Interferometry techniques for spaceborne gravitational wave detectors

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

Doktors der Naturwissenschaften
– Dr.rer.nat –

genehmigte Dissertation von

Dipl.Phys. Frank Steier

geboren am 1. August 1978
in Dorsten, Deutschland

2008

Zusammenfassung

Im Kontext dieser Arbeit wurden spezielle Aspekte der Interferometrie der beiden Weltraummissionen LISA und LISA Pathfinder behandelt. Die Laser Interferometer Space Antenna (LISA) ist ein zukünftig geplanter Gravitationswellendetektor, während LISA Pathfinder eine Mission ist, bei der bestimmte Technologien getestet werden, die in LISA ihre Anwendung finden.

Die Interferometertopologie von LISA beinhaltet mehrere Elemente, die in anderen Interferometern nicht notwendig sind. Eines dieser Elemente ist der Phasenlock zweier Laser, von denen einer eine sehr geringe Leistung in der Größenordnung von 100 pW aufweist. Dies wird durch die Divergenz der Laserstrahlen hervorgerufen, die über die Distanz der 5 Millionen Kilometer langen Arme aufgeweitet werden. Daher wird ein lokaler Laser phasenstarr an den einfallenden Strahl mit einem Frequenzversatz von 2 MHz bis 20 MHz gekoppelt und ausgesandt. Für diese Arbeit wurde ein Phasenlock mit Lichtleistungen vergleichbar zu denen in LISA demonstriert. Die erreichte Phasenstabilität liegt maximal um einen Faktor 5 über den für LISA geforderten Spezifikationen. Der experimentelle Aufbau wurde außerdem verwendet, um verschiedene Photodioden-Verstärker bezüglich ihrer Rauscheigenschaften zu untersuchen.

Ein weiteres besonderes Element der LISA Interferometrie ist eine optische Faser, die sich im optischen Pfad befindet. In ihr laufen zwei Laserstrahlen in entgegengesetzten Richtungen. Die angestrebte Empfindlichkeit von LISA kann nur erreicht werden, wenn beide Strahlen annähernd gleiche Weglängenänderungen bei ihrer Propagation durch die Faser erfahren. Die maximale Abweichung darf höchstens $1 \, \mathrm{pm}/\sqrt{\mathrm{Hz}}$ oberhalb von 3 mHz betragen. Zu diesem Zweck wurde ein quasi-monolithisches Interferometer mit Hilfe der sogenannten *hydroxide-catalysis bonding* Technik gebaut, das alle wesentlichen Bestandteile der Faserverbindung beinhaltet. Durch Anwendung dieser Technik war es möglich eine genügend hohe mechanische und thermische Stabilität zu erzielen. In diesem Experiment wurden Rückreflexionen an den Faserinjektoren als dominierende Rauschquelle iden-

tifiziert, die unvermeidbar sind, da sie von der Topologie des Interferometers hervorgerufen werden: Alle Rückreflexionen an den Faserinjektoren liegen automatisch auf derselben Strahlachse wie der engegengesetzt laufende Strahl. Durch entsprechende Nachverarbeitung der Messdaten konnte dieser Effekt jedoch minimiert werden, wodurch die Sensitivität im gesamten LISA Messband verbessert wurde. Die Verbesserung entspricht maximal einem Faktor von 30 und führt zu einer Empfindlichkeit von etwa $10\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ bei $10\,\mathrm{mHz}$. Dies ist das beste Ergebnis im Vergleich zu verschiedenen anderen Experimenten ähnlicher Zielsetzung weltweit.

Das Faserinterferometer ist eines der komplexesten Interferometer, die je mit der Technik des *hydroxide-catalysis bonding* gebaut wurden. Es beinhaltet zwei Faserinjektoren für die Eingangsstrahlen und zwei zusätzliche, die bezüglich fest vorgegebener Strahlen, die aus dem Interferometer kommen, justiert werden. Insgesamt besteht es aus drei Interferometern, von denen zwei einen optischen Pfad durch die Faser enthalten. Anstelle der Faser kann ein stabiler Referenzpfad auf der Basisplatte vermessen werden, ohne dass das Interferometer nachjustiert werden muss.

Ein zweiter wesentlicher Schwerpunkt dieser Arbeit ist die Definition und der Test der Prozessierung der interferometrischen Daten an Bord von LISA Pathfinder. Zu diesem Zweck wurde das bereits bestehende LISA Pathfinder Experiment erweitert, indem der Prozessierung spezielle Besonderheiten hinzugefügt wurden. Dazu gehören die Mittelung von nominellen und redundanten Datenkanälen, ein asynchroner Datentransfer der wissenschaftlichen Ausgangsdaten des Interferometers und ein Fehlerdetektions und -fortpflanzungsschema, das die Qualität der Ausgangsdaten anhand der Anzahl und Kombination der verwendeten Eingangssignale autonom klassifiziert. Dies ist notwendig, da diese Daten zur Kontrolle von Testmassen und Satellit verwendet werden. Die eben genannten Prozessierungsschritte, abgesehen von der Fehlerdetektion, wurden in das Experiment integriert, das somit eine Testumgebung des kompletten optischen Messsystems an Bord von LISA Pathfinder darstellt.

Mit den neuen Software-Elementen konnte die geforderte Empfindlichkeit des Interferometers experimentell demonstriert werden. Desweiteren wurde das Interferometer eingehend auf den Effekt von Fehlern in Datenkanälen und seine Justage charakterisiert. Die dadurch gewonnen Ergebnisse stellen wichtige Referenzen für die in naher Zukunft geplanten Tests von Engineering- und Flug-Modellen dar.

Schlagworte: Gravitationswellen, LISA, Phasenmessung

Abstract

This thesis deals with the interferometry for two space missions, LISA and LISA Pathfinder. LISA, the Laser Interferometer Space Antenna, is a planned gravitational wave detector in space, and LISA Pathfinder is a technology demonstration for LISA.

The LISA interferometry includes a number of new technologies which are usually not needed in other interferometers. One of those is the phase lock of two lasers, where one of the lasers has a power of only about $100\,\mathrm{pW}$. Due to the 5 million kilometres long interferometer arms, the beam divergence causes an attenuation of the laser beams which travel between the LISA spacecraft. Instead of direct back reflection of the laser light, a second laser is locked to the incoming beam with an offset frequency between $2\,\mathrm{MHz}$ and $20\,\mathrm{MHz}$ and sent out. An experimental demonstration of such a phase lock was performed. The achieved sensitivity is, at maximum, a factor of 5 above the required LISA performance. The experimental setup was used to test different photodiode transimpedance amplifiers which were specially made for the application in LISA.

Another feature of the LISA interferometry is the use of a fibre within the optical beam path in which two beams are counter-propagating. LISA will only reach its design sensitivity if the optical pathlength in the fibre is reciprocal to a level of about $1\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$, i.e. the pathlength fluctuations are identical for both directions. Therefore, a quasi-monolithic interferometer was built that includes all important features of the LISA fibre link. A technique called *hydroxide-catalysis bonding* was applied to ensure the thermal and mechanical stability of the interferometer. The fibre link was characterised, and ghost reflections in the setup were identified as an important noise source. These are inevitable in this interferometer topology: All direct back reflections of one beam automatically have the same beam axis as the counter-propagating beam. A post-processing correction method was developed to reduce the noise introduced by such ghost beams. This led to a sensitivity improvement in the entire LISA measurement band. The noise is reduced at most

by a factor of 30 down to a level in the $10\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$ range around $10\,\mathrm{mHz}$, which is at present the best result among several similar efforts worldwide.

The fibre interferometer is one of the most complex interferometers that was ever built using the hydroxide-catalysis bonding technique. It has two input fibre injectors and, in addition, two fibre injectors that are aligned with respect to beams coming from the baseplate. It includes three interferometers, while two of them include a path through the optical fibre. Instead of this fibre, length variations of a stable optical (reference) path can be measured without realignment of the interferometer.

For the LISA Pathfinder mission, the flight software for the interferometric data processing was defined as part of this thesis. A complete end-to-end testbed of the optical system on-board LISA Pathfinder was set up in order to test the data processing. In contrast to earlier breadboard experiments it includes averaging of redundant readout channels and asynchronous data transfer of the main data channels between two separated programs that simulate the two computers involved on-board LISA Pathfinder. A failure detection and propagation scheme was developed within this software. It automatically chooses the valid data channels for the data calculation and classifies the quality of the science output data which is used for the drag-free control of the LISA Pathfinder test masses and spacecraft. The end-to-end optical model was used to show the performance of the system including all of the features, apart from the error handling, described above. In addition, the interferometer was characterised in detail in terms of channel failures and alignment. The achieved results on the LISA Pathfinder end-to-end optical hardware model are an important reference for the upcoming tests of engineering and flight hardware and software.

Keywords: gravitational waves, LISA, phase readout

Contents

List of Figures

List of Tables

Glossary

AA	Anti-aliasing
AC	Alternating current
ADC	Analogue-to-Digital Converter
AEI	Albert Einstein Institute
AMP	amplifier
AOM	Acousto-Optical Modulator
BC	Beam Combiner
BFI	Backside Fibre link Interferometer
BS	Beamsplitter
c	cut-off
CMOS	Complementary Metal-Oxide-Semiconductor
corr	corrected
DAC	Digital to Analogue Converter
DAQ	Data Acquisition
DC	Direct current
DFACS	Drag-Free and Attitude Control System
DMU	Data Management Unit
DWS	Differential Wavefront Sensing
EADS	European Aeronautic Defence and Space Company
EM	Engineering Model
EOM	Electro-Optical Modulator
err	error signal
F1,F2,F3	Fix point components on the bonded interferometer
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FM	Flight Model
FPGA	Field Programmable Gate Array
GSFC	Goddard Space Flight Center

het	heterodyne
H/K	Housekeeping
ifo	interferometer
IIR	Infinite Impulse Response
InGaAs	Indium-Gallium-Arsenide
IRS	Inertial Reference Sensor
jn	Johnson noise
JPL	Jet Propulsion Laboratory
LISA	Laser Interferometer Space Antenna
LO	Local Oscillator
LP	low-pass
LTP	LISA Technology Package
LUT	Look-Up Table
m	modulation
M	Master laser
M‖	Mirror on the bonded interferometer
M1	Interferometer measuring the LISA fibre link in one direction
M2	Interferometer measuring the LISA fibre link in other direction
max	maximum
MB	Measurement Beam
meas	measured
Mil-BUS	BUS using military standards
mix	mixer
MX	Mixer
NASA	National Aeronautics and Space Administration
Nd:YAG	neodymium-doped yttrium aluminium garnet
NEP	Noise Equivalent Power
Ny	Nyquist
OBC	On-board Computer
OLG	Open Loop Gain
OMS	Optical Metrology System
op	Operational Amplifier
OpAmp	Operational Amplifier
OPD	Optical Pathlength Difference
p	parasitic
PAAM	Point Ahead Angle Mechanism
PANDA	Polarisation-Maintaining and Absorption-Reducing
PC	Personal Computer
PD	photodiode
PM	Phasemeter
PMS	Phase Measurement System
PT	Phase-tracking
PZT	Piezoelectric Transducer

Q	Quadrant
QPD	Quadrant Photodiode
RAM	Random access memory
RB	Reference Beam
rec	received
ref	reference
rms	root mean square
S	Slave laser, sometimes Straylight
samp	sampling
SMART2	former name of LISA Pathfinder
sn	shot noise
S/W	Software
T	Transit frequency
TM	Test mass
TNO	Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek
u	update
UGL	University of Glasgow
X1	Interferometer measuring test mass 1 position in the LTP
X12	Interferometer measuring differential test mass positions in the LTP
Zerodur	Glass ceramic from Schott

Part I

LISA Interferometry

The LISA mission

The Laser Interferometer Space Antenna (LISA) is a planned interferometric gravitational wave detector in space. Its measurement band lies between $100\,\mu\mathrm{Hz}$ and $1\,\mathrm{Hz}$ with a peak sensitivity of $10^{-23}/\sqrt{\mathrm{Hz}}$. The launch is planned to be in 2018 [1].

LISA consists of three satellites following the Earth in its orbit around the Sun. The formation of this triangle is illustrated in Figure 1.1. The three satellites are situated in the corners of an equilateral triangle with an arm length of $5{\cdot}10^6\,\mathrm{km}$. The triangle is tilted by 60 degrees to the ecliptic and trails the Earth on its orbit around the Sun by 20 degrees. The satellite orbits are chosen such that the formation of the equilateral triangle is as stable as possible in its constellation without application of additional forces [2].

In order to measure gravitational waves, the distances between drag-free masses have to be measured. The satellites are exposed to many external forces like solar radiation, external electro-magnetic forces etc. Each of the spacecraft carries two of those free-falling bodies called *test masses* or *proof masses* and protects those against the external forces. They are situated at the end of each interferometer arm and act as end mirrors of the interferometer as shown in Figure 1.2. The test masses are surrounded by a capactive sensor and actuator, the Inertial Reference Sensor (IRS), which measures and controls their position inside the spacecraft. In the sensitive direction the test masses are kept drag-free in the LISA measurement band, but outside this frequency band and along the insensitive axes, they are controlled to stay in the centre of the IRS [4].

The longitudinal measurements of the test masses along the sensitive axis, and the length measurements of the long arms are performed interferometrically. The beam transmitted to and the beam received from the far spacecraft are both routed through a single telescope with a diameter of $40\,\mathrm{cm}$ (also shown in Figure 1.2) [5].
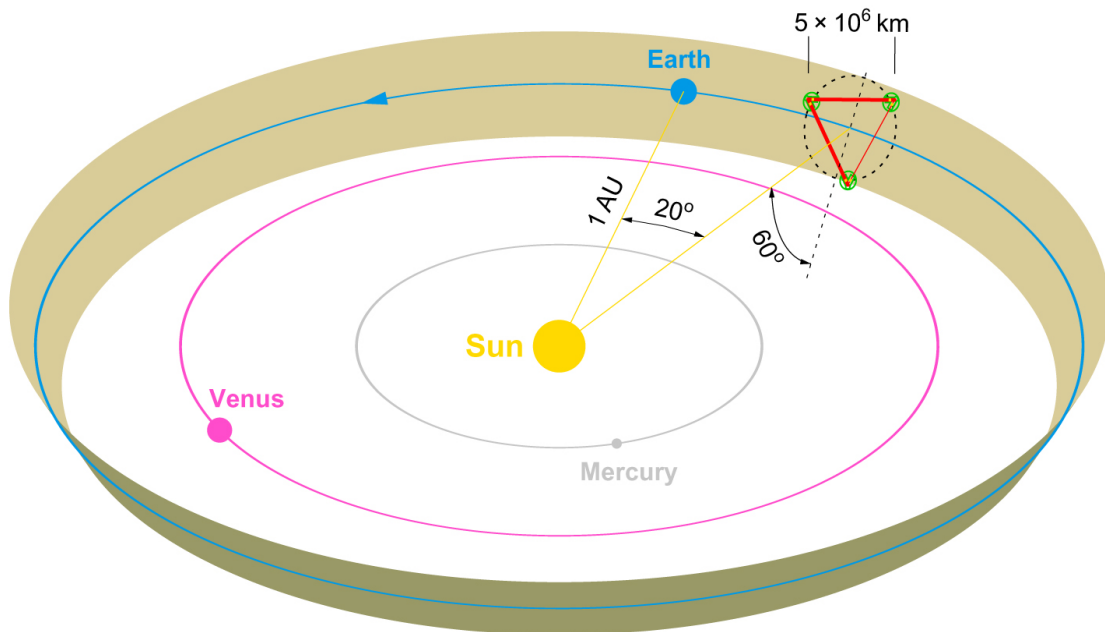
Figure 1.1: Orbits of the LISA satellites [3].
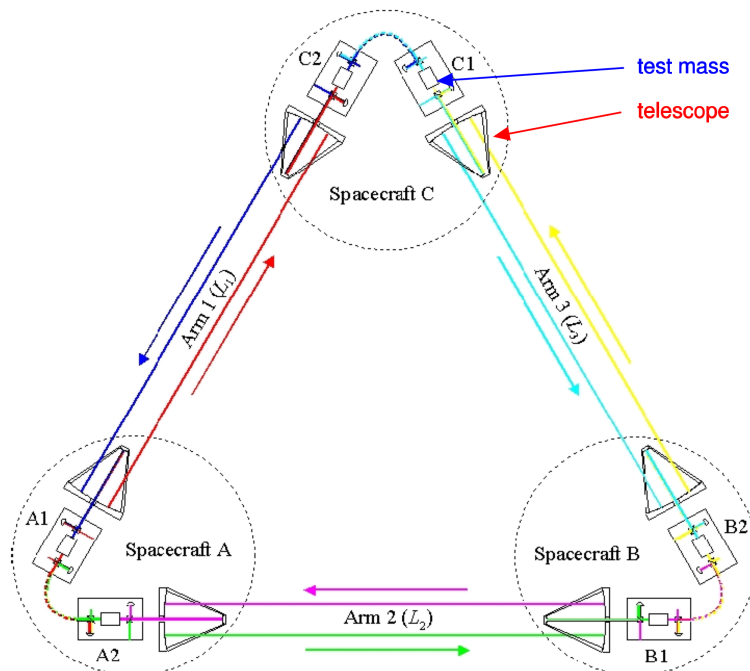
Figure 1.2: The six laser links of LISA [3].

Due to the long arms of the interferometer and the limited size of the telescope, the beam divergence causes an attenuation of the 1 W outgoing laser beam to a received laser beam power of about 100 pW at the far spacecraft. Therefore, the beam cannot be directly back reflected. Instead, a second laser in the far spacecraft is phase locked to the incoming beam with an offset frequency of about 10 MHz before being re-transmitted. A simplified schematic view of the whole LISA configuration is shown in Figure 1.2.

On-board each spacecraft there are two lasers, one for each arm. They are also phase locked with respect to each other. Each belongs to one of the optical benches, which carry the optics of the interferometer and are situated at the end of each interferometer arm. A picture of the optical bench baseline design is shown in Figure 1.3. The injected laser (*TX Laser Delivery*) is sent to the three main interferometers:

- In the *Reference Interferometer* the beam interferes with the second laser from the same satellite. Here both local lasers are phase locked with respect to each other.

- In the *Proof Mass Optical Readout* both local lasers are also interfered, but here the beam path of the laser from the other optical bench includes a reflection at the proof mass. Therefore, the differential measurement with the *Reference Interferometer* detects proof mass movement with respect to the optical bench.

- The *Science Interferometer* is used to interfere the incoming beam from the far spacecraft with the local laser. On one end of the interferometer arm both are phase locked and on the other end the length change is measured.

The topology of the interferometer is called *strap-down* [7, 8]. In earlier LISA designs the length changes between two test masses were measured directly, whereas in the current baseline design the measurement of the arm length is separated from the test mass readout. Both the arm and the test mass are measured with respect to the optical bench. The inter-spacecraft measurement is called *long arm interferometry* and test mass and optical bench to optical bench measurements are called *local interferometry*. The advantage of this design is the decoupling of both. They can be developed almost independently from each other and lead to an important simplification in the development, integration and testing.

Figure 1.4 shows the readout of one of the six test masses. It is illustrated for the second optical bench on the first satellite, OB 1/2. The setups on each of the six optical benches are symmetrical. Here one of the benches is chosen in order to introduce a nomenclature. The laser L 1/2 belongs to this optical bench. It is used as measurement beam for the test mass readout. The second laser from the same spacecraft L 1/1, but from the other optical bench, is routed to this bench via a fibre. It is used as reference beam in the reference interferometer REF 1/2 and the proof mass readout PM 1/2. In the differential measurement of both interferometers, the
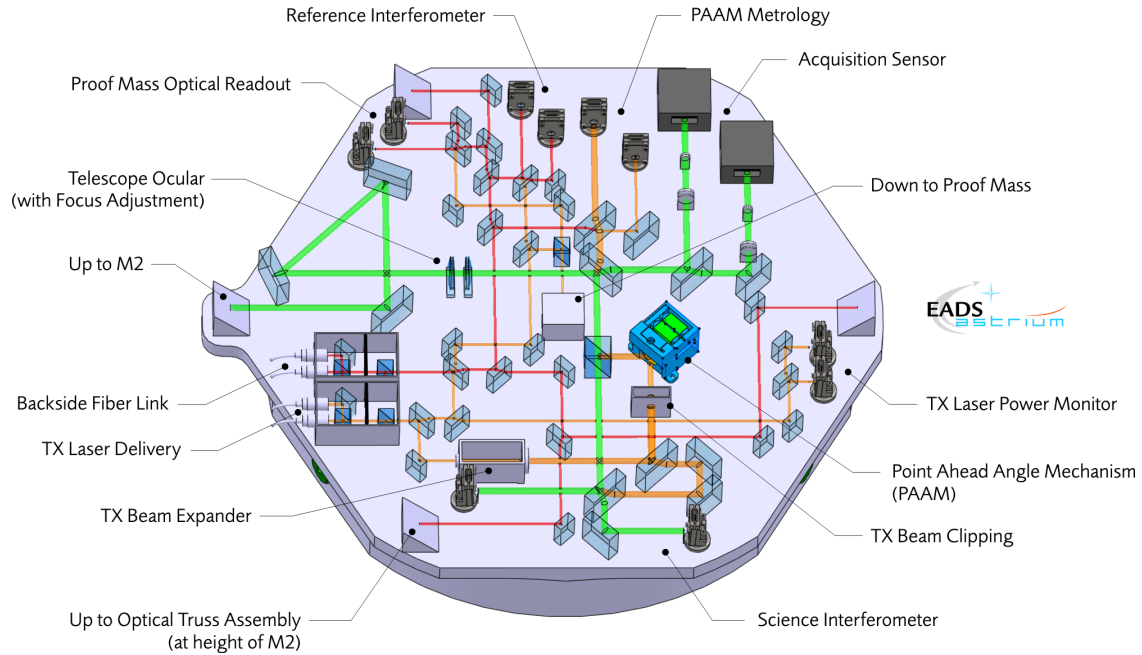
Figure 1.3: Schematic view of the LISA optical bench [6].

test mass position is measured with respect to the optical bench. It corresponds to the LISA Pathfinder interferometry discussed in Chapter 6.

The setup on the second optical bench, OB 1/1, on the same spacecraft is symmetric as shown in Figure 1.5. Here the laser L 1/2 acts as reference beam and the laser L 1/1 as measurement beam. In this figure only the two reference interferometers are highlighted. The measurements in both reference interferometers, REF 11 and REF 12, are sensitive to the optical pathlength changes of the fibre connecting the two optical benches. This connection is called *backside fibre link*. It plays an important role in the LISA interferometry and is discussed in Chapter 3.

The long LISA arms are measured as shown in Figure 1.6. The laser L 1/2 from bench OB 1/2 is transmitted to the remote spacecraft and received on bench OB 2/1 where it is routed to the science interferometer SCI 21. The laser L 2/1 from bench OB 2/1 is phase locked to the incoming laser and transmitted back to the bench OB 1/2. There, it is also routed to the interferometer SCI 12, where the phase of the measured output signal carries the information of the arm length changes.

An additional element, which is usually not necessary in other interferometers, is a piezo actuated mirror in the optical path of the incoming and outgoing beam. It is called the *Point Ahead Angle Mechanism* (PAAM) and corrects for the difference in the angle between the incoming and outgoing beam, which is changing during a year [8]. The PAAM is not part of the work presented in this thesis.

The backside fibre link and the test mass optical readout belong to the local interferometry, whereas the measurement of the arm is called long arm interferometry of LISA. Important aspects of both were investigated within this thesis.
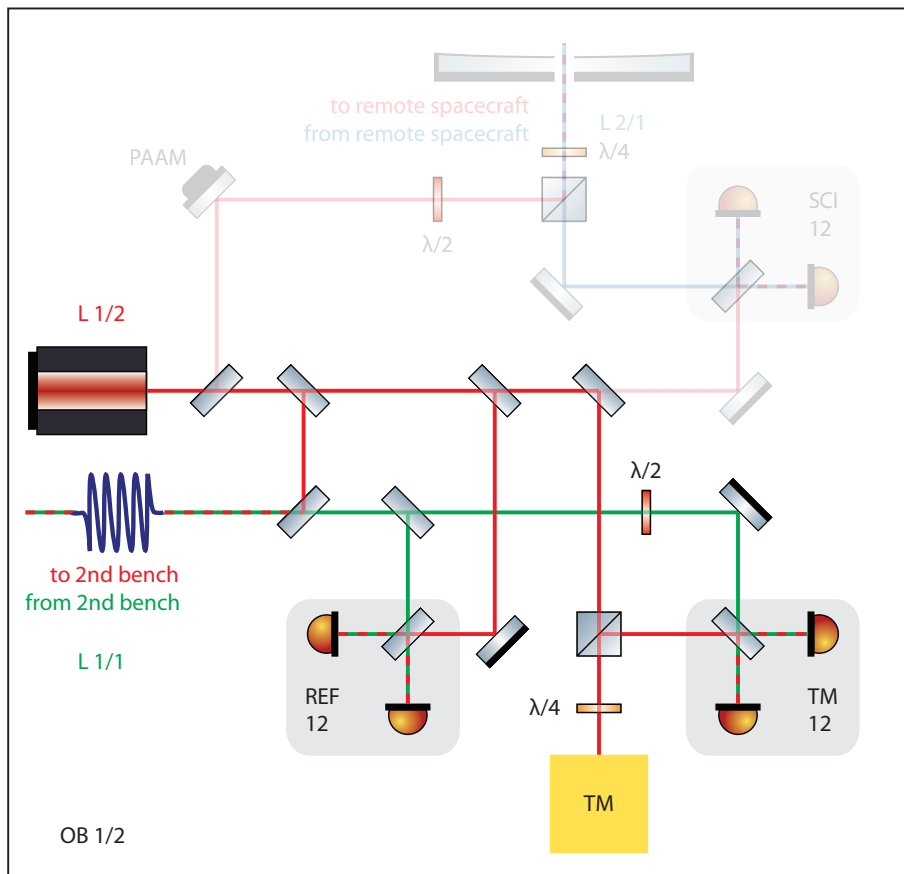
Figure 1.4: Read-out of the local test mass.
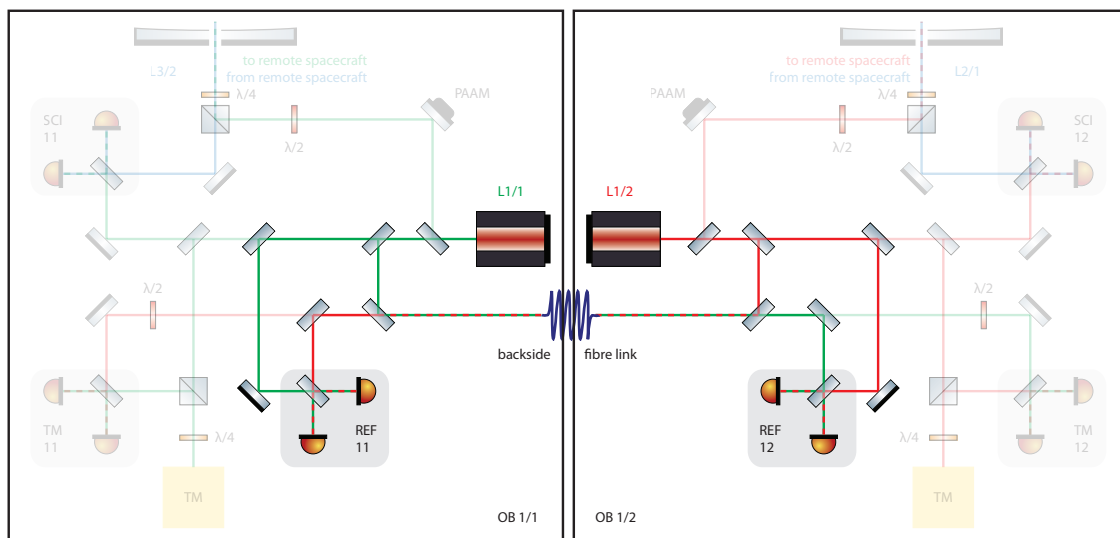
Figure 1.5: The two optical benches on-board one satellite.
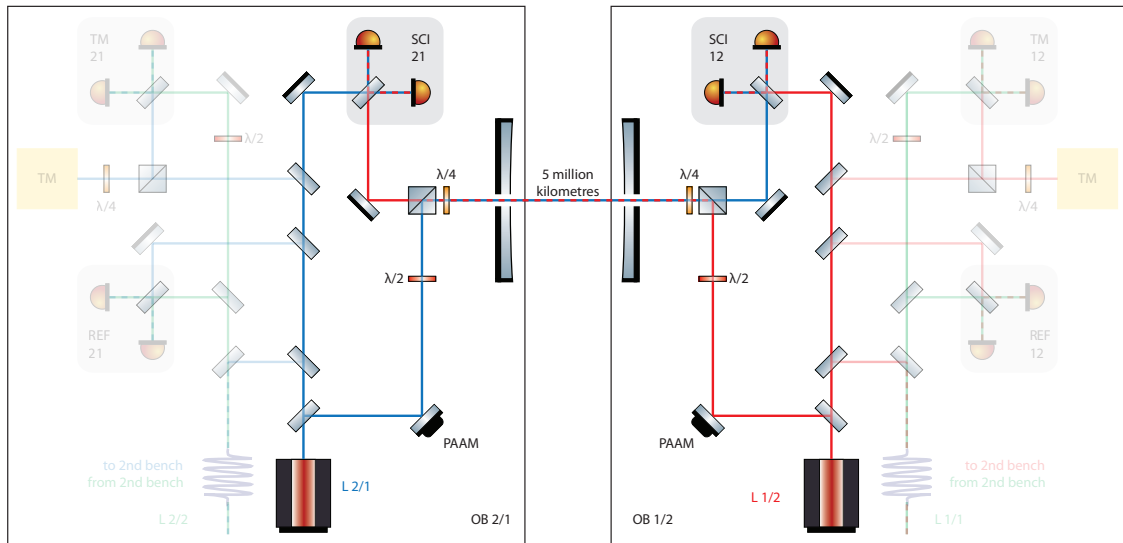
Figure 1.6: The LISA long arm interferometry.

One of the most important techniques to be developed for LISA is a shot noise limited phase lock between the local laser and the received weak beam. To the author, no experiment of a such phase lock is known that is sufficient for LISA.

The local interferometry for the test mass readout can in principle be copied from the technology demonstration mission for LISA: LISA Pathfinder. However, also in the local interferometry, challenging new features are included in the interferometer design. One of the most critical new elements is an optical fibre connecting both optical benches on-board each satellite. As already described, both local lasers interfere in the reference interferometers on both optical benches. In one of the interferometers both lasers are phase locked with respect to each other. One laser has the role of a frequency reference, and therefore, the beat in the second reference interferometer is free-running. This concept works only if the differential pathlength fluctuations in both reference interferometers are well below the required LISA sensitivity. This should in principle be the case, since the optical benches are symmetric and both beams are counter-propagating in the fibre link. However, the fibre is a birefringent element in the optical path and might have different effects on the two beams resulting in excess noise in the free-running beat note.

The optical bench itself is made of a low-expansion material in order to reduce pathlength changes introduced by thermal variations. For mechanical stability reasons, the optical components are fixed to the baseplate using *hydroxide-catalysis bonding*. By applying this technique the components and the baseplate form a quasi-monolithic body [9, 10].

In the following chapters, experiments investigating the weak light phase lock and the backside fibre link interferometer will be introduced and discussed. For the fibre link, an interferometer was built using the bonding technique mentioned.

Phase locking at picowatt power levels

One challenge in LISA is the phase lock of the outgoing lasers to the incoming weak light beams. Due to the long interferometer arms of LISA, the received beam is attenuated to about $100\,\text{pW}$ [11, 6]. In this chapter, first experiments in which a phase lock is performed with power levels comparable to LISA will be presented. The phase lock is implemented using analogue electronics. This will be different in LISA, but the outcome of these experiments is important, since the entire analogue front-end of the phase measurement system can be tested in this setup.

2.1 Phase lock for LISA

Each spacecraft has two lasers, whose light is sent out to the far spacecraft via telescopes as shown in Figure 1.2. At the same time these telescopes are used to receive the incoming light from the far spacecraft. Due to the large distance of $5 \cdot 10^9\,\text{m}$ between the satellites, it is not feasible to reflect the light directly back, since it is attenuated due to the beam divergence. The received power on one of the detectors on the far spacecraft can be approximated as

$$P_{\text{rec}} = \frac{\pi^2}{16}\frac{D^4}{\lambda^2 L^2}\eta_{\text{opt}}P_{\text{L}} \tag{2.1}$$

assuming that the beam radius at the transmitting telescope corresponds to the telescope radius [6]. D is the telescope diameter, λ the laser wavelength, L the LISA arm length and P_{L} the power of the outgoing beam. η_{opt} is the optical efficiency and accounts for all losses in the optical path. The laser wavelength is $1064\,\text{nm}$ due to the choice of the laser. The output power will be approximately $1\,\text{W}$. The received power scales with the 4th order of the telescope diameter. In

principle, the telescope should therefore be as large as possible in order to improve the LISA design sensitivity. However, a compromise must be found to keep the cost of the mission on a reasonably level. As a compromise, a 40 cm telescope is forseen in the LISA baseline design. This design also includes an optical efficiency of 23 % on the optical bench. This leads to a detected power,

$$P_{\text{rec}} \approx 127 \, \text{pW}, \tag{2.2}$$

of the received beam after it has passed the recombination beamsplitter in the science interferometer. This is an approximation based on the assumption, that two Gaussian beams interfere perfectly on the photo detector.

Due to the low power level of the incoming beam, a second (local) laser is phase locked with an offset frequency between 2 MHz and 20 MHz to the incoming beam. The two local lasers on each of the satellites are also phase locked with respect to each other. In contrast to the long arm, this lock is easier to implement, since the photodiode amplifiers do not require a huge amplification, and therefore, the electronic noise of the amplifier contributes less to the measured phase noise [3].

2.1.1 Shot noise limit

The biggest challenge in the phase lock for LISA is to handle the low power level of the incoming laser light. The available light power, calculated in Equation (2.2), yields to a fundamental limit in the residual phase fluctuations on the weak beam, which is called the *shot noise* limit. It is caused by the statistical phase distribution of the detected photons. Assuming a certain power, the phase noise cannot be measured better than

$$\tilde{\varphi}_{\text{sn}} = \sqrt{\frac{hc}{\lambda \eta P_{\text{rec}}}}, \tag{2.3}$$

where c is the speed of light and h the Planck constant [6]. The efficiency of the photodiode, η, is approximately 0.68 A/W. In LISA two quadrant photodiodes are used as main detectors. Therefore, the shot noise for each individual quadrant $\tilde{\varphi}_{\text{sn}}^{\text{Q}}$ can be calculated. Assuming a power of 32 pW per quadrant, the shot noise is

$$\tilde{\varphi}_{\text{sn}}^{\text{Q}} \approx 87 \, \mu\text{rad}/\sqrt{\text{Hz}}. \tag{2.4}$$

LISA is designed such that the pathlength measurement is almost shot noise limited at high frequencies – above 3 mHz. The overall sensitivity budget for one single arm in LISA is 12 pm/$\sqrt{\text{Hz}}$ relaxing with f^{-2} below 3 mHz. It is plotted as red curve in Figure 2.1 [3, 6]. The overall shot noise contributes with 8.6 pm/$\sqrt{\text{Hz}}$ (green) to the measurement sensitivity. As calculated in Equation (2.4), each quadrant has a higher shot noise level, since less power is impinging. When the signals from the four quadrants are combined, the effective shot noise in the combined signal corresponds to the total power on all 4 quadrants. For the laboratory experiments

presented in this thesis, the shot noise level for each quadrant is most important, since the first simple approach of the phase lock experiment includes only a single element photodiode. In this case the level of the blue curve in Figure 2.1 has to be reached in order to show a sufficient sensitivity.
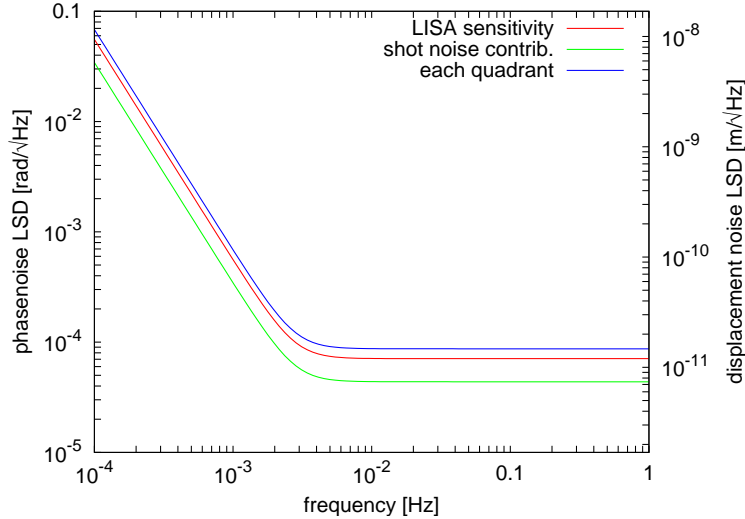
Figure 2.1: Phase lock related goals for LISA.

2.1.2 Phase lock alternatives

As already mentioned, direct back reflection of the laser light is not possible due to the strong attenuation of the laser beam. Therefore, the lasers are phase locked with respect to each other with variable offset frequencies of the order of 15 MHz. The necessity to use such high offset frequencies is mainly due to the Doppler shift present in LISA [12, 13].

The three satellites have elliptical orbits around the Sun. The orbits are chosen such that the spacecraft naturally remain in the formation of an equilateral triangle. This can be realised if the aphelion (the point with largest distance to the sun) is at different times in the year for each satellite. The aphelion of the second satellite is shifted by 60 degrees (one third of a year) with respect to the first spacecraft [2, 13]. The third satellites aphelion is again shifted by 60 degrees. Due to Kepler's second law, this results in different velocities of the satellites, which change during the year. In addition the Earth, the Moon and the other planets apply different forces on each satellite. The differential acceleration of the satellites results in a varying Doppler shift, which is different in each arm. The maximum relative velocity of the spacecraft is 15 m/s [3] which corresponds to a frequency change of about 15 MHz.

The implementation of the phase lock for LISA must be compatible with the readout of Doppler shifted frequencies that vary slowly (compared to the measurement band) in frequency during one year. In the following, three possible implementations will be discussed.

DC readout

The simplest solution for a phase lock implementation is a DC phase lock. The beatnote at the output of the photodiode amplifier can be mixed with an electrical reference oscillator. Both have the same frequency and the output of the mixer has contributions at the sum frequency and the difference frequency (DC) of the input signals. The sum signal is not important for the phase readout and can be filtered out. The signal at DC is the difference of both frequencies, which is the integrated differential phase of both oscillators. The phase lock is performed by adjusting the phase of one laser so that the output of the mixer is kept at zero.

For LISA, DC readout is not considered a viable option. The electrical oscillator would have to follow the Doppler induced laser frequency change. Initial lock acquisition would be difficult, because the mixer output is only proportional to the phase difference as long as the lasers have almost the same frequency, which is not the case when they are initially turned on in space.

An additional difficulty is the sensitivity of DC or homodyne readout to temperature variations. Amplifiers or active filters, which must be used to extract the phase signal from the mixer output and prepare it for the data acquisition, have usually temperature sensitive input offset voltages and currents. These are amplified and produce additional spurious signals, which disturb the measurement. This is not a fundamental problem but practically difficult to overcome. Heterodyne readout schemes usually have lower sensitivity to DC voltage drifts.

Intermediate heterodyne readout

An existing LISA Pathfinder phasemeter applicable for frequencies around $1\,\mathrm{kHz}$ is already available as an Engineering Model and will be available as flight hardware soon. One could make use of it also for LISA. One possible approach is to use an intermediate frequency and mix the photodiode output signal to $1\,\mathrm{kHz}$ and measure the phase of that signal instead of the directly sampled MHz signal. The performance of the LISA Pathfinder phasemeter would be sufficient for LISA. Nevertheless, more features would have to be included for LISA which would change both the soft- and hardware of the existing phasemeter. Therefore, the advantage of an available flight unit is lost. An additional intermediate frequency mixing step would have to be implemented including a kHz oscillator that is phase stable with respect to the reference clock on the satellite. The implementation of such a heterodyne readout scheme on-board LISA is currently not forseen.

