which is sent in regular intervals (usually the decimation rate) to the PC with separate UDP packets. The packet size was varied between 256 and 512 bytes with readout rates of up to 100Hz.

- DSP block: The actual signal processing is done in a separate block. The content of this block varies greatly between the different phasemeter types. It is controlled via the register interface and it usually creates the frame packets that are sent to the PC via the streaming interface.

- Hardware interface: Each peripheral connected to the ML605 required a representation in the FPGA that interfaces the specific signal lanes, clocks and command interfaces to the inner structure of the FPGA. Depending on the hardware used these blocks can have a large complexity and their correct implementation is very time consuming. Since they are part of the overall signal dynamics, by influencing delays and sampling rates, they have to be designed with great care to not spoil any system performance.

A.2. PC program

The program used to control and readout the ML605 was implemented under Linux using either C or C++. The use of command line flags, configuration files and hard coded parameters allowed us to control each system and to generate helpful meta data, like time and date of each measurement. The data received from the streaming interface is converted here back into separate signals, some of which are either used in further processing or simply stored in text files. The relevant signals are displayed on the control terminal to monitor the experiments.

A.3. Digital-to-analogue converter FMC

Several experiments using the ML605 infrastructure required an active control of analogue components via the phasemeter. Since no suitable digital-to-analogue converters were commercially available a dedicated FMC card was constructed together with I. Bykov. The card carries eight DACs with a sampling rate of approximately 2.5 MHz. Such a high rate was chosen to be able to construct control loops with high bandwidth and to generate modulation signals required for deep phase modulation (see Chapter 11). The schematics of the analogue back-end of one DAC channel are shown in Figure B.8.

Appendix B

Schematics

Schematics of circuits designed specifically for experiments and hardware in this thesis are presented here.

Figure B.1.: Schematic of two channel 1-bit digital-to-analogue converter used in Chapter 7. The power supply is not shown.
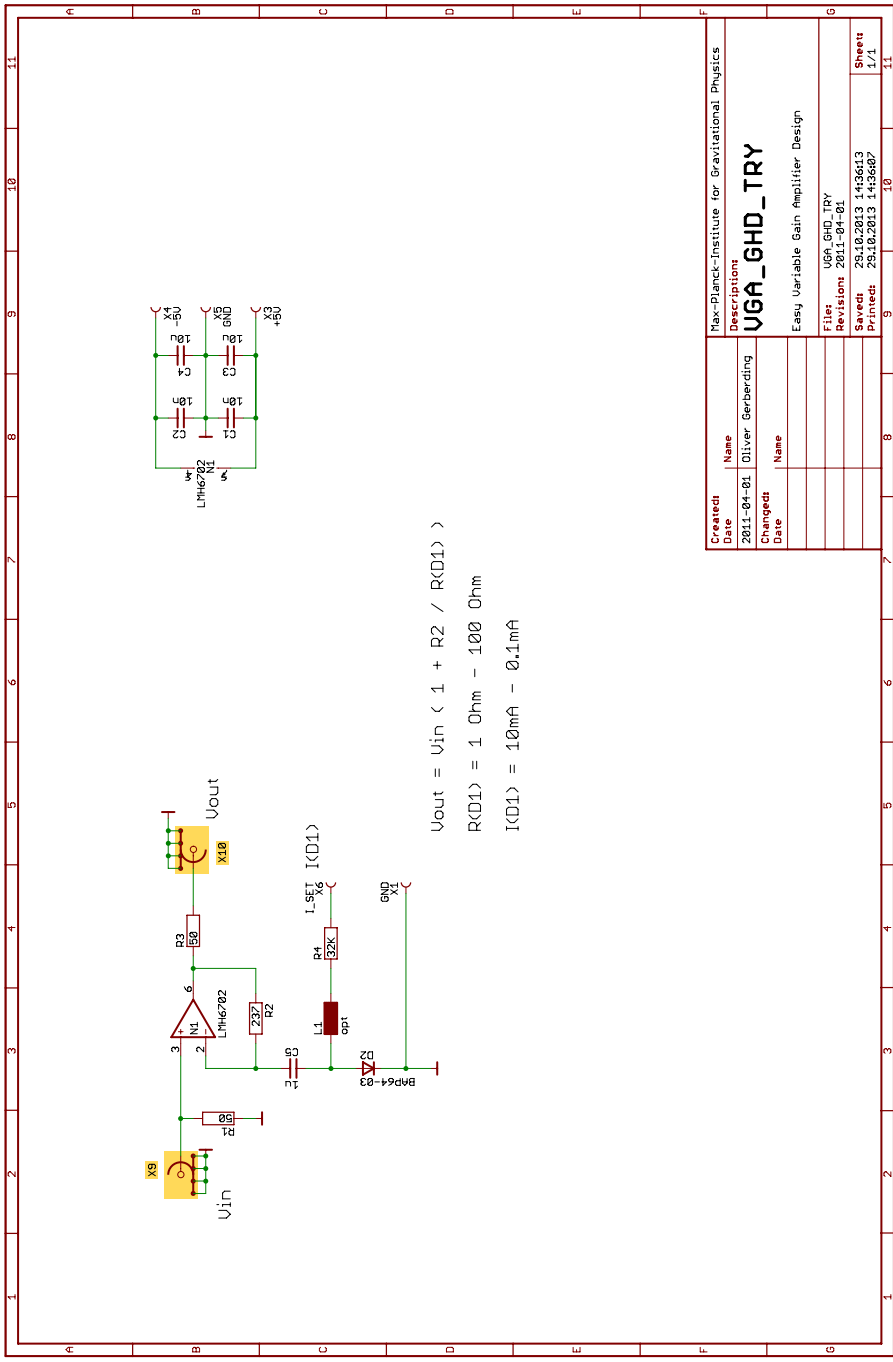
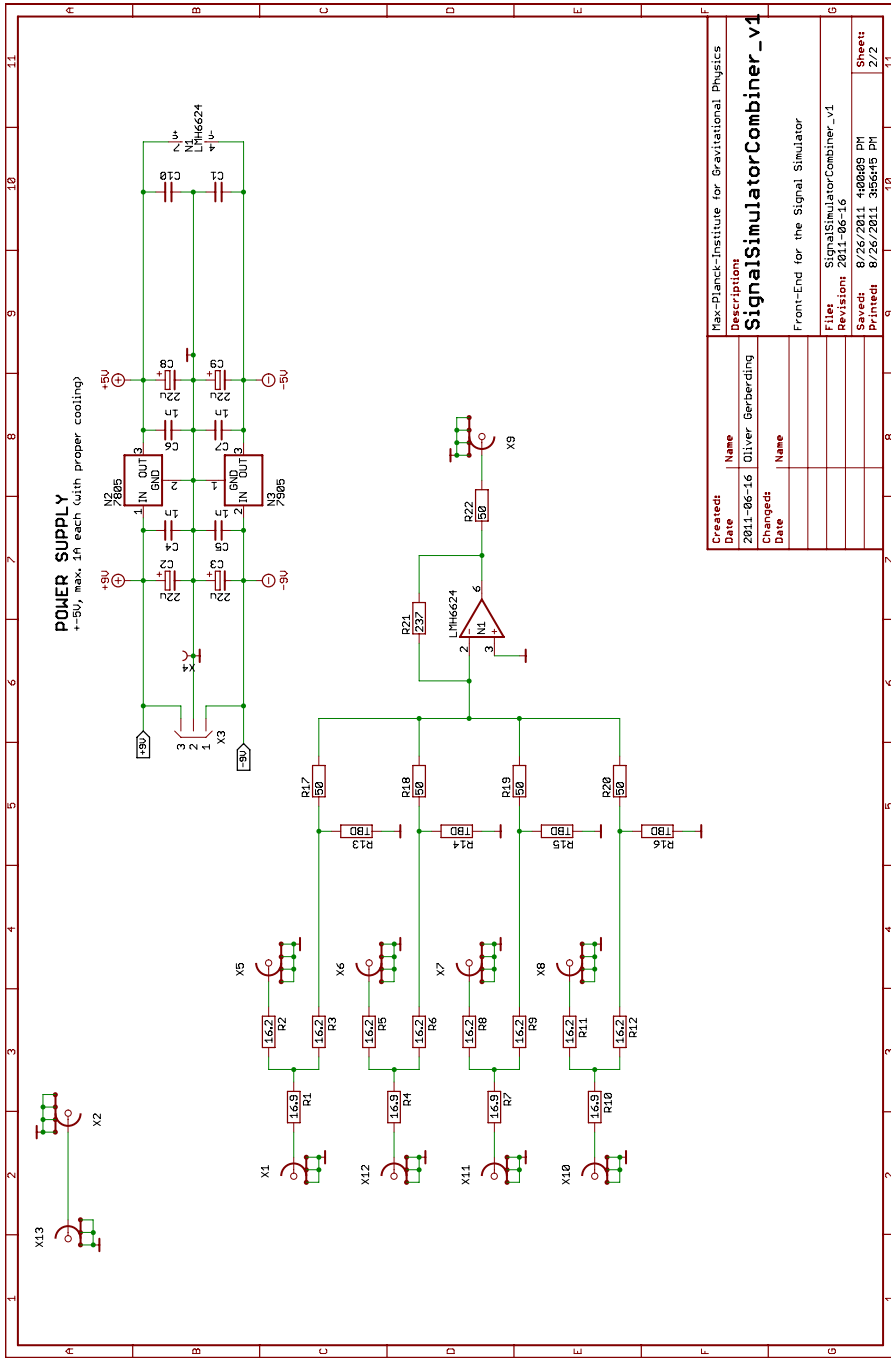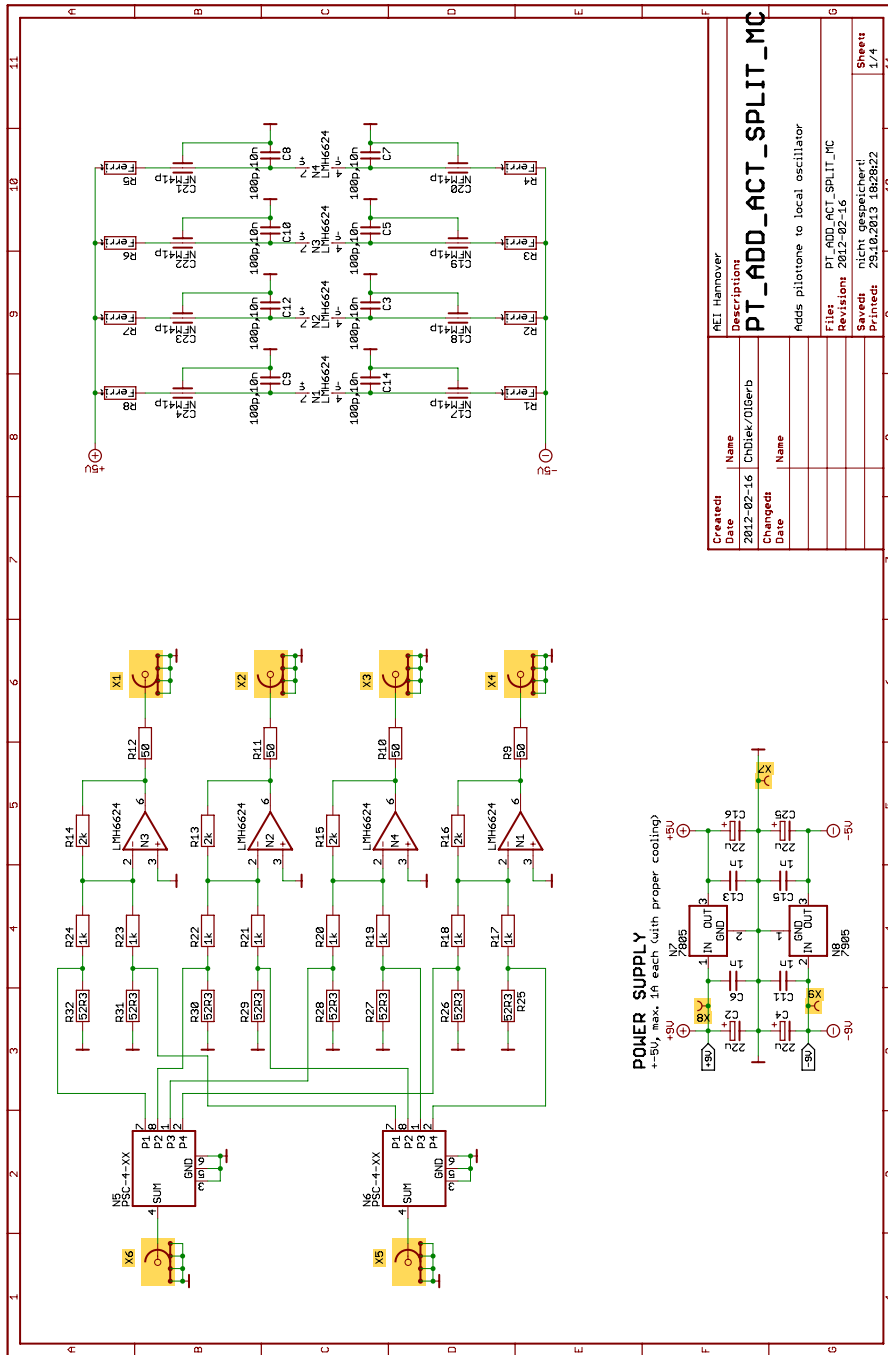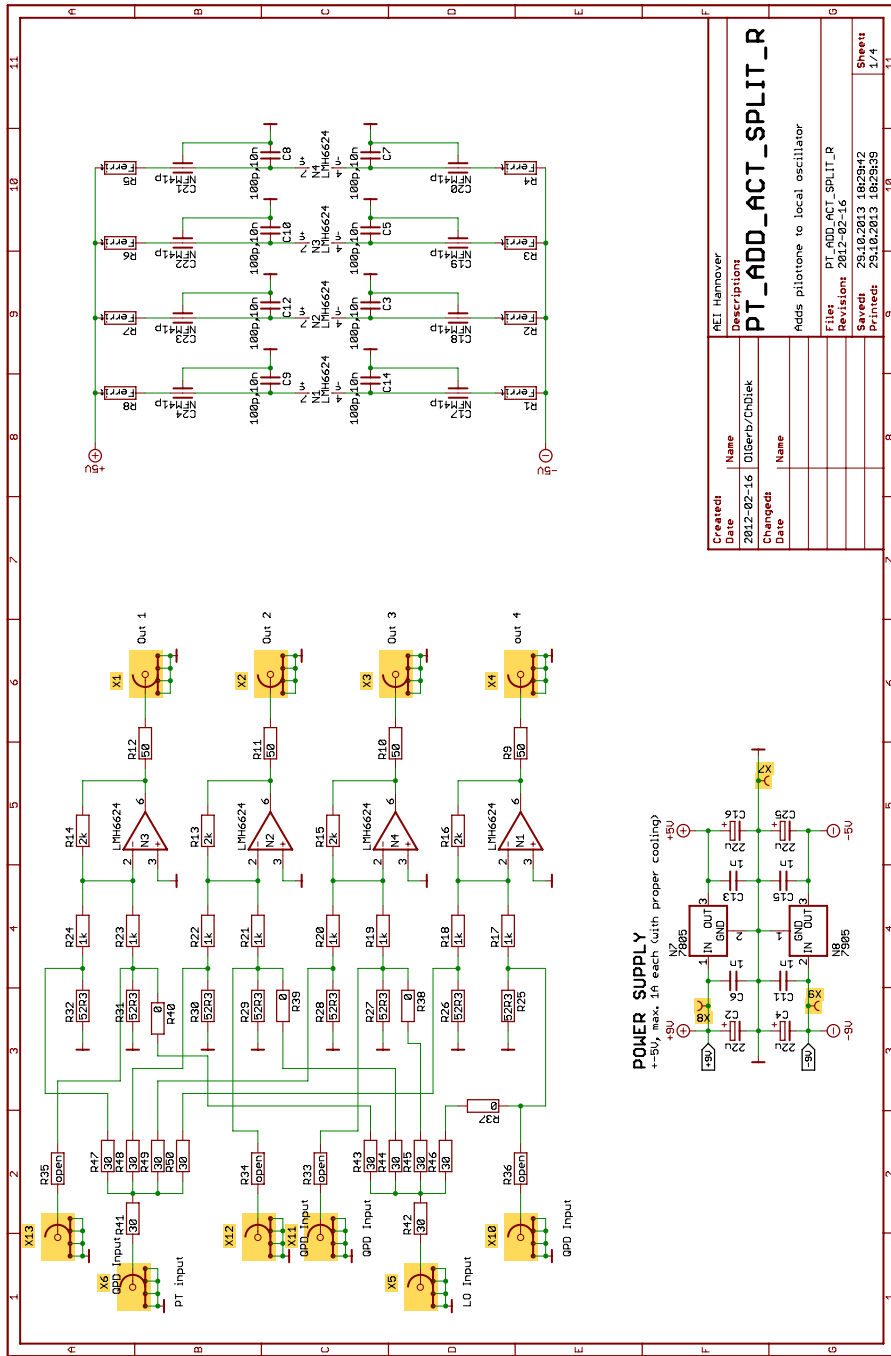**Figure B.2.:** Schematic of the variable-gain amplifier prototype used in Chapter 8.