Direct heterodyne readout

The most promising readout scheme is the direct heterodyne readout. Here the photodiode signal is directly sampled, and the phase estimation is performed directly on the MHz signal. In terms of phase measurement this detection method is theoretically simple but technically challenging. Such a phasemeter is of the same kind

as the LISA Pathfinder phasemeter, but operating at MHz-frequencies. It requires fast sampling with a frequency between 50 MHz and 80 MHz. A phasemeter, useful for laboratory experiments cannot be implemented on a standard PC. Due to the high sampling rate, the main calculations have to be performed in hardware. Even for first tests, a hardware phasemeter must be available with at least 3 analogue input channels (for 2 photodiodes and 1 electrical oscillator) and 1 analogue output channel, which are necessary to perform a phase lock between two lasers and have an independent out-of-loop phase measurement. Additionally, at least one analogue or digital output must be implemented in order to get the phase information out of the phasemeter.

An additional advantage of a digital phasemeter is the comparatively simple implementation of advanced modulation techniques for clock and data transfer between the satellites and tracking of the Doppler shift. This option is the current baseline design for the LISA mission.

2.2 Laboratory implementation of the phase lock

This section deals with the phase lock of two lasers under LISA-like conditions in a laboratory experiment. Due to the absence of a fast phasemeter, as it will be used for LISA, the lock is performed using a DC readout scheme. The scope of the presented experiment is the development of photodiode transimpedance amplifiers sufficient for LISA and the development of an independent readout scheme, which can be used for comparison with a digital phasemeter in the near future. The analogue detection chain including the photodiodes, their amplifiers and the analogue-to-digital conversion is susceptible to add noise to the phase measurement. Once a signal is digitised, the additional noise introduced by the phasemeter itself can be well characterised. Therefore, the test of such photodiodes and amplifiers is an important step for the LISA design verification.

Investigations on weak light phase lock experiments have already been performed: At the University of Glasgow an intermediate heterodyne readout scheme was developed focussing on electronics development [10]. For this setup a phase lock with weak light was not published up to now, but the electronics developed is sufficient for the LISA sensitivity. At Goddard Space Flight Center (NASA), a phase lock experiment with light powers comparable to LISA was performed [11]. The sensitivity reached is at high frequencies (above 100 mHz) almost at the shot noise level and decreases to lower frequencies with $f^{-0.5}$.

Groups at the Jet Propulsion Laboratory [14], Goddard Space Flight Center [15], University of Glasgow [10], University of Florida [16] and also at the AEI in Hannover [17] are working on a digital phasemeter for LISA, but this work is, at the present time, not focused on a phase lock with weak light power levels comparable to LISA.

This section focusses on the description and characterisation of the whole exper-

imental setup and readout. In the last part the results of the weak light phase lock are presented.

2.2.1 Optical setup

The aim of the investigations presented here is mainly the characterisation of a DC phase detection scheme and suitable photodiodes and amplifiers to be used for LISA. A simple Mach-Zehnder interferometer, shown in Figure 2.2, was built for this purpose using *hydroxide-catalysis bonding* (see also Chapter 4). Two beams are injected onto the optical bench and a fraction of both is split off (BSA1 and BSA2) for power stabilisation of the laser light. After splitting at BS3 and BS4, the beams travel through different paths in the interferometer. They are recombined with the beamsplitters BC1 and BC2. The difference of the two output phases is the differential pathlength change of both beams starting at the beamsplitters BS3 and BS4.

This optical bench was also designed to accommodate a frequency noise detection interferometer built by the components M1 and BC3 in Figure 2.2(b). The frequency interferometer would use the pathlength mismatch in the interferometer arms as frequency detector like in LISA Pathfinder (see Section 6.5.2). Its implementation was not necessary for the experiments presented here, but may be useful if the optical bench is used for other experiments. In this case only two additional components would be needed.

In order to characterise the readout electronics usually 'π-measurements' are performed. The interference signals leave the output of the recombination beamsplitter with a fixed phase difference of 180 degrees. This is due to energy conservation and offers the opportunity to test two independent detectors including their entire readout chain by comparing them to each other. This measurement reduces the experimental effort compared to 'real' interferometric length measurements and avoids the additional unknown pathlength fluctuations. In the LISA measurement band, experiments cannot usually be performed in air and must therefore be set up inside a vacuum chamber. For interferometric measurements with an accuracy of about $10 \, \mathrm{pm}/\sqrt{\mathrm{Hz}}$, setups on aluminium baseplates are not thermally stable enough (see Section 3.3.2) and interferometers with low thermal expansion have to be built. However, π-measurements can be performed in air on aluminium breadboards, since no pathlength changes are measured.

Even in π-measurements, noise sources might appear which are produced by the modulation or locking scheme of the input beams. From experiences in the LISA Pathfinder experiment, it is known that noise is introduced into the measurement by additional signals at the heterodyne frequency. In LISA Pathfinder both beams are modulated by Acousto-Optical Modulators (AOMs). Spurious signals are generated due to electrical cross-talk and pickup on their input signals, which are shown in Figure 2.3. A copy of these spectra is therefore also present on the two interfering beams. A detailed description of this noise source can be found in [18, 19, 20, 21].
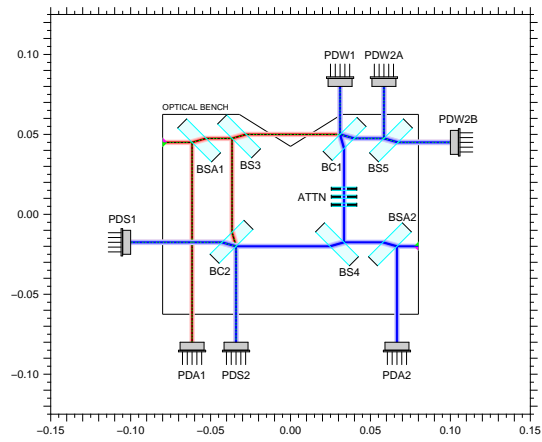
(a) Schematic layout.
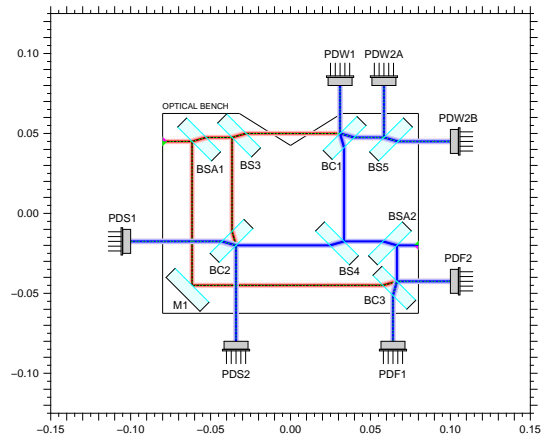
(b) Frequency interferometer.
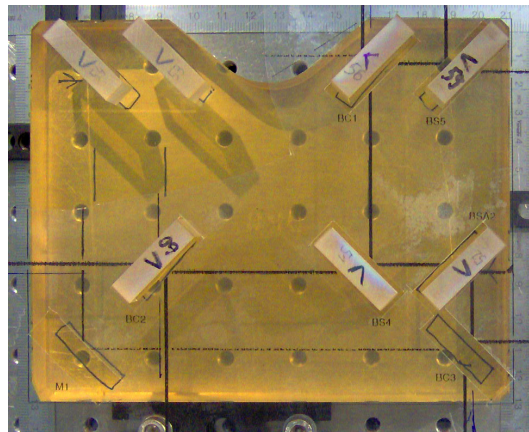
(c) Bonded interferometer.

Figure 2.2: Mach-Zehnder interferometer designed for the phase lock experiment.
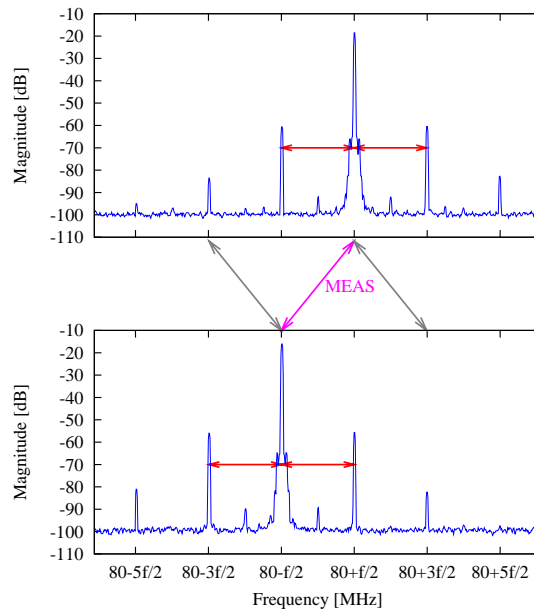
Figure 2.3: Typical spectra of the AOM input signals in the LISA Pathfinder experiment.

The phasemeter measures signals only at the heterodyne frequency. The signal carrying the 'real' length information is indicated by the magenta arrow in Figure 2.3. But additional spurious signals are also measured by the phasemeter. One type of those is indicated by the red arrows. It is an interference of the carrier signal with its own sideband. This modulation is already on the beam when it arrives at the recombination beamsplitter and therefore split with the same phase at both output ports. This introduces a spurious signal on the 'real' interference signal, which does not cancel in a π-measurement. The reason is that the spurious signal leaves the beamsplitter with the same phase, whereas the desired signal leaves with 180 degrees phase difference.

A second type of spurious signal is indicated by the grey arrows in Figure 2.3. It is the interference of the carrier of one beam with a second order sideband of the second beam. Since this signal is produced at the recombination beamsplitter, it leaves the two output ports with 180 degrees phase difference as the 'real' signal does. Therefore, the estimated phase might be calculated slightly wrong, but noise introduced by this signal cancels in the π-measurement.

The spectra of two lasers which are phase locked with respect to each other with a fixed offset frequency have a different shape compared to beams modulated by AOMs. Nevertheless, similar spurious signals might also be produced in the case of phase locked lasers. In order to avoid these complications in the characterisation of the readout and photodiode electronics, an additional beamsplitter (BS5) is forseen, which splits the interference signal of one interferometer for diagnostic purposes. At its two output ports, both beams always leave with the same phase, and in the difference only the electronic readout noise can be observed.

In Figure 2.2 an attenuator is also shown. It attenuates only one beam in the

upper right interferometer. It can be used to decrease the light power of one beam to pW level. In that case the upper right interferometer can be used to phase lock both lasers under conditions comparable to LISA. The second interferometer can then be utilised to perform an out-of-loop measurement of the phase stability. The advantage obtained is a beatnote in the lower left interferometer with two bright beams. In this interferometer the complications of a weak light phase lock do not have to be taken into account.

In the experiments presented here, this attenuator was not implemented. It is not possible to perform long-term measurements with this interferometer using DC readout with analogue mixers. On typical time scales for measurements (about 10 h) the thermally driven length change in the interferometer is of the order of a few cycles, where the analogue mixer leaves its linear range. However, it could be used with intermediate heterodyne readout. When a digital phasemeter using direct heterodyne readout is built, this feature of the setup can be used.

The Mach-Zehnder interferometer was built with similar methods as described in Chapter 4. However, due to the simplicity of the layout, the interferometer alignment is less critical and is not discussed further in this thesis.

2.2.2 Phase detection

As already mentioned, the phase lock experiment presented here is implemented using DC readout and analogue loop filters. The loop must have a high suppression at low frequencies, since the free-running phase noise of a laser is high. It can be calculated from the free-running frequency fluctuations of the laser which are shown in Figure 2.4 (green curve).
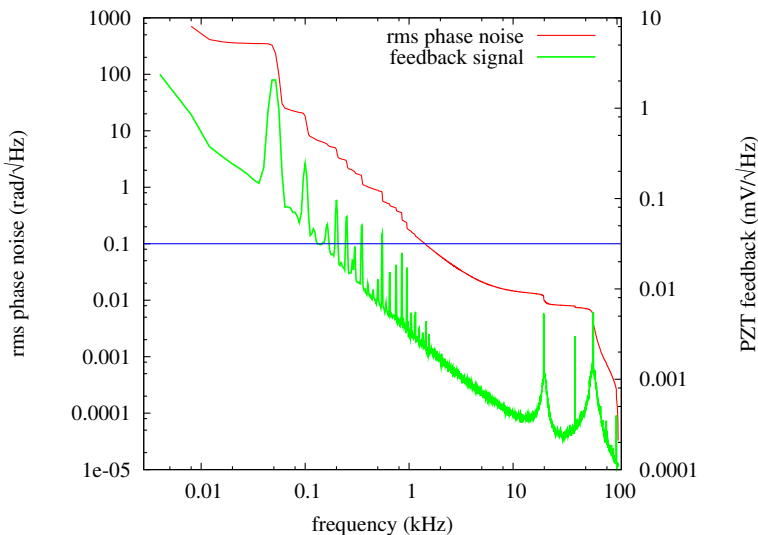
Figure 2.4: Free-running frequency and phase noise of a Nd:YAG laser.

The angular laser frequency ω_L is given by the time derivative of the laser phase φ_L:

$$\omega_L = \frac{d}{dt}\varphi_L. \tag{2.5}$$

In the Fourier domain, the derivative corresponds to multiplication with $i\omega$,

$$\frac{d}{dt} \curvearrowright i\omega, \tag{2.6}$$

with the angular Fourier frequency ω. Equation (2.5) can be written in the Fourier domain as

$$\mathcal{F}(\omega_L) = i\omega\mathcal{F}(\varphi_L), \tag{2.7}$$

where $\mathcal{F}$ is the applied Fourier transform. Taking into account that $\omega = 2\pi f$ for the Fourier frequency and $\omega_L = 2\pi\nu_L$ for the laser frequency, the linear spectral density of the phase noise can be expressed as

$$\tilde{\varphi}_L(f) = \frac{\tilde{\nu}_L(f)}{f}. \tag{2.8}$$

The laser is locked to a cavity that is used as frequency reference. Since the frequency fluctuations are suppressed by the loop, the free-running frequency fluctuations are calculated via the laser feedback signal, which is calibrated. The scaling factor is $2.7\,\text{MHz/V}$. From this measurement the rms of the free-running phase fluctuations can be calculated using the relation

$$\varphi_{\text{rms}}(f) = \left[\int_\infty^f df'\frac{\tilde{\nu}_L{}^2(f')}{f'^2}\right]^{1/2}. \tag{2.9}$$

This value is integrated from the high frequency end to lower frequencies. This can be done since the fluctuations at high frequencies can be assumed to be small. The result gives the residual phase fluctuations at the unity gain frequency and is illustrated as red curve in Figure 2.4.

The response of an analogue mixer to linear differential phase changes of its input signals is sinusoidal. Around a phase difference of 90 degrees, the output is around a zero-crossing. In a range of about $0.1\,\text{rad}$, it can be assumed to be linear (blue line). From this spectrum it can be seen that the loop bandwidth has to be well above $1\,\text{kHz}$. Otherwise the mixer leaves its linear range and the lasers may lose their lock. Sufficient suppression at low frequencies can be realised with an appropriate amount of integrators in the servo. A schematic view of the control loop for the phase lock is shown in Figure 2.5.

Two lasers interfere at a beamsplitter. Their electric fields can be written as

$$E_M = A_M \exp\left[i\left(\omega_M t + \varphi_M\right)\right], \tag{2.10}$$
$$E_S = A_S \exp\left[i\left(\omega_S t + \varphi_S\right)\right], \tag{2.11}$$
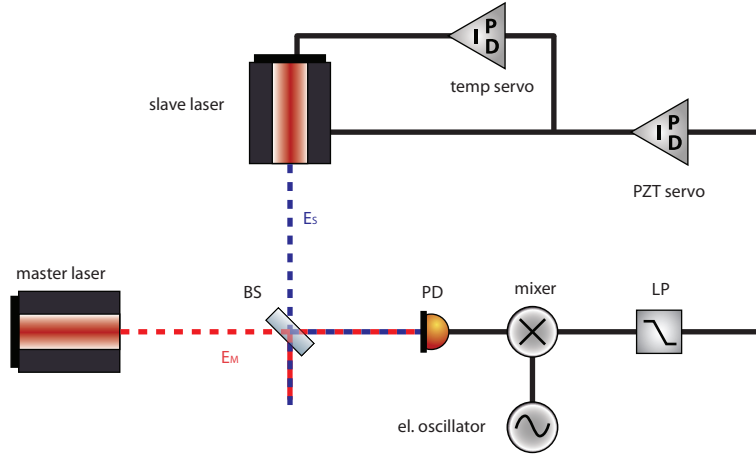
Figure 2.5: Stabilisation loop of the weak light phase lock experiment.

where the indices M and S refer to the incoming weak master and the local slave laser. ω is the laser frequency, φ its phase and A the amplitude. The formalism used here is similar to the discussions in [22]. The interfering field is the sum of both input fields and the photodiode detects the power

$$P = |E_{\mathrm{M}} + E_{\mathrm{S}}|^2. \tag{2.12}$$

This leads to the interfering power

$$P = A_{\mathrm{M}}^2 + A_{\mathrm{S}}^2 + 2cA_{\mathrm{M}}A_{\mathrm{S}}\cos\left[\left(\omega_{\mathrm{M}} - \omega_{\mathrm{S}}\right)t + \left(\varphi_{\mathrm{M}} - \varphi_{\mathrm{S}}\right)\right] \tag{2.13}$$

that is detected by the photodiode. The first two terms appear as DC power in the readout. The second term is oscillating at the frequency difference,

$$\omega_{\mathrm{het}} = \omega_{\mathrm{M}} - \omega_{\mathrm{S}}, \tag{2.14}$$

and has a phase of

$$\Delta\varphi = \varphi_{\mathrm{M}} - \varphi_{\mathrm{S}}. \tag{2.15}$$

The index 'het' indicates that ω is the heterodyne frequency. Additionally, the contrast c is added in order to take imperfections in the interference into account. The oscillating signal is the important one for the phase lock. The photodiode converts it into a current with an efficiency of η. The resulting AC photo current can then be written as

$$I_{\mathrm{AC}} = 2\eta cA_{\mathrm{M}}A_{\mathrm{S}}\cos\left(\omega_{\mathrm{het}}t + \Delta\varphi\right). \tag{2.16}$$

Transimpedance amplifiers convert the photo current into a voltage by measuring the voltage across the transimpedance resistor. The output of the amplifier is

$$U_{\mathrm{AC}} = 2c\eta RA_{\mathrm{M}}A_{\mathrm{S}}\cos\left(\omega_{\mathrm{het}}t + \Delta\varphi\right), \tag{2.17}$$

where R is the resistance of the transimpedance resistor. Here it is assumed that the response of the amplifier is flat. For simplification all prefactors are merged in

$$U_{\mathrm{PD}} = 2c\eta R A_{\mathrm{M}} A_{\mathrm{S}} \tag{2.18}$$

such that the output voltage can be written as

$$U_{\mathrm{AC}} = U_{\mathrm{PD}} \cos\left(\omega_{\mathrm{het}} t + \Delta\varphi\right). \tag{2.19}$$

This signal is fed into an analogue mixer and compared with an electrical reference oscillator running at the same frequency. Its signal can be written as

$$U_{\mathrm{ref}} = U_{\mathrm{el}} \cos\left(\omega_{\mathrm{het}} t + \varphi_{\mathrm{el}}\right) \tag{2.20}$$

The mixer multiplies both input signals. Using standard trigonometric identities, its output is

$$U_{\mathrm{mix}} = \frac{U_{\mathrm{PD}} U_{\mathrm{el}}}{2} \left[\cos\left(\Delta\varphi - \varphi_{\mathrm{el}}\right) + \cos\left(2\omega_{\mathrm{het}} + \Delta\varphi + \varphi_{\mathrm{el}}\right)\right]. \tag{2.21}$$

It consists of two terms. The first one is the DC term, which depends on the phase difference between the measured signal and the electrical reference oscillator. The second term is at the sum frequency and not used for the phase lock. Therefore, the signal is low-pass filtered in order to keep only the information of the phase difference, which is used as error signal for the phase lock and fed into a servo. The actuator is a cavity length stabilising piezo inside the laser head. The dynamic range of this actuator is quite small (a few 10's of MHz). Since typical low frequency fluctuations of the laser frequency over the duration of a measurement are higher, the piezo feedback signal is used as error signal for a second servo that actuates on the temperature of the laser head. This way the feedback signal for the piezo is kept around the centre of its dynamic range.

From Equation (2.5) it can be seen that the frequency is the time derivative of the phase, or equivalently, the phase is the frequency integral. For the loop presented here, this means that it has an intrinsic integrator, since the phase is measured, but the laser frequency is actuated. As a consequence the mixer output is not the true error signal, but already integrated. In these kind of loops, the frequency error signal is not directly accessible.

2.2.3 Data acquisition

The data acquisition is done by sampling the output of the low-pass filtered analogue mixer as indicated in Figure 2.6. An additional amplifier is needed to optimally use the dynamic range of the data acquisition (DAQ) system. Two signals are sampled. The in-loop signal is also used as servo input. A second, out-of-loop, signal is sampled in order to characterise the photodiodes, mixers and low-pass

filters. For the sampling system, a commercially available sampling card from National Instruments (NI6014) is used. It has 16 analogue input channels with a total maximum sampling rate of 200 kilosamples/second. The input channels are multiplexed and sampled by a single 16 bit Analogue-to-Digital Converter (ADC). The signals are adequately anti-alias filtered before being fed into the ADC of the DAQ.
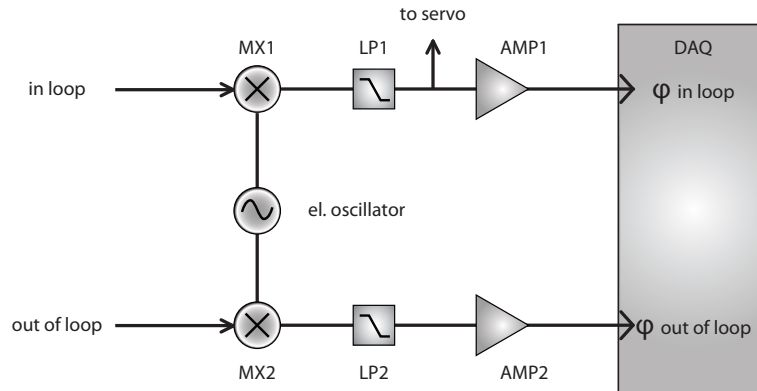
Figure 2.6: Data acquisition for the phase measurement at DC.

The in-loop signal is always derived from a photodiode output signal that is used to produce the error signal. The so-called out-of-loop signal on the other hand has its origin from different sources depending on the measurement performed. For characterisation of the readout electronics – including mixer, low-pass filter and amplifier – the output of one photodiode is split. For tests of the photodiode amplifier, two independent photodiodes deliver the input signals for the mixers in a π-measurement.

In order to convert the measured voltages to phase, the low-pass filtered mixer output must be scaled. In principle the scaling factor can be calculated from a combination of many parameters, but usually for each measurement the photodiode signals are temporarily substituted by synthetic signals from frequency generators. The amplitude of the photodiode is synthetically re-generated, and the phase is varied with respect to the electrical oscillator by a few degrees. The phase variation divided by the output voltage gives the scaling factor, which is directly measured with an accuracy only limited by the precision of the frequency generator. In the laboratory experiment this was the model DS345 from Stanford research.

2.2.4 Sensing noise of the Phase Measurement System

The initial step of the investigations presented here was the implementation of a Phase Measurement System (PMS) with sufficient sensitivity. Initial noise hunting and improvement of the electronics were done with high light power levels (of the order of 1 mW for both interfering beams) and simple photo diode pre-amplifiers.

For later weak light experiments, these photodiodes were replaced by other devices with dedicated amplifiers.

The servo used is a 2-stage controller using five integrators as shown in Figure 2.7. The mixer output is filtered with two second order active low-pass filters in order to remove the 2f-component and produce a 'clean' error signal for the controller. The first servo stage consists of the intrinsic integrator (see Equation (2.5)) plus two additional damped integrators for the path that actuates on the laser frequency via the PZT. The second stage includes three damped integrators and the laser temperature controller is used as actuator. The corner frequencies of the filter elements are also indicated in this figure.
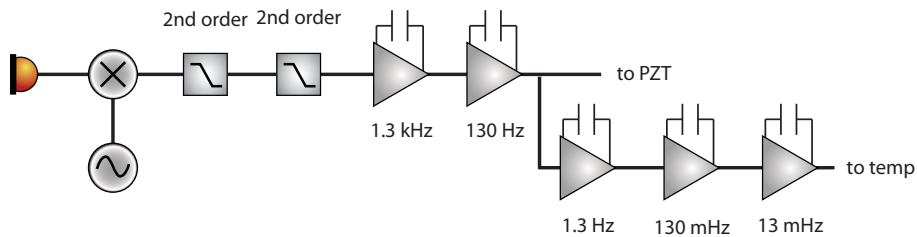
Figure 2.7: Phase lock controller.

With such a servo the in-loop phase stability shown as the red curve in Figure 2.8, can be reached. The first out-of-loop sensitivity, measured with a second independent photodiode (green curve), is reached on a simple tabletop setup. This noise level improves when both electronics and optical setup are covered by a box made of Plexiglas that reduces thermal variations and air flow around the components (blue curve). A time series of the measured in- and out-of-loop phase noise is shown in Figure 2.9.
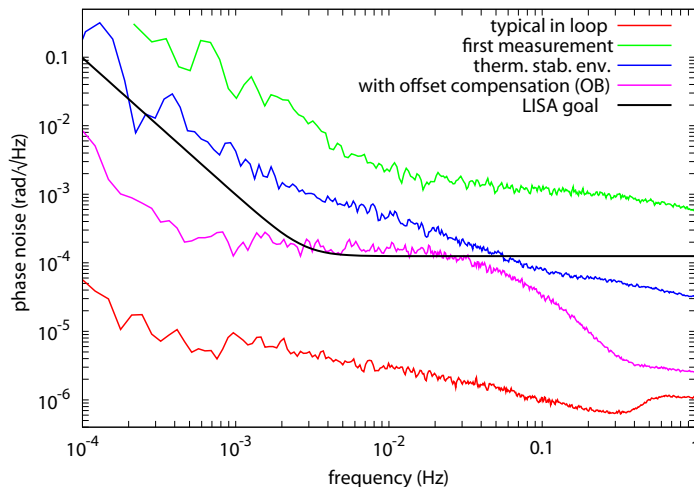
Figure 2.8: Improvements in sensitivity of the phase measurement system.
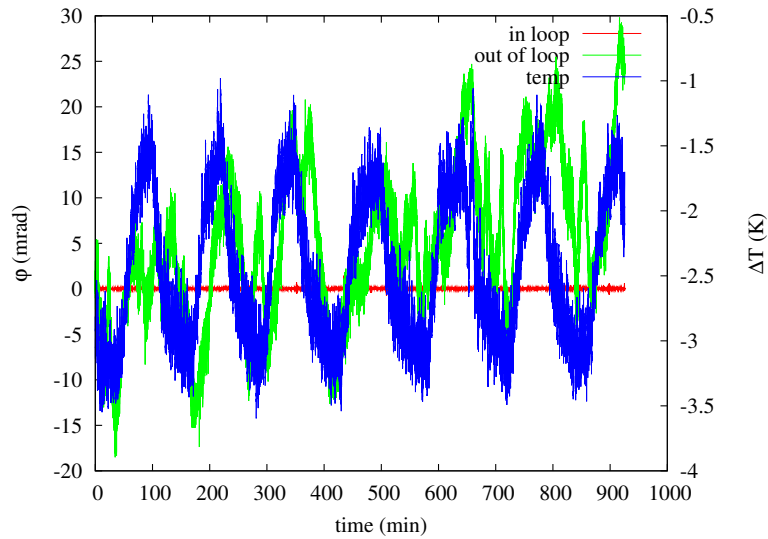
Figure 2.9: Temperature dependency of the measured phase.

The next step was the identification of the component where these fluctuations couple. These turned out to be the low-pass filters used to suppress the double frequency of the mixer output. The final low-pass filter design, which will be introduced in the following section, leads to a sensitivity of about $2 \cdot 10^{-4} \, \text{rad}/\sqrt{\text{Hz}}$ appearing as almost white noise between $1 \, \text{mHz}$ and $100 \, \text{mHz}$ (magenta curve).

2.2.5 Low-pass filter design

In this section the low-pass filter and its development is presented, since it turned out to be the dominant noise source in the analogue phase lock experiment.

Filter design aspects

Many aspects have to be taken into account when choosing the cut-off frequency of the low-pass filter. A low cut-off frequency reduces the loop bandwidth of the phase lock, and if the cut-off frequency is too high, it is impossible to find suitable amplifiers with sufficiently low noise.

The most critical part for the LISA phase lock is expected to be the photodiode amplifier. In the first amplification stage the gain should be as high as possible in order not to introduce additional amplifier noise in subsequent stages to the output signal. On the other hand this amplifier must have enough bandwidth. In this respect the test of the phase lock is performed under worst case conditions – meaning performance of the experiment at the upper end of heterodyne frequencies. For the results presented here, it is $18 \, \text{MHz}$ or $20 \, \text{MHz}$.

The function of the filter is the suppression of the double offset frequency of the phase lock, which is around $40 \, \text{MHz}$ in this experiment. In principle, the double

frequency does not have an influence on the phase lock as long as it does not exceed unity gain in the open loop gain. However, it is practically necessary to sufficiently suppress this contribution in order to minimise additional noise introduced by non-linear effects in the amplifiers. For this reason, an active low-pass filter with two second order stages was chosen.

The cut-off frequency of the filter has to be high compared to the unity gain frequency of the loop. If it is not, the phase loss at unity gain is too high and the lock cannot be performed. Assuming a cut-off frequency of 1 MHz and a fourth order low-pass filter as used here, the phase lag at 100 kHz is already 20 degrees. As discussed in Section 2.2.2 a unity gain frequency of about 10 kHz is necessary to perform the phase lock. Therefore, a cut-off frequency of 1 MHz is reasonable. Also it is worthwhile to have gain in the filter which improves the signal-to-noise ratio.

Filter implementation

The issues discussed above determine the choice of the amplifiers due to their gain-bandwidth-product. An additional important specification is their input current and voltage noise, but that is only defined at high frequencies and is of the order of $1\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$ and $1\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$. At low frequencies, the output voltage noise is dominated by the temperature sensitivity of the input offset current and voltage. The initial implementation of the low-pass filter was a Sallen-Key topology [23], which is realised using two different operational amplifiers (OpAmps) – OPA690 and AD797. The temperature coefficients of the offset current and voltage are of the order of $10\,\mathrm{nA/K}$ and $10\,\mu\mathrm{V/K}$.

None of these filters leads to a sufficient performance of the phase lock. The solution to this problem is a *composite amplifier*. Special chopper amplifiers that have a minimal offset voltage drift exist. They are usually quite slow and unsuitable for MHz applications. However, it is possible to use the trim inputs of fast operational amplifiers to adjust their offset voltage. The chopper operational amplifier measures the offset and feeds back to the fast amplifier. This leads to technical problems, since there are only a few fast amplifiers available that have trim inputs. Furthermore, trimming the amplifiers is done by feeding a certain current to their inputs at a voltage level near one of the supply voltages. The estimation of this current is only possible if the internal circuitry of the amplifier is known, which is usually not the case. However, application notes are available explaining how composite amplifiers can be built.

The input offset voltage is the voltage difference between the inverting and non-inverting input of the operational amplifier. In order to measure this difference, it is worthwhile to change the topology from Sallen-Key to *multiple feedback* [23]. The advantage of the latter is that its non-inverting input has ground potential. The voltage at the other input can therefore be directly measured against ground and its offset can be compensated. The final filter implementation is shown in Figure 2.10.

As the fast OpAmp, an LT1028 was chosen, since it has suitable characteristics and an application note, describing how to use it as composite amplifier, is available [24]. The offset compensation is done by an AD8628, which has an offset voltage drift of $30\,\mathrm{nV/K}$ and a supply voltage of $\pm 2.5\,\mathrm{V}$. The unity gain frequency of the feedback is chosen such that the output voltage noise is dominated by the noise of the fast amplifier at high frequencies and by the voltage noise of the chopper OpAmp at low frequencies. The lowest out-of-loop phase noise in Figure 2.8 was achieved using the described implementation of the composite amplifier.
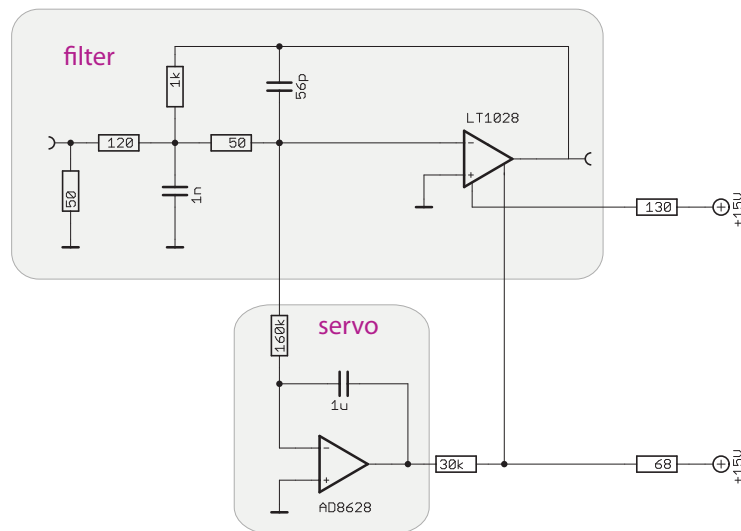
Figure 2.10: Low-pass filter built as composite amplifier.

Measurement sensitivity

The sensitivity is limited by the photodiode amplifiers that were used for the first experiments with mW power level. This can be seen in Figure 2.11. The green curve is the best out-of-loop measurement from Figure 2.8, which was achieved by a π-measurement. The blue curve shows a measurement where only one photodiode is used instead of two. Its output signal is split. One part is used to produce an error signal for the phase lock, and the second part is used as out-of-loop measurement signal. The out-of-loop signal is not completely independent from the in-loop signal, since a common photodiode is used for both paths. However, this measurement characterises the electronic noise of the phase measurement system (PMS) including the entire readout chain, apart from the photodiode. The sensitivity of the PMS is well below the required readout noise for LISA.
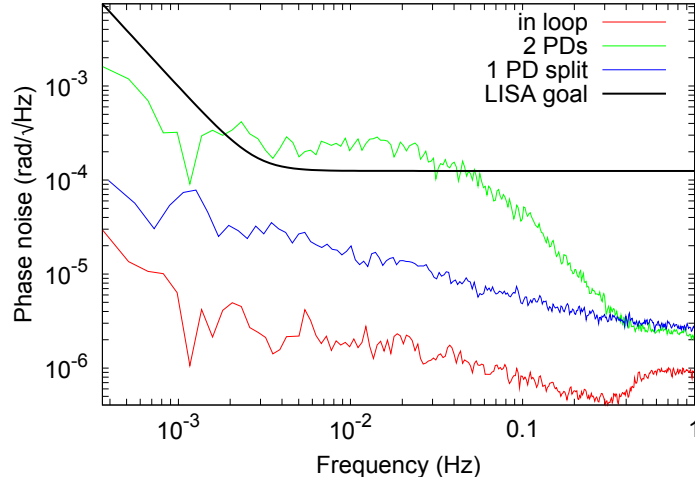
Figure 2.11: Measurement sensitivity of the phase lock electronics.

2.3 Weak light phase lock

For the measurements at low light power levels, the same readout electronics as described in the previous section is used. The only differences are the photodiode amplifiers that are specifically made for weak light phase lock applications. In this section the additional noise that appears if a weak light phase lock is performed is analytically described and experimental results are presented. The phase lock is performed with light powers of the order of a few 10's of pW. This work was done at the AEI Hannover in collaboration with Christian Diekmann, Benjamin Sheard and Gerhard Heinzel. The results are already published in a diploma thesis [25].

2.3.1 The photodiode

Initially a suitable photodiode had to be found for the weak light phase lock. The preferable semi-conductor material is Indium-Gallium-Arsenide (InGaAs). Photodiodes made of this material usually have a quantum efficiency of more than 90 % at 1064 nm. This type of photodiode is the only convenient option for LISA. The chosen model has a diameter of 0.5 mm. Its properties are summarised in Table 2.1. It is from PerkinElmer with the serial number C30619G.

Different suppliers provide photodiodes with similar properties, e.g., [27, 28]. The *noise equivalent power* (NEP) is small enough and not limiting the measurement. The responsivity is not well specified, and given by all suppliers with about 0.68 A/W, which corresponds to a quantum efficiency of about 80 %. Most important is the capacitance of the photodiode, which scales with the sensitive area. It limits the bandwidth of the sensor. A diameter of 0.5 mm is reasonable for the purpose of a weak light phase lock with an offset frequency of about 20 MHz. Small photodiodes would simplify the choice of the amplifier, but for LISA large photo-

Table 2.1: Summarised photodiode properties [26].

property	value
diameter (mm)	0.5
responsivity (A/W)	> 0.68
dark current (nA)	5
noise equivalent power (@1300 nm) (pW/$\sqrt{\text{Hz}}$)	< 0.1
capacitance (@ 5 V bias voltage) (pF)	8
maximum bias (V)	5

diodes are desirable, since quadrant detectors are needed and their slit size should be small compared to their sensitive area. For the capacitance, all supplier specifications are quite similar. The major difference between the various photodiodes is the maximum bias voltage that can be applied. Increasing the bias voltage decreases the capacitance and thereby the bandwidth of the detector increases. The combination of specified maximum bias voltage and parasitic capacitance was the reason for the choice of the model used for the weak light phase lock experiment.

2.3.2 The amplifier

The main challenge in the phase lock experiment is the design of a suitable amplifier with sufficient low output voltage noise compared to the shot noise. This amplifier produces an additional current that contributes to the measured signal. The relation of the input noise sources to the output voltage noise of the photodiode amplifier will be worked out in this section.

The shot noise current spectral density $\tilde{I}_{\text{sn}}$ can be calculated from the DC photo current I as

$$\tilde{I}_{\text{sn}} = \sqrt{2eI} = \sqrt{2ec\eta\left(P_{\text{S}} + P_{\text{M}}\right)} \approx \sqrt{2ec\eta P_{\text{S}}}, \tag{2.22}$$

using the approximation that the master laser power is considerably smaller than the slave laser power. The input current noise of the amplifier is required to be small compared to the shot noise current. The influence of the shot noise current can be reduced by increasing the slave laser power, since the signal scales with the slave laser amplitude as can be seen from Equation (2.13). Therefore, the signal-to-noise ratio

$$SNR \propto \frac{\sqrt{P_{\text{S}}P_{\text{M}}}}{\sqrt{P_{\text{S}}}} = \sqrt{P_{\text{M}}} \tag{2.23}$$

depends only on the master laser power. However, the slave laser power should not be too high for two reasons: Firstly, only a small fraction of the beam should be used to perform the phase lock, since most of the light has to be transmitted to the far spacecraft. Secondly, the power consumption on the optical bench should

be as small as possible in order to reduce the coupling of thermal effects into the measurement.

A reasonable assumption for the beam powers of the slave laser is $1\,\mathrm{mW}$ on one of the photo detectors. The power of the master laser is fixed to about $100\,\mathrm{pW}$ at the same position. For a typical responsivity of InGaAs photodiodes, this leads to a DC photo current of $0.68\,\mu\mathrm{A}$, which is dominated by the slave laser. The AC photo current is about $43\,\mathrm{nA_{peak}}$ and the shot noise of the local oscillator beam can be calculated using Equation (2.22) to be $4.7\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$. The bandwidth of the amplifier should be sufficiently larger than the Doppler shift of about $15\,\mathrm{MHz}$. In conclusion the chosen amplifier must have the following properties [29, 30]:

- a minimum bandwidth from $2\,\mathrm{MHz}$ to $20\,\mathrm{MHz}$,

- an equivalent input noise sufficiently smaller than $4.7\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$,

- low power dissipation and

- low phase dependency to temperature fluctuations.