217

Figure B.3.: Schematic of the signal combiner for the jitter corrected digital signal simulator used in Chapter 8.

Figure B.4.: Schematic of the analogue front-end for jitter correction used in Chapter 8 with active pilot tone adder and distribution via a power splitter.

Figure B.5.: Schematic of the analogue front-end for jitter correction used in Chapter 8 with active pilot tone adder and distribution via a resistive splitter.

Figure B.6.: Schematic of the analogue front-end for jitter correction used in Chapter 8 with power splitter ADP-2-1W used for pilot tone distribution and addition.
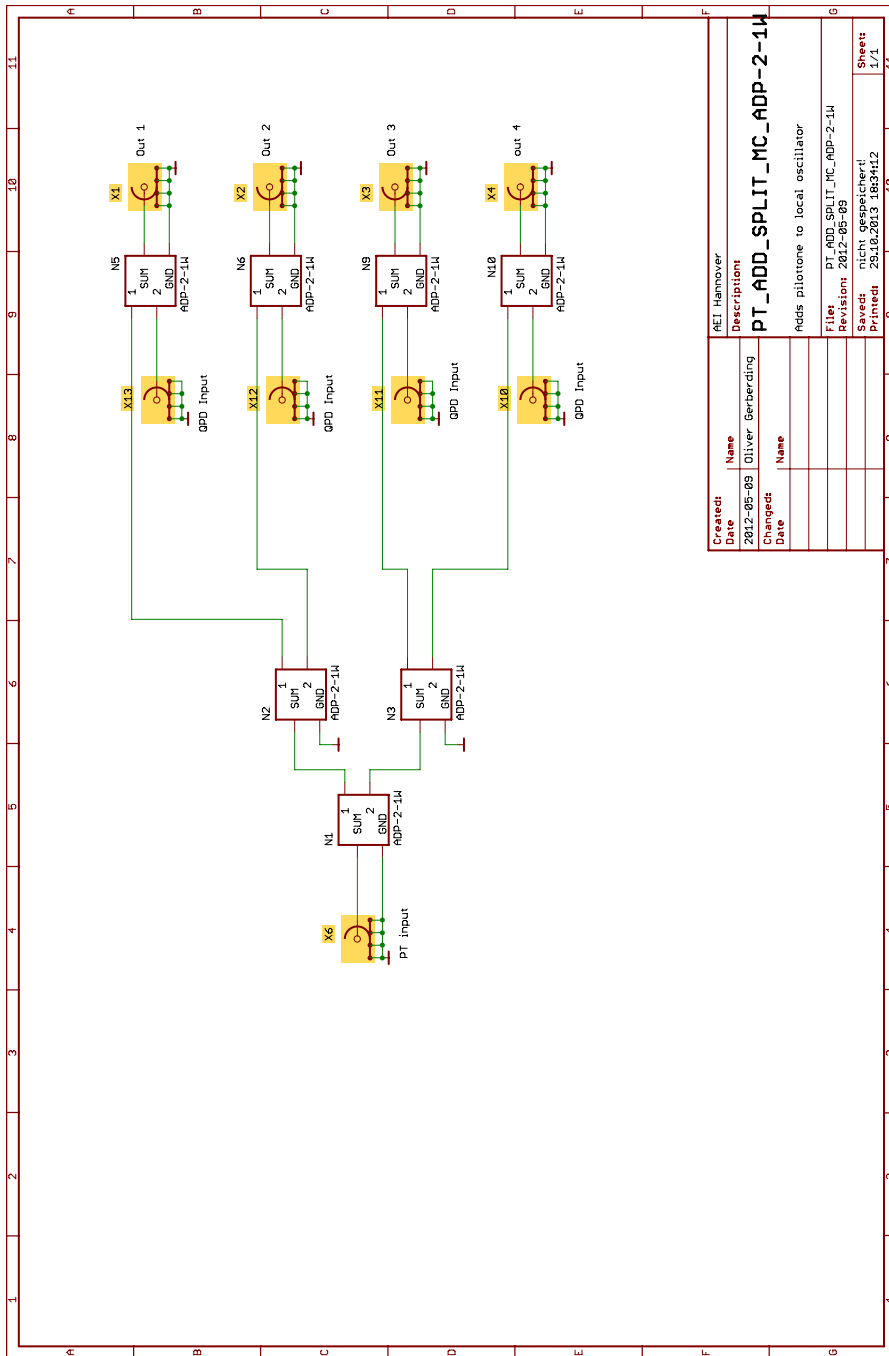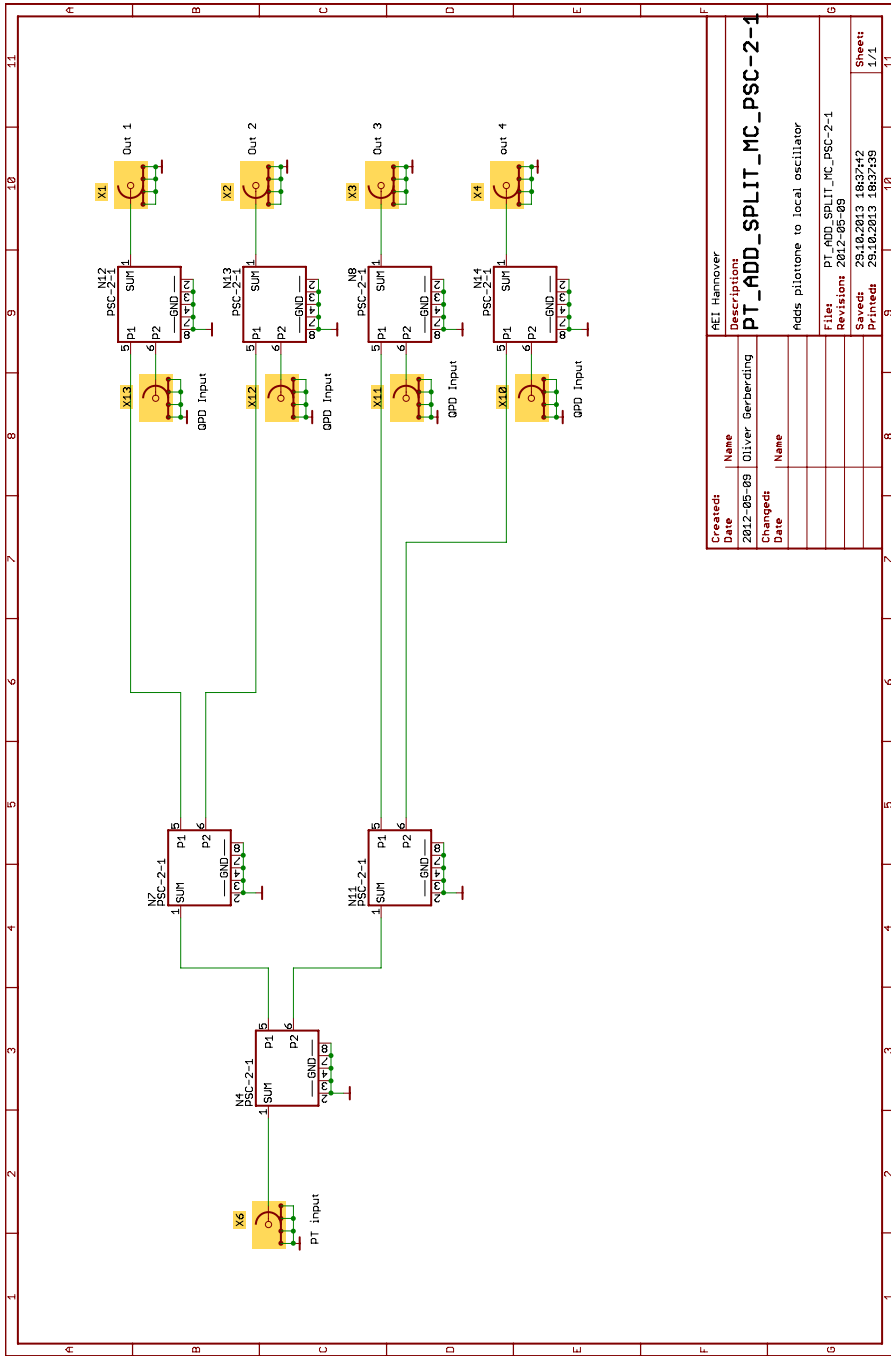
Figure B.7.: Schematic of the analogue front-end for jitter correction used in Chapter 8 with power splitter PSC-2-1 used for pilot tone distribution and addition.
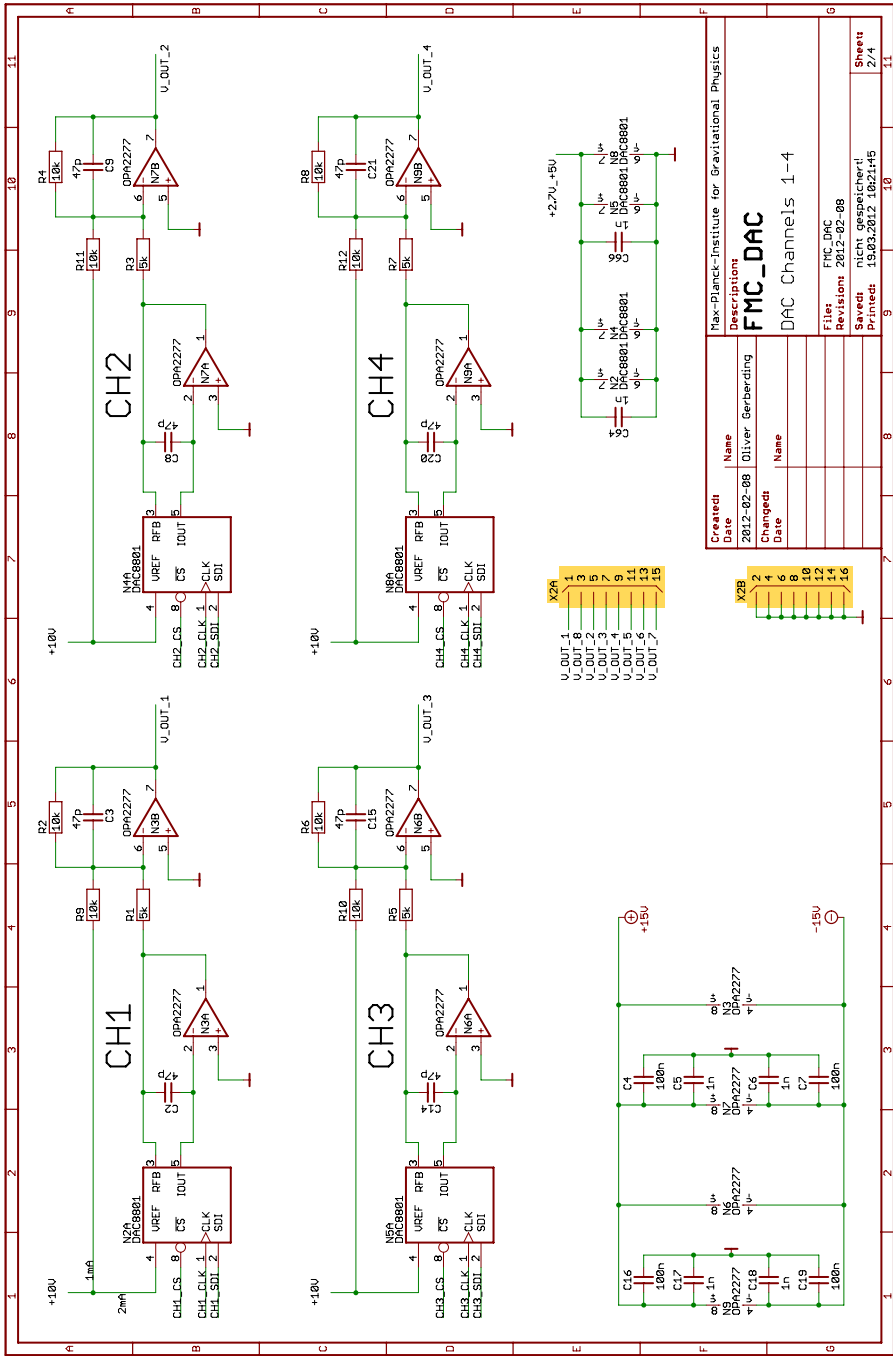
Figure B.8.: Schematic of the eight channel digital-to-analogue converter card designed for all set-ups using the ML605 (see Appendix A).

Figure B.9.: Schematic of the analogue phase shifter used in Chapter 12. The power supplies are not shown.