A transimpedance amplifier as shown in Figure 2.12 was chosen as basic configuration.
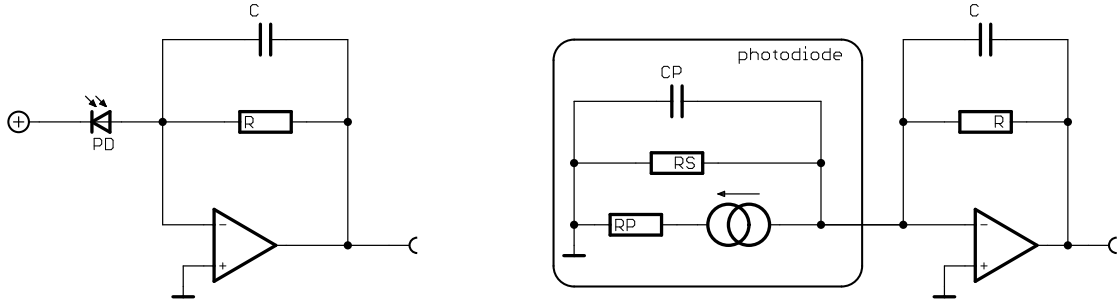
Figure 2.12: Simple transimpedance amplifier design.

The inverting input of the amplifier has a high impedance. Therefore, the photo current, I_p, flows through the transimpedance resistor, R. Since the inverting input has the same potential as the non-inverting input, which is grounded, an output voltage,

$$U_\mathrm{PD} = RI_\mathrm{p}, \tag{2.24}$$

proportional to the photo current is produced. This equation is an approximation that can be used in the flat part of the frequency response of the transimpedance amplifier. The capacitor, C, reduces the quality, Q, of the low-pass characteristic of the frequency response and is tuned such that the amplifier bandwidth is flat and maximised. The equivalent circuit diagram is shown on the right hand side of

the figure, which takes the parasitic photodiode properties into account. Most important here is the parasitic capacitance C_p, which limits the amplifier bandwidth. The maximum bandwidth of such a circuit with a flat response is

$$\omega_{\mathrm{max}} = \frac{\omega_{\mathrm{T}}}{\sqrt{RC_p\omega_{\mathrm{T}} + \sqrt{2RC_p\omega_{\mathrm{T}} - 1}}}, \tag{2.25}$$

where ω_{T} is the transit frequency of the operational amplifier [31].

Different noise sources add up at the amplifier output: The *Johnson noise*, the input current and voltage noise ($\tilde{I}_{\mathrm{op}}$ and $\tilde{U}_{\mathrm{op}}$) of the operational amplifier and the shot noise from Equation (2.22) have to be taken into account. They appear as output voltage noise which refers to an equivalent input current noise of

$$\tilde{I}_{\mathrm{in}} = \frac{\tilde{U}_{\mathrm{PD}}}{R} = \sqrt{\frac{4kT}{R} + \tilde{I}_{\mathrm{op}}^2 + \tilde{I}_{\mathrm{U}}^2 + 2eI_p}. \tag{2.26}$$

The input current noise, I_{U}, is equivalent to the input voltage noise, $\tilde{U}_{\mathrm{op}}$, and can be calculated as

$$\tilde{I}_{\mathrm{U}} = \frac{\tilde{U}_{\mathrm{op}}}{R}\left(1 + 2\pi f_c RC_p\right), \tag{2.27}$$

as described in [29]. For the properties assumed here, the input voltage noise has to be $0.75\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$ for an equivalent input current noise of $1\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$.

The Johnson noise is a current produced by Brownian motion of the electrons inside a resistor R and can be expressed as

$$\tilde{I}_{\mathrm{jn}} = \sqrt{\frac{4kT}{R}}, \tag{2.28}$$

where k is the Boltzmann constant and T the temperature [31]. The Johnson noise is shown for room temperature in Figure 2.13(a). This is the temperature at which the laboratory experiment and LISA operate. Slightly above $10\,\mathrm{k\Omega}$, it is about $1\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$. Therefore, the transimpedance resistor should be at least $15\,\mathrm{k\Omega}$. Using this value and a parasitic capacitance of $10\,\mathrm{pF}$, the closed-loop bandwidth of the transimpedance amplifier can be calculated using Equation (2.25). The photodiode capacitance is assumed to be slightly higher than specified. It is a conservative estimate that accounts for additional capacitances on the circuit board due to wiring and the input capacitance of the operational amplifer. The minimum transit frequency must therefore be $400\,\mathrm{MHz}$ for a $20\,\mathrm{MHz}$ cut-off frequency, f_c.

As a conclusion it is found that a *state of the art* operational amplifer is needed to reach the required noise level. The one that was chosen for the first approach is the LMH6624. Other amplifier topologies are also possible and may be suitable for the application. One of those is an amplifier using discrete transistor amplification stages, which was also used for the phase lock experiment presented here. The results of the measurements presented in the following section were achieved with these two different topologies. They were developed by Gerhard Heinzel and Benjamin Sheard, and their schematic layout is shown in Appendix A.

(a) Johnson noise.

(b) Amplifier bandwidth.
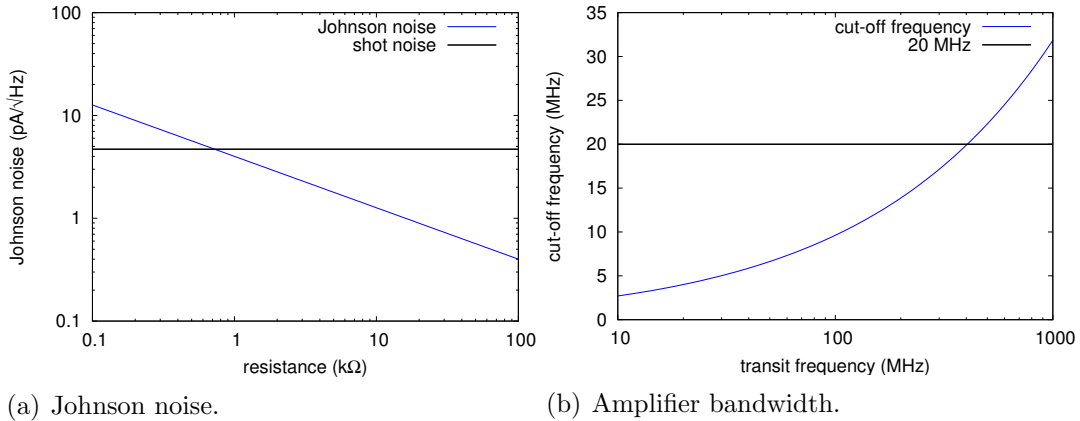
Figure 2.13: Effect of transimpedance resistor value on current noise and bandwidth.

2.3.3 Experimental results of the weak light phase lock

For the two photodiode designs introduced in Section 2.3.2, a phase lock was established with weak light power of the master laser. The parameters of the experiment are summarised in Table 2.2. The master laser power is higher than the expected power for LISA. Therefore, the phase noise equivalent shot noise limit from Equation (2.4) is reduced in the experiment. The transistor amplifier has slightly lower input current noise. However, for both amplifier designs it is of the order of a few $pA/\sqrt{Hz}$. For compensation of this mismatch, the measurements for the OpAmp design are performed with higher local oscillator power.

Table 2.2: Experimental parameters of the weak light phase lock.

	OpAmp design	Transistor design
master laser power (pW)	100	30
slave laser power (mW)	1	0,2
amplifier input current noise ($pA/\sqrt{Hz}$)	4	2.5
shot noise equivalent ($\mu rad/\sqrt{Hz}$)	47	84

The resulting phase noise is shown in Figure 2.14.   Both in-loop signals are below the required LISA read out noise, whereas the out-of-loop signals are slightly above. In the worst case, at $4\,mHz$, they are a factor of 5 worse. Conspicuously both amplifiers show the same performance. This can be understood by looking at Figure 2.15. The blue curve is achieved by splitting one photodiode signal for the in- and out-of-loop measurement path (as already shown in Figure 2.11 for the high power phase lock). Below $50\,mHz$ the measurement is limited by the readout sensitivity of the measurement system. It decreases to higher frequencies with $1/f$ until it almost reaches the in-loop noise.
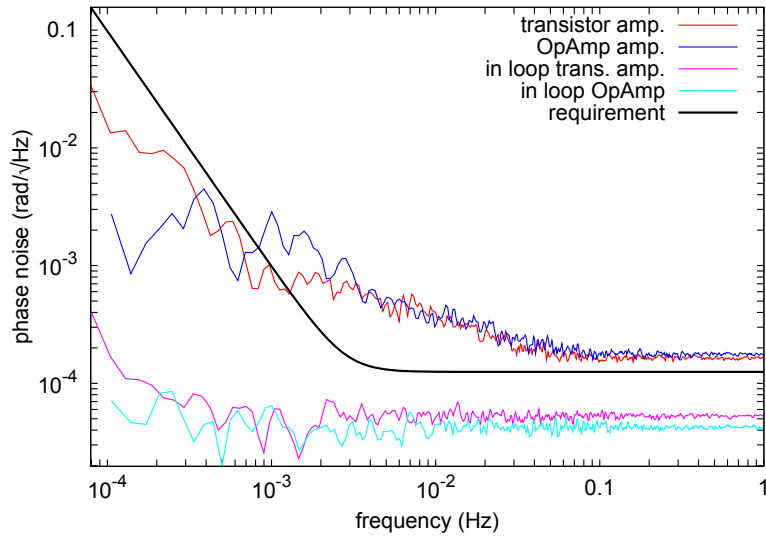
Figure 2.14: Comparison between the phase noise of the photodiode amplifiers.
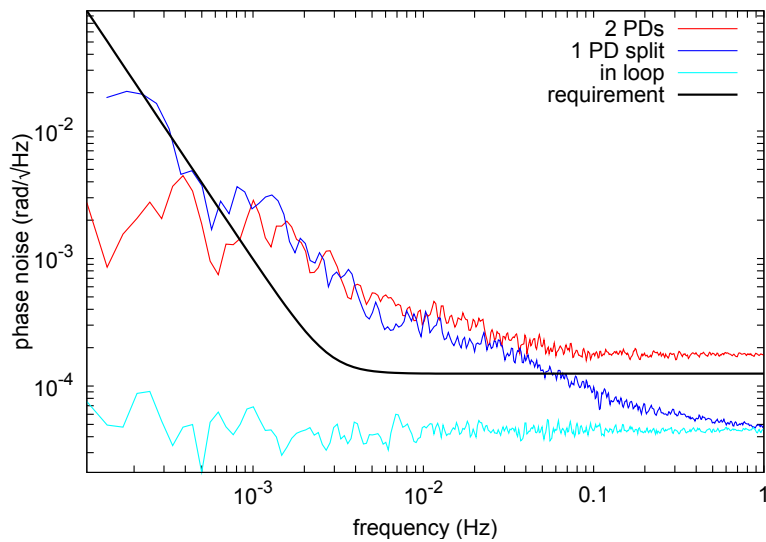
Figure 2.15: Readout noise of the phase lock experiment.

For the experiments at low and high power level, the amplitude of the phase measurement input signal is of the same order of magnitude. Therefore, the decreased performance is not a direct effect of the small signal. Concerning the phase measurement system only the signal-to-noise ratio has changed compared to the high power experiments.

As discussed in Section 2.3.2, the signal amplitude is about $43\,\mathrm{nA_{peak}}{=}30\,\mathrm{nA_{rms}}$. The shot noise is of the order of $5\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$ and is white. The photodiode output signal has a flat response of about $30\,\mathrm{MHz}$. Therefore, the rms noise of the photo current can be calculated to be $27\,\mathrm{pA_{rms}}$, which carries the same power as the signal itself. For high performance phase measurements, the mixer input is almost saturated in order to minimise electronic noise from the following components in the chain. The amplifiers of the low-pass filter are used also close to their maximum output range and may therefore introduce noise by non-linear internal processes.

2.4 Conclusion

For this thesis, a phase measurement system that uses DC readout to perform a phase lock under LISA-like conditions was developed. It is not the preferred implementation for LISA, but allows to test the whole analogue front-end of the phase measurement system. In addition to the experiment, a noise analysis of the system was done. Such a readout scheme will not be used for LISA. However, it can be used as a testbed for characterisation of photodetector prototypes. For mW power levels of both beams, the performance of the experiment is better than the required LISA sensitivity. For application to measurements with realistic light power levels, the performance is above the LISA requirement – in the worst case a factor of 5.

The low-pass filter suppressing the double output frequency of the mixer was identified to be the limiting element of the phase measurement. Improvements of the sensitivity were not further investigated, since the composite amplifier used for the low-pass filters are only rarely documented, and correct circuitry without documentation can only be done by trial and error. Furthermore, it was clear that a fast digital phasemeter which directly samples the photodiode amplifier output signal, will be available soon [17]. It will be a breadboard design of the LISA phasemeter and can be used to characterise the weak light phase lock with the amplifiers developed for this experiment. Since these measurements will be much closer to the LISA baseline design of the phase measurement system, the development of the analogue readout scheme was only continued up to the status presented here.

The LISA backside fibre link

The Laser Interferometer Space Antenna (LISA) is a complex interferometer. It uses 6 lasers that are phase locked to each other with offset frequencies in the MHz-range. The previous chapter dealt with aspects of the *long arm interferometry*, whereas in this chapter the focus lies on the *local interferometry* on each satellite. In contrast to the long arm interferometry, the local interferometry can be tested completely on ground except for the readout of free-falling test masses.

One of the most critical components in the interferometer is the optical fibre connecting the two optical benches on-board each spacecraft. In this chapter the important properties of the fibre and the setup of the optical bench interferometer concerning the fibre is investigated. This work was done in corporation with Roland Fleddermann.

3.1 LISA baseline design

In the LISA baseline design, an optical fibre connecting the two optical benches on-board each satellite is forseen. Figure 3.1 shows a schematic view of one of the LISA satellites with the backside fibre link connecting both benches.

The fibres to be used are polarisation-maintaining, single-mode fibres [6]. A commercially available fibre of this kind is shown in Figure 3.2(a). This fibre and all fibres that were used in this work are so-called PANDA fibres. These fibres have a fast and a slow axis. Therefore, the fibre is birefringent. A cross section of such a fibre is shown in Figure 3.2(b).

The birefringence is introduced by stress-applying material in the cladding. If additional stress is introduced from outside the fibre, the polarisation extinction ratio decreases. Thereby the light is converted into a different polarisation state
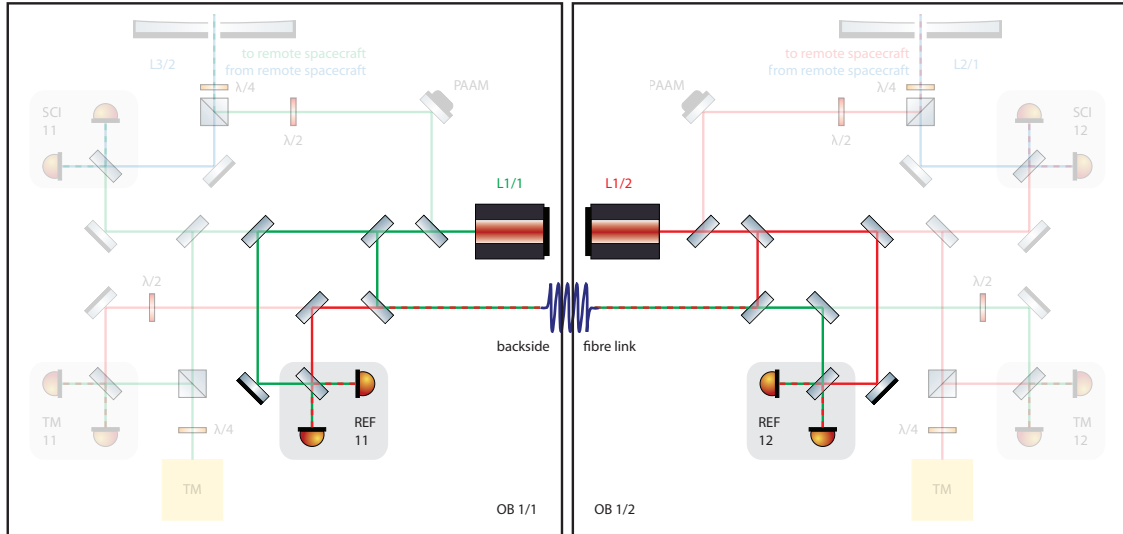
Figure 3.1: Fibre connecting the two optical benches on-board each LISA satellite.
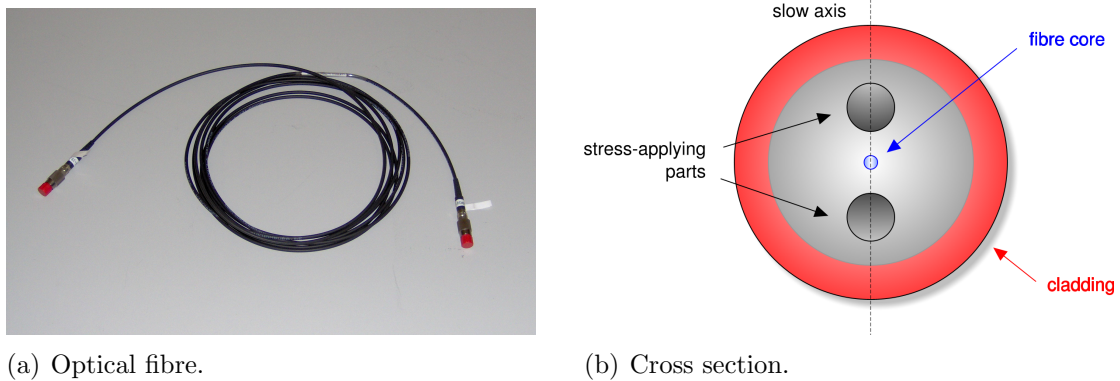
(a) Optical fibre.   (b) Cross section.

Figure 3.2: A PANDA fibre and its cross section.

and can introduce noise in the interferometric readout. Also thermal variations introduce stress due to differential expansion of the fibre core and the outer stress-applying parts, which also changes this ratio.

The fibre front surface has an angle of 8 degrees with respect to the fibre core, which reduces direct back reflection from the injected beam and internal reflections at the front surfaces. This is important in interferometric applications, since those beams produced by such back reflections also travel on the same optical path as the main (measurement) beam and impinge on the photodiodes. Therefore, they produce spurious signals in the interferometer that may spoil the measurement sensitivity.

Also other types of polarisation-maintaining single-mode fibres are available. However, the PANDA-type fibres are furthest developed [32]. Therefore, the chosen

fibre is the best solution of commercially available fabricates for the application in LISA.

In LISA all lasers are phase locked with respect to each other. This is not only the case for the long arm, but also for the two local lasers. Each optical bench has a master laser, which is sent to the remote spacecraft, but a fraction of it is used to sense the position of the local test mass and also to send it to the second optical bench on the same spacecraft. The phase lock between the two local lasers must be performed on the ultra-stable optical bench, since phase fluctuations accumulated after the lock, would couple directly into the longitudinal measurement of the test mass or the LISA arm.

In Figure 3.1 a simplified view of the LISA baseline optical bench can be seen. The local beam L 1/2 on the optical bench OB 1/2 is injected and split. One part is reflected by the test mass and travels to the interferometer PM 12 that measures the test mass position. Another part travels directly to the reference interferometer REF 12. A fraction of this laser light is coupled into the backside fibre link. Out of this fibre comes the local oscillator beam L 1/1 from the other optical bench, OB 1/1. It propagates to the reference interferometer REF 12 and Proof Mass Optical Readout PM 12 directly. It is used as reference beam in these interferometers. Thus, the local beam L 1/2 is used to sense the test mass position, and the beam from the second bench, L 1/2, acts as local oscillator beam in the interferometers. The setup of the second optical bench is identical [33], but the two laser beams swap their roles.

The LISA science signal is extracted by combination of all beat notes in the different interferometers. This concept works only if the differential phase noise in the two reference interferometers, REF 11 and REF 12, on both optical benches, shown in Figure 3.1, is well below the required LISA sensitivity. For this reason the fibre link is an essential element in the LISA local interferometry. Its contribution to the differential phase noise in both reference interferometers must be sufficiently lower than the required LISA sensitivity. The differential phase noise in the reference interferometers is only small enough if the pathlength fluctuations for both beams travelling in opposite directions through the fibre are equal.

The measurement of these differential pathlength changes starts at the beam-splitter directly in front of the backside fibre link. Apart from the path through the fibre, it only includes paths on the ultra-stable optical bench which introduce sufficiently low pathlength noise. Differential pathlength noise for two counter-propagating beams on the same optical path is called *non-reciprocity*.

Since the non-reciprocal phase noise directly couples into the science measurement, which has a total noise budget of $12\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$, the allocated contribution to noise introduced by the fibre is $1\,\mathrm{pm}/\sqrt{\mathrm{Hz}}$, equivalent to $6\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$. The requirement is shown in Figure 3.3. It relaxes with f^{-2} at frequencies below $3\,\mathrm{mHz}$ [34].
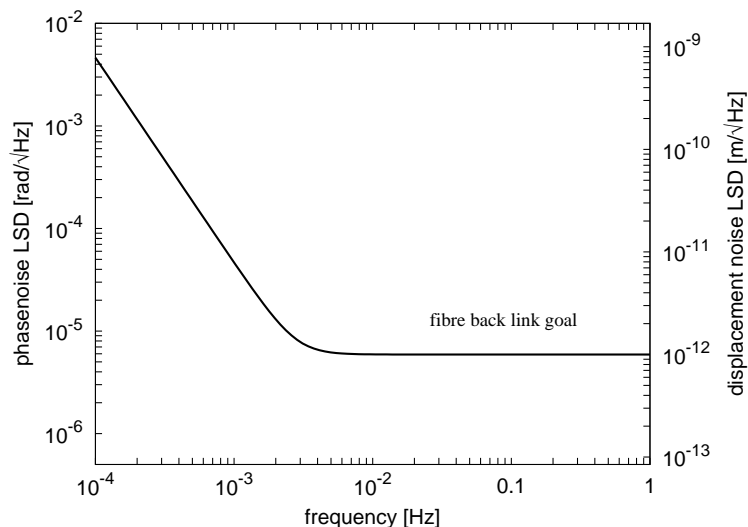
Figure 3.3: Required reciprocity of the backside fibre link.

3.2 Laboratory implementation of the fibre link

The LISA optical bench is complex and includes components that have not yet been tested to the accuracy needed for LISA. Before first breadboards of this complex optical bench are built, all critical components have to be tested in advance. This is challenging because the components have to be characterised interferometrically with picometre accuracy in the mHz measurement band. Usually interferometers made of ultra low-expansion material have to be built, which is a delicate task, but nevertheless, it turned out to be necessary for an appropriate investigation of the backside fibre link. In this section the characterisation of the fibre is presented starting with a Sagnac interferometer, which did not evolve from the work of this thesis, but will be described for completeness. A detailed description can be found in the diploma thesis of Roland Fleddermann [35]. An improved setup, which is very similar to the LISA fibre link interferometer, will be introduced. First experiments were performed on an aluminium breadboard with standard laboratory equipment. The experience and knowledge obtained from this setup was used to realise a quasi-monolithic interferometer using the *hydroxide-catalysis bonding* technique (see Chapter 4).

3.2.1 Previous approaches

First experiments have been performed on a Sagnac setup. This is reasonable, since it directly measures the non-reciprocity. Travel time effects of the light might couple into the measurement, but due to the short fibre, which will have a length of the order of 2 m, these effects are assumed to be small.

Other groups also started investigations on the backside fibre link. At the Uni-

versity of Florida, a polarising Sagnac interferometer was used. At the Jet Propulsion Laboratory (JPL) a fibre was inserted into a LISA testbed. However, this LISA testbed focuses on different aspects of the LISA interferometry and was not specifically made for the purpose of non-reciprocity measurements. The Group at University of Glasgow (UGL) will start an experiment soon and is in the interferometer design phase. The groups at JPL and UGL do not use Sagnac topologies, but more LISA-like configurations of the backside fibre link. At the time of writing this thesis, none of them have demonstrated the required sensitivity.

A simplified setup of the Sagnac interferometer, which was used in earlier experiments in Hannover is shown in Figure 3.4. A laser is split into two beams. One part is transmitted and the other is reflected by the input beamsplitter. After propagating in both directions through the fibre, both beams interfere at the same beamsplitter where they had been split. A technical difficulty arising in this configuration is the fact that, due to energy conservation, one output port of the interferometer is always dark and the other one is always bright at the usual operating point. Therefore, non-reciprocal effects result in tiny variations around a flat maximum or a minimum. In addition, any DC readout at spectral frequencies in the mHz-regime is very challenging and is usually avoided if possible. In order to get a useful error signal slope and to avoid the DC readout, an electro-optical modulator (EOM) is inserted. The modulation frequency and fibre length are chosen such that the output signal is optimally modulated at the modulation frequency. By mixing the electrical oscillator signal with the photodiode output signal, the result is proportional to non-reciprocal length changes.
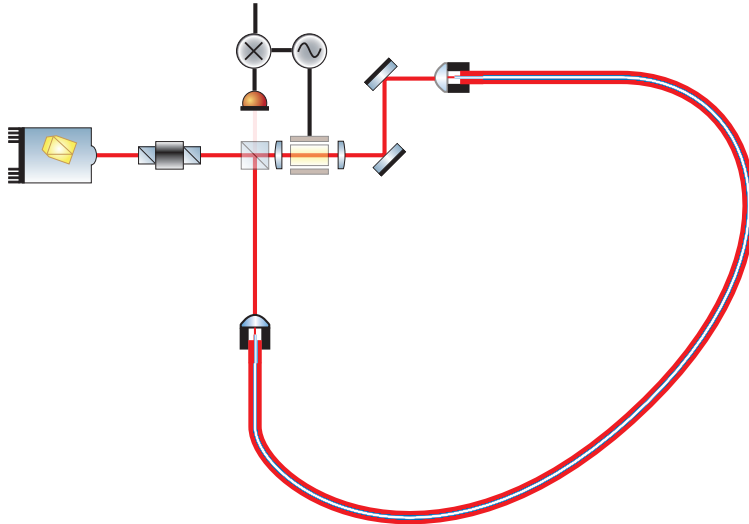
Figure 3.4: Sagnac setup for the back link fibre characterisation.

The non-reciprocity measured with the Sagnac setup is shown in Figure 3.5. The flat part of the spectrum is at about $3\,\mathrm{mrad}/\sqrt{\mathrm{Hz}}$ and thus far away from the

required stability in the $10\,\mu$rad range. Much time was spent in noise hunting on this setup which is summarised in [35]. Most probably the setup was limited by the EOM, which itself might introduce non-reciprocal effects. The Sagnac setup was a good first approach due to its promising sensitivity to non-reciprocal pathlength changes. However, a LISA-like setup is different to the Sagnac and since additional problems may appear, it was decided to investigate the fibre in an experiment which is more comparable to the actual optical setup on-board LISA.
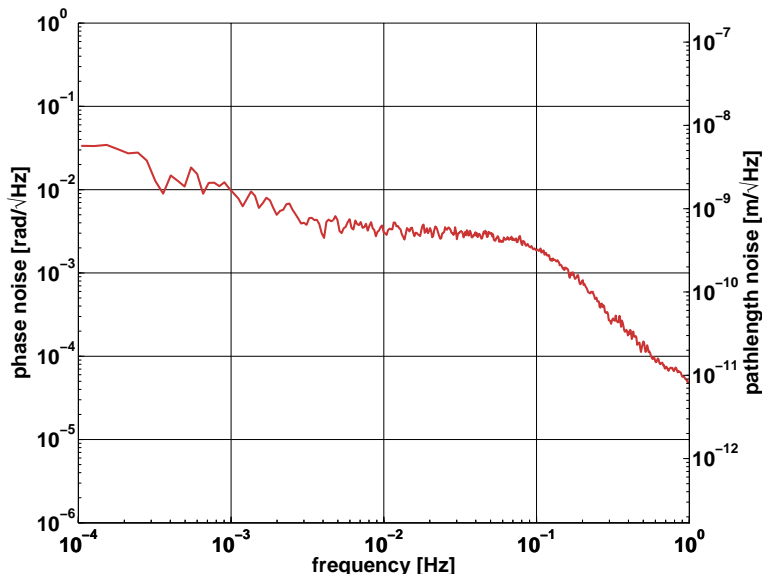
Figure 3.5: Non-reciprocity measured with the Sagnac setup.

3.2.2 The backside fibre link interferometer

Interferometer layout

With the first experiences from the Sagnac interferometer and the knowledge on the LISA baseline design, an interferometer layout that covers the important aspects of the LISA backside fibre link was developed. For its layout, a ray-tracing software called *OptoCad* was used [36]. A schematic of the interferometer layout is shown in Figure 3.6.

The setup consists of three main interferometers, shown in Figure 3.7. It is similar to the LISA Pathfinder interferometer. Two beams are injected onto the optical bench. In the measurements with lowest non-reciprocal noise level, discussed in Section 3.3, a copy of the LISA Pathfinder modulation bench was used (see Chapter 6). This bench is comparatively unstable, since it is set up in air, and in all interferometers the pathlength fluctuations of this modulation bench are measured as common-mode signal. Like in the LISA Pathfinder interferometer, the reference interferometer in Figure 3.7(a) is used to measure the pathlength fluctuations of the
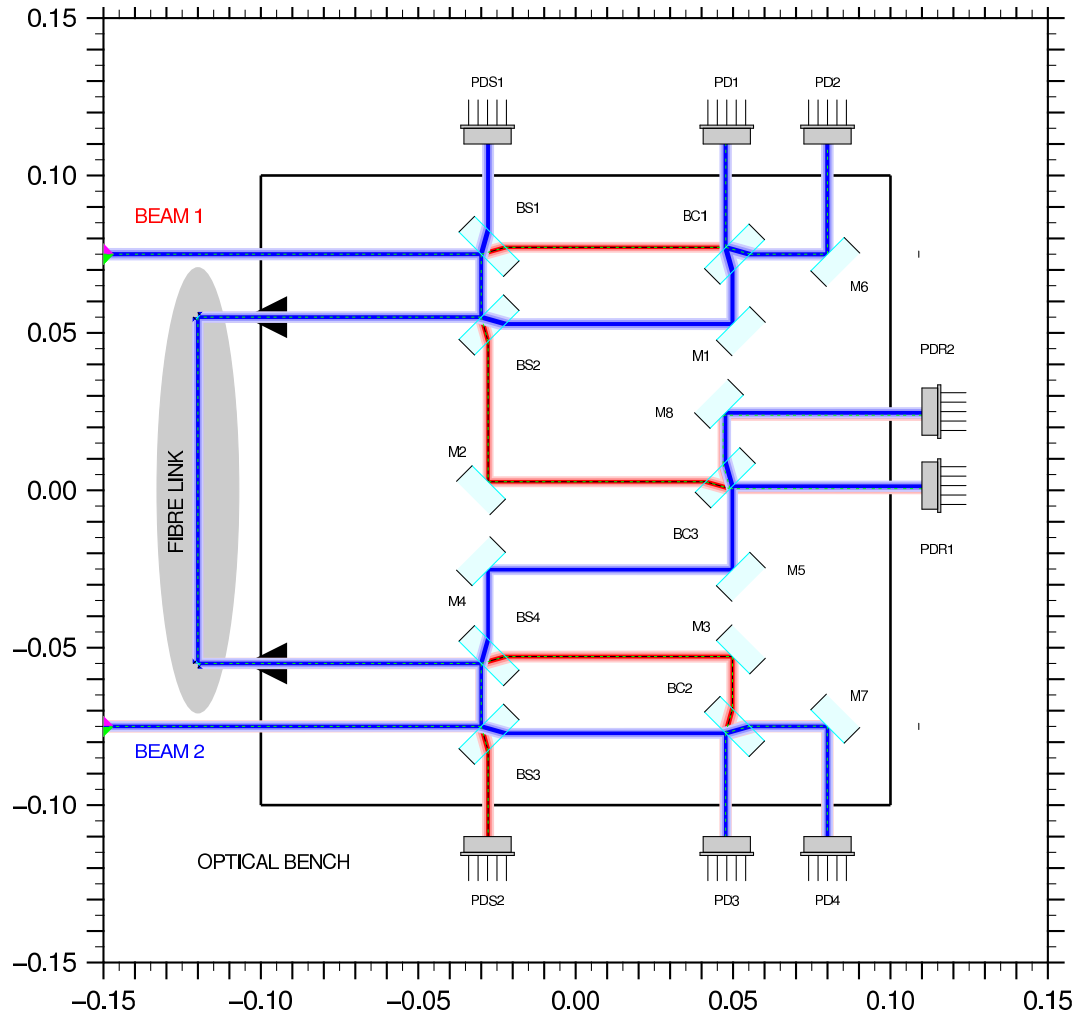
Figure 3.6: Layout of the backside fibre link interferometer.

unstable modulation bench. In the difference with the measurement interferometers, shown in Figure 3.7(b) and 3.7(c), they cancel and thereby only the pathlength changes on the interferometer baseplate are measured.

Non-reciprocity

The measurement interferometers are sensitive to the common-mode phase noise introduced by the fibre, since the optical path in the fibre changes, mainly due to temperature fluctuations. Now both measurement interferometers – M1 and M2 – measure the phase noise of the same fibre in opposite directions at the same time. The combination of both measurements leads to the the non-reciprocal phase noise.

In this section the output of the backside fibre link interferometer will be described analytically. It will be shown how the non-reciprocity is extracted. Furthermore, the effect of laser frequency coupling is pointed out.
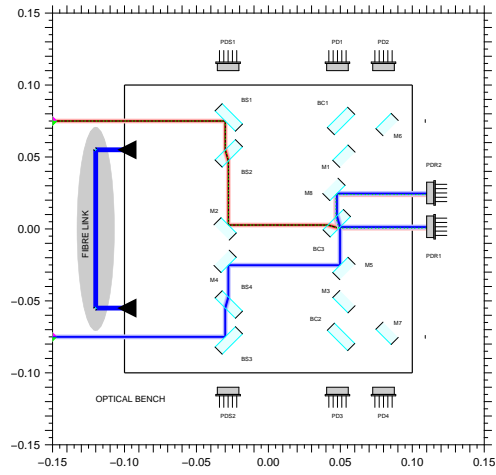
(a) Reference interferometer
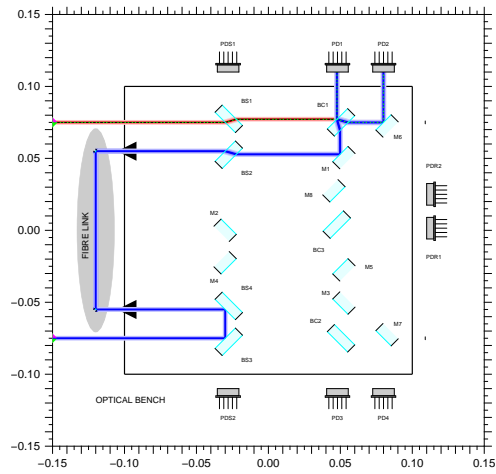
(b) M1 interferometer
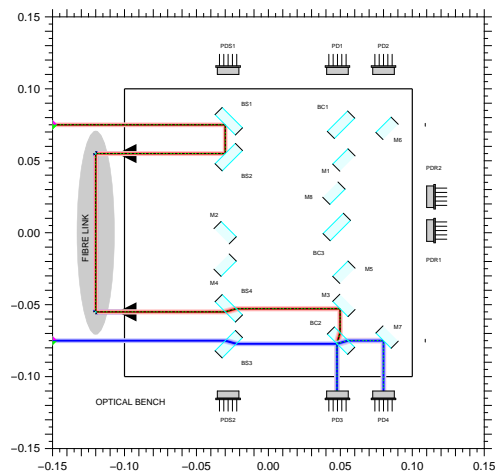
(c) M2 interferometer

Figure 3.7: Individual interferometers of the backside fibre link breadboard.

Without loss of generality, the normalised amplitude of both injected beams can be written as

$$A_1(t) = \exp\left[\mathrm{i} \cdot (\omega_{\mathrm{het}}t + \varphi_1(t))\right] \tag{3.1}$$

for the electric field component of *beam 1* and

$$A_2(t) = c \cdot \exp\left[\mathrm{i} \cdot \varphi_2(t)\right] \tag{3.2}$$

for *beam 2*. This approach is similar to the discussion in Section 2.2.2 with the simplification that only the frequency difference between both beams is taken into account. Each field has a phase φ with a subscript indicating the beam. Since the interferometer is used in a heterodyne scheme, the original (laser) frequency of the second beam is neglected, and the frequency difference between both beams ω_{het} is assigned to the first beam. c is a factor that takes imperfect interference and deviations of the light powers into account.

The non-reciprocity can be extracted from measurements of three photodiodes. Those are situated in the M1 and M2 interferometer and the reference interferometer. The amplitude of the electric field at the output of BC1 on photodiode PD1 is the sum of both fields and can be written as

$$A_{\mathrm{PD1}}(t) = A_1(t) \exp\left(\mathrm{i}\delta\right) + A_2(t - \tau) \exp\left(\mathrm{i}\phi\right). \tag{3.3}$$

Here the phase delay δ is introduced which accounts for the pathlength difference of both beams on the modulation bench. Beam 1 impinges on the photodiode, but beam 2 is delayed by τ, because it has travelled through the fibre. τ is the macroscopic delay and of the order of the fibre length. It can later be used to analyse the coupling of laser frequency noise into the phase measurement. ϕ is the microscopic common-mode length change of the fibre.

At the output of the M2 interferometer, the interfering fields lead to an amplitude of

$$A_{\mathrm{PD3}}(t) = A_1(t - \tau) \exp\left(\mathrm{i}\delta\right) \exp\left(\mathrm{i}(\phi + \Delta)\right) + A_2(t) \tag{3.4}$$

on PD3. Here the first beam accumulates the delay. Additionally, it has a phase $(\phi + \Delta)$ that accounts for the microscopic fibre length change plus the non-reciprocal pathlength change Δ which is only assigned to beam 1.

In the reference interferometer the detected power on PDR1 can be calculated as

$$A_{\mathrm{PDR1}}(t) = A_1(t - \tau^*) \exp\left(\mathrm{i}\delta\right) + A_2(t - \tau^*) \tag{3.5}$$

Both fields have a delay τ^*, which is different compared to the two measurement interferometers. However, since the pathlengths of both beams match in the reference interferometer, the two fields experience the same delay.

The photodiodes PD1, PD3 and PDR1 do not detect the field amplitude but the power,

$$P = |A|^2, \tag{3.6}$$

of the interference pattern which is proportional to the produced photo current. For the length measurement, only the oscillating term at the heterodyne frequency is taken into account, which is indicated by the index 'AC'. The detected AC power on each photodiode can be written as

$$
\begin{align}
P_{\text{PD1}}^{\text{AC}}(t) &= 2 \cdot c \cdot \cos\left(\omega_{\text{het}}t + \varphi_1(t) + \delta - \varphi_2(t - \tau) - \phi\right), \tag{3.7} \\
P_{\text{PD3}}^{\text{AC}}(t) &= 2 \cdot c \cdot \cos\left(\omega_{\text{het}}(t - \tau) + \varphi_1(t - \tau) + \delta - \varphi_2(t) + \phi + \Delta\right), \tag{3.8} \\
P_{\text{PDR1}}^{\text{AC}}(t) &= 2 \cdot c \cdot \cos\left(\omega_{\text{het}}(t - \tau^*) + \varphi_1(t - \tau^*) + \delta - \varphi_2(t - \tau^*)\right). \tag{3.9}
\end{align}
$$

For the readout of the photodiodes, a phasemeter similar to the LISA Pathfinder phasemeter is used (see Section 6.3). It measures the phase of the signals exactly at the heterodyne frequency. The common-mode fibre length changes, Ψ_1 and Ψ_2, can be measured in both directions by subtracting the phase of the measurement interferometers from the reference interferometer phase:

$$
\begin{align}
\Psi_1 &= \varphi_1(t) - \varphi_1(t - \tau^*) - \varphi_2(t - \tau) + \varphi_2(t - \tau^*) - \phi, \tag{3.10} \\
\Psi_2 &= \varphi_1(t - \tau) - \varphi_1(t - \tau^*) - \varphi_2(t) + \varphi_2(t - \tau^*) + \phi + \Delta. \tag{3.11}
\end{align}
$$

Here the fluctuations in the unstable part of the modulation bench cancel. For the reciprocal phase noise, changes in the stable delays τ and τ^* are caused only by laser frequency fluctuations, according to

$$
\frac{\tilde{\delta l}}{\Delta L} = \frac{\tilde{\delta \nu}}{\nu}. \tag{3.12}
$$

ΔL is the pathlength difference between one of the measurement interferometers and the reference interferometers, and ν is the laser frequency. $\tilde{\delta l}$ and $\tilde{\delta \nu}$ are their fluctuations. The fibre length is of the order of $10\,\text{m}$. A typical laser has a free-running frequency noise of about $3\,\text{kHz}/\sqrt{\text{Hz}}$ at $1\,\text{Hz}$ increasing with $1/f$ to lower frequencies. This limits the achievable sensitivity to about $100\,\text{nm}/\sqrt{\text{Hz}}$ at $3\,\text{mHz}$. For a sensitivity of $1\,\text{nm}/\sqrt{\text{Hz}}$, the laser frequency must be stabilised to $3\,\text{MHz}/\sqrt{\text{Hz}}$, which is easily achievable with the stabilisation scheme used for the experiments presented here.