Figure B.10.: Schematic of the analogue front-end of the ADC module of the LISA phasemeter elegant breadboard (see Chapter 9).

Appendix C

FIR filters for the EBB

The modelled and measured transfer functions of the finite-impulse response (FIR) filters used in the elegant breadboard (EBB) model of the LISA phaseme-ter are shown here (see Chapter 9).

Figure C.1.: Shown are the simulated and modelled transfer function of the first FIR filter (FIR1) in the EBB decimation chain. The filter has 54 coefficients and decimates from 610 Hz to 102 Hz (see Section 9.3.3).

Figure C.2.: Shown are the simulated and modelled transfer function of the second FIR filter (FIR2) in the EBB decimation chain. The filter has 104 coefficients and decimates from 102 Hz to 10 Hz (see Section 9.3.3).
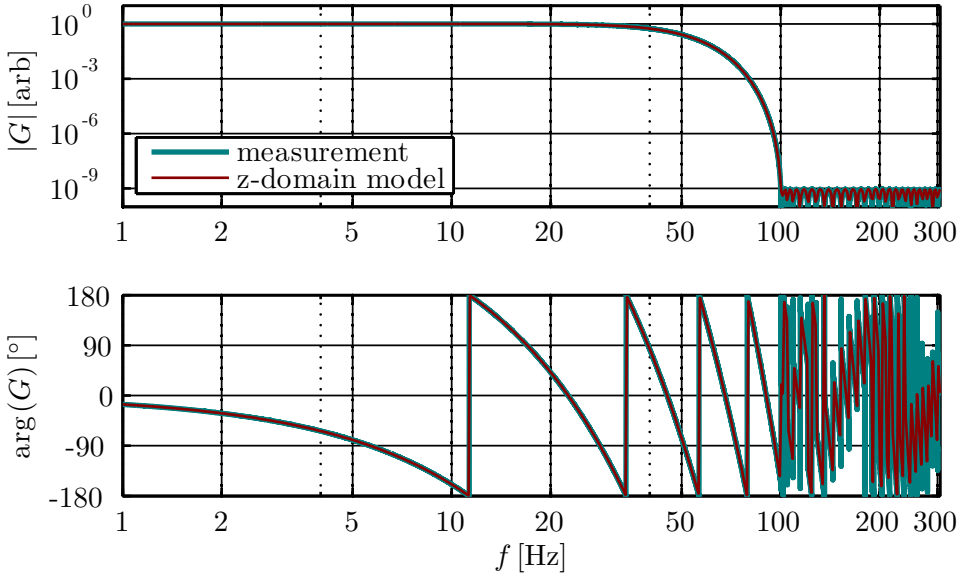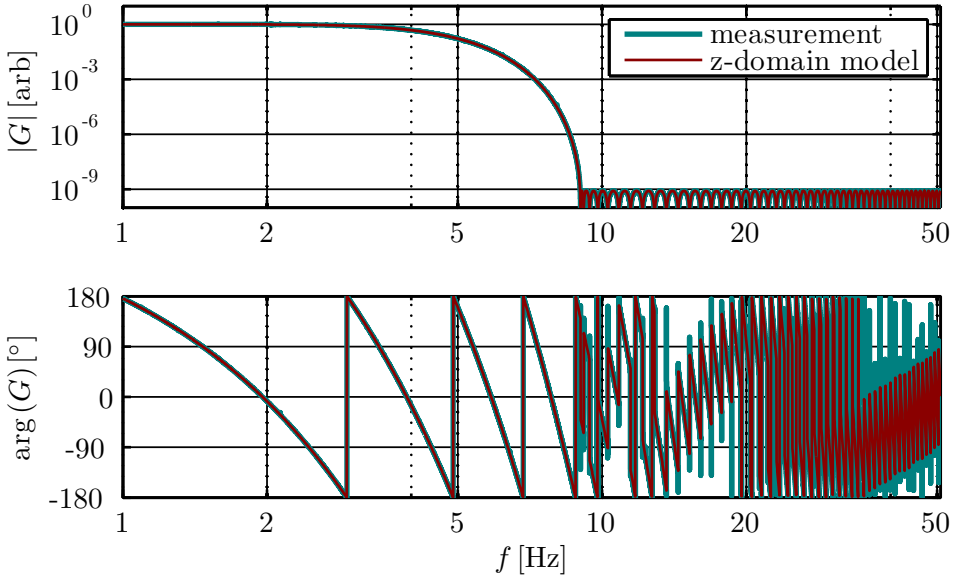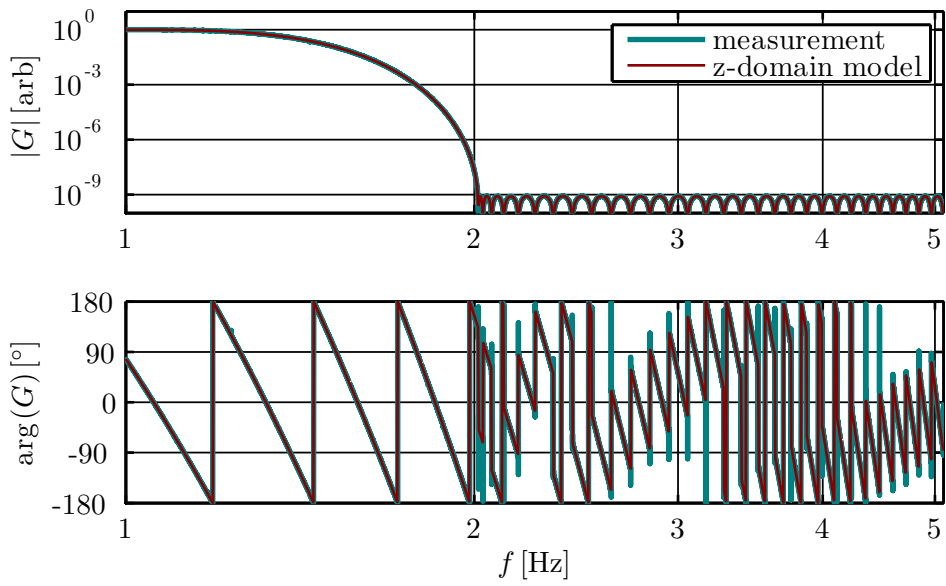
Figure C.3.: Shown are the simulated and modelled transfer function of the thrid FIR filter (FIR3) in the EBB decimation chain. The filter has 77 coefficients and decimates from 10 Hz to 3.4 Hz (see Section 9.3.3).

Appendix D

Additional LISA EBB performance measurements

Shown here are additional performance measurements of the elegant breadboard (EBB) model of the LISA phasemeter. All of these measurements were done using the means of temperature stabilisation described in Section 9.5.4.

The first two measurements were performed at low signal to noise ratio with a $C/N_0 = 75$ dBHz, and an input frequency of 15 MHz (Figure D.1) and 6.5 MHz (Figure D.2).

The following two measurements were performed at higher signal to noise ratio with a $C/N_0 = 95$ dBHz, and an input frequency of 15 MHz (Figure D.3) and 6.5 MHz (Figure D.4).

Figure D.1.: 15 MHz temp stab with ADC fans, pilot via 2-way, signal via 8 way, from DSS with frequency noise, Doppler, sidebands, low SNR
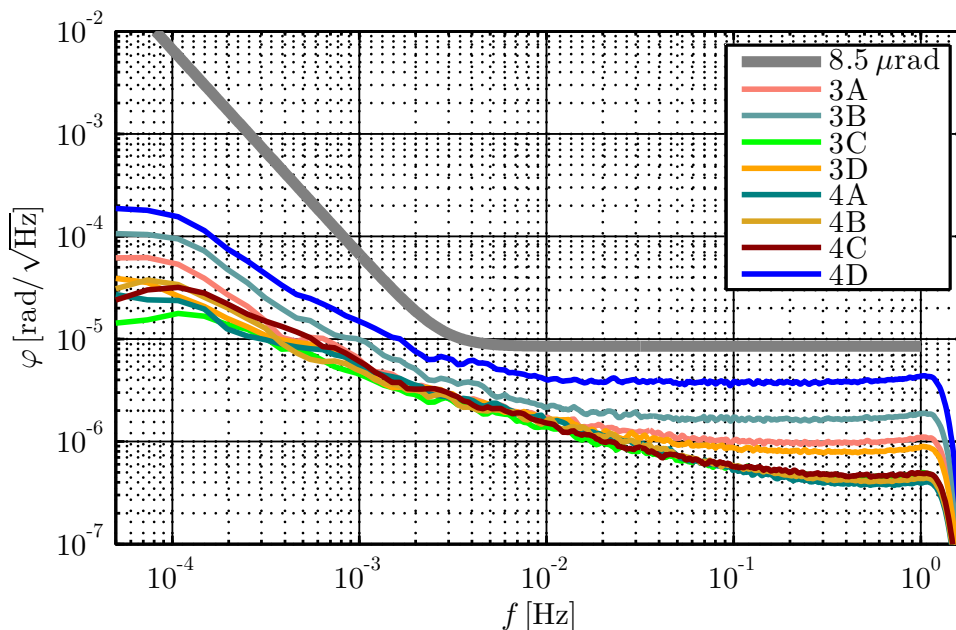
Figure D.2.: 6.5 MHz temp stab with ADC fans, pilot via 2-way, signal via 8 way, from DSS with frequency noise, Doppler, sidebands, low SNR
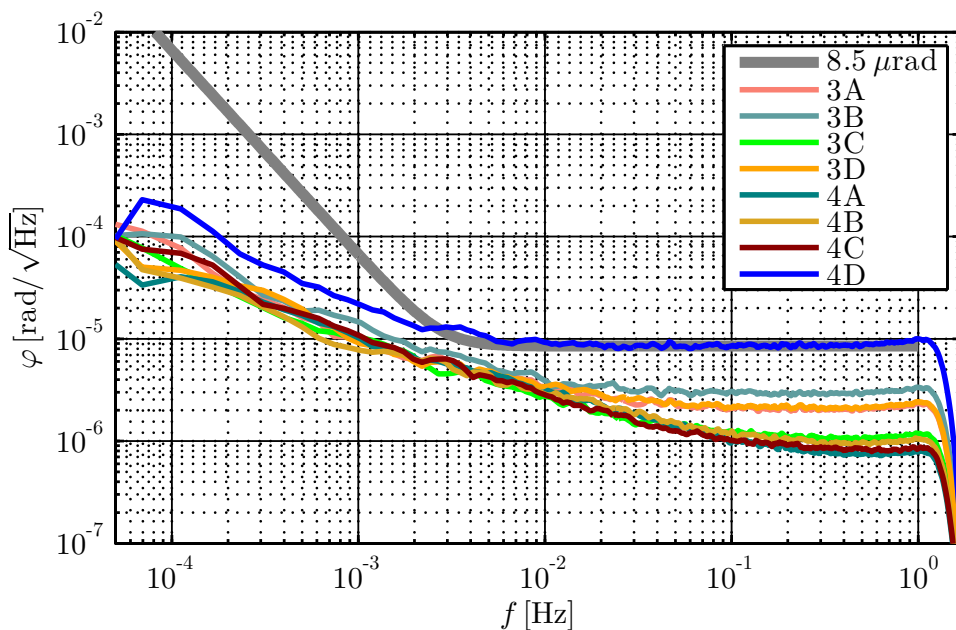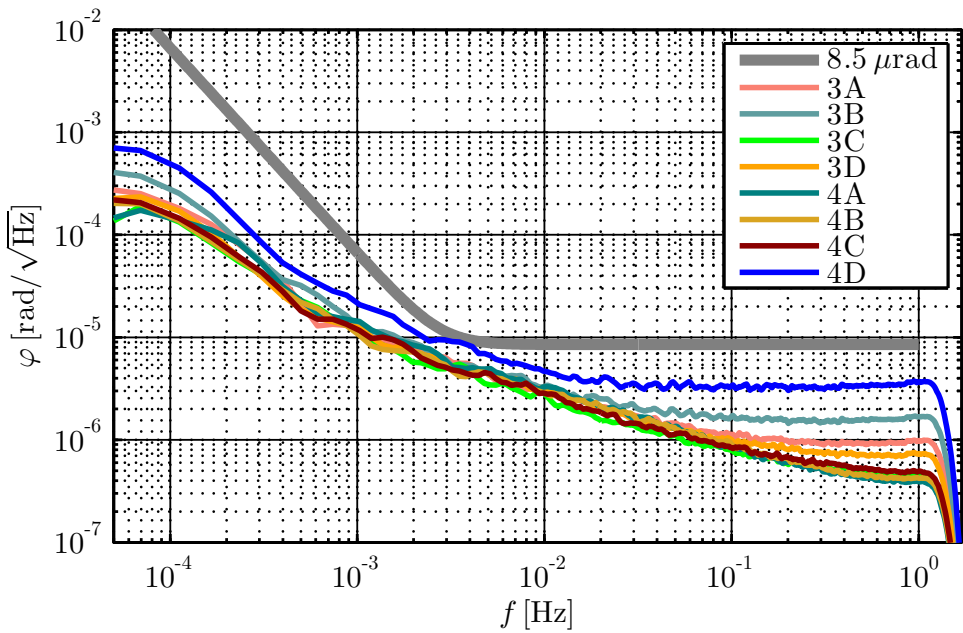
Figure D.3.: 15 MHz temp stab with ADC fans, pilot via 2-way, signal via 8 way, from DSS with frequency noise, Doppler, sidebands, high SNR
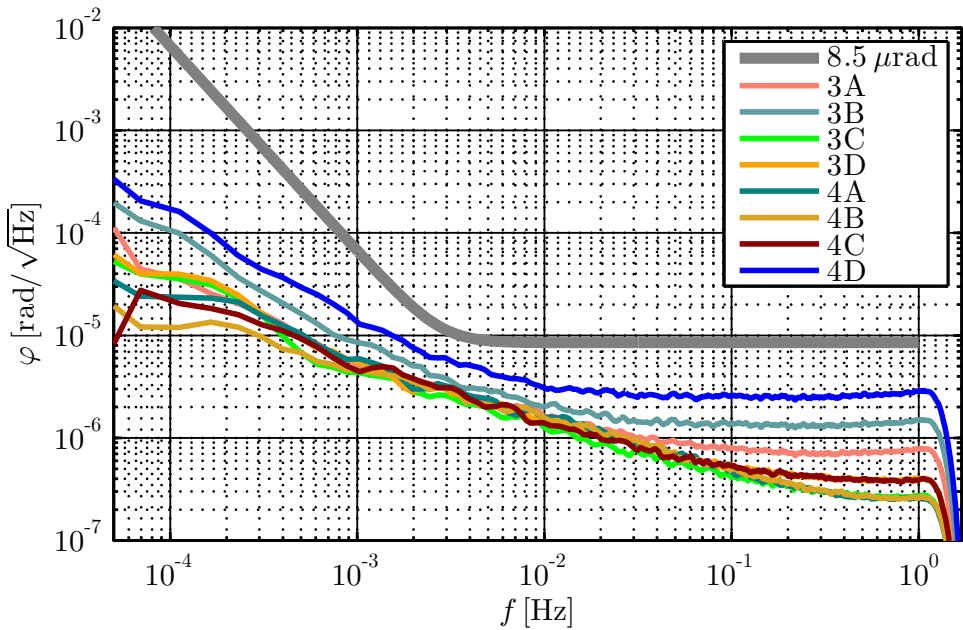
Figure D.4.: 6.5 MHz temp stab with ADC fans, pilot via 2-way, signal via 8 way, from DSS with frequency noise, Doppler, sidebands, high SNR