Therefore, the delays τ and $\tau*$ from Equations (3.10) and (3.11) can be neglected, which leads to

$$
\begin{align}
\Psi_1 &\simeq -\phi, \tag{3.13} \\
\Psi_2 &\simeq \phi + \Delta. \tag{3.14}
\end{align}
$$

One can see from these formulae that the sum is exactly the non-reciprocal phase change inside the fibre if the frequency fluctuations of the laser are neglected. However, it can again be calculated from Equations (3.10) and (3.11) taking the

laser frequency fluctuations into account:

$$
\begin{aligned}
\Psi_1 + \Psi_2 \;=\;& \varphi_1(t) - \varphi_1(t - \tau^*) \\
-\;& \varphi_2(t) + \varphi_2(t - \tau^*) \\
+\;& \varphi_1(t - \tau) - \varphi_1(t - \tau^*) \\
-\;& \varphi_2(t - \tau) + \varphi_2(t - \tau^*) \\
+\;& \phi + \Delta - \phi.
\end{aligned}
\tag{3.15}
$$

The first term is proportional to the laser frequency changes in the first beam, whereas the second term is proportional to the frequency changes in the second beam – both exactly at the same time. During the time τ^*, it can be assumed that the AOMs situated on the modulation bench do not introduce differential frequency noise in both beams and, therefore, the first and the second term cancel. The same argument holds for the third and the fourth term. Thus, the sum of Ψ_1 and Ψ_2 is the non-reciprocity

$$
\Psi_1 + \Psi_2 = \Delta.
\tag{3.16}
$$

Its measurement is insensitive to laser frequency fluctuations, as long as the pathlengths in the different interferometers are matched. This is important: In contrast to the reciprocal length changes, the non-reciprocal length changes are small.

Interestingly, both phases of the measurement interferometer have to be added instead of subtracted in order to extract the non-reciprocity. The reason for this is that both beams are counter-propagating in this setup.

Imbalanced power

Conventional interferometers are designed to have interfering beams with equal power. Due to its special topology, this is not the case for the backside fibre link interferometer. Detected signals of beams with unequal power lead to practical complications in the readout, since a high DC signal with a small AC amplitude has to be measured. In analogue readout schemes the photodiode output signals have to be high-pass filtered. In case of digitisation of the photodiode signals the dynamic range of analogue-to-digital converters cannot be used completely.

The situation in the present setup can be described as follows: Assuming a total input power of $P = P_1 = P_2$ for both beams and taking the attenuation of the beamsplitters in the optical paths into account, the power of beam 1 on PD1 is

$$
P_1^{\mathrm{PD1}} = \frac{1}{4} P,
\tag{3.17}
$$

whereas beam 2 has a power of

$$
P_2^{\mathrm{PD1}} = \frac{1}{16} \varepsilon P,
\tag{3.18}
$$

where ε is the coupling efficiency of beam 2 into the fibre injector. Taking into account that $E = \sqrt{P}$, the maximum reachable contrast in this setup is

$$c_{\text{max}} = 2 \cdot \frac{E_1 E_2}{E_1^2 + E_1^2} = \frac{4 \cdot \sqrt{\varepsilon}}{4 + \varepsilon}. \tag{3.19}$$

For an ideal (100 %) coupling efficiency, the maximum contrast is 80 %. In reality a well aligned fibre coupling has an efficiency of about 70 to 80 %, which results in a contrast of 70 to 75 %. The setup is symmetric: In the second interferometer the calculation is the same but the two beams swap their powers.

For the monolithic setup introduced in Section 3.4, two additional beamsplitters are in the path of the beams, which are going through the fibre. They implement a null-measurement path that can be used instead of the fibre and allows a characterisation of the instrument sensitivity. Therefore, Equation (3.18) becomes

$$P_2^{\text{PD1}} = \frac{1}{64} \varepsilon \cdot P \tag{3.20}$$

and Equation (3.19)

$$c_{\text{max}}^{\text{Zerodur}} = \frac{8 \cdot \sqrt{\varepsilon}}{16 + \varepsilon}. \tag{3.21}$$

In this case the maximum contrast is typically about 40 %.

Ghost beams

From the beginning it was suspected that noise introduced by ghost reflections in the setup is one of the dominant noise sources in the phase readout. This turned out to be partially true: In a conventional interferometer readout scheme it limits the sensitivity, but a correction can be applied in the data post-processing to increase the performance (see Section 3.4.1).

The most likely source of such ghost beams is the back reflection of the beam that is injected into the backside fibre link. This effect is illustrated in Figure 3.8. Since the interferometer setup is symmetric, the described effects and measurements in the M1 interferometer can also be applied to the M2 interferometer.

Beam 2 is transmitted through the fibre and travels to the M1 interferometer photodiodes. A part of it also impinges on photodiode PDS1. Additionally to the part of beam 1 going directly to the M1 photodiodes, there is a second part impinging on the photodiode, which is back reflected from the fibre coupling at the backside fibre link injector. This back reflection can have its origin on the lens in front of the fibre injector, the fibre front surface or even from inside the fibre due to back scattering. All these effects produce a spurious beat signal at the heterodyne frequency, which is measured in addition to the 'real' signal. The contribution to the measurement signal depends on the amount of light that is in the right spatial mode for the interference. Therefore, e.g diffuse scattered light contributes less
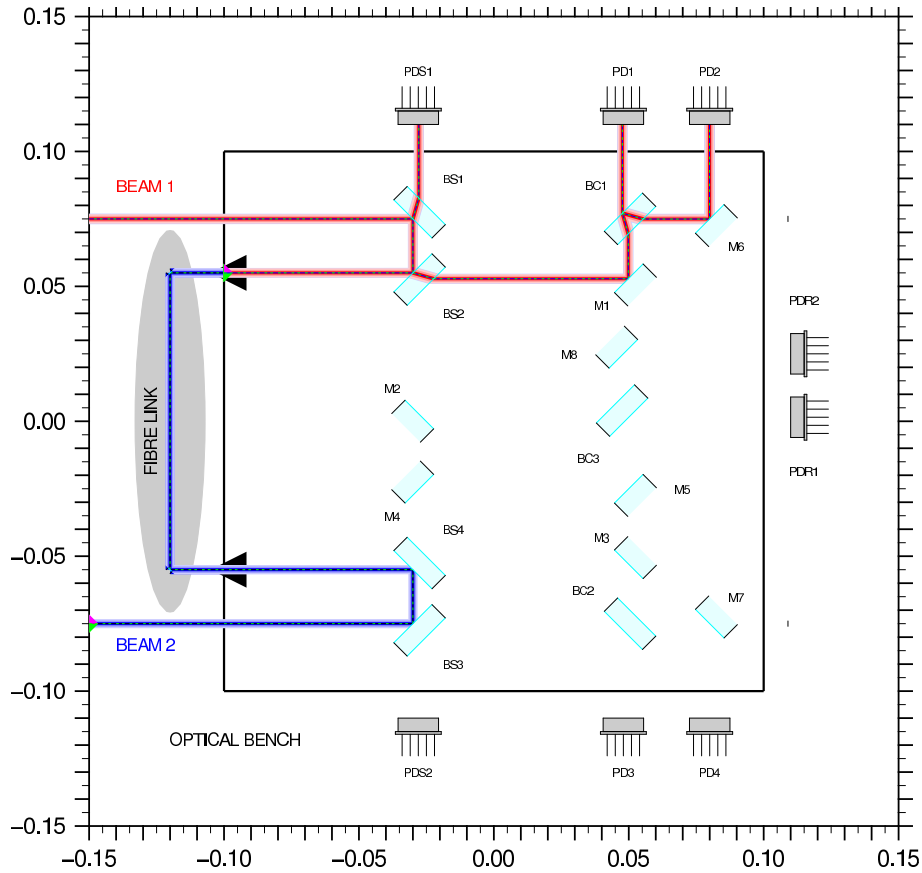
Figure 3.8: Straylight introduced by back reflections from the fibre injector.

than light that is directly back reflected from the lens or fibre surface in the right mode.

The interferometer baseplate is stable, since it is quasi-monolithic and made of low-expansion material. This is not true for the fibre injectors, which have metal parts and are situated in aluminium mounts. Thermal effects and vibrations introduce much more length and position changes in these materials. If the produced back reflections would contribute to the measurement with constant phase, the error would also be constant and therefore not spoil the interferometer performance, but this is not expected to be the case due to the material of the fibre injectors. Monolithic fibre injectors are being built at the University of Glasgow for the Flight Model of the LISA Pathfinder optical bench [37]. Those injectors have not been tested in an interferometer setup comparable to LISA in terms of interferometer sensitivity. In contrast to the fibres used in the presented setup here, the fibre output surface is not angled in the monolithic injector design. For LISA Pathfinder this is a minor problem, since there are no direct beams that are back reflected from the front surfaces of the fibre injectors. In a LISA-like setup, however, these injectors are expected to produce more back reflection in the right spatial mode for

the interference.

Straylight effects scale with the amplitude and not with the power of the light. This means if one wants to reduce the error introduced by back reflections by a factor of 10, the amplitude has to be reduced by a factor of 10 and thereby the power by a factor of 100.

In order to characterise the effect of all straylight contributions that produce spurious signals in the setup, an additional interferometer is forseen, which is read out by PDS1. Beam 2 and the reflected part of beam 1 have the same power ratio and phase relation on this photodiode as in the measurement interferometer M1.

The appearance of directly back reflected spurious beams discussed here is unavoidable in LISA as long as the backside fibre link with two counter-propagating beams is used. This problem will also appear in a free-beam setup, when both beams propagate in opposite directions on the same optical path. The influence on the interferometer performance and methods to minimise or subtract the spurious signal in the data post-processing are discussed in Section 3.3.

3.2.3 Modulation and readout of the interferometer

The interferometer readout during this measurement campaign was done in two different ways. In the beginning a LISA-like configuration was chosen, where two lasers are phase locked with an offset frequency of about 20 MHz with respect to each other. This is very similar to the phase lock presented in Section 2.2. The reference interferometer is used to lock the phase of both lasers to an electrical oscillator as shown Figure 3.9. This is implemented using a low-pass filtered output of an analogue mixer as error signal. The outputs of the measurement interferometers M1 and M2 are mixed down to 1.6 kHz in order to use a well characterised LISA Pathfinder-like phasemeter (see Section 6.3). Here, the phasemeter was not a hardware phasemeter using FPGAs, but running on a commercial PC with PCI AD-card like early breadboard designs of the LISA Pathfinder phasemeter [38].

Mixing the main measurement signals to an intermediate frequency and not to DC has important advantages: DC readout on long time scales is usually sensitive to thermally driven changes in the output voltages of amplifiers. A heterodyne readout is, in first order, not sensitive to those fluctuations. Moreover, the output phase of the M1 and M2 interferometer has a huge range, of the order of 100 rad, during a typical measurement period. Analogue mixers have a limited linear output of much less than 1 rad and for phase changes of more than half a cycle they change the sign in the output signal. Therefore, it is not possible to track the phase over a whole measurement when DC readout is used.

This readout scheme was used in the first approaches, but a lot of technical problems appeared. The photodiodes must be small in order to have a reasonable response to signals in the 20 MHz range. Therefore, the interference pattern has to be focused on the photodiode. Additional components have to be introduced in the setup and the photodiode alignment is very sensitive. A phase lock of two
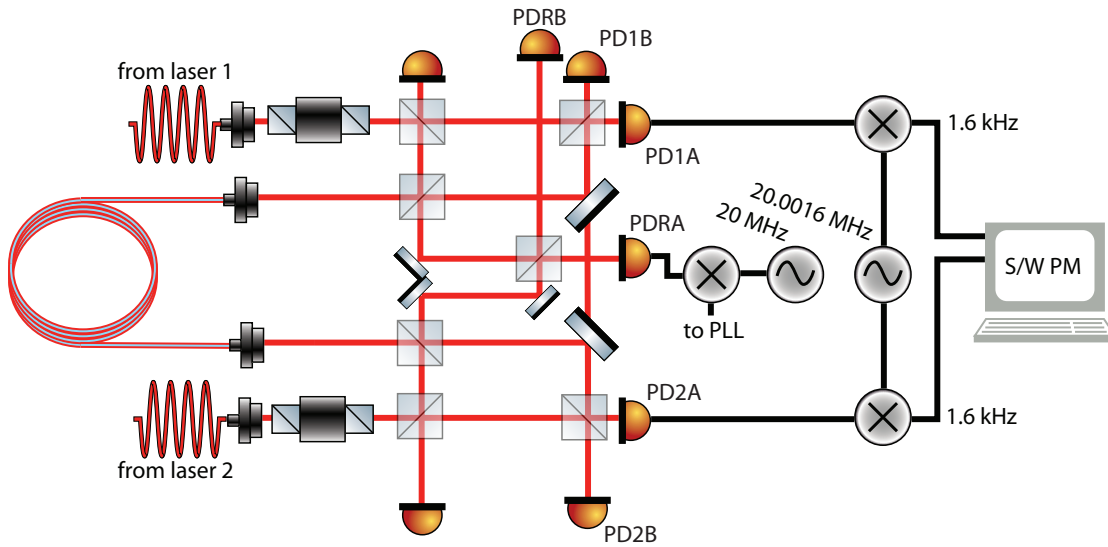
Figure 3.9: 20 MHz modulation scheme.

lasers is sometimes unstable over many hours. Additionally to this lock, a laser frequency stabilisation has to be used in order to reduce the coupling into the phase measurement that is caused by the residual pathlength mismatch in the different interferometers. The pathlength difference between one of the measurement interferometers and the reference interferometer causes a conversion of laser frequency noise into pathlength noise (see Equation (3.12)). Since many sensitive loops are necessary to be active, it is impractical to do measurements with this modulation scheme on long time scales.

In order to overcome these difficulties, a LISA Pathfinder modulation and readout scheme was used. It is shown in Figure 3.10 (see also Chapter 6). In this scheme only one laser is needed. Of course the setup also requires loops to be active in order to ensure a stable frequency difference of the two injected laser beams, but these loops only include electronic components. Therefore, they are not as sensitive as a phase lock between two lasers. Larger (silicon) photodiodes can be used which make focusing of the beams unnecessary. No additional mixers are needed in the photodiode readout paths and the photodiode amplifier output signals can be directly sampled with a kHz-phasemeter.

Due to the comparatively high noise floor in the first measurements, the laser frequency stability was not limiting the sensitivity and a stabilisation was not used. Later, a commercially available laser system with an internal iodine stabilisation was used.

Going to a heterodyne measurement at kHz is expected to be representative also for MHz frequencies in terms of non-reciprocity. There may be dispersion effects inside the fibre that could introduce different effects in the MHz regime. However, once the fibre performance is demonstrated with the actual setup at kHz, one can
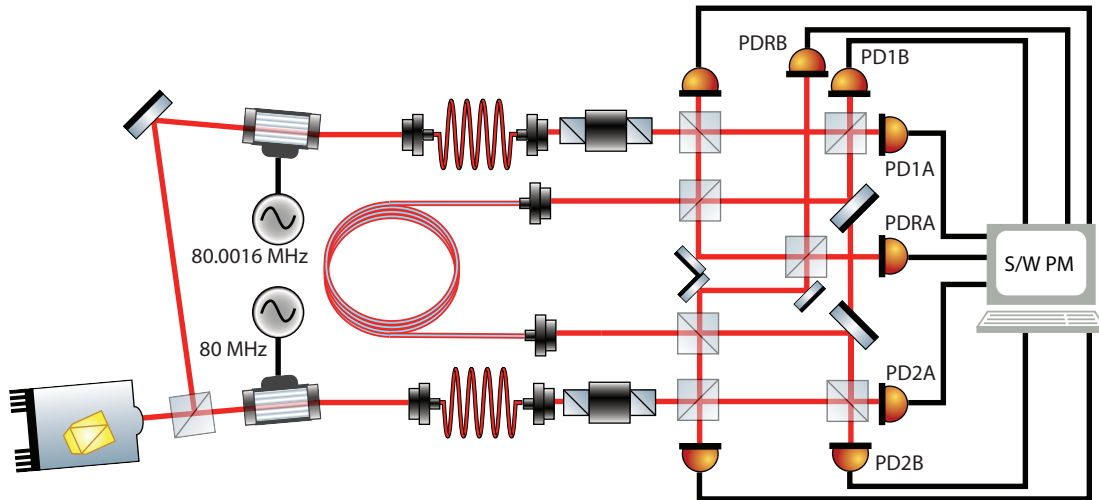
Figure 3.10: LISA Pathfinder modulation scheme.

use the MHz readout scheme to investigate the performance also at frequencies even more representative for LISA. The results achieved at kHz represent a lower limit for measurements in the MHz regime.

3.3 Tabletop experiments

3.3.1 Tabletop interferometer

From the very beginning it was clear that the fibre measurement on a LISA sensitivity level can only be performed on a thermally very stable setup. Nevertheless, the first breadboard of the backside fibre link interferometer was realised on an aluminium breadboard for several reasons. Although it was planned to build an interferometer using hydroxide-catalysis bonding techniques, a lot of parts of the setup could be tested on an aluminium breadboard to reasonable accuracy. This includes mainly the electronics and the data acquisition. Delivery times for bonding components and the preparation of the process took a long time, which was used to understand the difficulties of the interferometer topology in parallel. Initially the results from the breadboard experiment were expected to be meaningful at least at high frequencies around a few 10's of Hz, where thermal fluctuations usually couple much less into the pathlength measurement.

The aluminium breadboard is shown in Figure 3.11. The size of the breadboard is $45 \times 45\,\mathrm{cm}^2$. In front of the input fibre injectors Faraday isolators are inserted for two reasons: Firstly, the path of the second beam matches into the fibre injector of the first beam and vice versa. Especially in early experiments where two phase locked lasers were used, this path goes directly back to the laser head, which has to be avoided in order to prevent laser instabilities. Secondly, it was intended

to prevent at least the back reflection at the input fibre injectors from producing spurious signals.

Such Faraday isolators cannot be put in front of the back link fibre, because in that optical path two beams travel in opposite directions, and a properly used Faraday isolator acts as an optical diode. For LISA even the Faraday isolators in front of the beam injectors will not be used due to their strong magnetic filed. This would lead to additional forces on the test masses and thereby to a degradation of their drag-free state.

The performance of this setup is shown in Figure 3.12. After the initial measurements (red curve) various improvements were implemented:

- The measurement was performed in vacuum.

- The pathlength differences were adjusted such that the coupling of laser frequency noise was minimised.

- A laser frequency stabilisation was implemented.

- The effect of input polarisation alignment into the backside fibre was investigated.

- The modulation scheme was changed from two phase locked lasers to a LISA Pathfinder scheme.

- Electronic and digitisation noise was reduced.

Due to all these changes, the resulting performance reached a level shown by the blue curve in Figure 3.12. The peak sensitivity is $3 \cdot 10^{-4}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$ at a few Hz. Below that frequency, the noise increases with 1/f to low frequencies.

3.3.2 Limitations in the aluminium setup

It later turned out that none of the noise sources described above really limits the sensitivity of the aluminium breadboard backside fibre link interferometer. In order to characterise the noise produced by the baseplate itself, two simple Mach-Zehnder interferometers were compared, one on an aluminium breadboard and a second one on a Zerodur baseplate, which is shown in Figure 3.13. A typical non-reciprocity measurement is also shown in this figure.

For the readout of the interferometers, the same modulation bench, photodiodes and data acquisition system as for the backside fibre link interferometer were used. The resulting performance is shown in Figure 3.14. The pathlength noise of the Zerodur interferometer is at least one order of magnitude lower than that of the aluminium interferometer. The noise spectrum of the Zerodur interferometer is flat at low frequencies and decreasing to higher frequencies. This typical *shoulder*-shape is usually produced by the so-called *small vector* noise and not by 'real' pathlength
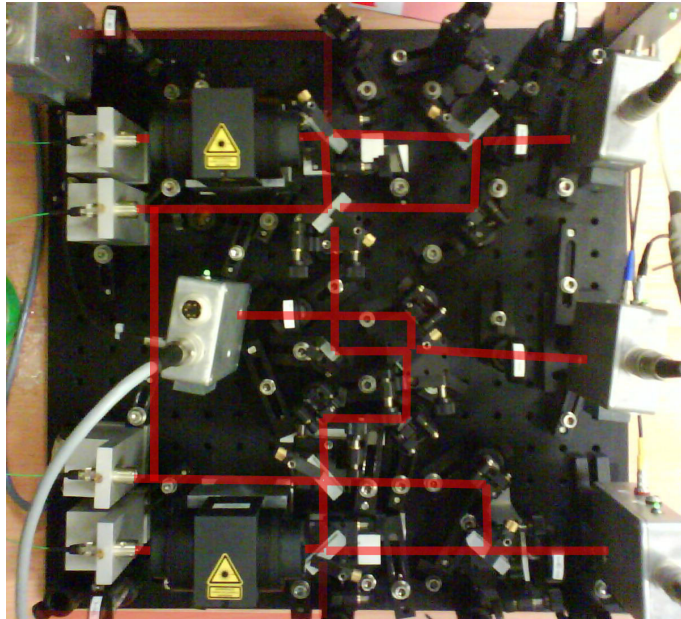
Figure 3.11: Aluminium breadboard of the backside fibre link interferometer.
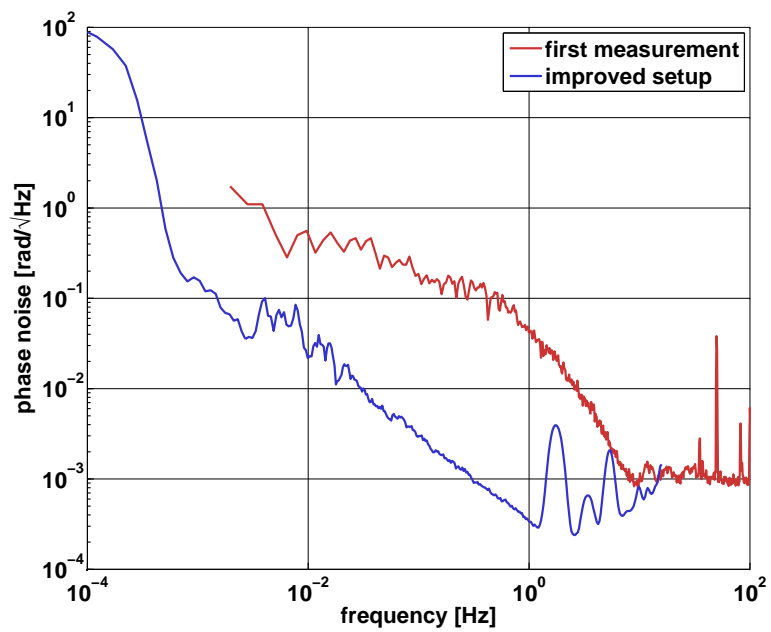
Figure 3.12: Non-reciprocity observed using the aluminium breadboard.

(a) Schematic view.   (b) Zerodur interferometer.
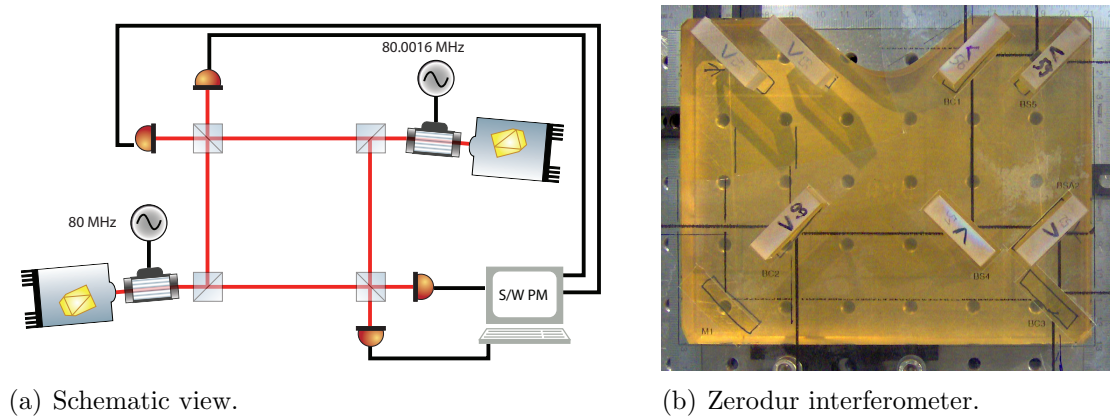
Figure 3.13: Mach-Zehnder interferometers.

fluctuations [18, 39]. For comparison, the sensitivity of the fibre interferometer on the aluminium baseplate is shown. Its shape is similar to the one of the aluminium Mach-Zehnder, but is about a factor of 3 better. This is probably due to the more compact setup of the fibre interferometer.
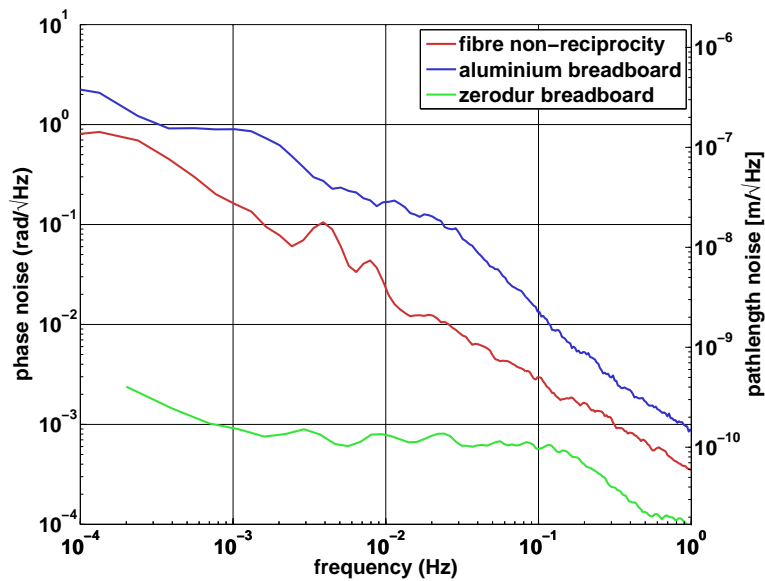
Figure 3.14: Comparison between interferometers based on aluminium and Zerodur.

The result obtained is very important. Initially it was expected that at least at high frequencies in the Hz-range a setup on an aluminium baseplate would be sufficient to reach a sensitivity of the order of the LISA sensitivity. This is obviously not true. Even at high frequencies the aluminium baseplate limits the sensitivity at a level of about $2 \cdot 10^{-4}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$. The mounts for the optical components might also play a role for the performance level reached.

3.4 The backside fibre link interferometer

In order to reach a better sensitivity at high frequencies and to get reasonable results also in the LISA measurement band, an interferometer was bonded on a baseplate made of low-expansion material. The bonding process is described in Chapter 4. In this section the important features of the interferometer and the data post-processing which lead to the final performance will be introduced. Apart from the straylight subtraction (Section 3.4.1), these features were already tested on the aluminium breadboard. However, due to the sensitivity reached at that time, the performance did not improve, but some of these effects do influence the sensitivity of the Zerodur interferometer.

The bonded backside fibre link interferometer used for the investigations presented here is shown in Figure 3.15. The setup of the interferometer is the same as introduced in Section 3.2.2 apart from two additional beamsplitters in front of the backside fibre link injectors. They were inserted to make a null-measurement possible, where the two beams do not travel through the back link fibre but on the ultra-stable optical bench. For a fibre measurement, the null-path must be blocked, and vice versa, if a null-measurement is performed, the fibre path must be blocked.
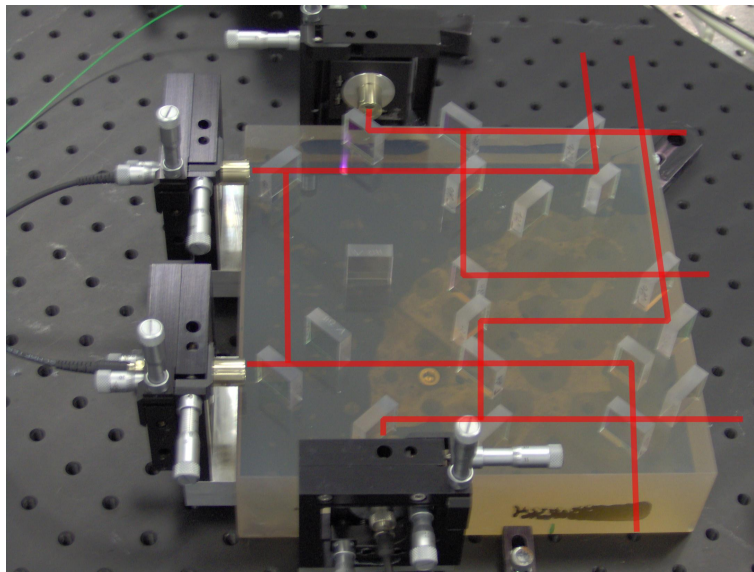
Figure 3.15: Bonded backside fibre link interferometer

3.4.1 Straylight subtraction

A huge improvement of the sensitivity can be achieved by subtracting the straylight effect. A typical spectrum of the non-reciprocity measured with the Zerodur interferometer is shown in Figure 3.16 (magenta).
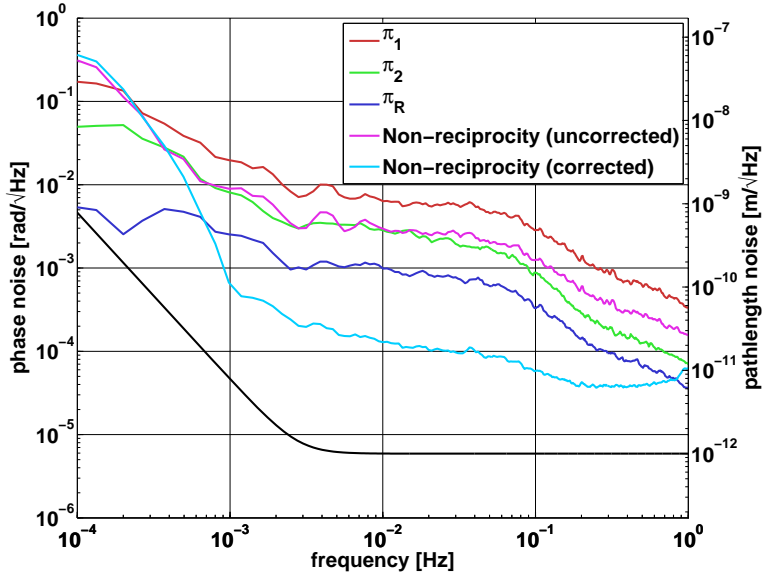
Figure 3.16: Initial performance of the Zerodur fibre interferometer.

A common technique in interferometry is to identify the differential readout noise of photodiodes including their entire signal chain by performing so-called π-measurements as already described in Section 2.2.1. In these π-measurements one makes use of the constant phase difference of 180 degrees at the two output ports of a recombination beamsplitter. The spectrum of the phase difference of the two photodiodes,

$$\pi_x = \varphi_A^x - \varphi_B^x, \tag{3.22}$$

should be constant. A and B indicate the two photodiodes at the output ports of the beamsplitter and x indicates the interferometer. Usually 'real' pathlength measurements cannot have higher sensitivity than the π-measurements which are also shown in Figure 3.16 for the three interferometers. The sensitivity of the reference interferometer is of the order of $10^{-3}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$ around $10\,\mathrm{mHz}$, which is better than the M1 ($6 \cdot 10^{-3}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$) and M2 interferometer ($3 \cdot 10^{-3}\,\mathrm{rad}/\sqrt{\mathrm{Hz}}$).

The uncorrected non-reciprocity is calculated as

$$\Delta = \varphi_1 - \varphi_R + \varphi_2 - \varphi_R = \varphi_1 + \varphi_2 - 2 \cdot \varphi_R, \tag{3.23}$$

where the index of the measured phase, φ, indicates the interferometer shown in Figure 3.7. For the uncorrected non-reciprocity, only one photodiode signal from each interferometer is used, since the second photodiode nominally delivers only redundant information. The non-reciprocity is obviously a factor of 2 better than the π_1-measurement. This seems contrary to the typical low noise level of π-measurements compared to 'real' length measurements and can only be explained if the outputs of the two photodiodes at the beamsplitter of the M1 interferometer have anti-correlated noise, which is high compared to the M2 and the reference

interferometer. In this case the π-measurement noise from Equation (3.22) doubles compared to the non-reciprocity from Equation (3.23), since it includes two noisy phase measurements, φ_{A}^1 and φ_{B}^1, whereas the non-reciprocity includes only one, φ_{A}^1 or φ_{B}^1. The effect results from the appearance of ghost beams which is different in the M1 and M2 interferometer. Fortunately, this effect can be compensated as discussed in Section 3.2.2.

The heterodyne signal is illustrated in Figure 3.17 as a complex vector. The signals of two redundant photodiodes to be measured (left: red vector and blue vector) leave the recombination beamsplitter with 180 degrees phase difference due to the conservation of energy in the interference process. The additional (ghost beam) signals (left: black vectors) interfere at the fibre injector of the backside fibre link. Therefore, those spurious signals leave the recombination beamsplitter with the same phase. Subtracting the resulting measured vectors (orange and light blue) yields twice the signal while the spurious signals automatically cancel (right: violet vector).
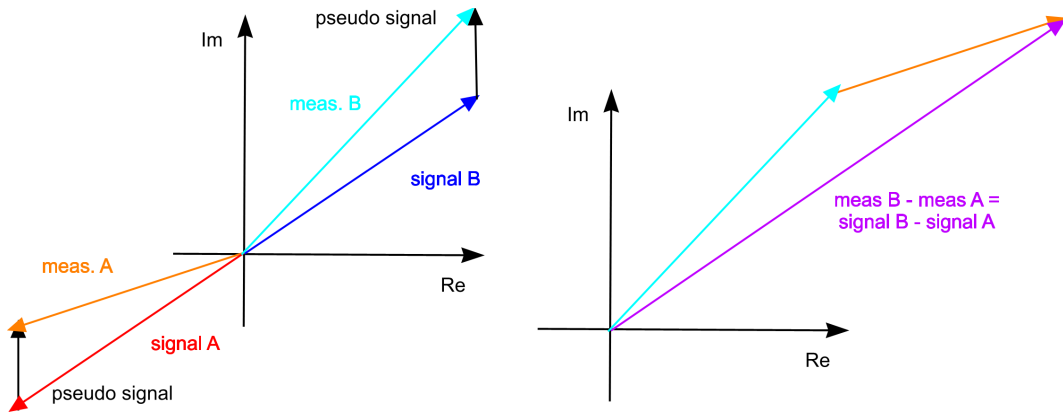
Figure 3.17: Compensation of pseudo signals by averaging.

In this case the individual calculation of each interferometer output phase can be re-written as

$$\varphi_{\mathrm{x}} = \arctan \frac{A_{\mathrm{A}}^{\mathrm{x}} \cos \varphi_{\mathrm{A}}^{\mathrm{x}} - A_{\mathrm{B}}^{\mathrm{x}} \cos \varphi_{\mathrm{B}}^{\mathrm{x}}}{A_{\mathrm{A}}^{\mathrm{x}} \sin \varphi_{\mathrm{A}}^{\mathrm{x}} - A_{\mathrm{B}}^{\mathrm{x}} \sin \varphi_{\mathrm{B}}^{\mathrm{x}}}. \tag{3.24}$$

Here 'x' is an index for the interferometer. The index A indicates one interferometer output port and B the complementary one. A denotes the amplitude of the heterodyne signal. The calculation includes a simple vector sum of the two measured signals and a recalculation of its angle in polar coordinates. Of course it is also possible to take the direct output of the Fourier transform, which are the real and imaginary part and calculate the vector sum directly in cartesian coordinates. However, the phasemeter used for the measurement in the experiment delivered the output in polar coordinates.

The calculation from Equation (3.24) corresponds to a subtraction of the two photo currents and could also be done using analogue electronics. However, when

applying the correction in the data post-processing no additional electronics is needed and mismatches in the gain of the photodiodes readout chains can easily be compensated.

The uncorrected non-reciprocity from Equation (3.23) is calculated using only one photodiode signal from each interferometer. For the corrected non-reciprocity the phase of each interferometer is re-calculated with both redundant photodiode signals using Equation (3.24), before the calculation from Equation (3.23) is performed. By applying this correction the sensitivity increases by more than one order of magnitude, which is shown as the light blue curve in Figure 3.16.

Limitation of the correction method

In the first approach of the straylight correction both photodiode outputs were normalised with respect to the 'real' signal amplitude which is not the measured amplitude, since this includes the spurious signal. This 'real' signal amplitude is difficult to estimate. It may be possible to choose the best ratio by weighting signal A differently to signal B such that the pathlength noise is minimised, but this was not further investigated. Additionally, drifts in photodiode amplifier gains could slowly change the ratio of both signals, which is difficult to model. One further improvement might be to include the imperfect reflection-transmission ratio of the beamsplitter in the model.

In Figure 3.18 the electric fields at the beamsplitter input are shown. E_1 and E_2 on the left hand side refer to those fields that produce the 'real' signal without any spurious signal. E_S is the spurious field produced by back reflection from the fibre injector. Assuming an amplitude reflectivity, ρ, and transmittance, τ, of the beamsplitter, the power of the interfering *signal* fields on the photodiode can be written as

$$P_\mathrm{A}^\mathrm{real} = (\rho \cdot E_1)^2 + (\tau \cdot E_2)^2 + 2(\rho E_1)(\tau E_2)\cos\left(\frac{f_\mathrm{het}}{2\pi} + \varphi_\mathrm{sig}\right), \qquad (3.25)$$

$$P_\mathrm{B}^\mathrm{real} = (\tau \cdot E_1)^2 + (\rho \cdot E_2)^2 + 2(\tau E_1)(\rho E_2)\cos\left(\frac{f_\mathrm{het}}{2\pi} + \varphi_\mathrm{sig} + \pi\right). \quad (3.26)$$

f_het is the heterodyne frequency, i.e. the frequency difference between the two incoming laser beams. Both beams have an arbitrary common phase φ_sig. The additional phase of π in P_2 accounts for the 180 degrees phase difference between the two outgoing interfered beams. One can see in these equations that the DC output power scales with the reflectivity and transmittance of the beamsplitter and may be different if the two incoming beams have different powers. This is not true for the heterodyne signal, which has always the same amplitude in both output ports independent of a potentially imbalanced reflection-transmission and input power ratio of the beamsplitter.

The situation is different for the straylight amplitude, since the interference has taken place at the fibre injector (see Section 3.2.2). The beam is just split at the
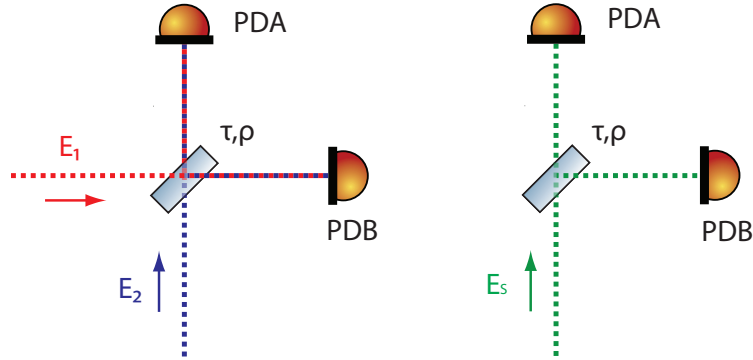
Figure 3.18: Electro-magnetic fields at the beamsplitter output.

beamsplitter with the same phase φ_S in both outputs, but different amplitudes. The AC power generated by the back reflected light can be written for both detectors as

$$P_A^S = \tau^2 P_S \cos\left(\frac{f_{\text{het}}}{2\pi} + \varphi_S\right), \tag{3.27}$$

$$P_B^S = \rho^2 P_S \cos\left(\frac{f_{\text{het}}}{2\pi} + \varphi_S\right). \tag{3.28}$$

The DC part of the back reflected light can be neglected, since the phasemeter only measures the phase of the signal at the heterodyne frequency.