Appendix E

Optical set-up digital interferometry experiment

Shown here is a detailed overview of the optical layout used for the digitally enhanced heterodyne interferometry experiments described in Chapter 12. The set-up has been modified throughout the experiments. Initially, the heterodyne beat note was generated from a single beam that was split and frequency shifted by an acusto-optic modulator (not shown). Later on, two lasers were used, which were frequency offset-phase locked to each other. One of these lasers was stabilised to an Iodine reference, which significantly reduced laser frequency noise.
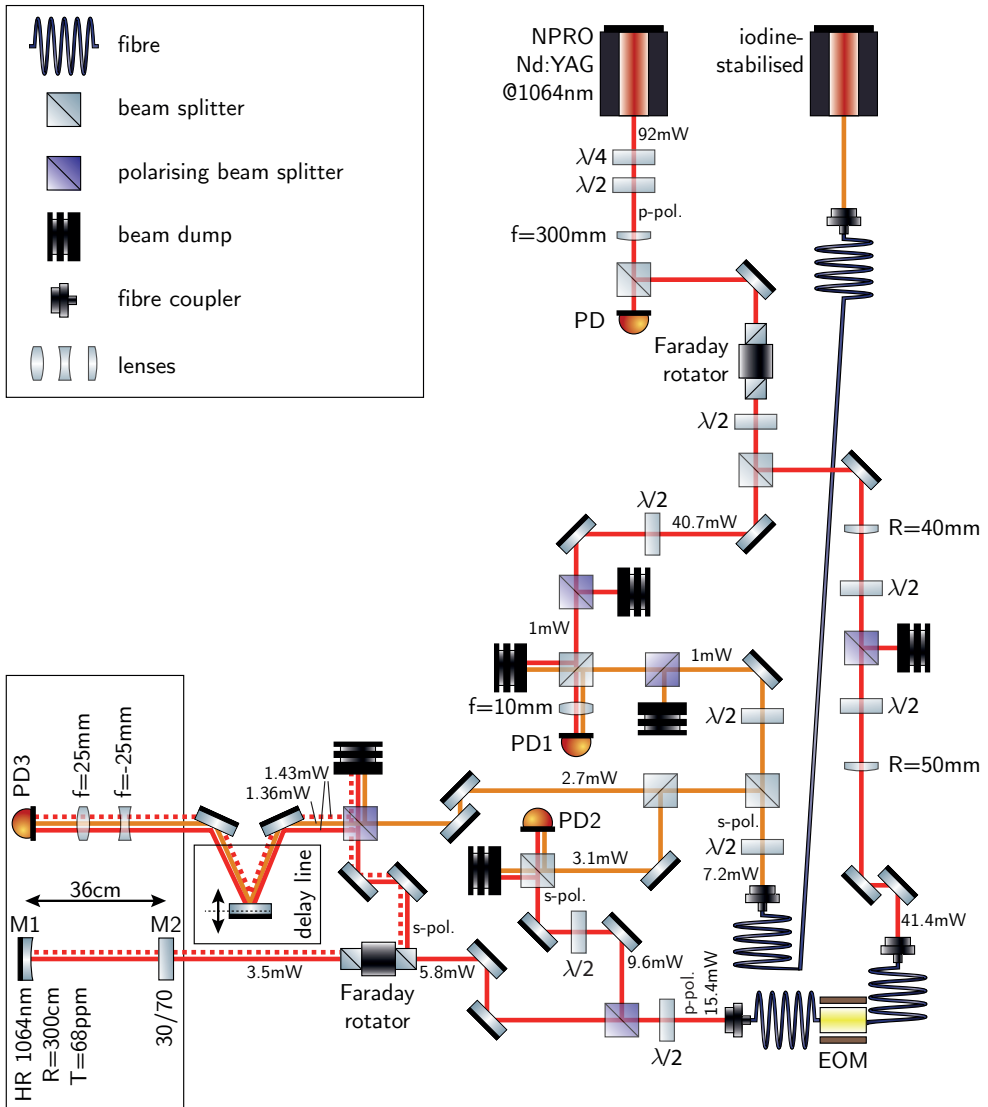
Figure E.1.: Detailed overview of the optical set-up used for the digital interferometry experiments, from K.-S. Isleif [145].

Bibliography

[1] B.S. Sathyaprakash and B. F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Reviews in Relativity*, 12(2):1, 2009.

[2] Albert Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–822, 1916.

[3] Joel M Weisberg, Joseph H Taylor, and Lee A Fowler. Gravitational waves from an orbiting pulsar. *Scientific American*, 245:74–82, 1981.

[4] Andreas Freise and Kenneth A. Strain. Interferometer Techniques for Gravitational-Wave Detection. *Living Reviews in Relativity*, 13(1):1, 2010.

[5] Gregory M Harry and the LIGO Scientific Collaboration. Advanced LIGO: the next generation of gravitational wave detectors. *Classical and Quantum Gravity*, 27(8):084006, 2010.

[6] F Acernese et al. Advanced Virgo baseline design. Technical report, VIRGO Collaboration, 2009.

[7] Kentaro Somiya. Detector configuration of KAGRA-the Japanese cryogenic gravitational-wave detector. *Classical and Quantum Gravity*, 29(12):124007, 2012.

[8] G Hobbs, A Archibald, Z Arzoumanian, D Backer, M Bailes, NDR Bhat, M Burgay, S Burke-Spolaor, D Champion, I Cognard, et al. The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Classical and Quantum Gravity*, 27(8):084013, 2010.

[9] Karsten Danzmann et al. The Gravitational Universe: Whitepaper for the ESA L2/L3 selection. arXiv preprint arXiv:1305.5720, 2013.

[10] LISA International Science Team. LISA Unveiling a hidden Universe. Assessment Study Report ESA/SRE(2011)3, ESA, February 2011. available at http://sci.esa.int.

[11] NGO science working team. NGO Revealing a hidden Universe: opening a new chapter of discovery. Assessment Study Report ESA/SRE(2011)19, ESA, December 2011. available at http://sci.esa.int.

[12] F Antonucci et al. From laboratory experiments to LISA Pathfinder: achieving LISA geodesic motion. *Classical and Quantum Gravity*, 28(9):094002, 2011.

[13] F Antonucci et al. The LISA Pathfinder mission. *Classical and Quantum Gravity*, 29(12):124014, 2012.

[14] Günter Seeber. *Satellite geodesy: foundations, methods, and applications.* Walter de Gruyter, 2003.

[15] B. D. Tapley, S. Bettadpur, M. Watkins, and C. Reigber. The gravity recovery and climate experiment: Mission overview and early results. *Geophysical Research Letters*, 31(9), 2004.

[16] MR Drinkwater, R Floberghagen, R Haagmans, D Muzi, and A Popescu. GOCE: ESA's first Earth Explorer Core mission. In *Earth Gravity Field from Space: from Sensors to Earth Sciences*, pages 419–432. Springer, 2003.

[17] Mark R Drinkwater, R Haagmans, D Muzi, A Popescu, R Floberghagen, M Kern, and M Fehringer. The GOCE gravity mission: ESA' s first core Earth explorer. In *Proceedings of the 3rd international GOCE user workshop*, pages 6–8, 2006.

[18] DN Wiese, RS Nerem, and FG Lemoine. Design considerations for a dedicated gravity recovery satellite mission consisting of two pairs of satellites. *Journal of Geodesy*, 86(2):81–98, 2012.

[19] Peter L Bender, John L Hall, Jun Ye, and William M Klipstein. Satellite-satellite laser links for future gravity missions. *Space science reviews*, 108(1-2):377–384, 2003.

[20] R. Pierce, J. Leitch, M. Stephens, P. Bender, and R. Nerem. Intersatellite range monitoring using optical interferometry. *Appl. Opt.*, 47(27):5007–5019, Sep 2008.

[21] I. Coddingtion, W. C. Swann, L. Nenadovic, and N. R. Newbury. Rapid and precise absolute distance measurements at long range. *Nature Photonics*, 3:351, 2009.

[22] P Silvestrin, M Aguirre, L Massotti, B Leone, S Cesare, M Kern, and R Haagmans. The Future of the Satellite Gravimetry After the GOCE Mission. In *Geodesy for Planet Earth*, pages 223–230. Springer, 2012.

[23] M Tröbs, P Weßels, C Fallnich, M Bode, I Freitag, S Skorupka, G Heinzel, and K Danzmann. Laser development for LISA. *Classical and Quantum Gravity*, 23(8):S151, 2006.

[24] M Tröbs, S Barke, J Möbius, M Engelbrecht, D Kracht, L d'Arcio, G Heinzel, and K Danzmann. Lasers for LISA: Overview and phase characteristics. *Journal of Physics: Conference Series*, 154(1):012016, 2009.

[25] G. Wanner. *Comples optical ssystem in space: numerical modelling of the heterodyne interferometer of LISA Pathfinder and LISA*. PhD thesis, Leibniz University Hannover, Albert-Einstein Institute Hannover, 2010.

[26] Gudrun Wanner, Gerhard Heinzel, Evgenia Kochkina, Christoph Mahrdt, Benjamin S. Sheard, Sönke Schuster, and Karsten Danzmann. Methods for simulating the readout of lengths and angles in laser interferometers with Gaussian beams. *Optics Communications*, 285(24):4831 – 4839, 2012.

[27] Vitali Müller. Simulations for LISA & GRACE-Follow-On: Satellite constellations at lagrangian points for LISA-like missions interferometer simulations for the GRACE Follow-On mission. Master's thesis, Leibniz Universität Hannover, 2013.

[28] Evgenia Kochkina. *Stigmatic and Astigmatic Gaussian Beams in Fundamental Mode*. PhD thesis, Leibniz Universität Hannover, 2013.

[29] Evgenia Kochkina et al. Simulating and Optimizing Laser Interferometers. In G. Auger, P. Binetruy, and E. Plagnol, editors, *ASP Conference Series; 9th LISA Symposium*, volume 467, page 291, 2012.

[30] Evgenia Kochkina, Gudrun Wanner, Dennis Schmelzer, Michael Tröbs, and Gerhard Heinzel. Modeling of the general astigmatic Gaussian beam and its propagation through 3D optical systems. *Appl. Opt.*, 52(24):6030–6040, Aug 2013.

[31] Gerhard Heinzel, Claus Braxmaier, Roland Schilling, Albrecht Rüdiger, David Robertson, M te Plate, Vinzenz Wand, K Arai, Ulrich Johann, and Karsten Danzmann. Interferometry for the LISA technology package (LTP) aboard SMART-2. *Classical and Quantum Gravity*, 20(10):S153, 2003.

[32] Matthew Pitkin, Stuart Reid, Sheila Rowan, and James Hough. Gravitational Wave Detection by Interferometry (Ground and Space). *Living Reviews in Relativity*, 14(5):1, 2011.

[33] H Audley et al. The LISA Pathfinder interferometry - hardware and system testing. *Classical and Quantum Gravity*, 28(9):094003, 2011.

[34] K Dahl, G Heinzel, B Willke, K A Strain, S Goßler, and K Danzmann. Suspension platform interferometer for the AEI 10 m prototype: concept, design and optical layout. *Classical and Quantum Gravity*, 29(9):095024, 2012.

[35] V. Wand. *Interferometry at low frequencies: Optical phase measurement for LISA and LISA Pathfinder*. PhD thesis, Leibniz Universit(a)t Hannover, 2007.

[36] D Shaddock, B Ware, P Halverson, R E Spero, and B Klipstein. Overview of the LISA Phasemeter. *AIP Conf. Proc.*, 873:689–696, 2006.

[37] Floyd M. Gardner. *Phaselock Techniques*. (New Jersey: John Wiley & Sons, Inc.), 3rd edition, 2005.

[38] Muthu Jeganathan and Serge Dubovitsky. Demonstration of nanometer-level active metrology for long-range interferometric displacement measurements. In *Interferometry in Optical Astronomy*, volume 4006, pages 838–846, 2000.

[39] 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, 2012.

[40] C Diekmann, F Steier, B Sheard, G Heinzel, and K Danzmann. Analog phase lock between two lasers at LISA power levels. *Journal of Physics: Conference Series*, 154:012020, 2009.

[41] P. W. McNamara. Weak-light phase locking for LISA. *Classical and Quantum Gravity*, 22:S243, 2005.

[42] Brent Ware and Glenn de Vine, Jet Propulsion Laboratory, California Institute of Technology. Personal communication.

[43] Joachim Kullmann. *Development of a digital phase measuring system with microradian precision for LISA*. PhD thesis, Leibniz Universität Hannover, 2012.

[44] J. J. Esteban. *Laser ranging and data communication for the laser interferometer space antenna*. PhD thesis, Leibniz University Hannover, 2012.

[45] Patrick Kwee. *Laser Characterizatioon and Stabilization for Precision Interferometry*. PhD thesis, Leibniz Universität Hannover, 2010.

[46] Gerald Hechenblaikner. Common mode noise rejection properties of amplitude and phase noise in a heterodyne interferometer. *J. Opt. Soc. Am. A*, 30(5):941–947, May 2013.

[47] R. Fleddermann. *Interferometry for a space-based gravitational wave observatory Reciprocity of an optical fiber*. PhD thesis, Leibniz Universität Hannover, 2012.

[48] C. Diekmann. *Development of core elements for the LISA optical bench - Electro-optical measurement ssystem and test devices*. PhD thesis, Leibniz Universität Hannover, 2013.

[49] Gerald Hechenblaikner. Measurement of the absolute wavefront curvature radius in a heterodyne interferometer. *J. Opt. Soc. Am. A*, 27(9):2078–2083, Sep 2010.

[50] H.I. Campbell, A.H. Greenaway, H.I. Campbell, and A.H. Greenaway. Wavefront sensing: From historical roots to the state-of-the-art. *EAS Publications Series*, 22:165–185, 1 2006.

[51] Thilo Schuldt, Martin Gohlke, Dennis Weise, Ulrich Johann, Achim Peters, and Claus Braxmaier. Picometer and nanoradian optical heterodyne interferometry for translation and tilt metrology of the LISA gravitational reference sensor. *Classical and Quantum Gravity*, 26(8):085008, 2009.

[52] Gerhard Heinzel, Vinzenz Wand, A Garcia, Oliver Jennrich, Claus Braxmaier, David Robertson, Kevin Middleton, David Hoyland, Albrecht Rüdiger, Roland Schilling, et al. The LTP interferometer and phasemeter. *Classical and Quantum Gravity*, 21(5):S581, 2004.

[53] Euan Morrison, Brian J. Meers, David I. Robertson, and Henry Ward. Automatic alignment of optical interferometers. *Appl. Opt.*, 33(22):5041–5049, Aug 1994.

[54] F Ales, P. Gath, U. Johann, and C Braxmaier. Modeling and simulation of a laser ranging interferometer acquisition and guidance algorithm. In *AIAA Guidance, Navigation, and Control (GNC) Conference*. American Institute of Aeronautics and Astronautics, 2013.

[55] W. M. Folkner, G. de Vine, W. M. Kleipstein, K. McKenzie, D. A. Shaddock, R. E. Spero, R. Thompson, D. Wuchenich, N. Yu, M. Stephens, J. Leitch, M. Davis, J. deCino, C. Pace, and R. Pierce. Laser frequency stabilization for GRACE-II. *Earth Science*, 2010.

[56] W. M. Folkner, G. de Vine, W. M. Kleipstein, K. McKenzie, R. E. Spero, R. Thompson, N. Yu, M. Stephens, J. Leitch, R. Pierce, T. T.-Y. Lam, and D. A. Shaddock. Laser frequency stabilization for GRACE-2. *Earth Science*, 2011.

[57] M Tröbs, L dArcio, S Barke, J Bogenstahl, I Bykov, M Dehne, C Diekmann, ED Fitzsimons, R Fleddermann, O Gerberding, et al. Testing the LISA optical bench. In *Proceedings of the 9th International Conference on Space Optics, Ajaccio*, 2012.

[58] S Mitryk, V Wand, and G Mueller. Verification of time-delay interferometry techniques using the University of Florida LISA interferometry simulator. *Classical and Quantum Gravity*, 27(8):084012, 2010.

[59] Massimo Tinto and Sanjeev V. Dhurandhar. Time-delay interferometry. *Living Reviews in Relativity*, 8(4):1, 2005.

[60] Glenn de Vine, Brent Ware, Kirk McKenzie, Robert E. Spero, William M. Klipstein, and Daniel A. Shaddock. Experimental Demonstration of Time-Delay Interferometry for the Laser Interferometer Space Antenna. *Phys. Rev. Lett.*, 104:211103, May 2010.

[61] J J Esteban, I Bykov, A F Garcia Marin, G Heinzel, and K Danzmann. Optical ranging and data tranfer development for LISA. *Journal of Physics: Conference Series*, 154:012025, 2009.

[62] J J Esteban, A F Garcia Marin, J Einholz, A M Peinado, I Bykov, G Heinzel, and K Danzmann. Ranging and phase measurment for LISA. *Journal of Physics: Conference Series*, 228:012045, 2010.

[63] J J Esteban, A F Garcia Marin, S Barke, A M Peinado, F Guzman Cervantes, I Bykov, G Heinzel, and K Danzmann. Experimental demonstration of weak-light laser ranging and data communication for LISA. *Optics Express*, 19(17):15937, 2011.

[64] A. Sutton, K. McKenzie, B. Ware, and D. A. Shaddock. Laser ranging and communications for LISA. *Optics Express*, 18:20759, 2010.

[65] A. J. Sutton, K. McKenzie, B. Ware, G. de Vine, R. E. Spero, W. Klipstein, and D. A. Shaddock. Improved optical ranging for space based gravitational wave detection. *Classical and Quantum Gravity*, 30:075008, 2013.

[66] Massimo Tinto, Michele Vallisneri, and J. W. Armstrong. Time-delay interferometric ranging for space-borne gravitational-wave detectors. *Phys. Rev. D*, 71:041101, Feb 2005.

[67] Kirk McKenzie, Robert E. Spero, and Daniel A. Shaddock. Performance of arm locking in LISA. *Phys. Rev. D*, 80:102003, Nov 2009.

[68] Benjamin S. Sheard, Malcolm B. Gray, David E. McClelland, and Daniel A. Shaddock. Laser frequency stabilization by locking to a LISA arm. *Physics Letters A*, 320(1):9 – 21, 2003.

[69] B S Sheard, M B Gray, D A Shaddock, and D E McClelland. Laser frequency noise suppression by arm-locking in LISA: progress towards a bench-top demonstration. *Classical and Quantum Gravity*, 22(10):S221, 2005.

244

[70] J. I. Thorpe, P. Maghami, and J. Livas. Time domain simulations of arm locking in LISA. *Phys. Rev. D*, 83:122002, Jun 2011.

[71] Y Yu, S J Mitryk, and G Mueller. Experimental validation of dual/modified dual arm locking for LISA. *Classical and Quantum Gravity*, 28:094009, 2011.

[72] Andrew Sutton and Daniel A. Shaddock. Laser frequency stabilization by dual arm locking for LISA. *Phys. Rev. D*, 78:082001, Oct 2008.

[73] Antonio F García Marín et al. Phase locking to a LISA arm: first results on a hardware model. *Classical and Quantum Gravity*, 22(10):S235, 2005.

[74] Simon Barke. Inter-Spacecraft Clock Transfer Phase Stability for LISA. Diploma thesis, Leibniz Universität Hannover, jun 2008. Diploma Thesis.

[75] G Heinzel, J J Esteban, S Barke, M Otto, Y Wang, A F Garcia Marin, and K Danzmann. Auxiliary functions of the LISA laser link: ranging, clock noise transfer and data communication. *Classical and Quantum Gravity*, 28:094008, 2011.

[76] Shubhashish Datta, Abhay Joshi, and Jim Rue. Large-area InGaAs quad photoreceiver for Laser Interferometry Space Antenna. volume 7817, pages 78170F–78170F–12, 2010.

[77] F Guzman Cervantes, J Livas, R Silverberg, E Buchanan, and R Stebbins. Characterization of photoreceivers for LISA. *Classical and Quantum Gravity*, 28(9):094010, 2011.

[78] Abhay Joshi, Shubhashish Datta, Jim Rue, Jeffrey Livas, Robert Silverberg, and Felipe Guzman Cervantes. Ultra-low noise large-area InGaAs quad photoreceiver with low crosstalk for laser interferometry space antenna. volume 8453, pages 84532G–84532G–10, 2012.

[79] Vishay. Thin film Mini-MELF resistors SMM0204 data sheet. http://www.vishay.com/docs/20004/smm0204.pdf, 2011.

[80] Rodolphe Boudot and Enrico Rubiola. Phase noise in RF and microwave amplifiers. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 59(12):2613–2624, 2012.

[81] Paul Horowitz, Winfield Hill, and Thomas C Hayes. *The art of electronics*, volume 2. Cambridge university press Cambridge, 1989.

[82] Steven W Smith et al. *The scientist and engineer's guide to digital signal processing*. California Technical Pub. San Diego, 1997. http://www.dspguide.com/.

[83] C. M. Burnett. Development of an Ultra-precise Digital Phasemeter for the LISA Gravitational Wave Detector. Master's thesis, Lulea University of Technology Department of Computer Science, Electrical and Space Engieering, 2010.

[84] Stanley P Lipshitz, Robert A Wannamaker, and John Vanderkooy. Quantization and dither: A theoretical survey. *Journal of the audio engineering society*, 40(5):355–375, 1992.

[85] Peter J Ashenden. *Digital design (VHDL): an embedded systems approach using VHDL*. Morgan Kaufmann Publishers Inc., 2007.

[86] B. Widrow and I. Kollar. *Quantization Noise*. New York: Cambridge University Press, 2009.

[87] Mark R Santoro, Gary Bewick, and Mark A Horowitz. Rounding algorithms for ieee multipliers. In *Computer Arithmetic, 1989., Proceedings of 9th Symposium on*, pages 176–183. IEEE, 1989.

[88] Huirae Cho and Jin-up Kim. An efficient method for compensating the truncation dc-error in a multi-stage digital filter. In *Communications and Information Technology, 2004. ISCIT 2004. IEEE International Symposium on*, volume 1, pages 190–193. IEEE, 2004.

[89] Oliver Gerberding, Benjamin Sheard, Iouri Bykov, Joachim Kullmann, Juan Jose Esteban Delgado, Karsten Danzmann, and Gerhard Heinzel. Phasemeter core for intersatellite laser heterodyne interferometry: modelling, simulations and experiments. *Classical and Quantum Gravity*, 30(23):235029, 2013.

[90] Robert Bogdan Staszewski and Poras T Balsara. Phase-domain all-digital phase-locked loop. *Circuits and Systems II: Express Briefs, IEEE Transactions on*, 52(3):159–163, 2005.

[91] Stefan Mendel and Christian Vogel. A z-domain model and analysis of phase-domain all-digital phase-locked loops. In *Norchip, 2007*, pages 1–6. IEEE, 2007.

[92] Marina Dehne, Michael Tröbs, Gerhard Heinzel, and Karsten Danzmann. Verification of polarising optics for the LISA optical bench. *Opt. Express*, 20(25):27273–27287, Dec 2012.

[93] E.B. Hogenauer. An economical class of digital filters for decimation and interpolation. *IEEE ASSP Magazine*, 29(2):155, 1981.