As a consequence, the spurious signals indicated as black vectors in Figure 3.17 do not have the same amplitude ratio as the 'real' signals and thus cannot be compensated completely by normalising the signal amplitude. Therefore, the measured signals have to be scaled with the amplitude of the spurious signal and not (as assumed in an initial simple model) with the 'real' signal amplitude. For the measurements presented here, this was not done, since the amplitude of the spurious signal is not directly accessible in the measurement.

3.4.2 Polarisation

The backside fibre link interferometer uses optical fibres for the input beam injection. In addition, the fibre under test is a central part of the interferometer. Polarisation-maintaining optical fibres are used, and such fibres are well known to be sensitive to polarisation alignment of the input beams. Only if the input beam polarisation is well aligned to the fibre axis, the output polarisation is maximally linearised. Furtermore, the polarisation should be constant, since the fibres are birefringent. Therefore, it is desirable to include polarisers at the fibre inputs and outputs.

The Zerodur interferometer does not include any bonded polarising components. The reason for this is that standard polarisers cannot be used for bonding, since

they are coated. Furthermore, the fibres used are known to be sensitive to input polarisation alignment and hence the possibility of re-alignment is desirable.

Two types of polarisers are inserted in the setup as highlighted in Figure 3.19. The input polarisers prepare the polarisation of the injected beams. This way it is ensured that the local oscillator beam in the measurement interferometers has a well defined polarisation state. The second type of polarisers are the back link output polarisers. With those, it is possible to align the input polarisation to the fibre axis and also to 'clean' the polarisation of the outcoming beam.
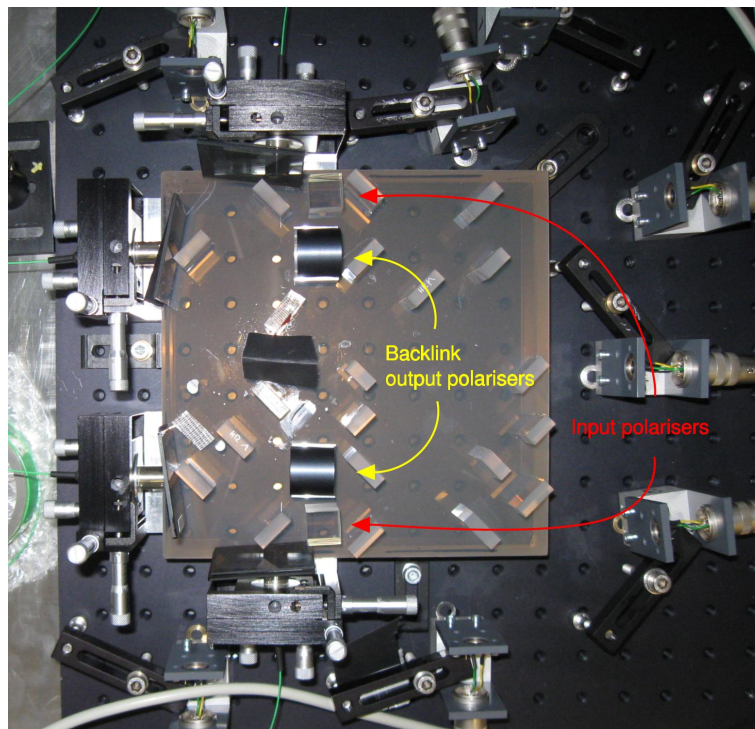
Figure 3.19: Inserted polarisers on the optical bench

The sensitivity that is reached, when polarisers are used, is shown in Figure 3.20. The straylight correction method (from Section 3.4.1) is also applied in these measurements. The sensitivity improves in the whole LISA measurement band; at some frequencies the improvement is up to one order of magnitude.

The input polarisers are polarising beamsplitter cubes. They are just put on the optical bench on one of their flat surfaces without any additional fine adjustment in rotation around the beam axis. Due to their design, their transmissive output delivers s-polarised light. The polarisers in front of the backside fibre are adjustable. However, their alignment is difficult, because it has to be done with respect to the fibre axis. The only way to optimise the alignment is to characterise the fibre output polarisation in terms of *Degree of Linear Polarisation* and orientation of the polarisation axis.
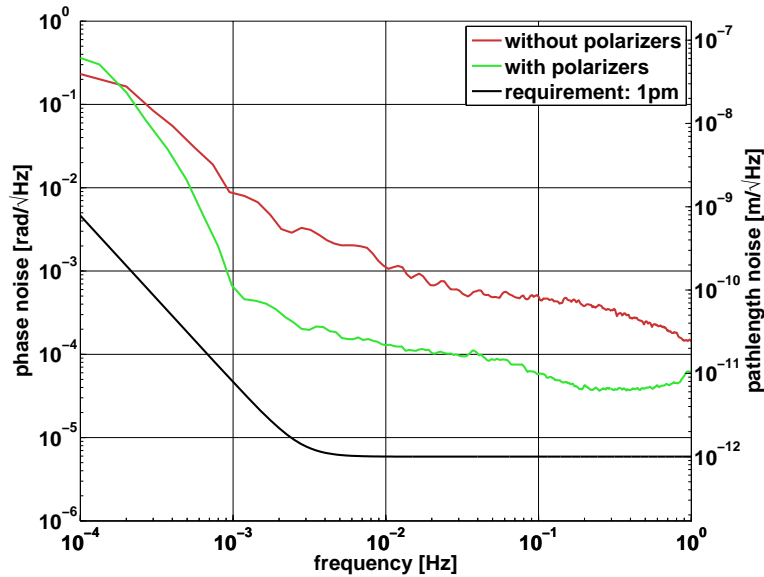
Figure 3.20: Interferometer performance with and without polarisers.

The sensitivity shown in Figure 3.20 is reached by only optimising the heterodyne amplitude in the measurement interferometers when changing the input beam polarisation. This method is not optimal, because by changing the input polarisation usually the alignment of the input beam changes and thereby the output power.

Determination of polarisation states requires complex experimental setups. A good option would be to use a commercially available polarimeter. Since such a polarimeter was not available at the time, the accuracy of the polarisation alignment of the incoming beam and the fibre axis is expected to be of the order of 1 degree. Once a polarimeter is available, the polarisation alignment can be further optimised, and the influence on the sensitivity will be investigated.

The reason for the performance decrease without polarisers is most probably caused by conversion of polarisation inside the fibre. The PANDA fibres used are known to only deliver polarisation-maintained light as long as the input polarisation matches the fibre axis. If this is not the case, the output polarisation contains not only light that has travelled on the right axis, but also light that has travelled on the wrong axis at some point. If the amount of light travelling along the wrong axis is different for the two counter-propagating beams, the measurement of the non-reciprocity is affected. In addition, the fibre acts as a retarder. If the fibre axes have different orientations at both ends, the two counter-propagating beams experience different pathlength changes inside the fibre depending on their polarisation angle.

3.4.3 Fibre stabilisation

In order to minimise non-reciprocal effects inside the fibre, a length stabilisation has been used whose actuator was developed as an OPD actuator for LTP [20].

A picture of it is shown in Figure 3.21. The fibre is wrapped around a ring-piezo
and glued to it. With this actuator it is possible to keep the (reciprocal) length
changes in the fibre constant by stabilising the phase of one of the measurement
interferometers to the reference interferometer phase. With such a length stabili-
sation the sensitivity shown in Figure 3.22 does not change. Interestingly, this is
different when no polarisers are inserted into the interferometer. Somehow the noise
introduced by polarisation conversion inside the fibre is suppressed by stabilising
the fibre length. However, at the sensitivity level of a few $10\,\mu\mathrm{rad}/\sqrt{\mathrm{Hz}}$ the length
stabilisation does not help any further.
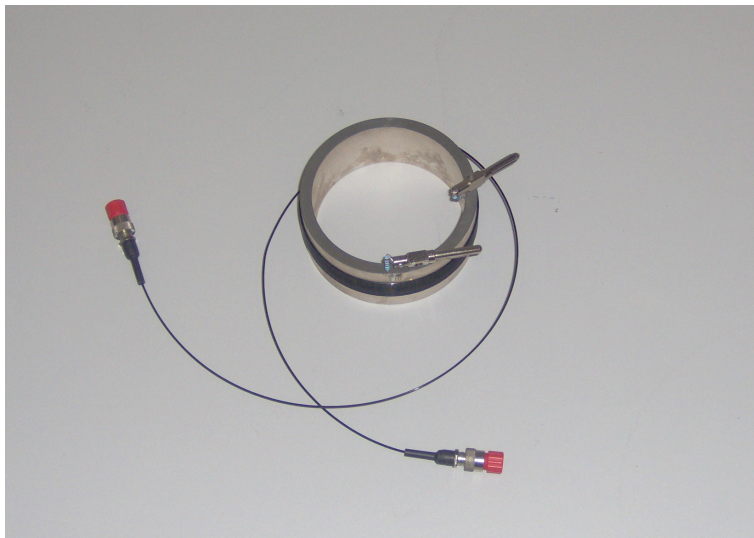
Figure 3.21: Actuators for fibre length actuator using a ring-piezo.

3.4.4 Limitations of the backside fibre link interferometer

The lowest non-reciprocal noise level achieved so far with the bonded backside fibre
link interferometer is shown in Figure 3.23. In order to reach the performance, a
LISA Pathfinder modulation scheme was used with a kHz-phasemeter running on a
PC. An OPD stabilisation scheme, which is also a copy from the LISA Pathfinder
setup (Section 6.1), is implemented. A commercially available frequency stabilised
laser with iodine stabilisation included is used. The polarisation of all beams was
linearised to an accuracy of at least 1 degree (Section 3.4.2), but further optimisa-
tion might be possible when a polarimeter is available. Length and temperature
stabilisation of the backside fibre is not necessary at the performance level shown
here. In order to reduce the noise caused by ghost beams, the subtraction method
(Section 3.4.1) was applied.

In the frequency range above $1\,\mathrm{mHz}$ the non-reciprocal phase noise is at most
about a factor of 20 above the $1\,\mathrm{pm}$ requirement. The null-measurement without
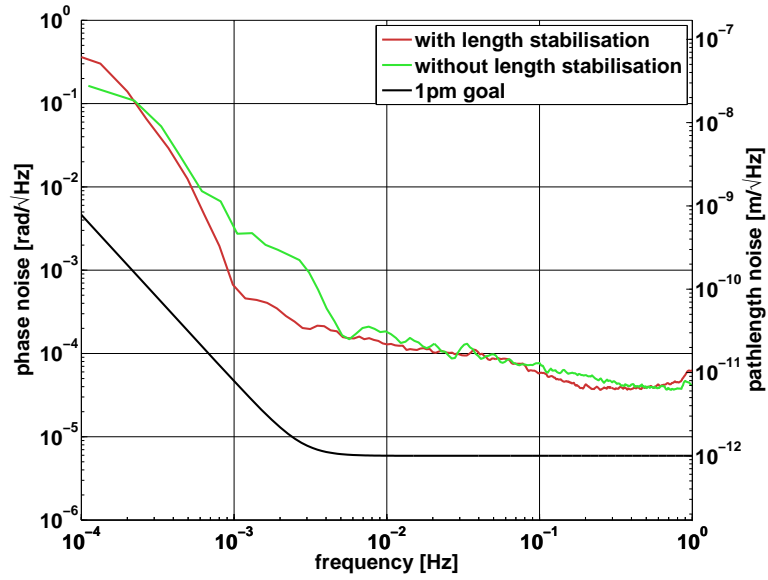
Figure 3.22: Interferometer performance with and without fibre length stabilisation.
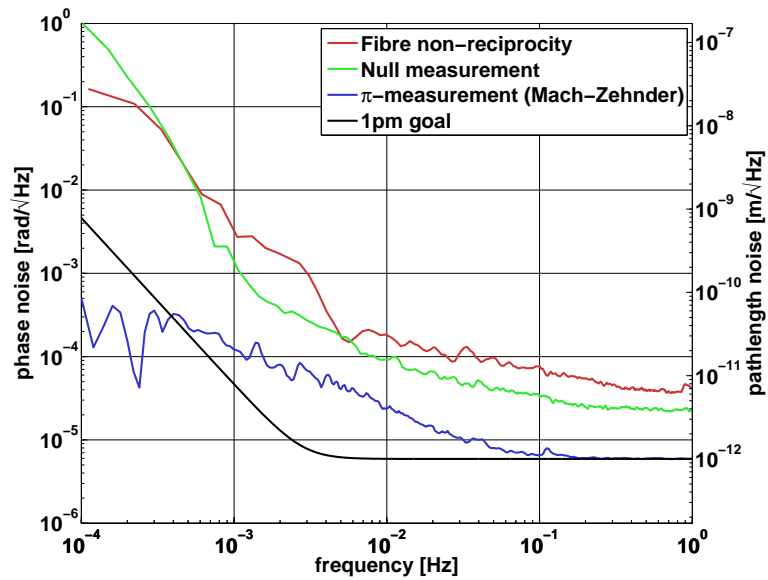
Figure 3.23: Comparison of performance reached with and without backside fibre.

the backside fibre, also shown in Figure 3.23, is at almost exactly the same level. This leads to the conclusion that the sensitivity is not limited by the fibre itself.

In order to assess the electronic readout noise of the photodiodes and the phasemeter, a spectrum of a π-measurement performed on a simple Mach-Zehnder interferometer is also shown. It is at least a factor of 4 below the fibre interferometer performance. This means that the setup is not yet limited by the readout noise either.

Laser frequency noise can be eliminated as a noise source, since it cancels to first order in the measurement of the non-reciprocity and the laser frequency noise is low compared to the LISA Pathfinder setup, where a higher sensitivity is achieved. Also from experience with LISA Pathfinder, it seems unlikely that pressure induced changes in the refractive index limit the performance. The thermal expansion coefficient of the baseplate material is also too low to have a significant influence on the noise level.

Mechanical vibrations couple into the measurement and decrease the interferometer sensitivity if the vacuum pumps are running. For high sensitivity measurements, these pumps were turned off. This coupling is part of further investigations, but also unlikely to be the limiting noise source, since the experiment is not very sensitive to daily laboratory business.

An optimised polarisation alignment in the interferometer might improve the sensitivity. However, the null-measurement should be less sensitive to polarisation misalignment, and at the current sensitivity level, the non-reciprocity of the fibre would not decrease.

A small unknown is the modulation electronics, which has not been tested yet with the LISA Pathfinder interferometer. However, it is an improved design and the characterisation of its output signals indicates a sufficient performance.

No amplitude stabilisation of the interferometer beams is implemented. Even if this was done, the required photodiodes detect spurious signals from back reflections when they are inserted after the input beam injectors. In the LISA Pathfinder experiment the amplitude stabilisation improved the sensitivity [20], and in the future the implementation will be one focus of the fibre investigations.

Noise introduced by the spurious signals from the ghost beams can limit the performance of the setup. The improvement in sensitivity, when applying the subtraction method, is already very high (about a factor of 10). One might still improve the model for the subtraction, which may lead to better results.

The ghost beams also produce signals in the null-measurement, even if the backside fibre can be removed in that case. This is due to the symmetry of the setup. *Beam 1* from Figure 3.6 has a direct optical path to the fibre injector of *beam 2* and vice versa. Ghost beams produced at the input fibre injectors cannot be avoided in the null-measurement. A solution to overcome this problem is a selection of the best possible fibre injectors for this purpose. An experiment has already been set up to characterise fibre injectors in a dedicated way.

Another possibility to reduce the ghost beam amplitude is to change the ratio

between the beam that is directly sent to the recombination beamsplitter, and the one that is coupled into the backside fibre. This is the case for LISA anyway, but due to the limited availability of bonding components, a 50-50 beamsplitter is used to separate the local oscillator beam from the back link beam (BS1 and BS3 in Figure 3.8). By increasing the local oscillator beam, the beat signal amplitude increases compared to the spurious beatnote and the error is therefore reduced. Nevertheless, the reduction of the spurious signal has to be huge, since the error produced scales with the amplitude of the signal. If the actual sensitivity was exclusively limited by the additional spurious signals, the local oscillator beam power has to be increased by a factor of about 100 to reach the required performance.

Assembly of monolithic interferometers

During the work for this thesis, two interferometers have been built using *hydroxide-catalysis bonding* techniques. One is a Mach-Zehnder interferometer that was mainly used for the weak light phase lock experiment. The other one is the backside fibre link interferometer. The former is shown in Figure 4.1.

Figure 4.1: Bonded Mach-Zehnder interferometer.

In this chapter the technique for building quasi-monolithic interferometers is described. Since the Mach-Zehnder interferometer is simpler than the backside

fibre link interferometer, only the construction of the latter one will be described in this chapter. The applied techniques are similar for both interferometers.

The hydroxide-catalysis bonding technique has already been demonstrated to be compatible for space application and performance for low-frequency measurements in the LISA measurement band [9, 10]. No fundamental development of the technique itself is part of this thesis, but special procedures for the construction were developed in order to build the complex backside fibre link interferometer (BFI). The work presented here was done in close collaboration with the Institute for Gravitational Research of the University of Glasgow.

The aim of the discussions presented in this chapter is to highlight all challenging aspects of the interferometer alignment. They are described as they appear in the build-up process. The preparation and the bonding process itself was done in collaboration with Roland Fleddermann with the support of Johanna Bogenstahl.

The alignment methods presented, the tools and the design of the fibre injectors can in general be used also for construction of other interferometers with similar stability requirements. Features and techniques that were not applied in other interferometers for similar applications are:

- the fibre injector mounting,

- alignment of fibre injectors with respect to fixed beams coming from the interferometer baseplate and

- construction of an interferometer that can be switched to measure one of two different optical paths (Here, these are the paths through the fibre and a null-measurement path on the stable interferometer baseplate.)

However, the main goal of the work presented in this chapter is the construction of an interferometer for the backside fibre link characterisation on breadboard level with a sensitivity sufficient for LISA.

4.1 The interferometer layout

From the experience with the aluminium breadboard of the BFI, the layout of the Zerodur interferometer was slightly changed. An *OptoCad* model of the backside fibre link interferometer is shown in Figure 4.2. All modifications are forseen for practical and alignment reasons. The only change in the setup, relevant for the measurement, is an additional path, which gives the possibility of performing a measurement restricted to the ultra-stable optical bench without the fibre.

In contrast to the initial setup shown in Figure 3.6, the following changes are forseen:

1. Beamsplitters BS5 and BS6 are inserted. The backside fibre can be blocked and the path between these two beamsplitters can be used as null-measurement on the ultra-stable optical bench.
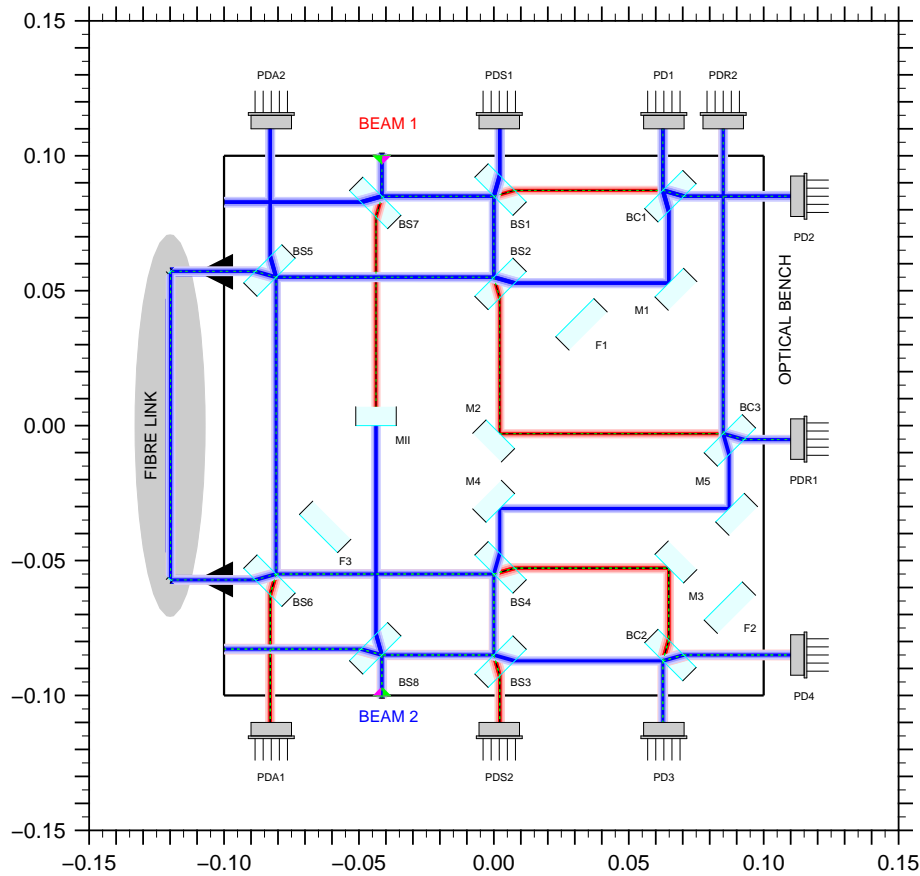
Figure 4.2: Layout of the backside fibre link interferometer.

2. The two input fibre injectors are put on different sides of the interferometer due to the small size of the baseplate ($20 \times 20\,\text{cm}^2$) and the need for room for the beam injectors.

3. The beamsplitters (BS7 and BS8) redirect the input beams such that the reflected beams have the same orientation as in the initial interferometer setup. The additional output beams are used to align the input beams to the optical bench.

4. The mirror M$\parallel$, which is close to the centre of the baseplate, reflects the input beam back into the fibre injector. It is used to align the input beam parallel to the baseplate.

5. The components F1, F2 and F3 are not used optically, but also fixed using the same bonding technique in order to benefit from the strength of the bond. They define reference points for a template that are well defined with respect to the interferometer baseplate.

4.2 Properties of the bonding material

Hydroxide-catalysis bonding is used to build mechanically and thermally highly stable assemblies. The reason for using this technique to build the backside fibre link interferometer is to have a material with low expansion coefficient as baseplate. This reduces pathlength changes introduced by thermal variations.

The baseplate of the BFI is made of Zerodur, which is a glass ceramic with a similar structure as fused silica. The components are made of pure fused silica. The simplified structure of both component and baseplate is shown in Figure 4.3. Both mainly consist of silicon dioxide (SiO_2), but at the edge of the material they are terminated by hydroxyl groups. A catalyst is used to build a crystal lattice between both components [9].
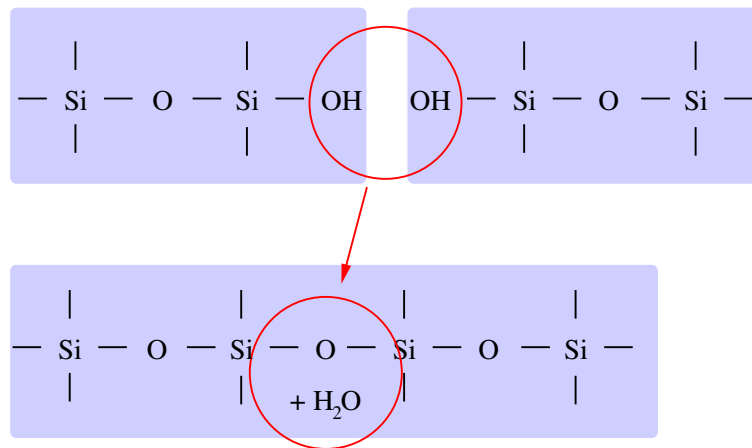
Figure 4.3: Basic chemical reaction of the bonding process.

Zerodur is designed to have a minimal index of thermal expansion at temperatures around room temperature. It is available with expansion coefficients between 10^{-7} and $2 \cdot 10^{-8}$/K [40].

The components can only be aligned along the surface normal of the baseplate, the vertical degrees of freedom are fixed and require very accurate machining of the components. In order to ensure an adequate bond strength, the baseplate surface and the bottom surface of the components have to be polished to an optical surface quality of $\lambda/10$. The baseplate is shown in Figure 4.4(a). It has a size of $20 \times 20 \, \text{cm}^2$. The perpendicularity of the optical surfaces to the bottom surface of the components is specified for all components to be less than 1 arcsecond.

An example of an optical component bonded on a Zerodur baseplate is shown in Figure 4.4(b). Its alignment is very challenging. Since the surfaces of the component and the baseplate are polished, they contact optically once the component is put on the baseplate [10]. This has to be avoided, since it cannot be taken apart without destroying the surface quality. In order to ensure a proper bond, a catalyst is used, which is put on the bottom surface of each component. The component can then

be placed onto the baseplate, where it floats on the fluid catalyst. It can be aligned for less than one minute until the bond hardens. It is possible to get components off the baseplate for a short while after this time if one accepts small damages of the surface. For critical alignment, a non-bonding organic buffer fluid can be used, which allows longer alignment times (of the order of 3 minutes). However, after the right position has been found, the component has to be taken off and both the component and the baseplate have to be cleaned again, before the catalyst can be used.

(a) Zerodur baseplate.   (b) Bonded optical component.
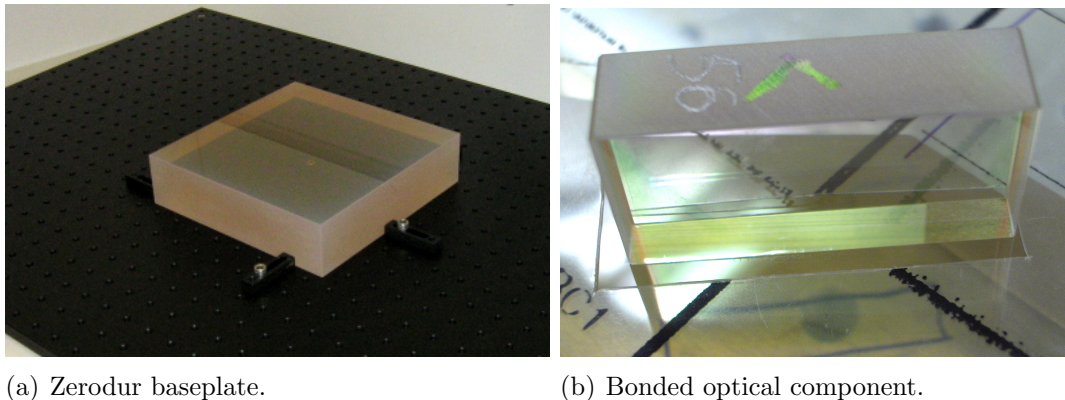
Figure 4.4: Picture of the Zerodur baseplate and a bonded component.

These constrains make the bonding process very challenging, and special tools have to be prepared to have fast and accurate alignment possibilities, as described in [10]. The following sections will show in more detail how the construction of the BFI was realised.

4.3 Alignment of non-critical components

Apart from the recombination beamsplitters, all components that only split or reflect the beams are less critical in terms of alignment. In order to place them on the baseplate, a template, shown in Figure 4.5(a) is used.

It is made of brass with a typical machining accuracy of $10..100\,\mu$m. In the figure it is mounted above the interferometer baseplate, which is covered in order to avoid contamination. Each component has a pocket which defines its position. The pockets have three small spheres as shown in Figure 4.5(b). They form a three point reference for the optical component. In order to keep the components touching the spheres, the baseplate is tilted by an angle of 5 degrees during the bonding procedure. Otherwise the components would randomly float on the bonding fluid within the pockets without having any well defined position. The overall accuracy of the template including the component reference spheres is assumed to be well below 1 mm. In Figure 4.5(b) the three different applications of the pockets are shown:
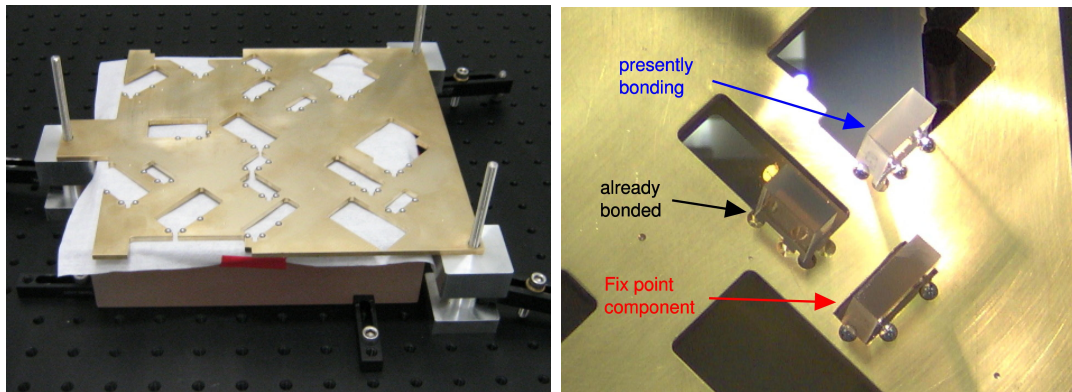
(a) Template for non-critical components.   (b) Pockets for components.

Figure 4.5: Template for the component alignment.

1. The fix point components define the position of the template with respect to the baseplate. Three of these fix points are forseen on the optical bench. They are bonded in the first step of the bonding procedure. Their distance to the pockets is 1 mm. The spheres stay in the pockets for the entire bonding procedure in order to be able to put the template in the same position before each bonding step.

2. The component which is presently bonding touches the spheres and has, therefore, a well defined position with respect to the template and the baseplate. The template stays for at least half an hour in its position until the component has settled and the bond has started to harden. Afterwards it is removed.

3. Since the bonding process happens in many steps with intermediate removal of the template and cleaning of components and baseplate, the spheres are removed for components that are already bonded in order to avoid misalignment of the template due to the presence of too many reference points.

The way the template is used, it is always ensured that the reference position is given by the bonded fix points on the optical bench. By using the small sized pockets for the fix point components, the template cannot touch and push optical components and break them as long as the fix point components are not broken off the baseplate.

4.4 Input beam alignment

One of the most critical alignment parts of the interferometer is the input beam alignment. The input beams define the location of all other optical components. Therefore, it was decided to glue all fibre injectors directly to the optical bench. One injector is shown in Figure 4.6(a) while it is prepared. In order to have a

solid connection between fibre injector and optical bench, an epoxy resin (ER2188 from Electrotube) is used. It has a low index of thermal expansion (40 ppm/K) and the water absorption during 10 days is less than 0.5 % [41]. For the alignment of the fibre injectors, a mount with 4 degrees of freedom is forseen: Tip, tilt and two translational stages, as shown in Figure 4.6(b), are necessary to align the input beams in all degrees that are predefined by the optical components. A compact mount with an overall width of less than 10 cm was chosen. The alignment actuators are micrometer-screws.
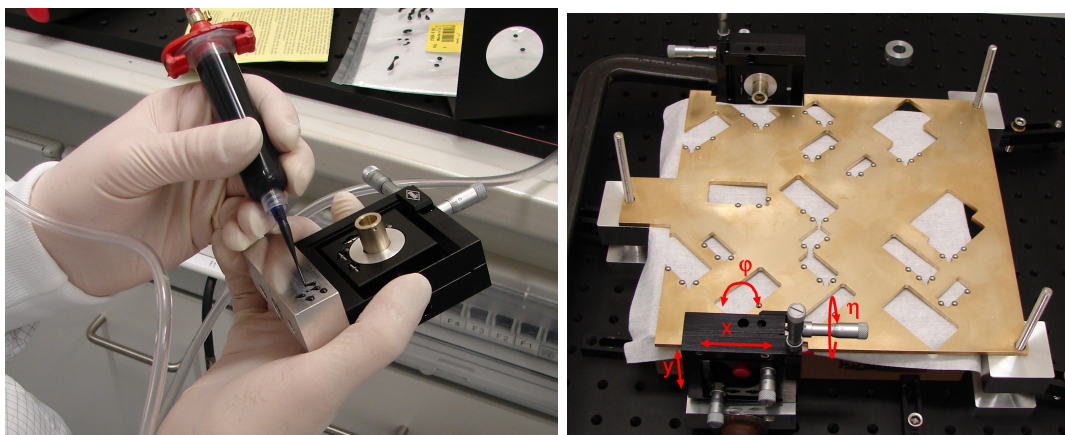
(a) Glueing the fibre injector. (b) Attached fibre injectors.

Figure 4.6: Attachment of fibre injector to the baseplate.

All other optical components are bonded to the baseplate using hydroxide-catalysis bonding. Before this is possible, the surfaces of components and baseplate have to be adequately cleaned. This happens in two stages. In the first – rough cleaning – stage organic substances are removed. This cleaning step is illustrated in Figure 4.7(a). In a second step the surfaces are cleaned more careful and at the same time prepared for bonding.

After this cleaning stage the surfaces have to be covered with clean water to avoid contamination and absorption of particles from the ambient air. The whole cleaning procedure has to be done as shortly as possible before each component is bonded. Therefore, it is desirable to bond only very few components at once.

The first components to be bonded on the baseplate are the fix point components and the mirror M∥. They are shown in Figure 4.7(b). Since the vertical angle between the beam and the baseplate cannot be changed by the alignment of the bonded components, the input beams have to be as parallel to the baseplate as possible. The first beam injected onto the optical bench defines the beam height in the interferometers. The mirror M∥ is used to reflect the injected light directly back into the fibre. This ensures the parallelism of the first beam with respect to the baseplate to the level of the mirror perpendicularity, which is less than 1 arcsecond.

All non-critical components are bonded using the template as reference. Only
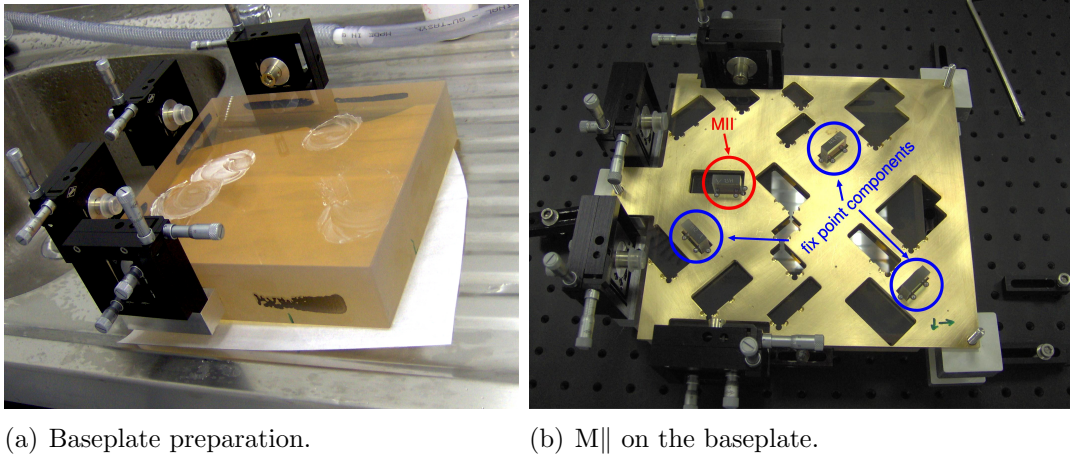
(a) Baseplate preparation.   (b) M∥ on the baseplate.

Figure 4.7: Preparation and bonding of first set of components.

the recombination beamsplitters need more accurate alignment. Due to the symmetry of the interferometric layout, it is possible to easily align the second injected beam with respect to the first one. For this purpose, the null-measurement path is used, which is indicated in Figure 4.8. This is the key point in the alignment of the interferometer, since then it is possible to optimise the fibre coupling and interferometer alignment for the fibre and for the null-measurement in parallel.
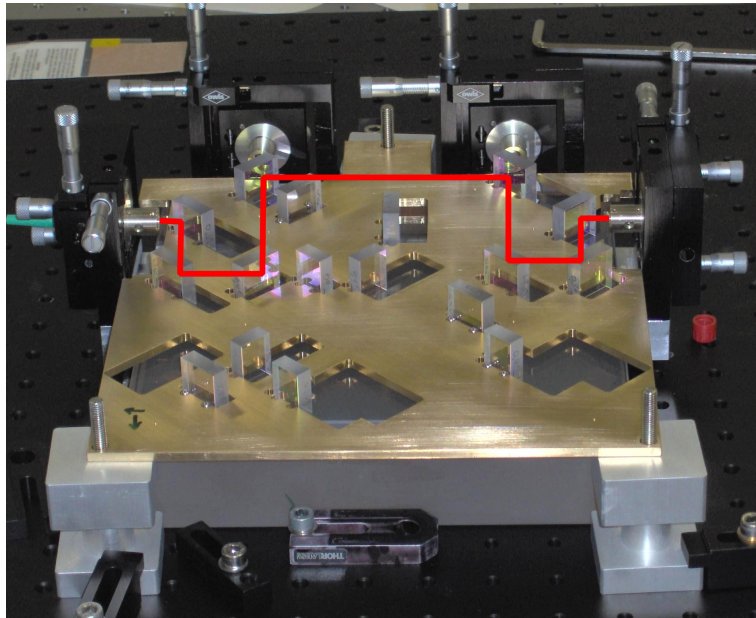
Figure 4.8: Null-measurement path on the optical bench.

By optimising the coupling into the the second injector, both input beams automatically get the same height and the same angle. Since the first beam is already

aligned to be parallel to the baseplate, the second beam is also parallel after this alignment step. After both backside fibre link injectors are aligned to the input beams, they are also parallel and at the same height as the input beam injectors. In addition the null-measurement path has automatically a completely common path with the fibre measurement path and the contrast in the measurement interferometers is optimised for both at the same time independent of the path (fibre or null-measurement path) that is used for the alignment.

When all components apart from the recombination beamsplitters are bonded, it is still possible to optimise the fibre injector alignments. Before bonding of the recombination beamsplitters, additional auxiliary interferometers were set up next to the baseplate in order to optimise the contrast by adjusting the fibre injector lenses. Thus, the modes of all four fibre injectors are matched with respect to each other such that all interferometers have optimal contrast and the coupling efficiency into the backside fibre link is sufficient.

4.5 Recombination beamsplitter alignment

The most challenging step in the interferometer construction procedure is the recombination beamsplitter alignment. All fibre injectors and the reference interferometers have to be optimally aligned before. There is exactly one optimal position for the beam combiner. It has to be found by aligning the component in one lateral and one angular degree of freedom. As already mentioned, the time to position the component when the bonding fluid is in place is less than one minute. During alignment the recombination beamsplitter output signals must be measured with photodiodes in order to be able to optimise the contrast in the interferometer.

For adjustment, an alignment tool, shown in Figure 4.9(a), was used. It consists of a bridge with a movable console. The console itself carries two micrometer screws with mechanical fingers attached to them. Due to the tilt of the baseplate, the recombination beamsplitter is pushed against the finger tips and aligned as illustrated in Figure 4.9(b). Lateral and angular alignment of the component is not disentangled with this tool. The position and angle of the component are iteratively varied in small steps until the contrast in the interferometer is optimised.

The first interferometer to be built was the reference interferometer, since it only needs two well aligned fibre injectors (input beam injectors). The achieved contrast in the final setup is about $80\,\%$. After this step, the fibre injector alignment is defined completely and cannot be changed without loosing contrast in the reference interferometer.

As final step in the bonding procedure, the recombination beamsplitters in the measurement interferometers were aligned. In order to check the alignment state of all fibre couplers during the bonding of the measurement interferometers, the transmitted power was continuously monitored during the process. Due to the alignment of the fibre injectors via the null-measurement path, the contrast of the
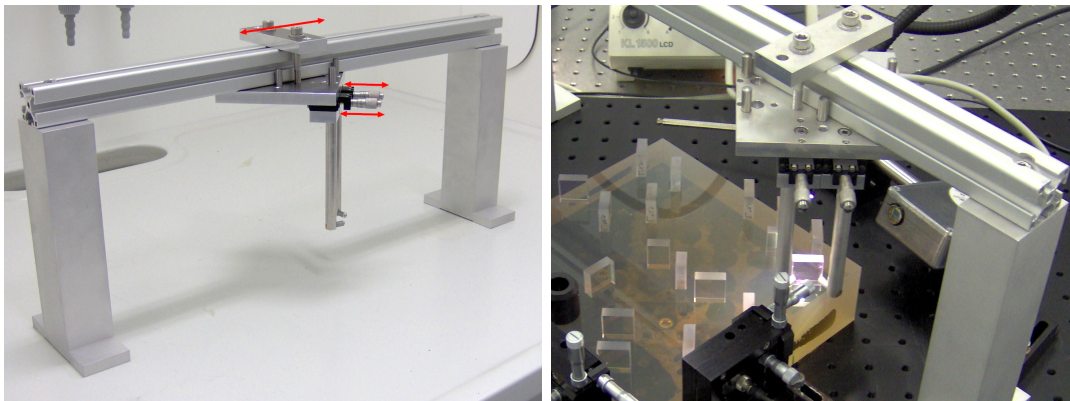
(a) Alignment tool.   (b) Recombination beamsplitter alignment.

Figure 4.9: Pictures of the alignment tool in use for recombination beamsplitter adjustment.

interferometers is optimal if the fibre and also if the null-measurement path is used. Both must be checked, before it can be decided whether or not to leave the component on the baseplate.

The complete setup of the interferometer is shown in Figure 4.10. The picture was taken while the last recombination beamsplitter bond settles.

4.6 Conclusion

The backside fibre link interferometer is one of the most complex interferometers that was ever built using hydroxide-catalysis bonding. Up to now only two other interferometers with similar complexity have been built using this technique. One of those is the first breadboard constructed at University of Glasgow (UGL) [10]. It was intended to demonstrate a sufficient performance for the LISA Pathfinder and LISA missions. Apart from this, the LISA Pathfinder Engineering Model optical bench was constructed [42] and its performance was demonstrated (see also Section 6.2) [21]. Presently other interferometers are built at UGL: One is the Flight Model for LISA Pathfinder, but also further interferometers for testing parts of the LISA interferometry are planned [43].

The interferometer presented here includes specific aspects of the LISA *local interferometry* that have not been investigated before. These are the alignment of fibre injectors with respect to fixed beams coming from the optical bench. Also a fibre is included in the interferometer path and an additional path on the ultra-stable bench is realised, which can be used for instrument characterisation. The fibre injectors, the backside fibre link interferometer and the null-measurement interferometer are aligned at the same time, which is technically difficult to implement. None of these features was realised in other quasi-monolithic interferometers that have been built up to now.
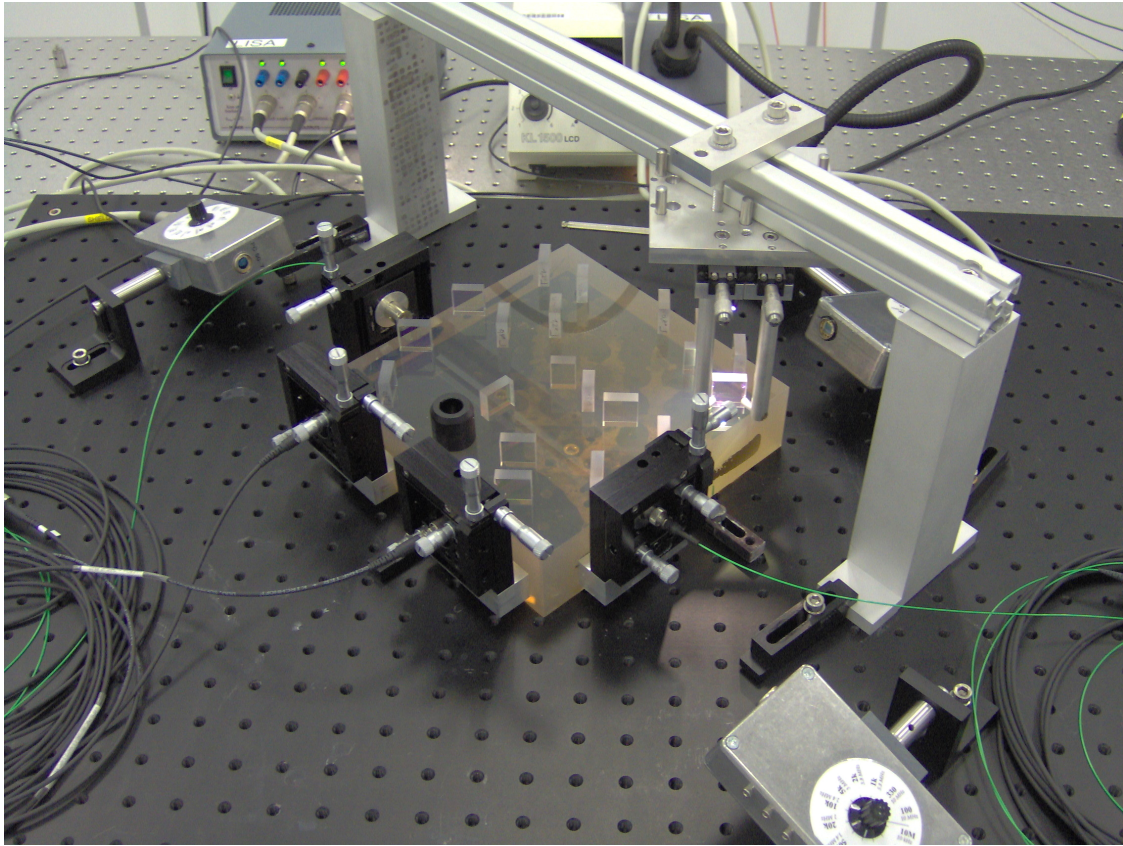
Figure 4.10: Complete backside fibre link interferometer.

Fibre injectors in such interferometers are a big issue, since mounts are needed to align them with respect to the optical bench, or alternatively the beam position and orientation of the outgoing beam has to be measured accurately. For the LISA Pathfinder Flight Model, the fibre injector will also be quasi-monolithic and their assembly will be fixed to the interferometer baseplate using hydroxide-catalysis bonding. It is technically not possible to build such injectors with standard laboratory equipment.

The fibre injector mounting used here seems to be adequate for LISA interferometry tests on breadboard level. It is a trade-off between a stable assembly fixed to the baseplate and the use of additional input beam alignment optics. It is realised using combined translational rotational stages with four degrees of freedom that are fixed to the interferometer baseplate. Even if this is not applicable for LISA, a sufficient performance on breadboard level could be shown down to $10 \, \mathrm{pm}/\sqrt{\mathrm{Hz}}$ at $10 \, \mathrm{mHz}$. This performance is probably limited by the appearance of spurious signals as discussed in Section 3.4.4. Therefore, it might also be useful if higher sensitivities are required.

Part II

LISA Pathfinder Interferometry

The LISA Pathfinder mission

LISA Pathfinder is a technology demonstration mission to be launched in 2010. The aim of the mission is to demonstrate a differential acceleration noise of two test masses of less than $3 \cdot 10^{-14}\,\mathrm{N}/\sqrt{\mathrm{Hz}}$ in the frequency range between 3 and 30 mHz [44]. This is close to the required acceleration noise for the LISA test masses. Both sensitivity goals are shown in Figure 5.1: On the left hand side the required differential pathlength noise of the two test masses is shown, and on the right hand side it is converted to acceleration noise.
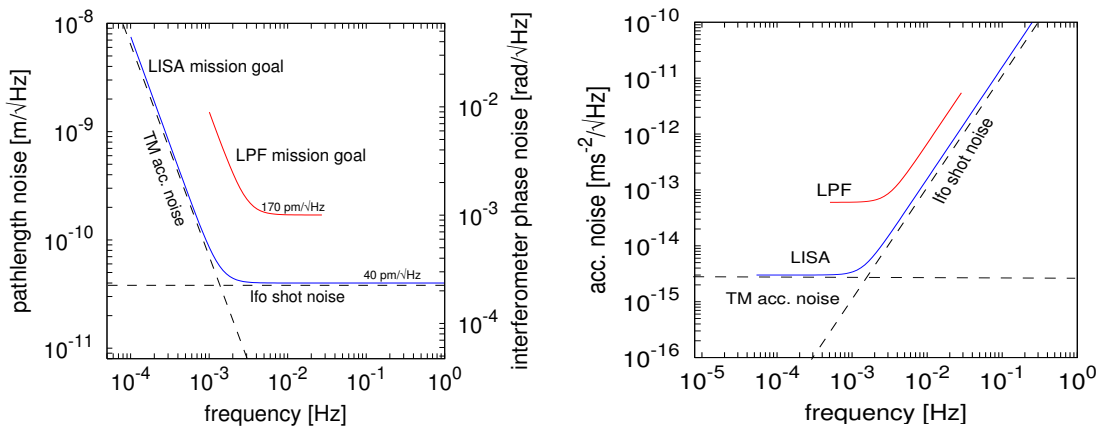
Figure 5.1: Sensitivity goals for LISA and LISA Pathfinder.

Some of the elements on LISA Pathfinder are impossible to test on earth. These are

- the Inertial Sensor,

- the Micro-Newton thrusters,

- the LISA *local interferometry* and

- the test mass and satellite control.

They will be introduced in more detail in the following. The tests of these techniques delivers essential information for LISA.

The differential test mass acceleration is measured along the main sensitive axis, which connects the two centres of the test masses and is called x. On-board LISA Pathfinder, the distance between the two test masses is reduced from $5 \cdot 10^6$ km (for LISA) down to about 30 cm. For the test mass sensing and control, an electrostatic sensor and actuator is used. It is called Inertial Reference Sensor (IRS). For the position control of the satellite, Micro-Newton thrusters are installed. Both systems will be similar for LISA.

Ideally both test masses would be drag-free, but due to effects like spacecraft internal gravity gradients, solar wind and electro-magnetic forces acting on the test masses, this is technically not possible and active control of both test masses and spacecraft in multiple degrees of freedom is needed [45]. The control is performed by the Drag-Free and Attitude Control System (DFACS).

The distance fluctuations between the two test masses along the sensitive axis are additionally read out with an interferometer, which is more accurate than the electrostatic readout. Without this additional readout, it would not be possible to measure a residual acceleration noise of $3 \cdot 10^{-14}$ N/$\sqrt{\text{Hz}}$ in the mHz-range. Furthermore, the interferometer is a test of the LISA local interferometry.

The LISA Pathfinder satellite will have a Lissajous orbit at the Lagrangian point L1. The mission duration is planned to be three months with a potential extension of another three months.

5.1 LISA Pathfinder subsystems and operation

The satellite protects the two test masses, shown in Figure 5.2(a), from external forces. Since this protection is not perfect, actuators for the test masses and the satellite are implemented. In order to compensate for the residual differential forces on the test masses, a capacitive sensor and actuator, which is shown in Figure 5.2(b), is used. The actuator is an integrated part of the Inertial Reference Sensor and can actively control, but also sense the test masses in all degrees of freedom. For precise control of the satellite, Micro-Newton thrusters, shown in Figure 5.2(c), are implemented, which can actuate the satellite accurately with forces of up to a few 100's of μN.

The inner part of the satellite includes the main part of the LISA Pathfinder experiment. It is called the LISA Technology Package (LTP) core assembly. It is shown in Figure 5.3. The test masses are surrounded by vacuum chambers (grey

(a) Test mass   (b) Inertial Reference Sensor   (c) Micro-Newton thruster
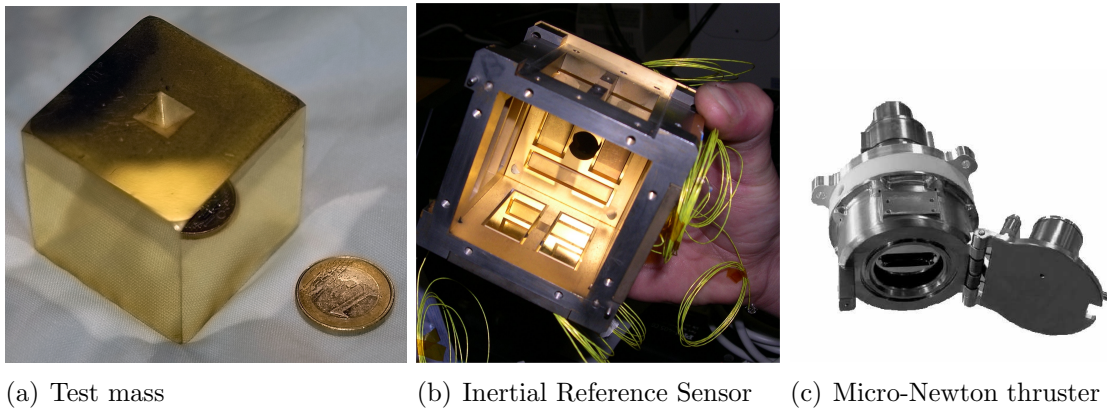
Figure 5.2: Subsystems of LISA Pathfinder [46].
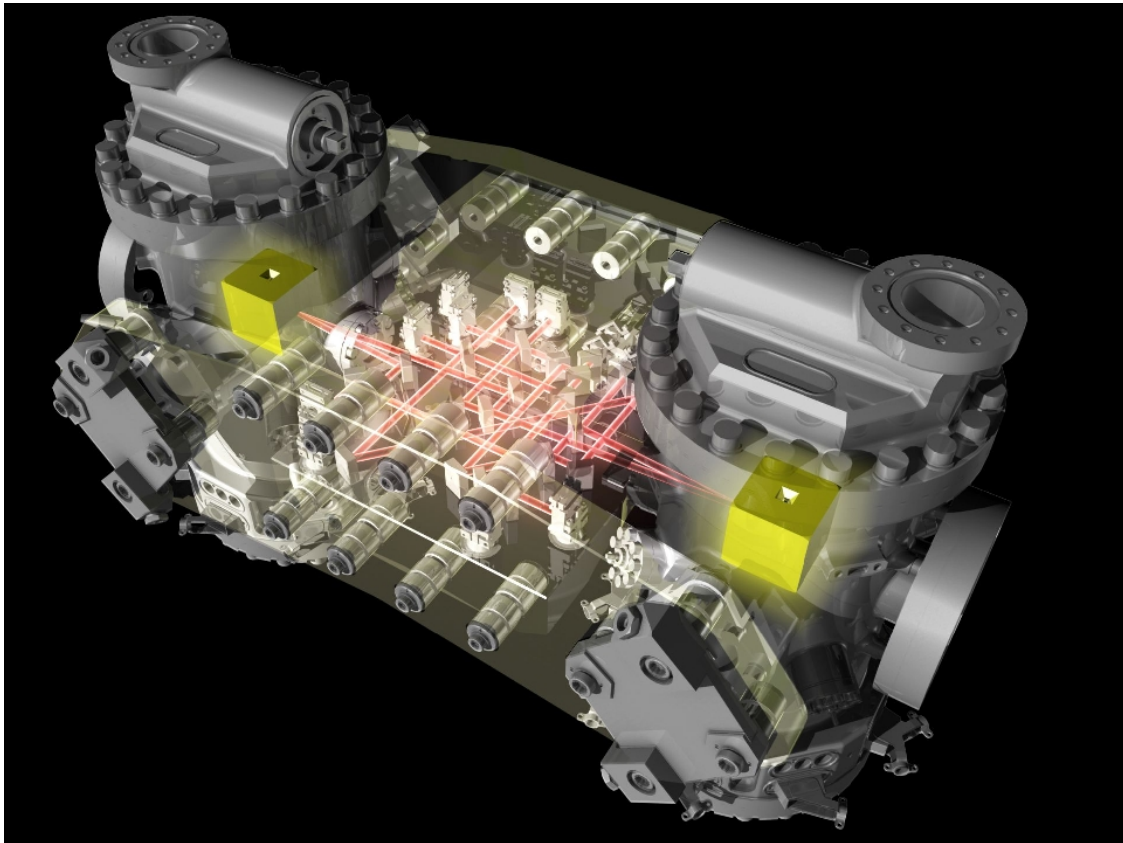
Figure 5.3: The LISA Technology Package core assembly [46].

units). Each of those chambers carries one inertial sensor with a test mass inside. The interferometer is indicated by the red rays in this figure. It measures the position of one test mass with respect to the satellite and the differential movement of both. Additionally, two angles of each test mass are read out [38]. The interferometer, its sub-units (e.g. Laser, Laser Modulator) plus its back-end, the so-called Data Management Unit (DMU) is called the Optical Metrology System (OMS) [47]. It will be described in more detail in the following sections.

The three systems introduced – Micro-Newton thrusters, Inertial Reference Sensor and the interferometer – deliver signals to, or receive signals from, the DFACS [47]. A schematic view of the satellite and its control in the standard science mode of LISA Pathfinder is shown in Figure 5.4. The DFACS uses interferometric and IRS data as input error signals and controls the satellite via the Micro-Newton thrusters, with the controller (C_{df}), and the test masses via the electrostatic actuator inside the IRS, with the controller (C_{sus}). The interferometer data is processed inside the Data Management Unit.
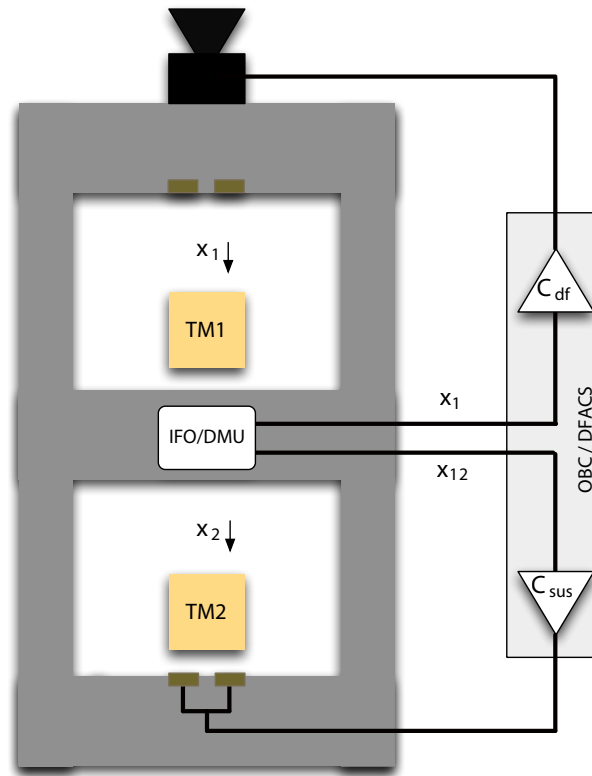
Figure 5.4: Schematic view of the satellite control.

The control must be performed carefully such that the differential test mass acceleration can be measured in a quasi drag-free mode. This is realised as follows: The satellite follows test mass 1 in order to keep it in the centre of the Inertial Reference Sensor, where the highest sensitivity is expected. For this feedback, the

interferometer data is used that measures the test mass 1 position x_1 with respect to the spacecraft. This actuation can be applied with high bandwidth, since only the satellite moves, but the differential displacement between the test masses stays in principle constant. The second test mass is made to follow the satellite in order to compensate for differential test mass acceleration [45] using the interferometer output that measures the differential displacement x_{12} of both test masses. The control of the second test mass is necessary, since the DC acceleration acting on the test mass must be compensated. The test mass 2 control is more critical, since it changes the x_{12} measurement. Therefore, the actuation takes place below the LISA Pathfinder measurement band.

The expected residual pathlength noise of both interferometers is shown in Figure 5.5. The test mass 1 position is measured with respect to the unstable satellite. Its noise is dominated by the Micro-Newton thrusters. The differential test mass displacement is below the required LISA Pathfinder mission goal.
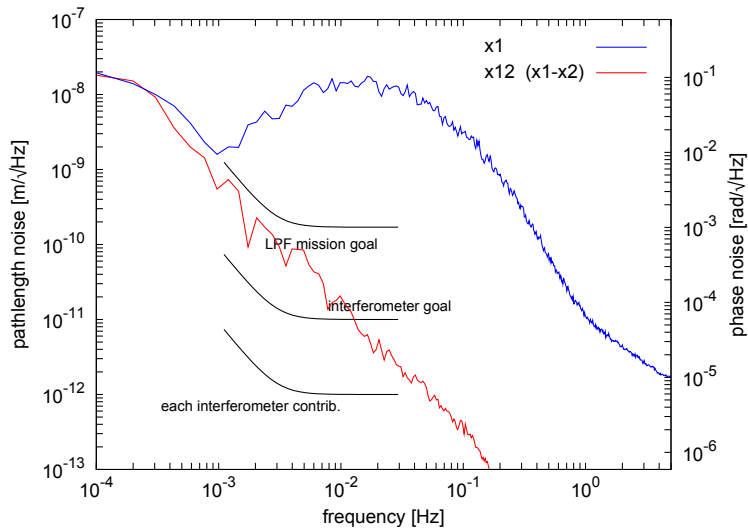
Figure 5.5: Expected residual test mass motion in science mode.

The main aim of LISA Pathfinder is to identify and understand all forces acting on the satellite, the (single) test masses and especially the differential force noise on both test masses. The second important point is to optimise the DFACS control scheme such that the sensitivity is optimised in order to develop a slightly different and improved control scheme also for LISA.

In the following chapters, the LISA Pathfinder interferometry is introduced. The focus lies on the interferometric readout performance with specific aspects of the LISA Pathfinder on-board data processing software. It includes the definition and testing of this processing and its influence on the DFACS.

5.2 LISA Pathfinder coordinate frame

For the following chapters, it is useful to understand the coordinate system on-board LISA Pathfinder. If not otherwise stated, the notation used refers to the one introduced here. The basic coordinate system of LISA Pathfinder is shown in Figure 5.6. The x-axis is the sensitive axis of the interferometer. It points from test mass 2 to test mass 1. The coordinate system is right-handed, and the y-axis is in the plane of the optical bench. The origin is between the test masses and therefore at the centre of the optical bench. All angles are indicated. The interferometer is sensitive to test mass rotation around the y- and z-axes, which are η and φ.
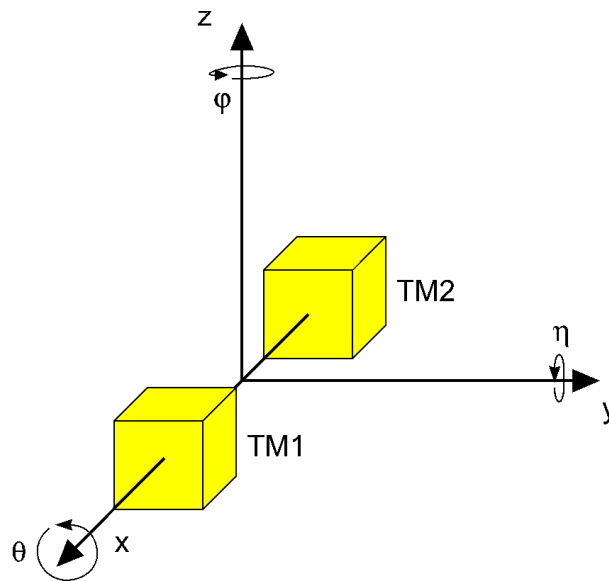
Figure 5.6: The LTP reference coordinate frame.

This notation is important mainly for the description of the interferometer data back-end processing (Chapter 8) and the experimental results of the software tests (Chapter 9).

CHAPTER 6

The LTP interferometer

In this chapter, the experimental setup of the LISA Pathfinder interferometer is introduced. A simplified block diagram is shown in Figure 6.1. The essential parts are the optical interferometer setup, which is the Engineering Model (EM) of the LISA Pathfinder interferometer, and the phasemeter that reads the detector output signals and calculates from those the desired signals. The modulation bench prepares the two input beams for the optical bench. Some of the data processing steps are implemented in a C-program running on a commercial PC. Two feedback loops are discussed here, the laser frequency and the *optical pathlength difference* (OPD) stabilisation. Their error signals are extracted from photodiodes on the optical bench.

The interferometer - also called optical bench (OB) - was designed at AEI Hannover and built at the Rutherford Appelton Laboratories with support of the LISA Pathfinder groups at AEI Hannover and the University of Glasgow. Environmental qualification tests for space applications were performed on the interferometer at TNO – an institute in Delft, The Netherlands – in 2004 and since then it was used at AEI for tests of performance for the longitudinal and angular test mass readout. Results were already published in [19, 20, 48].

The individual parts of the setup will be briefly described in this chapter. In particular the readout of the interferometer is important for the understanding of the back-end data processing inside the Data Management Unit (DMU) (Chapter 8).
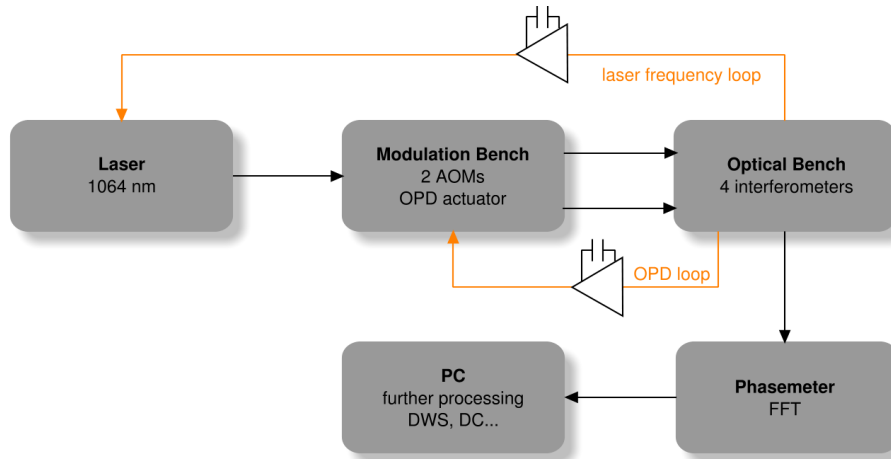
Figure 6.1: Simplified schematic view of the LTP laboratory experiment.

6.1 The length measurement

The LTP interferometer is of heterodyne Mach-Zehnder type. For the beam generation, a laser is split into two beams and transmitted through Acousto-Optical Modulators (AOMs). Both are frequency shifted by about 80 MHz with a slight frequency difference of the order of 1 kHz, as indicated in Figure 6.2. The recombination beamsplitter in this figure is not part of the modulation bench and here, used only to illustrate the generation of the interference signal.
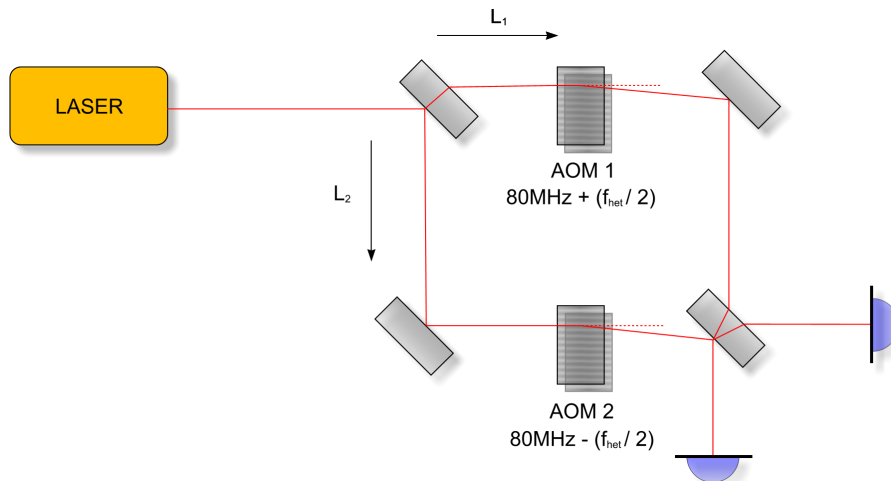
Figure 6.2: Generation of the heterodyne signal in the LTP interferometer.

Two beams interfering at this recombination beamsplitter lead to a photo current of

$$I_{\mathrm{PD}} = I_{\mathrm{DC}} + I_0^{\mathrm{AC}} \cos\left(\omega_{\mathrm{het}} t + \varphi\right) \tag{6.1}$$

on the photodiode. The AC photo current with amplitude I_{AC} oscillates with the frequency difference

$$f_{\mathrm{het}} = \frac{1}{2\pi} \cdot \omega_{\mathrm{het}} \tag{6.2}$$

of the two laser beams, which is of the order of $1\,\mathrm{kHz}$. For such a heterodyne scheme, the DC photo current I_{DC} is not important for the length measurement and can be omitted. The phase changes, φ, of this oscillation are proportional to the differential pathlength changes, ΔL, in the two interferometer arms, which start at the first beamsplitter. The relation can be written as

$$\varphi = \frac{2\pi}{\lambda} \cdot \Delta L, \tag{6.3}$$

where λ is the wavelength of the laser – in this case $1064\,\mathrm{nm}$.

In the LISA Pathfinder interferometer, the modulation bench and the ultra-stable optical bench are separated as illustrated in Figure 6.3. In addition, both beams are split on the optical bench and the phases φ_{R} and φ_{M} are measured in the two interferometers, where both injected beams are brought to interference. The index 'R' indicates the reference interferometer, whose entire optical paths lies on the stable baseplate. The index 'M' indicates the measurement interferometer which measures the fluctuations in its variable path. In Figure 6.3 these fluctuations are indicated as ΔM introduced by a moveable mirror. In LISA Pathfinder, this is one of the test masses. The measured phases include the optical paths starting at the
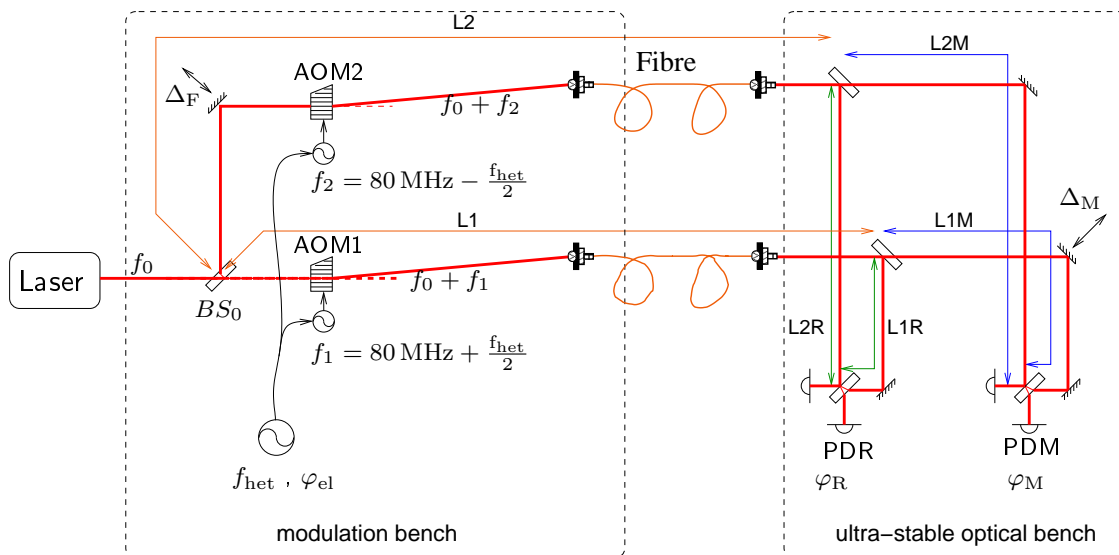
Figure 6.3: Separation of the modulation bench and the ultra-stable optical bench.

beamsplitter BS0 on the modulation bench. They can be written as

$$\varphi_{\mathrm{R}} \quad = \quad \frac{2\pi}{\lambda}\left(L1 + L1R - L2 - L2R\right), \tag{6.4}$$

$$\varphi_{\mathrm{M}} \quad = \quad \frac{2\pi}{\lambda}\left(L1 + L1M - L2 - L2M\right), \tag{6.5}$$

with the paths L indicated in Figure 6.3. The phase difference

$$\Psi = \varphi_{\mathrm{M}} - \varphi_{\mathrm{R}} = \frac{2\pi}{\lambda}\left(L1M - L1R - L2M + L2R\right) \tag{6.6}$$

includes only optical path on the ultra-stable optical bench, since the length fluctuations on the modulation bench, L1 and L2, are common-mode in both interferometers.

6.2 The optical bench

A schematic of the LTP Engineering Model optical bench is shown in Figure 6.4. The optical bench carries two single element photodiodes (PDA1 and PDA2) for an amplitude stabilisation of the injected laser beams. The other eight detectors are quadrant photodiodes (QPDs). They read out the interference of the four different interferometers. For each interferometer, two of these QPDs carry redundant information and can be used individually for the phase readout or be averaged as described in Section 8.1.3.

The two modulated beams are injected onto the optical bench via fibre output couplers from single-mode polarisation-maintaining fibres. The first beamsplitter for each beam (BS11 and BS16) samples some light for the amplitude stabilisation of the laser beams. The next beamsplitters are BS1 and BS4. They split the beams for the four interferometers situated on the optical bench. This has the consequence that all fluctuations in pathlength between the beamsplitter in front of the AOMs (on the modulation bench) and BS1 and BS4 are common-mode in all interferometers, which are individually illustrated in Figure 6.5.

The reference interferometer in Figure 6.5(a) has components only on the optical bench, which is made of Zerodur and therefore very insensitive to thermal fluctuations. It measures the common pathlength changes in the unstable part meaning the modulation bench. Those fluctuations are the same in all interferometers, and therefore, the reference interferometer can be used for subtraction. Its output is also used to produce the error signal for the OPD stabilisation. The frequency interferometer in Figure 6.5(b) has an intentional pathlength mismatch and is therefore sensitive to laser frequency fluctuations. Usually it is used to produce the error signal for the laser frequency stabilisation loop, but it can also just have a monitoring function or be used for noise subtraction. The X1 interferometer in Figure 6.5(c) measures the motion of test mass 1 with respect to the optical bench and the X12 interferometer in Figure 6.5(d) measures the differential motion of both test masses.
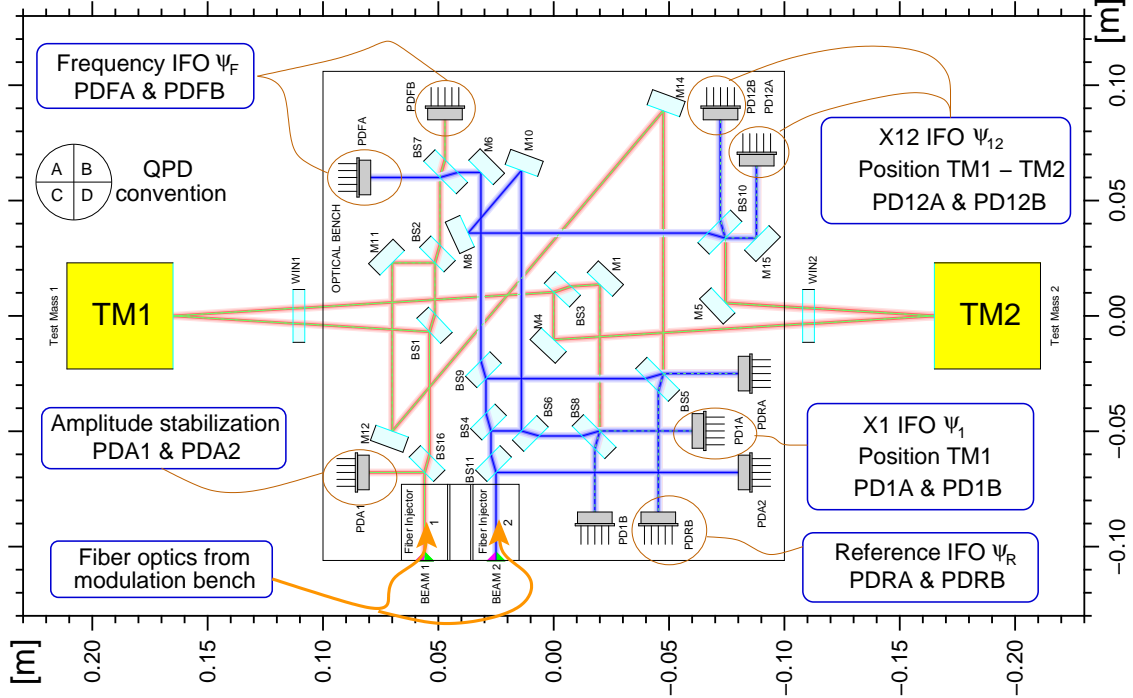
Figure 6.4: Schematic layout of the LISA Pathfinder optical bench [20].

6.3 Phase readout

The interference signal is measured with quadrant photodiodes which detect the heterodyne signal. They produce a photo current which is converted into a voltage using transimpedance amplifiers. The amplifer outputs are digitised and fed into the phasemeter. The phasemeter is used to calculate the phase by a *single bin discrete Fourier transform* for each quadrant of each photodiode [38]. 'Single bin' means that the Fourier transform of each segment of N samples is only performed at exactly the heterodyne frequency, f_{het}. The output of the photodiode pre-amplifier is transformed into a complex vector, F, which represents magnitude and phase of the f_{het}-component in the photo current. Additionally, the DC output voltage is averaged over N samples. With x_i being the amplifier output at time t_i, the DC and complex vector calculation can be performed using the following equations

$$DC = \frac{1}{N} \sum_{i=0}^{N-1} x_i(t_i), \tag{6.7}$$

$$\Re(F) = \frac{1}{N} \sum_{i=0}^{N-1} \cos\left(2\pi f_{\text{het}} \frac{i}{f_{\text{samp}}}\right) \cdot x_i(t_i), \tag{6.8}$$

$$\Im(F) = \frac{1}{N} \sum_{i=0}^{N-1} \sin\left(2\pi f_{\text{het}} \frac{i}{f_{\text{samp}}}\right) \cdot x_i(t_i). \tag{6.9}$$

(a) Reference interferometer

(b) Frequency interferometer

(c) X1 interferometer

(d) X12 interferometer
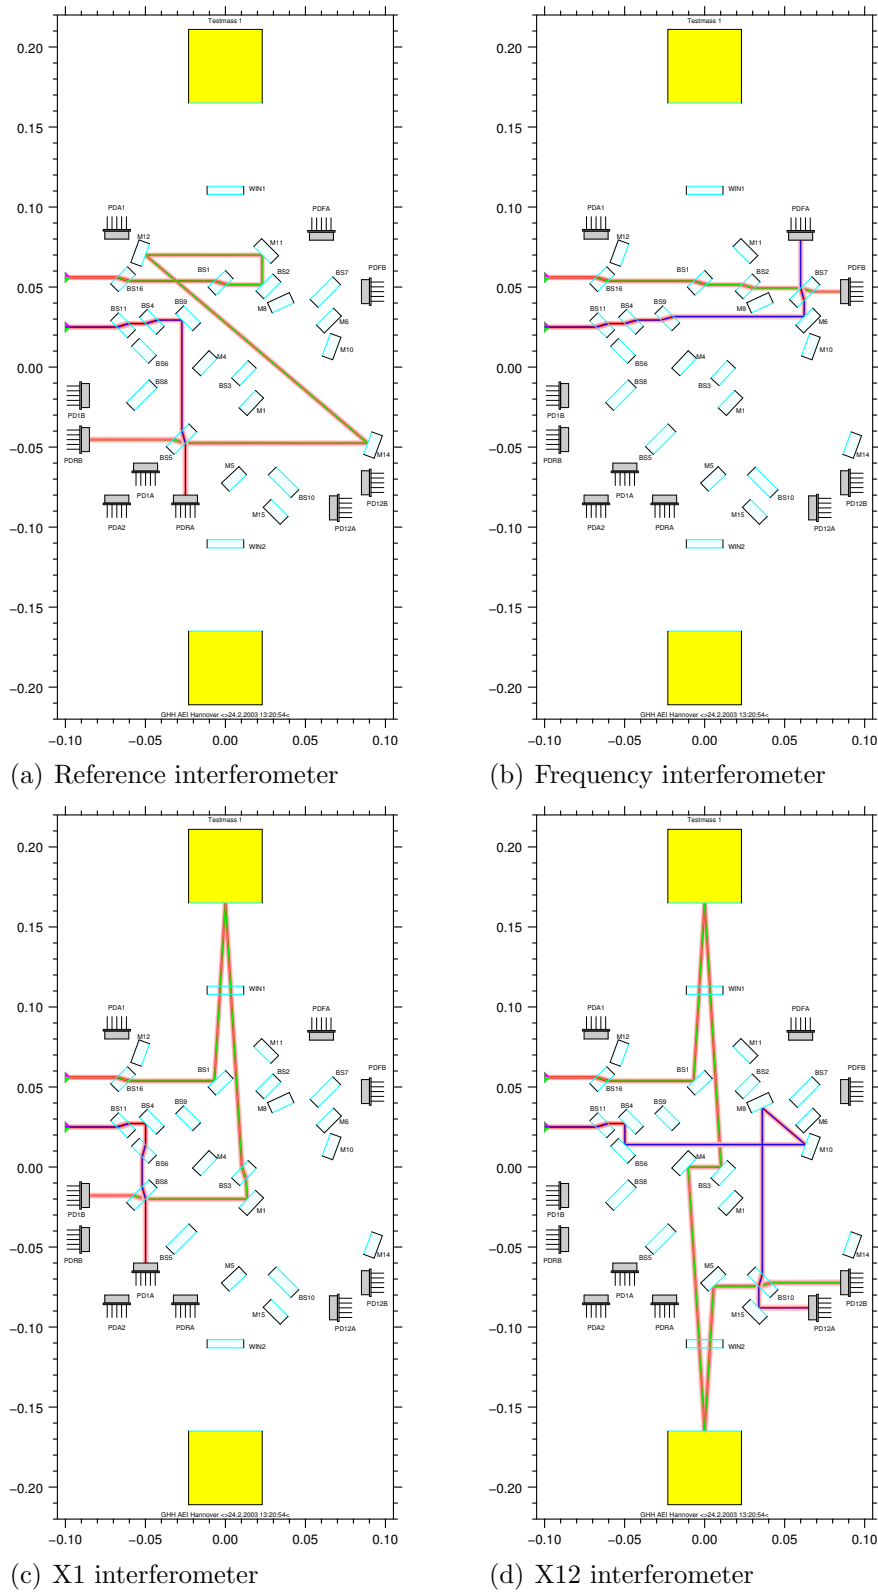
Figure 6.5: The four interferometers situated on the optical bench.