[94] Vesa Lehtinen and Markku Renfors. Truncation noise analysis of noise shaping DSP systems with application to CIC decimators. In *Proceedings of the European Signal Processing Conference*, pages 207–210, 2002.

[95] R. Spero, B. Bachman, G. de Vine, J. Dickson, W. M. Klipstein, T. Ozawa, K. McKenzie, D. A. Shaddock, D. Robison, A. Sutton, and B. Ware. Progress in Interferometry for LISA at JPL. *Classical and Quantum Gravity*, 28:094007, 2011.

[96] G. Ascheid and H. Meyr. Cycle slips in phase-locked loops: A tutorial survey. *IEEE Transactions on Communications*, Com-30(10):2228, 1982.

[97] A.J. Viterbi. Phase-locked loop dynamics in the presence of noise by fokker-planck techniques. *Proceedings of the IEEE*, 51(12):1737–1753, 1963.

[98] Pavel Hasan. Letter: Cycle slip statistics of a first-order phase-locked loop in the presence of additive noise and cochannel interference. *European Transactions on Telecommunications*, 5(5):619–624, 1994.

[99] G. J. Dick, M. Tu, D. Strekalov, K. Birnbaum, and N. Yu. Optimal phase lock at femtowatt power levels for coherent optical deep-space transponder. *IPN Progress Report*, 42-175:1–17, 2008.

[100] Reza Danesfahani, Mohammad Moghaddasi, and Mahmood Mahlouji. Enhanced acquisition and tracking in all digital phase-locked loops. *Circuits, Systems & Signal Processing*, 27(4):537–552, 2008.

[101] National Instruments. Manual, LabVIEW User. www.ni.com/labview/?, 1998.

[102] Jürgen Reichardt and Bernd Schwarz. *VHDL-Synthese: Entwurf digitaler Schaltungen und Systeme*. Oldenbourg Verlag, 2012.

[103] William T Scherer, Thomas A Pomroy, and Douglas N Fuller. The triangular density to approximate the normal density: decision rules-of-thumb. *Reliability Engineering & System Safety*, 82(3):331–341, 2003.

[104] 4DSP. FMC10x user manual. http://www.4dsp.com/FMC107.php, 2011.

[105] Daniel Schütze, Gunnar Stede, Vitali Müller, Oliver Gerberding, Christoph Mahrdt, Benjamin Sheard, Gerhard Heinzel, and Karsten Danzmann. LISA-like Laser Ranging for GRACE Follow-on. In *Astronomical Society of the Pacific Conference Series*, volume 467, page 285, 2013.

[106] G. Stede. Interferometrische Charakterisierung von Retroreflektoren für satellitengestützte Erdschwerefeldbestimmung. Diploma thesis, Leibniz Universität Hannover, 2011.

[107] Physik Instrumente (PI). M-824 Compact 6-Axis-Positioning System. http://www.physikinstrumente.com/en/pdf/M824_Datasheet.pdf.

[108] Marina Dehne. *Construction and noise behaviour of ultra-stable optical systems for space intererometers*. PhD thesis, Leibniz Universität Hannover, 2012.

[109] Texas Instruments. ADS6445 Datasheet. `http://www.ti.com/product/ads6445`, 2009.

[110] Texas Instruments. ADS64XX EVM User's Guide. `http://www.ti.com/product/ads6445`, 2007.

[111] 4DSP. FMC122 user manual. `http://www.4dsp.com/FMC122.php`, 2010.

[112] 4DSP. FMC204 user manual. `http://www.4dsp.com/FMC204.php`, 2010.

[113] NXP. PAB64-03 silicon PIN diode data sheet. `http://www.nxp.com/documents/data_sheet/BAP64-03.pdf`, 2004.

[114] Texas Instruments. LMH6702 1.7 GHz, Ultra Low Distortion, Wideband Op Amp. `www.ti.com/product/lmh6702`, 2005.

[115] Texas Instruments. LMH6624/LMH6626 Single/Dual Ultra Low Noise Wideband Operational Amplifier. `www.ti.com/product/lmh6624`, 2005.

[116] Mini-Circuits. Datasheet Power Splitter/Combiner ADP-2-1W. `www.minicircuits.com/pdfs/ADP-2-1W.pdf`.

[117] Mini-Circuits. Datasheet Power Splitter/Combiner PSC-2-1. `www.minicircuits.com/pdfs/PSC-2-1.pdf`.

[118] M Tröbs, L d'Arcio, S Barke, J Bogenstahl, C Diekmann, ED Fitzsimons, O Gerberding, J Hennig, FG Hey, H Hogenhuis, et al. LISA Optical Bench Testing. In *Astronomical Society of the Pacific Conference Series*, volume 467, page 233, 2013.

[119] Andreas Schreiber. Subpikometer-stabiles Frequenzkonvertierungssystem für die optische Bank von LISA. Bachelor thesis, Leibniz Universität Hannover, August 2011. Bachelorarbeit.

[120] 4DSP. FMC116/FMC112 user manual. `http://www.4dsp.com/FMC116.php`, 2011.

[121] Coherent. Datasheet Mephisto/Mephisto S: Ultra-Narrow Linewidth CW DPSS Laser . `http://www.coherent.com/downloads/Mephisto_DS_1013revA_2.pdf`, 2013.

[122] N. Brause. Bau eines Phasenmeters. Diploma thesis, Leibniz Universität Hannover, 2011.

[123] J Bogenstahl, L Cunningham, E D Fitzsimons, J Hough, C J Killow, M Perreur-Lloyd, D Robertson, S Rowan, and H Ward. LTP fibre injector qualification and status. *Journal of Physics: Conference Series*, 154(1):012011, 2009.

[124] B. S. Sheard, G. Heinzel, and K. Danzmann. LISA long-arm interferometry: an alternative frequency pre-stabilization system. *Classical and Quantum Gravity*, 27:084011, 2010.

[125] TT-Y Lam, Bram JJ Slagmolen, Jong H Chow, Ian CM Littler, David E McClelland, and Daniel A Shaddock. Digital laser frequency stabilization using an optical cavity. *Quantum Electronics, IEEE Journal of*, 46(8):1178–1183, 2010.

[126] R J Cruz, J I Thorpe, A Preston, Delgadillo R., M Hartmann, S Mitryk, A Worley, G Boothe, S R Guntaka, S Klimenko, D B Tanner, and G Mueller. The LISA benchtop simulator at the University of Florida. *Classical and Quantum Gravity*, 23:S751–S760, 2006.

[127] Stephen B Wicker and Vijay K Bhargava. *Reed-Solomon codes and their applications*. Wiley. com, 1999.

[128] Rhode & Schwarz. SMB100A RF and Microwave Signal Generator Specifications. http://www.rohde-schwarz.com/en/product/smb100a-productstartpage_63493-9379.html, 2013.

[129] Texas Instruments. LMH6551 Differential, High Speed Op Amp. http://www.ti.com/product/lmh6551, 2005.

[130] Tektronix. Arbitrary/Function Generators AFG 3102 Datasheet. http://de.tek.com/sites/tek.com/files/media/media/resources/AFG3000_Series_Arbitrary-Function_Generators_Datasheet_76W-18656-5.pdf, 2012.

[131] Laird Technologies. Datasheet AA PowerCool Series, AA-200-24-22 Thermoelectric Assembly. http://www.lairdtech.com/Products/Thermal-Management-Solutions/Thermoelectric-Assemblies/, 2010.

[132] Laird Technologies. Datasheet PR-59, TC-XX-PR-59 Temperature Controller. http://www.lairdtech.com/Products/Thermal-Management-Solutions/Thermoelectric-Modules/, 2010.

[133] Linear Technology. LTSpice IV. http://www.linear.com/designtools/software/, 2012.

[134] PL Bender, RS Nerem, and JM Wahr. Possible future use of laser gravity gradiometers. *Space science reviews*, 108(1-2):385–392, 2003.

[135] Gerhard Heinzel, Felipe Guzmán Cervantes, Antonio F García Marin, Joachim Kullmann, Wang Feng, and Karsten Danzmann. Deep phase modulation interferometry. *Optics Express*, 18(18):19076–19086, 2010.

[136] Felipe Guzman Cervantes. *Gravitational wave observation from space: Optical measurement Techniques for LISA and LISA Pathfinder*. PhD thesis, Leibniz Universität Hannover, 2009.

[137] Joachim Kullmann. Aufbau einer interferometrischen Positionsauslesung für die LISA-Testmassen. Diploma thesis, Leibniz Universität Hannover, 2007.

[138] T. S. Schwarze. FPGA-based modulation signal synthesis and measurement system for high precision laser interferometry. Masters thesis, Leibniz Universität Hannover, 2013.

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

[140] Malcom J. King. An internally sensing active optical phased array using digital interferometry. Master's thesis, The Australian National University, 2010.

[141] Glenn de Vine, David S Rabeling, Bram JJ Slagmolen, Timothy TY 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, 2009.

[142] Danielle MR 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):672–674, 2011.

[143] O. P. Lay, S. Dubovitsky, D. A. Shaddock, and B. Ware. Coherent range-gated laser displacement metrology with compact optical head. *Opt. Lett.*, 32(20):2933–2935, Oct 2007.

[144] S. M. Köhlenbeck. Digital unterstützte heterodyne Interferometrie: Ein Sensor zur Abstandsstabilisierung mit Picometerauflösung im sub-Hertz Bereich. Diploma thesis, Leibniz Universität Hannover, 2012.

[145] Katharina Sophie Isleif. Investigation of Noise Sources in Digital Interferomerty. Diploma thesis, Leibniz Universtität Hannover, 2013.

[146] F Jessie MacWilliams and Neil JA Sloane. Pseudo-random sequences and arrays. *Proceedings of the IEEE*, 64(12):1715–1729, 1976.

[147] e2v semiconductors. EV8AQ160 QUAD ADC Datasheet. http://www.e2v.com/e2v/assets/File/documents/broadband-data-converters/EV8AQ160.pdf, 2010.

[148] Xilinx. Virtex-6 FPGA GTX Transceivers. http://www.xilinx.com/support/documentation/user_guides/ug366.pdf, 2011.

[149] ANALOG DEVICES. AD9517-3 Datasheet. www.analog.com/static/imported-files/data_sheets/AD9517-3.pdf, 2013.

[150] Photline. Datasheet NIR-MPX series 1000 nm band Phase Modulators. http://www.photline.com/product/view/53/, 2010.

[151] New Focus. 12-GHZ Amplified Free-Space Photoreceivers Models 1567-A 1577-A. http://www.newport.com, 2002.

[152] A. Sutton, O. Gerberding, G. Heinzel, and D. Shaddock. Digitally enhanced homodyne interferometry. *Optics Express*, 20(20):22195, 2012.

[153] Andrew Sutton. *Digital Techniques for Coherent Optical Metrology*. PhD thesis, Australian National Unviersity, 2013.

[154] Antonio F García Marin. *Minimization of optical pathlength noise for the detection of gravitational wave with the spaceborne laser interferometer LISA and LISA Pathfinder*. PhD thesis, Leibniz Universität Hannover, 2007.

[155] F Steier, R Fleddermann, J Bogenstahl, C Diekmann, G Heinzel, and K Danzmann. Construction of the LISA back-side fibre link interferometer prototype. *Classical and Quantum Gravity*, 26(17):175016, 2009.

[156] Xilinx. ML605 hardware user guide. http://www.xilinx.com/products/boards-and-kits/EK-V6-ML605-G.htm, 2011.

[157] Peter Fall. 1g eth UDP / IP stack. http://opencores.org/project,udp_ip_stack, 2011.