This calculation is done by a digital FPGA-based phasemeter. Equations (6.7) to (6.9) are simplified in order to illustrate the basic idea of the calculation. The implementation inside the phasemeter is slightly different and will be introduced in more detail in Section 8.1.1, where the scaling of the phasemeter output is discussed. The actual output data are three values d, y and z for each interferometer, quadrant photodiode and quadrant that correspond to $DC, \Re(F)$ and $\Im(F)$:

$$
\begin{aligned}
DC &\longrightarrow d_{i,\hat{j},k} & i &= 1, 2, \mathrm{R}, \mathrm{F} \\
\Re(F) &\longrightarrow y_{i,\hat{j},k} & \hat{j} &= \mathrm{A}, \mathrm{B}, \mathrm{C}, \mathrm{D} \\
\Im(F) &\longrightarrow z_{i,\hat{j},k} & k &= \mathrm{A}, \mathrm{B}.
\end{aligned}
\tag{6.10}
$$

The index i indicates the interferometer, $\hat{j}$ the quadrant and k the redundant quadrant photodiodes of the interferometer, from which the signal is derived. The 'hat' above the index j is irrelevant here and is described later on in Section 8.1.2.

A picture of the phasemeter breadboard developed in Hannover is shown in Figure 6.6. It has 20 input channels, which are usually sampled at $800\,\mathrm{kHz}$. The resolution of the Analogue-to-Digital Converters (ADCs) is $18\,\mathrm{bit}$. It has a parallel interface with a commercial PC that is used to program the phasemeter and also to read the output data. Software running on this computer does the data post-processing, which is described in detail in Chapter 8.
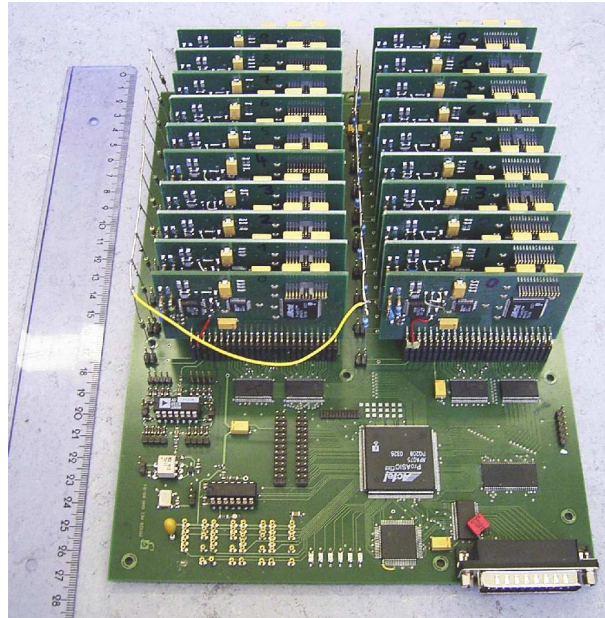
Figure 6.6: The FPGA-based phasemeter developed at AEI.

6.4 Differential Wavefront Sensing

For the test masses, not only the longitudinal motion is measured with the interferometer, but also the orientation in 2 degrees of freedom. In Chapter 8.1.6 a detailed description can be found, how this is realised. However, the basic idea is introduced here.

The angular motion of each test mass is measured in two different ways. One way is to estimate the beam position on a quadrant photodiode from the differential light power contribution at DC on the individual quadrants. This method is very robust and leads to reasonable results as long as at least a small fraction of the beam hits the photodiode.

A more sensitive method is Differential Wavefront Sensing (DWS) [49, 50]. Figure 6.7 shows two interfering beams on a quadrant photodiode. The reference beam is fixed and not moving. In the LISA Pathfinder setup, this is the beam that is confined to the ultra-stable optical bench. If the other (measurement) beam is tilted with respect to the reference beam, then its wavefront arrives earlier on the left half of the quadrant photodiode, whereas on the right half, it arrives later.

(a) Wavefront tilt   (b) Quadrant labelling
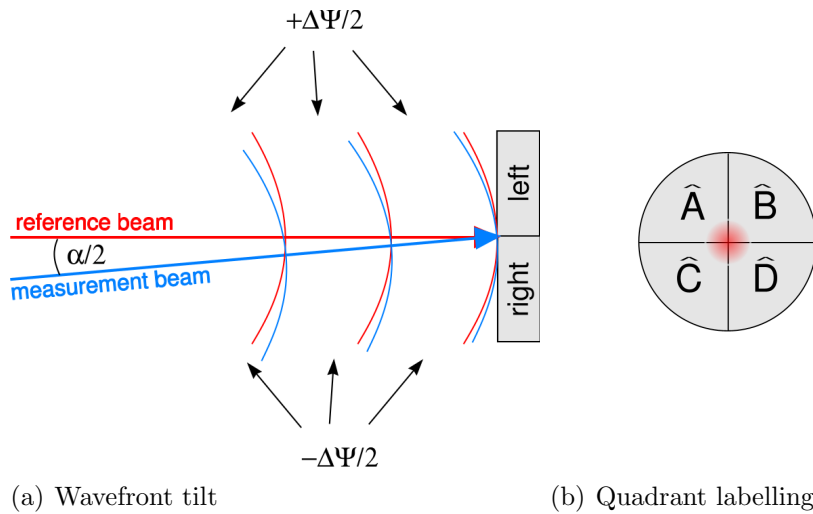
Figure 6.7: Two interfering beams on a quadrant photodiode.

The individual quadrants of the photodiode are read out as if they were single element photodiodes. The complex vectors for the different segments of the quadrant photodiode are added as follows:

$$\text{left} = A + C, \tag{6.11}$$
$$\text{right} = B + D, \tag{6.12}$$
$$\text{up} = A + B, \tag{6.13}$$
$$\text{down} = C + D. \tag{6.14}$$

Then the phase difference of the left and right half of the photodiode is calculated. A detailed description of this calculation can be found in Equation (8.36) and (8.37). The same is done for the upper and lower half respectively. The result is proportional to the angle of the measurement beam with respect to the reference beam. The scaling factor between physical test mass angle α and wavefront tilt, $\Delta\Psi$, is approximately given by

$$\frac{\Delta\Psi}{\alpha} = \frac{\lambda}{2\pi\omega_0},\tag{6.15}$$

where ω_0 is the beam waist [31]. The factor 2 accounts for the reflection of the beam at the test mass, which doubles the tilt of the reflected beam [31]. In case of the LISA Pathfinder Engineering Model optical bench, the coupling factor is of the order of $1/5000\,\mathrm{rad/rad_{DWS}}$ [18]. The index DWS indicates that the unit is not related to a physical angle, but to the spatial phase difference of the interference pattern.

The last calculation step is an $arg()$-function of the ratio of the complex vector elements. Its output is always between $-\pi$ and π, which limits the range for the DWS measurement contrary to the DC readout. However, it is more accurate and around zero it has the lowest noise level of about $10\,\mathrm{nrad}/\sqrt{\mathrm{Hz}}$ in the mHz range.

6.5 Stabilisation loops

In order to achieve the sensitivity necessary for LISA Pathfinder, four stabilisation loops are implemented in the interferometer. Two of them stabilise the laser intensity of both injected beams around the heterodyne frequency. Additionally, two other stabilisations are implemented: the laser frequency and optical pathlength difference (OPD) stabilisation loops.

The laser intensity stabilisation is of minor importance for the discussions in this thesis and therefore not further described. Investigations on this stabilisation can be found in [20].

6.5.1 OPD stabilisation

The reference interferometer from Figure 6.5(a) is intended to measure the phase fluctuations, φ_R, of the unstable modulation bench. By subtracting the measurement phase, φ_M, these fluctuation cancel under ideal conditions. It turns out that the cancellation is not perfect and limits the interferometer sensitivity at a level of about 300 pm, which can be reduced by stabilising the reference phase to the electronic phase of the reference oscillator [21, 39].

The reason was traced to electronic cross-talk between different circuit boards of the modulation electronics. Spurious sidebands on the AOM drivers produce additional signals on the laser light at the heterodyne frequency (see also Figure

2.3 in Section 2.2.1). These signals directly couple into the measurement of the phase with a random phase variation.

An illustration of the effect on the length measurement is shown in Figure 6.8. In order to show its influence, the effect is exaggerated. The 'real' pathlength change, in units of the laser wavelength, λ, is on the x-axis in this plot. Therefore, the ideal response of the interferometer should be the red curve. Due to the spurious signals produced by the sidebands, a periodic error is introduced into the measurement. It has two components, a 1f- and a 2f-component. The former comes from a beat of the carrier signal on one beam with a first order sideband on the same laser beam, whereas the latter comes from the beat of the carrier signal on one beam with a second order sideband on the other beam (see also Section 2.2.1). A detailed description of this effect can be found in [51, 39, 19, 18]. The 1f-component has one period and the 2f-component has two periods over one fringe.
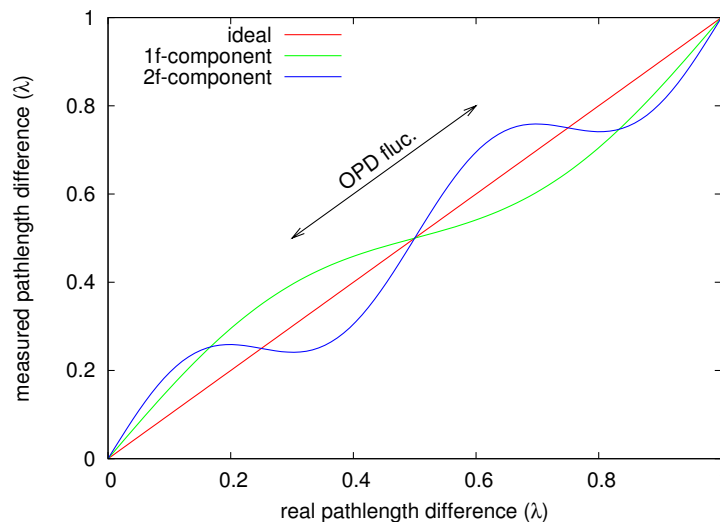
Figure 6.8: Effect of the non-linear error in the length measurement.

In all interferometers not only the length fluctuations on the stable interferometer baseplate are measured. Also the paths on the unstable modulation bench are included. Therefore, the measurements vary randomly over many fringes and the non-linear error converts into additional pathlength noise in the measurement band.

In order to overcome this problem, the reference phase, φ_R, is phase locked to the electronic reference signal, φ_{el}. This phase lock does not reduce the spurious signals, but converts the error to a small and constant offset on the measured phase difference.

6.5.2 Laser frequency stabilisation

A critical effect in highly sensitive interferometers is the coupling of laser frequency noise into the length measurement. It is present as soon as both interferometer

arms have unequal length. In the LTP interferometer a pathlength mismatch, ΔL, of about 1 cm is expected due to construction. The coupling of the frequency noise, $\tilde{\delta\nu}$, into the length measurement, $\tilde{\delta l}$, can be expressed as

$$\frac{\tilde{\delta l}}{\Delta L} = \frac{\tilde{\delta\nu}}{\nu}. \tag{6.16}$$

This leads to a required laser frequency noise of less than

$$\tilde{\delta\nu} = 28\frac{\text{kHz}}{\sqrt{\text{Hz}}} \tag{6.17}$$

per cm pathlength difference in the interferometers, which is about 100 times smaller at 3 mHz than the free-running laser frequency noise. Therefore, this noise must be suppressed. For this purpose, the phase measured in the frequency interferometer from Figure 6.5(b) is stabilised to the reference interferometer phase. The frequency interferometer has an intentional pathlength mismatch such that the coupling effect gets amplified. By closing the loop, the free-running laser frequency noise is suppressed.

Interferometer Characterisation

In this chapter the performance of the LTP interferometry with the present setup, described in Chapter 6, is presented. Instead of the real free-falling test masses, dummy mirrors were used that were mounted either on manual mechanical mounts or PZT-actuated mirrors as shown in Figure 7.1(a). The interferometer performance is characterised here by calculating the linear spectral density of the pathlength fluctuations of the X1 and X12 interferometer output signals in the LISA Pathfinder measurement band.

Fixed dummy mirrors were used to characterise the interferometer sensitivity itself. On-board LISA Pathfinder the residual test mass noise is much higher than the noise of the dummy mirrors, and therefore, especially the angular jitter can cause additional noise in the length measurement. In order to characterise the interferometer response to dynamic test masses, PZT-actuated mirrors were used, and here their long-term stability was investigated.

In addition, the angular test mass readout was characterised. The fibre injectors on the Engineering Model optical bench were identified to be unstable, and their noise contribution was subtracted from the angular measurement of the test masses.

Also, the differential photodiode alignment and its influence on the measurement is discussed.

7.1 Long-term stability of piezo actuated mirrors

The main interferometric readout is the monitoring of the differential longitudinal test mass motion. The Engineering Model optical bench, shown in Figure 7.1(b), is made of Zerodur.

For the performance tests of the interferometer, the dummy mirrors shown on

(a) Mounted and PZT actuated mirrors   (b) Engineering Model optical bench
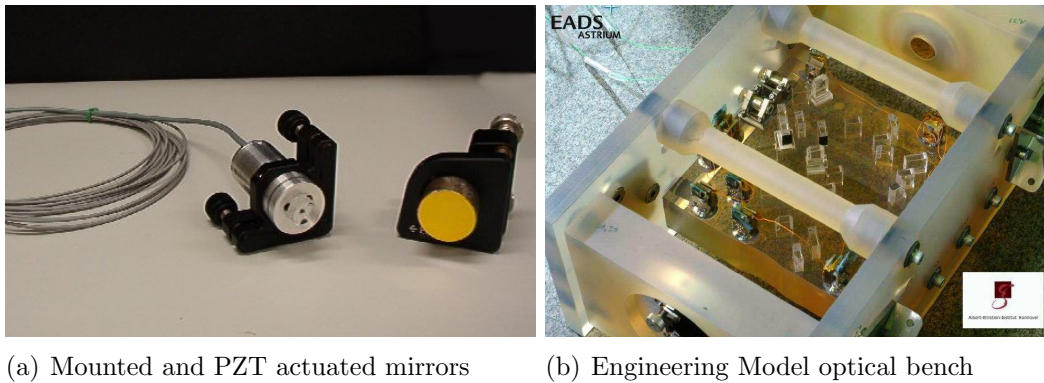
Figure 7.1: The Engineering Model of the LISA Pathfinder optical bench used with mounted mirrors or PZT actuated mirrors.

the right hand side in Figure 7.1(a) were used. They are made of Zerodur like the optical bench, but are inserted into manually adjustable mounts. These mounts were inserted into Zerodur sideplates in order to put them in the nominal test mass positions as shown in Figure 7.2.

(a) Zerodur dummy mirror   (b) PZT actuated mirror
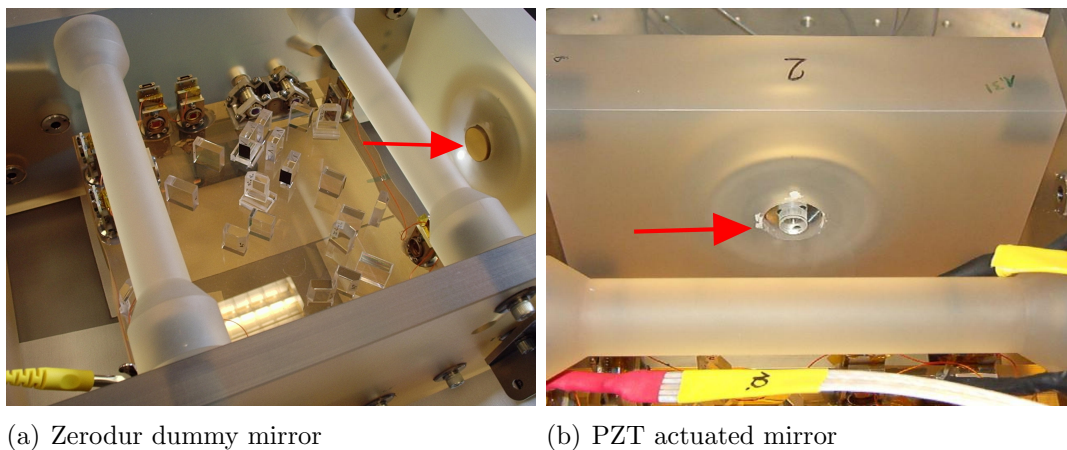
Figure 7.2: The test mirrors inserted in the Zerodur sideplates.

The interferometer performance with these mirrors is shown in Figure 7.3. It can also be seen in this figure that the mirrors mounted on piezo electric transducers lead to a sensitivity, which is slightly worse than with Zerodur dummy mirrors. It was very important to show this performance also with PZT actuated mirrors in order to test the interferometer with dynamic test mass behaviour. It is possible to reach the interferometer target, when additional angular noise is applied to the PZT actuated mirrors. The angular motion is simultaneously measured and can be subtracted from the longitudinal measurement [52]. Furthermore, the PZTs are used to test the on-orbit initial test mass alignment, which is described in Section 9.5.
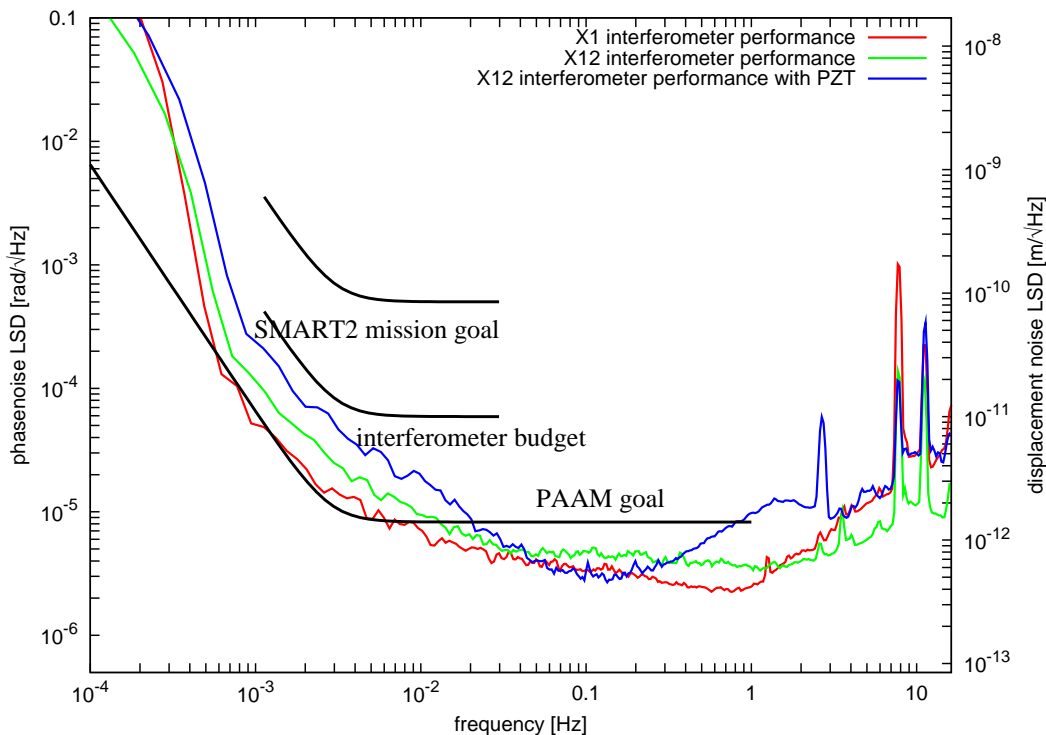
Figure 7.3: Interferometer performance with Zerodur mirrors and PZT actuated mirrors.

The long-term stability of the PZTs is also interesting for LISA. On the optical bench of LISA, it is planned to have PZT-actuated mirrors in the optical path. This is necessary, since a mechanism is needed that corrects for the small angle between the outgoing and incoming laser beam, which is varying periodically during the year. This device is called Point-Ahead Angle Mechanism (PAAM) [8]. The performance shown here sets an upper limit for the long-term stability of PZT actuated mirrors in the LISA measurement band. It was measured without a voltage applied to the PZT axes, which were, however, mounted in standard mounts made of aluminium. These mounts are probably the dominating noise source due to their thermal expansion.

The mirrors were glued on the PZTs using a standard two-component epoxy resin. Right after the glueing procedure, the PZT actuated mirrors were inserted into the interferometer and the longitudinal stability was monitored over seven weeks. Figure 7.4 shows a gradual improvement of the performance by more than one order of magnitude in the frequency range of interest. The improvement in the sensitivity is due to the hardening of the epoxy resin.

It can be concluded that it is possible to reach the interferometer sensitivity needed for LISA Pathfinder with PZT actuated mirrors. This is possible with and without small stabilised voltages applied to the PZT axes. Typical high voltage amplifiers, which were optimised for long-term stability and used for example for the alignment procedure (Section 9.5), lead to a performance, which is about one order of magnitude worse.
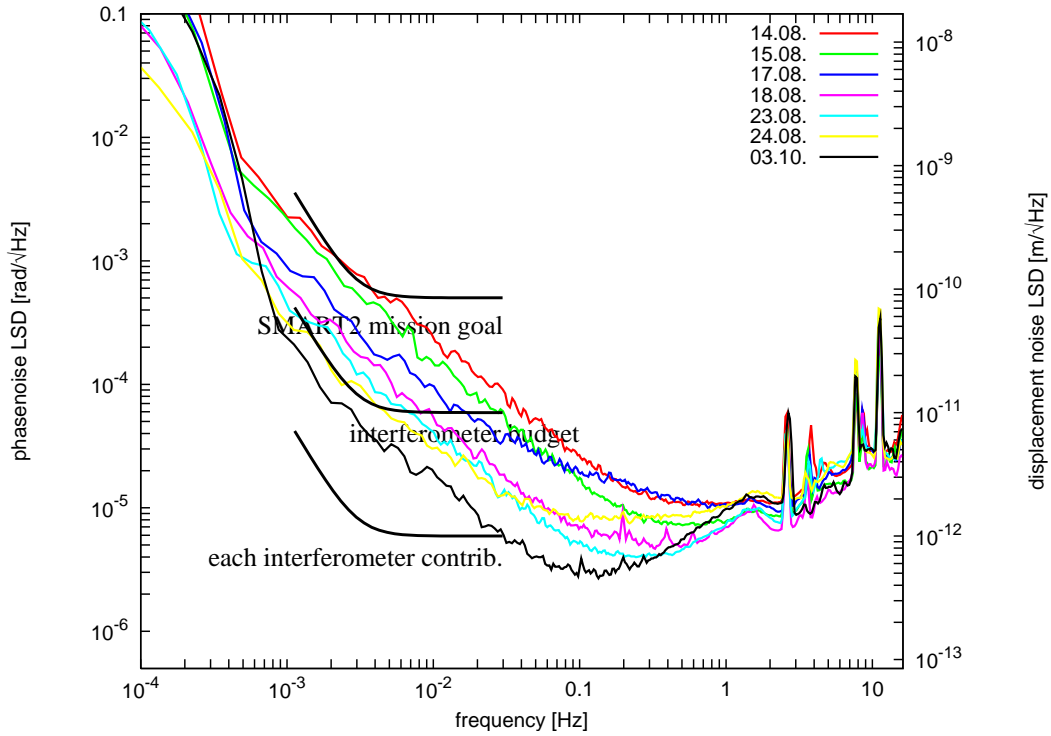
Figure 7.4: Improving PZT performance after glueing.

7.2 Angular sensitivity

The high sensitivity angular measurement of the test masses is achieved by Differential Wavefront Sensing (DWS). For each longitudinal measurement, the phase of the reference interferometer is subtracted from the measurement interferometer phase in order to cancel the common-mode noise of the unstable part of the interferometer [53]. This was not intended to be done for the angular measurement, since the reference beam, which is confined to the optical bench, is stable as long as the input fibre injectors are stable. Therefore, the DWS signals are calculated by subtracting the spatially differing phases on the photodiode surface of only the quadrant photodiode in the respective measurement interferometer.

The sensitivity directly derived from the DWS calculation is illustrated in Figure 7.5. It almost reaches the LISA Pathfinder specification. However, the angular noise in the measurement interferometers and the reference interferometers is correlated. This is an indication of a moving component in the common part of all interferometers on the optical bench. The angular noise in the reference interferometer can be subtracted, which leads to a performance shown in Figure 7.6. The sensitivity curves are shown for the corrected angles, $\varphi_{1,12}^{\mathrm{corr}}$ and $\eta_{1,12}^{\mathrm{corr}}$, that are
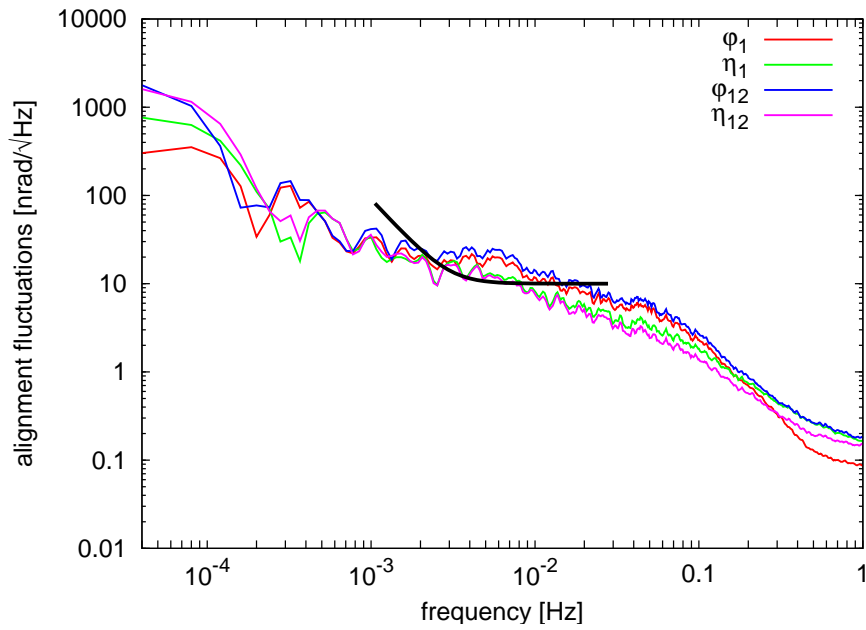
Figure 7.5: Test mass alignment fluctuations derived directly from the DWS signals of the measurement interferometers.
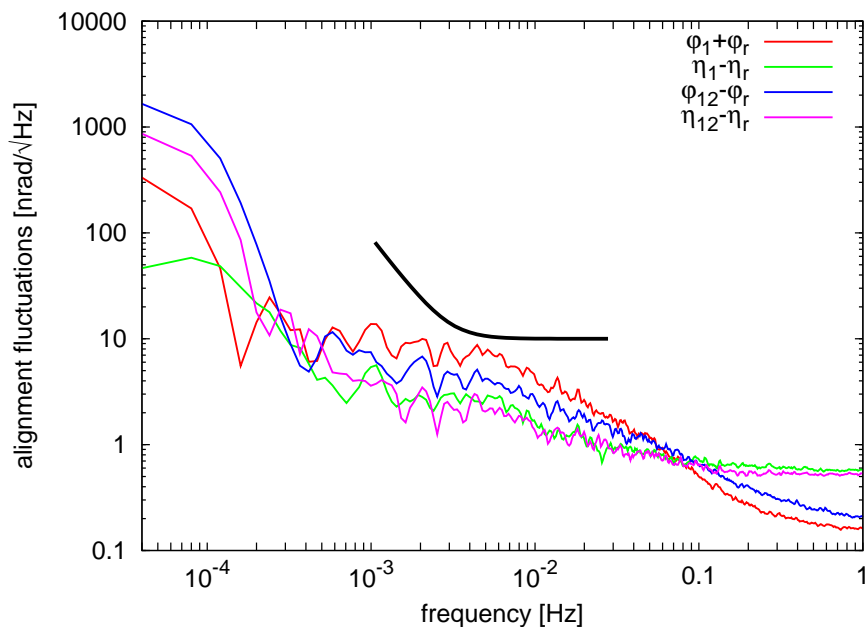
Figure 7.6: Test mass alignment fluctuations derived by subtraction of the angles of the reference interferometer.

calculated as

$$\varphi_1^{\text{corr}} \quad = \quad \varphi_1^{\text{meas}} + \varphi_R^{\text{meas}}, \tag{7.1}$$

$$\eta_1^{\text{corr}} \quad = \quad \eta_1^{\text{meas}} - \eta_R^{\text{meas}}, \tag{7.2}$$

$$\varphi_{12}^{\text{corr}} \quad = \quad \varphi_{12}^{\text{meas}} - \varphi_R^{\text{meas}}, \tag{7.3}$$

$$\eta_{12}^{\text{corr}} \quad = \quad \eta_{12}^{\text{meas}} - \eta_R^{\text{meas}}, \tag{7.4}$$

and derived either by addition or subtraction of the respective measured angles, $\varphi_{1,12,R}^{\text{meas}}$ and $\eta_{1,12,R}^{\text{meas}}$, of the reference and measurement interferometers.

In case of the Engineering Model optical bench, the subtraction of the angle of the reference interferometer improves the angular measurement, since the fibre injectors used are not very stable. They are commercially off-the-shelf fibre injectors produced by Schaefter&Kirchoff, which have a cylindrical shape. A special mount was made to bring them to the required beam height on the optical bench. The very sensitive fine adjustment of their height and vertical angle was done by putting small shims underneath the injectors [42]. The size and position of the shims was iteratively adapted to get an adequate beam alignment. A photograph of these injectors is shown in Figure 7.7. In order to fix them, they were clamped and glued into the custom made mount. For the Flight Model (FM) of the optical bench, a more advanced fibre injector design will be realised. It is a quasi-monolithic construction, which will be mounted to the optical bench using *hydroxide-catalysis bonding* in the same manner as for the other optical components [54, 37].
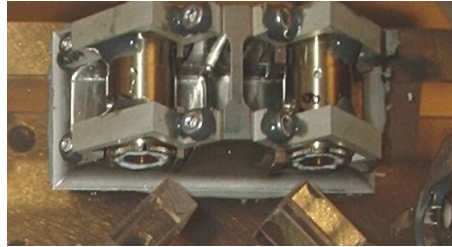
Figure 7.7: The fibre injectors on the Engineering Model optical bench.

In order to get a better alignment sensitivity, the measured horizontal angle of the X1 interferometer, φ_1, has to be added to the horizontal angle of the reference interferometer, φ_R, whereas all other angles (η_1, φ_{12} and η_{12}) have to be subtracted from the respective angles of the reference interferometer. This can be understood as follows:

When looking in the direction of beam propagation, the direction in which the beam is reflected changes with every reflection. This is equivalent to a change in the sign of the measured horizontal angle, φ. The vertical angle, η, is unaffected. By counting the number of reflections, the source of the jitter can be determined.

For the measurement shown in Figure 7.5, the set of 'B'-photodiodes from Figure 6.4 was used. Table 7.1 shows the number of reflections for each beam until it hits the photodiode.

Table 7.1: Number of reflections of both beams until hitting the photodiodes.

	PDRB	PD1B	PD12B
reference beam	2	2	5
measurement beam	4	3	6

The reference beam is the one that travels only on the stable optical bench, whereas the measurement beam path includes the test masses. An even number of reflections of a certain beam means a certain sign in the angular readout for the horizontal angle. An odd number of reflections produces the opposite sign. How can the sign in the correlation be understood?

The sign depends on the source of the jitter. One can either assume the reference beam or the measurement beam to move. Assuming that the former was varying in φ, it would produce the same sign in the angular readout of the reference and the X1 interferometer, but a different one in the X12 interferometer. This is obviously not the case, since the φ_R and φ_1 must be added in order to improve the sensitivity of the angular measurement. Therefore, the reference beam cannot be the source of the jitter.

Assuming on the other hand that the measurement beam was moving, the signs in the angular readout would be the same in the reference and the X12 interferometer, but not in the X1 interferometer. This is in agreement with the measurement, and therefore, the measurement beam injector is identified to produce the biggest angular noise.

7.3 Quadrant photodiode angular alignment

For the interferometer characterisation, the differential alignment of the quadrant photodiodes on the Engineering Model optical bench was investigated. Each quadrant photodiode has its own coordinate system, since the output channels refer to its physical orientation. Therefore, it is different from the LISA Pathfinder coordinate frame introduced in Section 5.2.

Each interferometer has a recombination beamsplitter, which has two output ports. In each of them a quadrant photodiode is situated in order to measure the interference signal. The quadrant photodiodes measure nominally redundant signals, but their slits may be rotated with respect to each other. This was tested by varying the angle of the PZT actuated test mass 1 dummy in two almost orthogonal angles (φ_1^{TM1} and η_1^{TM1}). Test mass 1 has its own coordinate system, which is defined on LISA Pathfinder by the electrostatic actuator for this test mass. In the laboratory setup it is defined by the PZT coordinate frame.

The variations of the TM 1 angle appear both in the X1 and X12 readout, since the X12 interferometer measures the common motion of both test masses. They are

shown in Figure 7.8. For all quadrant photodiodes in the X1 and X12 interferometer, the DWS signals were measured, while the angle of test mass 1 was modulated. The graph shows the direction of motion as observed by each photodiode in a vector representation.
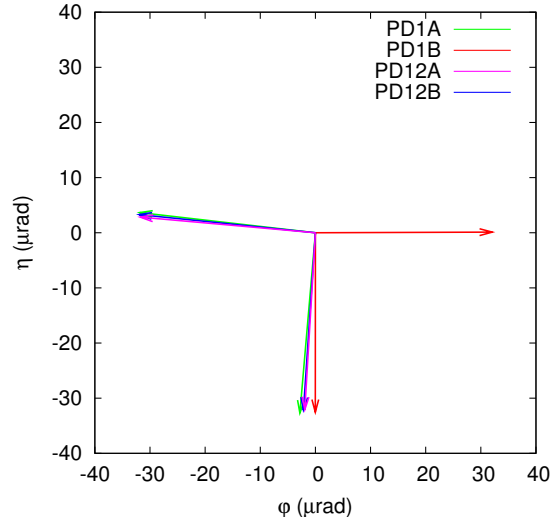
Figure 7.8: Measured tilt of the photodiodes with respect to each other.

The DWS signals are scaled with the well known scaling factors [48]. The plot shows the angle of the test mass in each frame defined by each quadrant photodiode. The results are summarised in Table 7.2. In the table, *roll* of the photodiodes means the rotation of the photodiodes around the axis orthogonal to their surface. Quadrant photodiode PD1B is taken as reference. The X12 axes differ by about 1 degree. The X12 differential alignment is about 1 degree. This is different for the X1 interferometer. PD1A differs by about 5 degrees from PD1B.

Table 7.2: Differential roll of the quadrant photodiodes in the measurement interferometers in degrees.

[deg]	PD1A	PD1B	PD12A	PD12B
φ	180+6.7	0	0.6	1.4
η	4.9	0	1.1	1.5

The roll of the quadrant photodiodes was aligned visually during the construction of the optical bench. Therefore, a deviation of about 5 degrees is a little high, but not totally unexpected. For the Flight Model of the interferometer, it is hoped to be at most of the same order of magnitude as long as no special alignment procedures are forseen for the compensation of the differential roll of the photodiodes.

The angular measurements for test mass 1 and test mass 2 can be assumed to be almost independent. The DFACS has an orthogonalisation matrix that can be used to decouple the actuation signals and error signals: e.g. the error signal $\varphi_{12}^{\mathrm{err}}$ can be experimentally calculated by a linear combination of the measured angles $\varphi_1^{\mathrm{meas}}$, η_1^{meas}, $\varphi_{12}^{\mathrm{meas}}$ and $\eta_{12}^{\mathrm{meas}}$, and the weighting factors can be measured on-board LISA Pathfinder (see also Section 8.1.7). If the X1 and X12 interferometer photodiodes are rotated with respect to each other, the orthogonalisation matrix changes, but it still delivers 'clean' error signals.

The differential rotation of the nominal/redundant quadrant photodiode pair cannot be compensated. In the averaging process as described in Section 8.1.4, the signals of redundant quadrant pairs (with differential rotation) are averaged before the angles are calculated. The differential rotation violates the redundancy and leads to an error in the calculated angle. However, the effect is assumed to be small, since in science operation of LISA Pathfinder, the test masses are nominally controlled to the DWS zero, where the influence is small.

Definition of the on-board data processing

In this chapter the processing of the interferometric data in the LTP on-board software of the Data Management Unit (DMU) is described. In addition to the phasemeter back-end processing, the DMU has other tasks such as controlling, diagnostics and sensing of other subsystems. Here only the phasemeter processing is considered. The main functionality of the DMU is the calculation of the longitudinal and angular signals of the test masses, which has been presented in Chapter 6. In contrast to the laboratory setup, the DMU has additional functionalities like averaging processes, failure detection and an asynchronous link to the On-board Computer (OBC), which calculates the feedback signals for the Drag-Free and Attitude Control System (DFACS) [55].

An on-orbit test mass alignment procedure was defined, demonstrated and already published [56]. The work presented here uses a more complex alignment procedure with wide range piezo actuated mirrors and an asynchronous link in the feedback loop.

Many aspects of the data processing on-board LISA Pathfinder are also interesting for LISA. Both missions use quadrant photodiodes for the main measurement. Therefore, many DMU functionalities can be copied from LISA Pathfinder. The phasemeter used in LISA is faster compared to the LISA Pathfinder phasemeter. However, an output data rate of the LISA science data is not necessarily higher than in LISA Pathfinder. The calculation of longitudinal and alignment signals can be performed as described in this chapter. The basic idea of the failure detection scheme, described in Section 8.4, can be adapted, and also downsampling as introduced here will have comparable constraints. Hopefully, the asynchronous data transfer, discussed in Section 8.2, and all problems that come with it are not of relevance for LISA.

8.1 Software definition

The phasemeter front-end calculates a complex vector $(y_{i,j,k}, z_{i,j,k})$ and a DC amplitude $d_{i,j,k}$ for each quadrant of each photodiode on the optical bench as described in Section 6.3 at a rate of $100\,\mathrm{Hz}$. This is the scientific input data of the Data Management Unit. The task of the DMU is to extract from those the main longitudinal measurements, x_1 and x_{12}, of the test masses, but also two angles φ and η for each test mass either derived from the Differential Wavefront Sensing (DWS) or from the DC readout of the photodiodes. The results are transferred to the On-board Computer (OBC) as illustrated in Figure 8.1 [38].
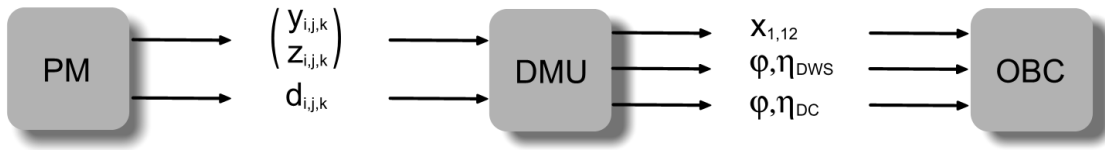
Figure 8.1: Input and output data of the DMU.