Acknowledgements

I am very grateful to Professor Karsten Danzmann for giving me the opportunity to undertake my PhD studies at the Albert Einstein Institute. This institute, its research and its people have been motivating and inspiring for me. I especially want to thank him for allowing and encouraging me to start my "own" research project with the Digital Interferometry experiment. Being able to turn my ideas into a real project is the dream of every scientist, and doing it without writing a proposal is something that only seems possible at a place like this.

PD Dr. Gerhard Heinzel has not only been my direct supervisor and in charge of the Space Interferometry Group, he has also shared his seemingly endless knowledge, motivation and fascination with me. He always gave me new perspectives on things and I really appreciate the trust he put in me. Thank you Gerhard!

A special thanks goes to two Post-Docs from the Space Interferometry Group whom I have been bugging the most, Benjamin Sheard and Michael Tröbs. Who would have thought that standing in front of a white board and talking can be such a joy.

A big thank you goes to the whole Space Interferometry Group and all its members over the years. I always enjoyed working in this group and with its highly motivated people and I like the fact that we all share the experience of working 'now' and 'down here' on something for 'at some point in the future' and 'up there', others just don't get it.

The time at the AEI preparing this thesis was simply a blast. The atmosphere here has turned many colleagues from all parts of the Institute into friends. I am very grateful to all the people behind the curtains, in the workshops, the IT and the administration, who keep this infrastructure alive and who sometimes keep us PhD students in check. A special thanks goes to my colleagues who have been with me since the beginning of my studies, especially to Christina Vollmer, who strangely always knows how I am by simply looking at me.

It has been pleasure to work with the phasemeter project team, including our Danish Industry partners from Axcon, the partners from DTU Space, the colleagues from the European Space Agency, and, of course, the team from the AEI. Especially I would like to thank Torben Rasmussen and Torben Vendt Hansen. 'Calling Torben' has been the solution to many of my problems with the hardware and it was always a pleasure to work with you.

My visit to the Australian National University was very motivating. This is thanks to Professor Daniel Shaddock, who has opened up a new perspective on

interferometry for me and who is also so kind as to referee this thesis. I also want to thank my "exchange student from the other side of the world" Andrew Sutton. Working with you was not only a lot of fun, but I also learned a lot from it. In addition to Andrew I also want to thank Danielle Wuchenich and Fritz Dengel for hosting me so often during my stay at the ANU. Your hospitality was simply "Bäm!".

I also want to express my gratitude to Professor Piet O. Schmidt for being a referee for this thesis, and to Professor Clemens Hammerer for being the chair of the defence committee.

I am very grateful to the students I have supervised: Nils Brause, Sina Köhlenbeck, Thomas Schwarze and Katharina-Sophie Isleif. It was a pleasure to work with you and I wouldn't want to have missed this experience.

I want to thank my long time office mates Henning Kaufer and Stefan Ast. I really don't know how to summarize this accurately, but the mixture of optomechanics, squeezing, phasemeters, general banter, movies, TV-shows, computer games and spoilers was just right for me. Thanks also to Melanie Meinders and Stefan Ast for the Game-of-Thrones breakfast tradition.

I also thank all the members of the Mensa-Gang. Nerdiness, banter, parties and sarcasm combined with mensa food have never been so delicious.

For proofreading this thesis I would like to thank, in alphabetical order, Stefan Ast, Heather Audley, Simon Barke Gerhard Heinzel, Katharina-Sophie Isleif, Sina Köhlenbeck, Moritz Mehmet, Daniel Schütze, Benjamin Sheard, Andrew Sutton, and Michael Tröbs. Your comments and corrections have improved the quality of this thesis a lot and all remaining errors are, of course, my own responsibility.

The outreach campaign we started in 2013 was an amazing experience. I would like to thank Simon Barke, Heather Audley, Benjamin Knispel, and Maike Lieser for this time. The excitement of reaching two million people with a meme and producing videos with more than 100k views is something I won't forget.

I especially like to thank my fellow PhD student Heather Audley, our friendship means a lot to me.

Of course, I would also like to thank all my friends from outside the Institute, especially Christian Lange, Maren Brenneke, Timko Dubielzig, Katrin Bürger, Felix Frost, Irene Geisler, and Florian Pini. You are simply the best. I also want to thank my teachers and friends from salsa dancing, especially my long time dancing partner Nadine Raudssus, for sharing this amazing obsession with me.

My family has always supported me, and their love and confidence in me made this thesis possible in the first place. I am forever grateful.

This thesis utilised the ComponentLibrary by A. Franzen:
(http://www.gwoptics.org/ComponentLibrary/)

Curriculum vitae

Oliver Gerberding
Schlosswenderstraße 13
30159 Hannover

Born on 11th of April 1986, in Hannover, Germany

1990 - 1998	Grundschule & Orientierunsstufe
1998 - 2005	Abitur, Otto-Hahn-Gymnasium, Springe
2005 - 2009	Diploma Studies of Technical Physics at Leibniz University Hannover
11/2009	Physic Diploma, thesis title: "Untersuchung fraktionaler Füllfaktoren durch Magnetotransport und Magnetokapazitätsmessungen"
2010 - 2013	Scientific Assistant at the Max-Planck Insitute for Gravitational Physics (Albert-Einstein-Institute), Hannover; Stipend holder of the International Max-Planck Research school for gravitational wave astronomy
04/2011-05/2011	Research visit to the Australian, National University, Canberra, Australia

Publications

Refereed journal articles

O. Gerberding, B. Sheard, I. Bykov, J. Kullmann. J. J. Esteban, K. Danzmann and G. Heinzel, *"Phasemeter core for intersatellite laser heterodyne interferometry: modelling, simulations and experiments"*
Classical Quantum Gravity, Vol. 30, 235029 (2013)
http://dx.doi.org/10.1088/0264-9381/30/23/235029

A. Sutton, O. Gerberding, G. Heinzel and D. Shaddock
"Digitally enhanced homodyne interferometry"
Optics Express, Vol. 20, Issue 20, pp. 22195-22207 (2012)
http://dx.doi.org/10.1364/OE.20.022195

Non-refereed articles

K. Danzmann et al.
"The Gravitational Universe"
white paper for the ESA L2/L3 science theme selection (2013)
http://arxiv.org/abs/1305.5720

Published proceedings

O. Gerberding, S. Barke, I. Bykov, K. Danzmann, A. Enggaard, J. J. Esteban, A. Gianolio, T.V. Hansen, G. Heinzel, A. Hornstrup, J. Kullmann, R. Ostergard, S.M. Pedersen, T. Rassmussen, J. Reiche and Z. Sodnik, *"Breadboard model of the LISA phasemeter"*
ASP Conference Series, Vol. 467, 9th LISA Symposium, pp. 271 (2012)
http://arxiv.org/abs/1208.6418

G. Heinzel, B. Sheard, N. Brause, K. Danzmann, M. Dehne, O. Gerberding, C. Mahrdt, V. Müller, D. Schütze, G. Stede, W. Klipstein, W. Folkner, R. Spero, K. Niklaus, P. Gath and D. Shaddock *"Laser Ranging Interferometer for GRACE follow-on"*
Proceedings of the 9th International Conference on Space Optics, Ajaccio (2012)
http://congrex.nl/icso/2012/papers/FP_ICSO-062.pdf

M. Tröbs, L. d'Arcio, S. Barke, J. Bogenstahl, I. Bykov, M. Dehne, C. Diekmann, I. Diepholz, E. D. Fitzsimons, R. Fleddermann, O. Gerberding, J.-S. Hennig, F. G. Hey, H. Hogenhuis, C. J. Killow, E. Kochkina,J. Kullmann, M. Lieser, S. Lucarelli, M. Perreur-Lloyd, J. Pijnenburg, D. I. Robertson, A. Shoda, A. Sohmer, A. Taylor, G. Wanner, H. Ward, D. Weise, G. Heinzel and K. Danzmann *"Testing the LISA Optical bench"*
Proceedings of the 9th International Conference on Space Optics, Ajaccio (2012)

http://congrex.nl/icso/2012/papers/FP_ICSO-110.pdf

M. Tröbs, L. d'Arcio, S. Barke, J. Bogenstahl, C. Diekmann, E. D. Fitzsimons, O. Gerberding, J.-S. Hennig, F. G. Hey, H. Hogenhuis, C. J. Killow, M. Lieser, S. Lucarelli, S. Nikolov, M. Perreur-Lloyd, J. Pijnenburg, D. I. Robertson, A. Sohme, A. Taylor, H. Ward, D. Weise, G. Heinzel and K. Danzmann *"LISA optical bench testing"*
ASP Conference Series, Vol. 467, 9th LISA Symposium, pp. 233 (2012)
http://www.aspbooks.org/a/volumes/table_of_contents/467

D. Schütze, G. Stede, V. Müller, O. Gerberding, C. Mahrdt, B. Sheard, G. Heinzel, and K. Danzmann *"LISA-like laser ranging for GRACE follow-on"*
ASP Conference Series, Vol. 467, 9th LISA Symposium, pp. 285 (2012)
http://www.aspbooks.org/a/volumes/table_of_contents/467

A. Taylor, L. d'Arcio, J. Bogenstahl, K. Danzmann, C. Diekmann, E. D. Fitzsimons, O. Gerberding, G. Heinzel, J.-S. Hennig, H. Hogenhuis, C. J. Killow, M. Lieser, S. Lucarelli, S. Nikolov, M. Perreur-Lloyd, J. Pijnenburg, D. I. Robertson, A. Sohmer, M. Tröbs, H. Ward and D. Weise *"Optical Bench Interferometer - From LISA Pathfinder to NGO/eLISA"*
ASP Conference Series, Vol. 467, 9th LISA Symposium, pp. 311 (2012)
http://www.aspbooks.org/a/volumes/table_of_contents/467

M. Lieser, L. d'Arcio, S. Barke, J. Bogenstahl, C. Diekmann, I. Diepholz, E. D. Fitzsimons, O. Gerberding, J.-S. Henning, M. Hewitson, F. G. Hey, H. Hogenhuis, C. J. Killow, S. Lucarelli, S. Nikolov, M. Perreur-Lloyd, J. Pijnenburg, D. I. Robertson, A. Sohmer, A. Taylor, M. Tröbs, H. Ward, D. Weise, G. Heinzel, K. Danzmann, *"LISA Optical Bench Testbed"*
ASP Conference Series, Vol. 467, 9th LISA Symposium, pp. 301 (2012)
http://www.aspbooks.org/a/volumes/table_of_contents/467

B. Sheard, M. Dehne, C. Mahrdt, O. Gerberding, V. Müller, G. Heinzel and K. Danzmann *"Laser Interferometry for future satellite gravimetry missions"*,
GEOTECHNOLOGIEN Science Report, Nr. 17, pp 160-166 (2011)
dx.doi.org/10.2312/GFZ.gt.17.24

Supervision

Supervision of Master and Diploma thesis projects (one year, full time):

N. Brause (2011)
Diploma Thesis
"Bau und Test eines Phasenmeters"

S. Köhlenbeck (2012)
Diploma Thesis
"Digital unterstützte heterodyne Interferometrie: Ein Sensor zur Abstandsstabilisierung mit Picometerauflösung im sub-Hertz Bereich"

T. S. Schwarze (2012)
Master Thesis
"FPGA-based modulation signal synthesis and measurement system for high precision laser interferometry"

K.-S. Isleif (2013)
Master Thesis
"Investigation of Noise Sources in Digital Interferometry"