8.1.1 Scaling the phasemeter output

The phasemeter data arrives as integer numbers and has to be converted to floating point numbers for the DMU internal data processing [57]. Therefore, the scaling inside the DMU is done as the first step in order to get the following variables:

$$DC_{i,\hat{j},k} = c_{i,\hat{j},k}^{\mathrm{DC}} \cdot d_{i,\hat{j},k} - c_{i,\hat{j},k}^{\mathrm{offset}} \tag{8.1}$$

$$\Re\left(F_{i,\hat{j},k}\right) = c_{i,\hat{j},k}^{\mathrm{Ry}} \cdot y_{i,\hat{j},k} + c_{i,\hat{j},k}^{\mathrm{Rz}} \cdot z_{i,\hat{j},k} + c_{i,\hat{j},k}^{\mathrm{Rd}} \cdot d_{i,\hat{j},k} \tag{8.2}$$

$$\Im\left(F_{i,\hat{j},k}\right) = c_{i,\hat{j},k}^{\mathrm{Iy}} \cdot y_{i,\hat{j},k} + c_{i,\hat{j},k}^{\mathrm{Iz}} \cdot z_{i,\hat{j},k} + c_{i,\hat{j},k}^{\mathrm{Id}} \cdot d_{i,\hat{j},k}. \tag{8.3}$$

The phasemeter output is a complex phase vector and can, after scaling, be written as

$$F_{i,\hat{j},k} = \Re\left(F_{i,\hat{j},k}\right) + \mathrm{i} \cdot \Im\left(F_{i,\hat{j},k}\right). \tag{8.4}$$

Additionally, the DC photodiode output $DC_{i,\hat{j},k}$ is scaled. The coefficients c can be chosen such that they are used for:

- Normalisation such that 1.0 represents the nominal value after both beams are switched on and the test masses are aligned.

- Removal of 180 degrees phase difference between each two redundant photodiodes.

- DC compensation: Since the phasemeter performs unsigned integer operations, a DC offset has to be removed from the sine and cosine components, y and z.

- Amplitude calibration.

- Phase calibration.

- Photodiode dark current compensation.

Determination of the scaling coefficients is quite complicated, since the simple form of Equations (6.7) to (6.9) is modified in the phasemeter. The ratio of the heterodyne frequency and the sampling frequency is related to the frequency bin B of the Fourier transform and the length N of the time series, to which the Fourier transform is applied:

$$\frac{f_{\text{het}}}{f_{\text{samp}}} = \frac{B}{N}. \tag{8.5}$$

Using this identity, the phasemeter output can be written as:

$$d \quad = \quad \sum_{n=0}^{N-1} x_n(t_n), \tag{8.6}$$

$$y \quad = \quad \sum_{n=0}^{N-1} \text{round}\left[A\left(1 + w(n)\cos\left(2\pi \cdot B\frac{n}{N}\right)\right)\right] x_n(t_n), \tag{8.7}$$

$$z \quad = \quad \sum_{n=0}^{N-1} \text{round}\left[A\left(1 + w(n)\sin\left(2\pi \cdot B\frac{n}{N}\right)\right)\right] x_n(t_n) \tag{8.8}$$

with the modified Fourier coefficients

$$c_n^* \quad = \quad \text{round}\left[A\left(1 + w(n)\cos\left(2\pi \cdot B\frac{n}{N}\right)\right)\right], \tag{8.9}$$

$$s_n^* \quad = \quad \text{round}\left[A\left(1 + w(n)\sin\left(2\pi \cdot B\frac{n}{N}\right)\right)\right]. \tag{8.10}$$

The input data $x_n(t_n)$ is already digitised and has a range R corresponding to the ADC that is used to convert the photodiode amplifier output voltage into an unsigned integer number. In case of the breadboard phasemeter, it has 18 bits. An additional factor of w(n) can be inserted in order to apply a time-domain window function (such as, e.g., Hanning, Blackman-Harris, flat-top). Not using a window means using a 'rectangular' window with all $w(n) = 1$.

For efficiency, the sine and cosine coefficients in the lookup table are directly multiplied with the window function before being uploaded to the phasemeter front-end. Therefore, the window function must be normalised such that its maximum value is unity in order to avoid an overflow in the integer arithmetic.

The coefficients which now would have a range between -1 and +1 are shifted by +1 so that the final coefficients, s_n^* and c_n^*, to be uploaded are always positive and scaled such that they fit in the range $0...2^{b-1}$. This requires the factor A to be

$$A = 2^{b-1} - 1. \tag{8.11}$$

A corresponds to the bit-depth, b, of the pre-computed sine and cosine coefficients which is chosen comparable to the ADC resolution. The values from Equation (8.7) and (8.8) are rounded to the next integer for use in the integer multiplier of the FPGA front-end.

The DC output of the phasemeter, d, is scaled such that '1' corresponds to a maximum input voltage. This means that it must be divided by the FFT length N and the ADC range R. Therefore, the scaling factors c^{DC} and c^{offset} are

$$
\begin{aligned}
c^{\mathrm{DC}} &= \frac{1}{RN}, & (8.12)\\
c^{\mathrm{offset}} &= 0. & (8.13)
\end{aligned}
$$

An offset, due to imperfections of the photodiode amplifier or the ADC, can be subtracted, but is here assumed to be zero.

The scaling of the real and imaginary part of the complex vector is more complicated. Most reasonable is a scaling of the vector length to approximately unity magnitude. The input signal can generally be written as

$$
x_n(t_n) = R \cdot (x_n^0 + o), \tag{8.14}
$$

where x_n^0 is a pure oscillation around zero at the heterodyne frequency and o an arbitrary offset. It relates to the phasemeter output data, d, via

$$
d = R \sum_{n=0}^{N-1} o \;\curvearrowright\; o = \frac{d}{RN}. \tag{8.15}
$$

When applying the generalised input signal from Equation (8.14) to Equation (8.7), it can be rewritten as:

$$
\sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right] \cdot Rx_n^0 = y - \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right] \cdot \frac{d}{RN} \quad (8.16)
$$

The left hand side of this equation is exactly the desired output of the Fourier transform. The first term, $(1+..)$, can be neglected, since it does not change the result of the summation.

For the normalisation, the left hand side must be divided by the same sum with a signal in the cosine quadrature:

$$
\begin{aligned}
\Im(F) &= \frac{R \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right] \cdot x_n^0}{R \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right] \cdot c_n} \\
&= \frac{1}{R \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right] \cdot c_n} \cdot y \\
&\quad - \frac{\sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right]}{RN \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1 + w(n)c_n\right)\right] \cdot c_n} \cdot d. & (8.17)
\end{aligned}
$$

From this equation one can read the coefficients, c, for $\Re(F)$. For $\Im(F)$ they can be calculated accordingly. The complete set of coefficients is

$$c^{\mathrm{Ry}} = \frac{1}{R \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1+w(n)c_n\right)\right] \cdot c_n}, \tag{8.18}$$

$$c^{\mathrm{Rd}} = \frac{\sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1+w(n)c_n\right)\right]}{RN \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1+w(n)c_n\right)\right] \cdot c_n}, \tag{8.19}$$

$$c^{\mathrm{Iz}} = \frac{1}{R \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1+w(n)s_n\right)\right] \cdot s_n}, \tag{8.20}$$

$$c^{\mathrm{Id}} = \frac{\sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1+w(n)s_n\right)\right]}{RN \sum_{n=0}^{N-1} \mathrm{round}\left[A\left(1+w(n)s_n\right)\right] \cdot s_n}. \tag{8.21}$$

The coefficients $c^{\mathrm{Ry}}_{i,\hat{j},k}$ and $c^{\mathrm{Iz}}_{i,\hat{j},k}$ can be used to multiply each channel individually by a phase factor of $e^{\mathrm{i}\varphi}$, e.g. to compensate different static phase shifts in the analogue front-ends. However, the default values are zero:

$$c^{\mathrm{Ry}}_{i,\hat{j},k} = 0, \tag{8.22}$$

$$c^{\mathrm{Iz}}_{i,\hat{j},k} = 0. \tag{8.23}$$

For all coefficients that scale the complex vector, the sign depends on which of the two redundant photodiodes is used. In the laboratory setup the convention was chosen to have negative $c^{\mathrm{Rz}}_{i,\hat{j},k}$ and $c^{\mathrm{Iy}}_{i,\hat{j},k}$ for the photodiode set A. This makes $c^{\mathrm{Rd}}_{i,\hat{j},k}$ and $c^{\mathrm{Id}}_{i,\hat{j},k}$ positive. For the set of B photodiodes, the sign changes for each of these coefficients. Thus, the 180 degrees phase difference of the two beamsplitter outputs is taken into account.

8.1.2 Channel labelling

Before the phasemeter channels are processed, they have to be relabelled. The reason is that due to an extra reflection, a moving beam ends up on different halves on the different quadrant photodiodes on the optical bench. This is illustrated in Figure 8.2(a): The blue beam moves to the left half (quadrants A/C) on the upper photodiode and to the right half (quadrants B/D) on the right photodiode if the incoming beam rotates counter-clockwise. The initial labelling of the photodiode quadrants refers to their physical orientation on the hardware device as seen from the direction of the incoming light, which is illustrated in Figure 8.2(b) [58].

In order to simplify the processing of the redundant information, coming from the two redundant photodiode sets on the optical bench, the quadrants are relabelled before they are processed. The non-ambiguous convention was chosen that the reference beam – the beam that does not leave the optical bench – moves per definition to quadrant A and C of each photodiode if it rotates in positive φ-direction as shown in Figure 8.3. The reference beam in the picture rotates to the

(a) Beam angle on photodiode pair.

(b) Physical quadrant labelling.
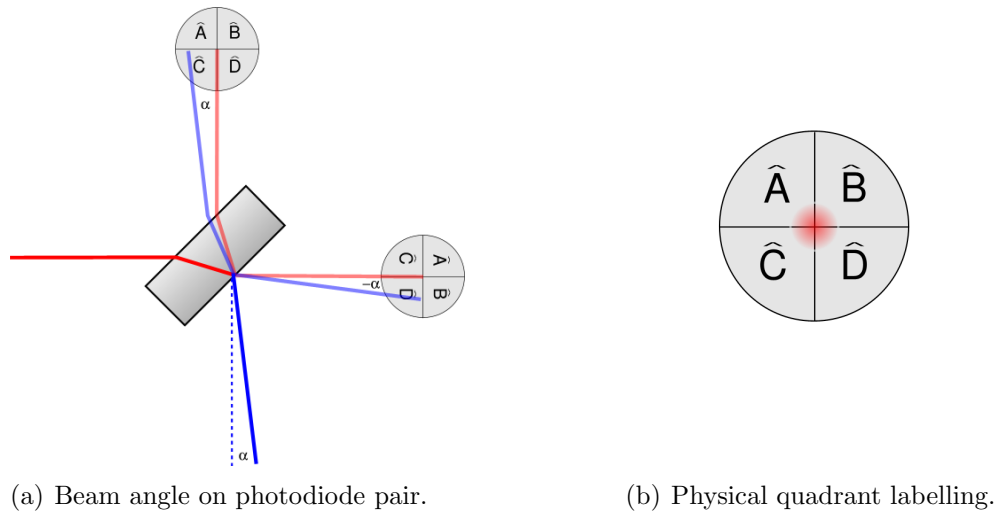
Figure 8.2: Labelling of redundant photodiode quadrants.
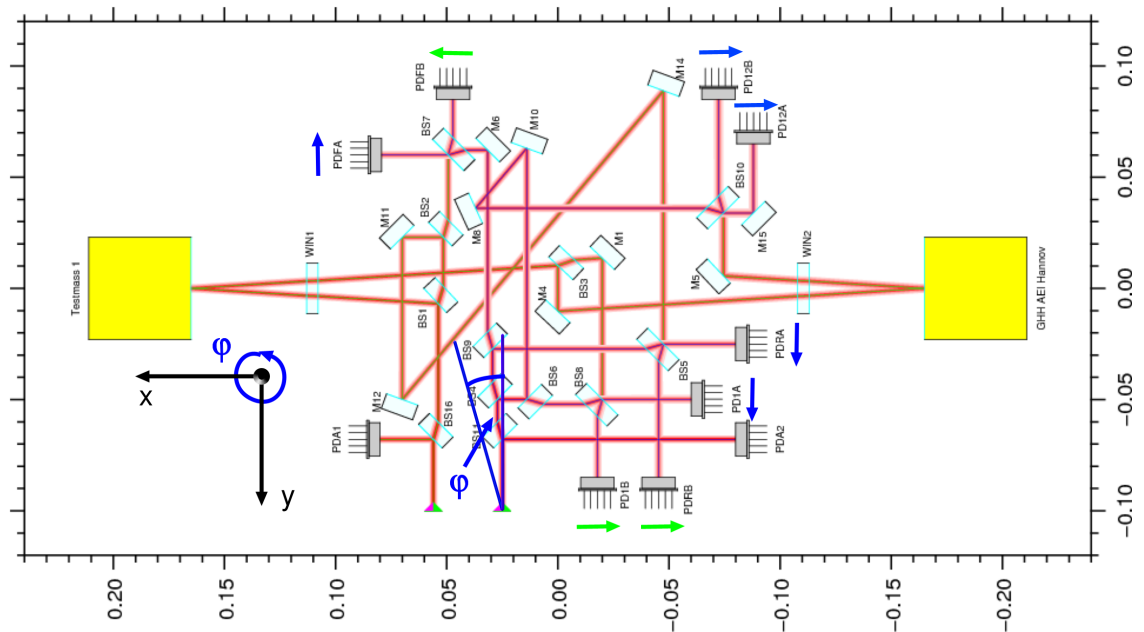
Figure 8.3: Illustration of the quadrant labelling convention.

left. Photodiodes where the beam moves to the left half are indicated by a green arrow. For photodiodes where the beam moves to the right half, a blue arrow is shown. This means in case of the Engineering Model optical bench: Quadrant A has to be exchanged with quadrant B, and quadrant C has to be exchanged with quadrant D for the photodiodes PD1A, PD12A, PD12B, PDFA and PDRA. For all other quadrant photodiodes, the relabelled quadrants correspond to the original label. In order to avoid confusion, a 'hat' is above the quadrant label, j, as long as it refers to the physical position. If no 'hat' is used, the quadrant is already relabelled.

8.1.3 Redundancy and averaging

One of the main differences in the realisation of experiments in space compared to laboratory experiments is the use of redundant systems. As far as possible, independent redundant systems are installed that can be used in case of failure of the nominal device. On-board LISA Pathfinder, a loss of the test masses, the inertial sensors or the optical bench would lead to a loss of the experiment, or at least to a strong performance degradation, since for these elements it was not possible to implement redundancy. However, at least the readout of each interferometer is redundant in the form of a set of two photodiodes behind each recombination beamsplitter and two redundant phasemeters. One set of photodiodes – namely all PDxA photodiodes, where 'x' can be any of the interferometers – is connected to phasemeter A, and the other set, PDxB, is connected to phasemeter B respectively. In normal operation both sets of photodiodes are used such that, as long as two redundant channels are available and failure-free, both are averaged for further calculations.

In principle two scenarios are possible for the averaging process:

1. The DMU output signals are computed for each individual photodiode, and afterwards the longitudinal and the 2×2 angular output signals (derived from DWS or DC readout) are averaged as illustrated in Figure 8.4(a).

2. The information of two redundant quadrants is averaged and afterwards handled as a single quadrant as shown in Figure 8.4(b).

Both options have advantages and disadvantages. In recent experiments the redundant photodiodes were not used mainly due to the limited amount of input channels of the breadboard phasemeter [19, 20]. It is known that the performance goal for the LISA Pathfinder interferometer can be reached by using either the nominal or redundant set of photodiodes alone. Therefore, this is also true for their average.

The second option from Figure 8.4(b) was chosen to be implemented for different reasons: It is more robust in case of loss of individual quadrants. In particular, reasonable angular signals can only be computed with information from 4 independent quadrants (A,B,C and D) of the same photodiode. In the case of a temporary
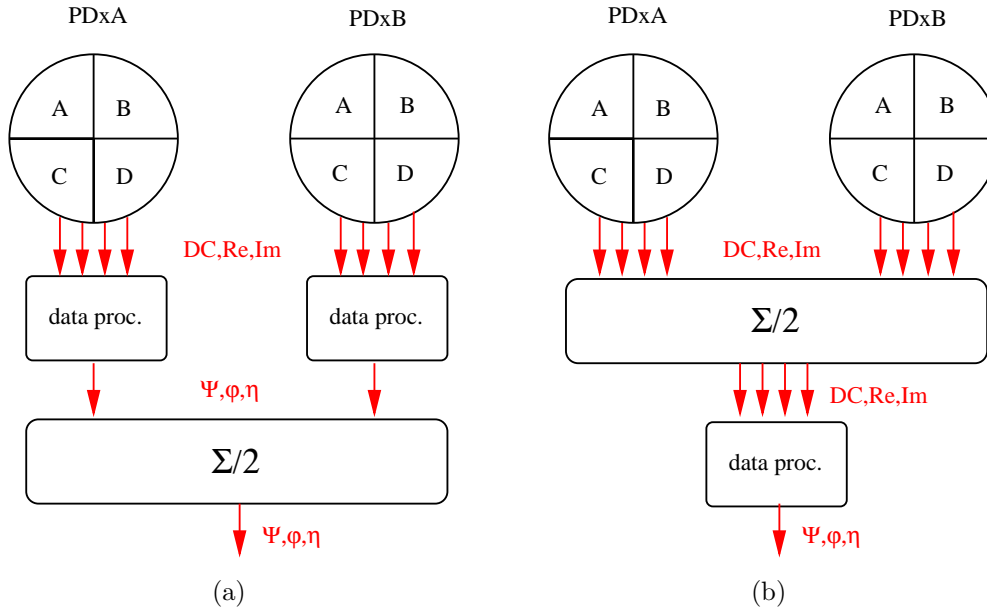
Figure 8.4: Scenarios for the averaging process inside the DMU.

or a continuous failure of a quadrant in PDxA and also one in PDxB in the same interferometer, those signals cannot be computed anymore. By using the chosen option, many more combinations of channel failures can be tolerated.

One recently discovered and important noise source in heterodyne interferometers are pseudo signals produced by sidebands at multiples of the heterodyne frequency on the laser beams [18, 39]. This is described in detail in Section 3.4.1 and used there to correct for the noise produced by spurious signals in the backside fibre link interferometer. The formalism for the phase calculation used here is different to the subtraction method introduced in that section, but the result is exactly the same. The combination of channel relabelling and averaging of the complex vectors of redundant channels has the same effect. This was not intended when the software was defined, but is a fortunate side effect. However, the influence of ghost beams and straylight in the LISA Pathfinder interferometer has no influence on the performance level achieved with the laboratory setup.

8.1.4 Calculating the averaged signals

The DC signals and the complex vectors for each quadrant are calculated as follows:

$$
DC_{i,j} = \begin{cases}
\frac{1}{2}\left(DC_{i,j,\mathrm{A}} + DC_{i,j,\mathrm{B}}\right) & \text{A and B failure-free} \\
DC_{i,j,\mathrm{A}} & \text{A failure-free} \\
DC_{i,j,\mathrm{B}} & \text{B failure-free} \\
\text{non-computable} & \text{A and B have failure}
\end{cases}
\tag{8.24}
$$

$$F_{i,j} = \begin{cases} \dfrac{1}{2}\left(F_{i,j,\mathrm{A}} + F_{i,j,\mathrm{B}}\right) & \text{A and B failure-free} \\ F_{i,j,\mathrm{A}} & \text{A failure-free} \\ F_{i,j,\mathrm{B}} & \text{B failure-free} \\ \text{non-computable} & \text{A and B have failure} \end{cases} . \tag{8.25}$$

Usually a channel and its redundant counterpart are averaged and used as input for the following calculations. If a failure in one of the channels appears, it will be excluded from further processing, and only in the case of errors in both the channel and its redundant counterpart the complex vector and the DC signal is non-computable. A description of how errors are detected and handled can be found in Section 8.4.

After this processing step, the interferometric output is calculated as if just one photodiode was used in each interferometer. This means that the nominal and redundant quadrant are already averaged, and further processing handles the averaged signals as if they were derived from only one quadrant. The following processing steps were already described and published in [19, 20]. However, for completeness these calculations will also be described here.

8.1.5 Calculating the longitudinal results

For the longitudinal signals, information from all 4 quadrants of one photodiode is summed (or averaged). The next step is to calculate the DC-signal, Σ_i, and a complex vector, F_i, for each quadrant photodiode pair as

$$\Sigma_i = \sum_j DC_{i,j} \tag{8.26}$$

$$F_i = \sum_j F_{i,j}. \tag{8.27}$$

This leads to the phase, ϕ_i, and the length of the complex vector, a_i,

$$\phi_i = \arg\left(F_i\right) \tag{8.28}$$

$$a_i = \left|F_i\right|. \tag{8.29}$$

The phase depends linearly on the pathlength change (see Equation (6.3)), and the vector length represents the amplitude of the heterodyne signal. Here it was chosen to add the complex vectors on each quadrant first and then calculate the phase, instead of calculating the phases first and do the averaging afterwards. This has the advantage that quadrants with bigger heterodyne amplitude contribute more to the resulting phase, which corresponds to a weighted averaging of the phase and improves the signal-to-noise ratio.

The phase differences of the X1 and X12 interferometer and the reference interferometer are proportional to the pathlength changes introduced by the movement

of the test masses, but all results from Equation (8.28) are within 2π since they are the output of the arg()-function. 2π corresponds to one interferometer fringe. This means that the interferometer test mass position readout is a relative displacement measurement. The interferometer cannot distinguish between a certain test mass position and the test mass shifted such that the optical pathlength differs by an integer multiple of the wavelength.

In order to track the test mass position over many fringes, a phase-tracking algorithm is applied. The algorithm compares the actual with the recent phase sample and assumes that phase changes between them are small. If their difference is larger than π an integer multiple of 2π is added or subtracted to the new phase sample such that the interferometer can follow a slow continuous drift of the test mass over many laser wavelength. Finally the subtraction of the reference phase is applied:

$$\Psi_1 \;=\; \mathrm{PT}\left(\arg\frac{F_1}{F_\mathrm{R}}\right) \tag{8.30}$$

$$\Psi_{12} \;=\; \mathrm{PT}\left(\arg\frac{F_{12}}{F_\mathrm{R}}\right) \tag{8.31}$$

$$\Psi_F \;=\; \mathrm{PT}\left(\arg\frac{F_\mathrm{F}}{F_\mathrm{R}}\right) \tag{8.32}$$

$$\Psi_R \;=\; \mathrm{PT}\left(\arg F_R\right). \tag{8.33}$$

The phase-tracking algorithm which is used in the test setup can be found in [59].

8.1.6 Calculating the angular results

As already mentioned, not just the longitudinal interferometric measurement is important, but also the angular measurements in order to provide error signals for the DFACS. The angles are either derived from the DC signals or from the DWS signals. The former estimate the centre of the beam power on the quadrant photodiode. The DWS signals on the other hand interferometrically measure the tilt of the two interfering wavefronts with respect to each other (see also Section 6.4). The two types of signals are calculated as follows:

$$DC_i^\phi \;=\; \frac{\Sigma_i^\mathrm{left} - \Sigma_i^\mathrm{right}}{\Sigma_i} = \frac{(DC_{i,\mathrm{A}} + DC_{i,\mathrm{C}}) - (DC_{i,\mathrm{B}} + DC_{i,\mathrm{D}})}{\Sigma_i} \tag{8.34}$$

$$DC_i^\eta \;=\; \frac{\Sigma_i^\mathrm{up} - \Sigma_i^\mathrm{down}}{\Sigma_i} = \frac{(DC_{i,\mathrm{A}} + DC_{i,\mathrm{B}}) - (DC_{i,\mathrm{C}} + DC_{i,\mathrm{D}})}{\Sigma_i} \tag{8.35}$$

$$DWS_i^\phi \;=\; \arg\left(\frac{F_i^\mathrm{left}}{F_i^\mathrm{right}}\right) = \arg\left(\frac{F_{i,\mathrm{A}} + F_{i,\mathrm{C}}}{F_{i,\mathrm{B}} + F_{i,\mathrm{D}}}\right) \tag{8.36}$$

$$DWS_i^\eta \;=\; \arg\left(\frac{F_i^\mathrm{up}}{F_i^\mathrm{down}}\right) = \arg\left(\frac{F_{i,\mathrm{A}} + F_{i,\mathrm{B}}}{F_{i,\mathrm{C}} + F_{i,\mathrm{D}}}\right). \tag{8.37}$$

Here 'left' indicates the sum of the DC signals or complex vectors of quadrants A and C and 'right' the sum of quadrants B and D respectively. In the same manner 'up' means the sum of quadrants A and B and 'down' of quadrants C and D.

For the calculation of the angular signals, the relabelling of the quadrants is most important. Only if the relabelling is performed before the above calculations, the nominal and redundant diodes produce the same signal.

The Differential Wavefront Sensing signals are very sensitive (see Equation (6.15)). However, the highest sensitivity is around zero. The noise performance degrades with increasing static misalignment. Furthermore, phase flips appear for angles above about $500\,\mu$rad, since the arg()-function used for the DWS calculations can only produce values between $-\pi$ and π. So the DC alignment signals are useful especially for larger angles, when not even a heterodyne signal is available. The DC signals can be used as long as a fraction of the beam hits the photodiode.

Both types of angular output signals are used for the on-orbit test mass alignment procedure introduced in Section 9.5. There, the contrast of the interferometers,

$$c_i = \frac{|F_i|}{\Sigma_i} \tag{8.38}$$

is also needed which is the ratio of the complex vector length and the DC signal on the photodiode with a range between 0 and 1. It is a quality criterion of the interferometer alignment and a useful diagnostic output. A high contrast indicates that the two interfering beams have comparable intensities and the overlap of their wavefronts is achieved.

One must be careful how the Fourier transform is performed. If using Equations (6.7) to (6.9), an additional factor of $1/2$ is needed for the contrast calculation if a range between 0 and 1 is desired. However, for the calculations introduced in Equations (8.1) to (8.3), the above calculation is correct.

8.1.7 Scaling the results

The main science data channels – the longitudinal and angular test mass alignment – transferred to the OBC are converted to engineering units, meaning [m] for the length measurement and [rad] for the angles. The conversion can be expressed by

$$\phi_1^{DWS} = g_1 \cdot DWS_1^\phi, \tag{8.39}$$

$$\eta_1^{DWS} = g_2 \cdot DWS_1^\eta, \tag{8.40}$$

$$\phi_2^{DWS} = g_3 \cdot DWS_1^\phi + g_4 \cdot DWS_{12}^\phi, \tag{8.41}$$

$$\eta_2^{DWS} = g_5 \cdot DWS_1^\eta + g_6 \cdot DWS_{12}^\eta, \tag{8.42}$$

for the DWS signals, and by

$$
\begin{aligned}
\phi_1^{DC} &= g_{11} \cdot DC_1^{\phi}, & (8.43) \\
\eta_1^{DC} &= g_{12} \cdot DC_1^{\eta}, & (8.44) \\
\phi_2^{DC} &= g_{13} \cdot DC_1^{\phi} + g_{14} \cdot DC_{12}^{\phi}, & (8.45) \\
\eta_2^{DC} &= g_{15} \cdot DC_1^{\eta} + g_{16} \cdot DC_{12}^{\eta}, . & (8.46)
\end{aligned}
$$

for the DC signals. Even without scaling, the DWS signals have the unit [rad], but this is the phase difference between the integrated phasefronts on different halves of a quadrant photodiode, whereas φ and η represent the geometric test mass orientation. These signals thus have a scaling factor with unit [rad/rad$_{\text{DWS}}$], where the index 'DWS' indicates the phase character of the unit. The longitudinal phases, Ψ_i, are converted to pathlength via

$$
x_1 = \frac{\lambda}{2\pi} \cos(3\,^{\circ}) \cdot \Psi_1, \tag{8.47}
$$

$$
(x_1 - x_2) = \frac{\lambda}{2\pi} \cos(3\,^{\circ}) \cdot \Psi_{12}. \tag{8.48}
$$

The cosine accounts for the non-perpendicular incidence of the beam on the test masses. The differential pathlength measurement, $(x_1 - x_2)$, is the main interferometer output. It corresponds to the differential motion of the two test masses. The x_1 measurement is used as error signal for the DFACS. The measured angles in the X12 interferometer are also differential measurements, but for the transfer to the OBC, the TM 1 angular motion is subtracted. The scaling factors g that convert interferometric angles into test mass angular displacement, depend on the interferometer design and were experimentally determined using calibrated piezoelectric transducers. For the DC signals, they are of the order of $1/500$ [rad] and for the DWS of the order of $1/5000$ [rad/rad$_{\text{DWS}}$] [48]. They depend on the precise beam parameters and must thus be determined experimentally for every optical bench model.

8.2 Asynchronous data transfer and downsampling

At this point the Drag-Free and Attitude Control System (DFACS) has to be introduced in more detail. A simplified scheme is shown in Figure 8.5. The DFACS is part of the OBC software and controls the test masses using the electrostatic actuators. The position of the test masses is read out by the interferometer in the x-axis and in two angles for each test mass [60]. The phasemeter does a Fourier transform of the photodiode signals and the output is post-processed by the DMU. The interface between the DMU, which produces the DFACS error signals, and the OBC has an asynchronous data link.

Additionally, the bandwidth of the data transfer between DMU and OBC is limited. The DMU calculates the data at 100 Hz data rate, but for the DFACS a data
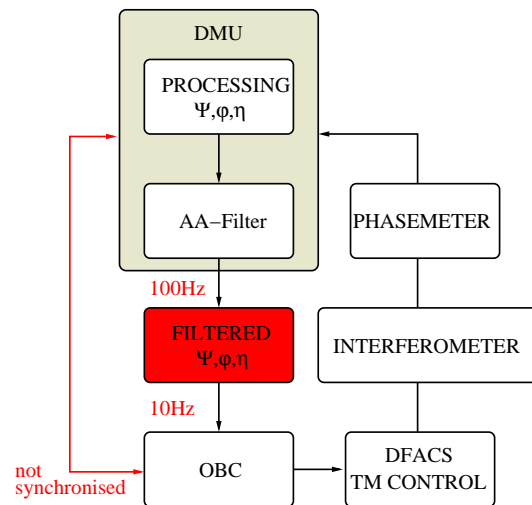
Figure 8.5: Simplified schematic view of the Drag-Free and Attitude Control System.

rate of about 10 Hz is sufficient, since its control loops are quite slow. The reason for the high data rate inside the DMU is the calculation of error signals for the optical pathlength difference (OPD) stabilisation and the laser frequency stabilisation described in [20]. Therefore, the data has to be adequately downsampled before being transferred to the OBC.

On-board LISA Pathfinder, the DMU and OBC are connected via a Mil-BUS, which cannot be simulated in our laboratory setup. The roles of the DMU and the OBC are taken over by two different C-programs. The calculations of the OMS output are done on a DMU-simulator and written into a shared memory with about 32 Hz, which is read out by an OBC-simulator with about 3.2 Hz. The written data at high frequency is always an average of the last ten samples and the process of writing and reading the shared memory with different programs is intrinsically asynchronous. This way the asynchronous data link including averaging is realised in the laboratory setup.

The actual implementation on-board LISA Pathfinder is different. However, also on-board LISA Pathfinder the OBC receives an average of the last 10 samples. Two situations can occur: Either the OBC has a faster clock than the DMU, or its clock is slower. The averaging for both cases is illustrated in Figure 8.6 [61]. In the first case, when the OBC samples faster, the first 10 samples are averaged. Due to the early polling of the OBC for the second average, occasionally the latest sample of the first 10 samples and the following 9 new samples are averaged. In the second case, when the OBC clock is slower, one sample is occasionally left out.

Many additional mechanisms had to be implemented onboard LISA Pathfinder in order to compensate for the asynchronism between the DMU and OBC. For LISA the complications and hidden extra cost will hopefully be remembered and a synchronisation will be implemented in hardware.
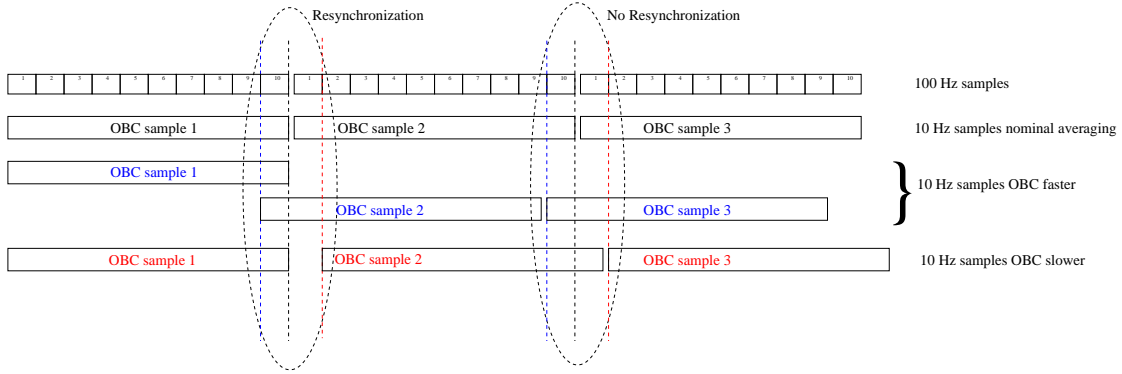
Figure 8.6: Averaging of DMU samples with respect to the OBC clock.

8.2.1 Anti-aliasing filters

Data reduction has to be done carefully. A well known effect appearing in down-sampling processes is aliasing. Noise at frequencies above the sampling frequency of the reduced data appears at low frequencies in the measurement band. After downsampling, this noise is indistinguishable from the 'real' signal at the original frequency. Therefore, it has to be ensured that the original data has no significant signals above half the sampling frequency, the so-called Nyquist frequency. For this purpose, the data has to be low-pass filtered using so-called anti-aliasing filters.

For the implementation of the anti-aliasing filters, two types of digital filters are discussed here: FIR (Finite Impulse Response) and IIR (Infinite Impulse Response) are characterised. Both compute their output values using a linear combination of recent input data. IIR filters, in addition, have feedback, since they use the recent output data as well. Mathematically a filter can be described by its transfer function $H(s)$ or as recursive formula with input data x_n and output data y_n

$$
\begin{aligned}
y_i &= a_0\, x_i + a_1\, x_{i-1} + \cdots + a_n\, x_{i-n} \\
&- b_1\, y_{i-1} - b_2\, y_{i-2} - \cdots - b_m\, y_{i-m}.
\end{aligned}
\tag{8.49}
$$

The index i indicates a certain time sample, $i-1$ means the previous time sample etc. For FIR filters, all b_i are zero. Using these coefficients the transfer function can be written as

$$
H(s) = \frac{a_0 + a_1 z^{-1} + \cdots + a_n z^{-n}}{1 + b_1 z^{-1} + b_2 z^{-2} + \cdots + b_m z^{-m}},
\tag{8.50}
$$

where

$$
z = \exp(sT),
\tag{8.51}
$$

$$
s = \mathrm{i}\omega = 2\pi i f_{\mathrm{samp}}.
\tag{8.52}
$$

The relation between sampling time, T, and the sampling frequency, f_{samp}, is

$$
T = \frac{1}{f_{\mathrm{samp}}}.
\tag{8.53}
$$

The main differences between the two types of filters are:

- FIR filters are always stable. IIR filters may be unstable depending on the choice of coefficients.

- For downsampling, IIR filters must be applied for each value at high (input) data rate, whereas FIR filters can be applied at slow (output) data rate.

- For a given shape of a typical transfer function, IIR filters usually need much less coefficients than FIR filters.

- IIR filters are more sensitive to rounding errors introduced by finite precision arithmetic.

8.2.2 Filter design aspects

A lot of publications on aliasing can be found under the keywords Shannon- and Nyquist-theorem, e.g. in [62, 63]. It plays an important role in digital data processing. By taking discrete samples of a continuous time series, signals at frequencies higher than the Nyquist frequency

$$f_{\mathrm{Ny}} = f_{\mathrm{samp}}/2, \tag{8.54}$$

contribute to the spectrum below the Nyquist frequency. All signals or even noise at frequencies

$$f_n = f \pm n \cdot f_{\mathrm{samp}}, \tag{8.55}$$

for integer n, appear in the spectrum at frequency f. These signals are indistinguishable from the 'real' signal at this frequency, once the time series is decimated to the sampling frequency. An illustration of aliasing in the frequency domain is shown in Figure 8.7. In order to reduce the effect of aliasing, low-pass filters are needed, which are usually also called anti-aliasing (AA) filters. These filters reduce all signals above the Nyquist frequency to a negligible level.

In the case of the data processing inside the DMU the data is calculated at 100 Hz and downsampled to 10 Hz. For this purpose, a suitable AA-filter was designed taking the following aspects into account:

- *Flat frequency response with 0dB gain between DC and 1 Hz:* This is the important measurement band for LISA Pathfinder, where the signal has to be as 'clean' as possible.

- *Maximal group delay of 100 ms between DC and 1 Hz:* The DMU output signals are used for DFACS. Due to constraints on the unity-gain frequencies of the controllers, the overall delay from test mass motion to provision of error signals has to be less than 150 ms (50 ms are allocated for the processing inside the phasemeter and DMU).
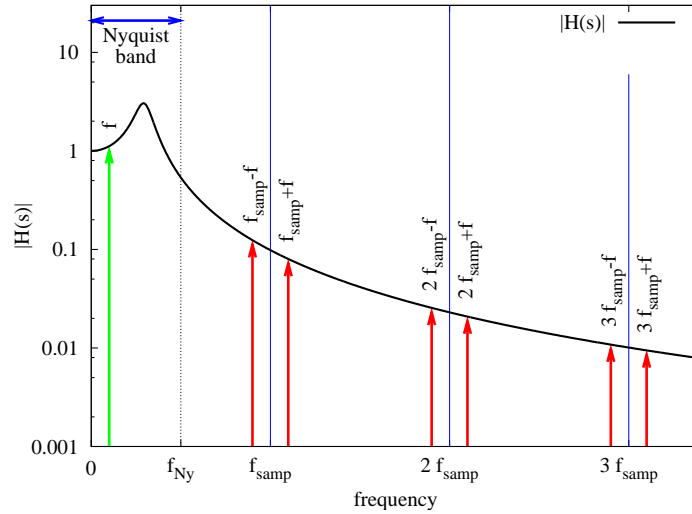
Figure 8.7: Frequency contributions causing aliasing in the measurement band [64].

- No significant peaks above 0 dB for frequencies above 1 Hz.

- *Not too much overshoot or ringing in the time domain*: This makes the filter less sensitive to peaks or steps in the time series of the signal to be filtered.

- *Best possible suppression at 10 Hz and its multiple frequencies*: Signals at those frequencies would be aliased close to DC and therefore directly in the measurement band.

- *Reasonable suppression at all frequencies above 7 Hz*: Signals at these frequencies may be aliased into the measurement between 0 . . . 3 Hz.

- *Smallest possible computational effort*, since the DMU processor load has to be kept small.

Those requirements on the filter were put together in a target response shown in Figure 8.8. The LISO plot shows the targets to be met as cross ($+$), and maximum values not to be exceeded as arrows ($\downarrow$). The DMU processor load reduction turned out to be very important, since not just the main science channels, but also additional channels (see Section 8.3) that are filled with intermediate calculation steps, have to be downsampled before being transferred to the OBC, and the capacity of the DMU is quite small. Therefore, the number of coefficients was limited to the order of 20 ($10\,a+10\,b$) for IIR filters and 30 for FIR filters for the investigations described here. For the optimisation of the filters, a software tool called LISO was used to match the coefficients to the behaviour of the target function [65].

The filter characteristics of the most reasonable candidates are shown in Figure 8.9 and 8.10. The graphs read line by line show the following filter characteristics:
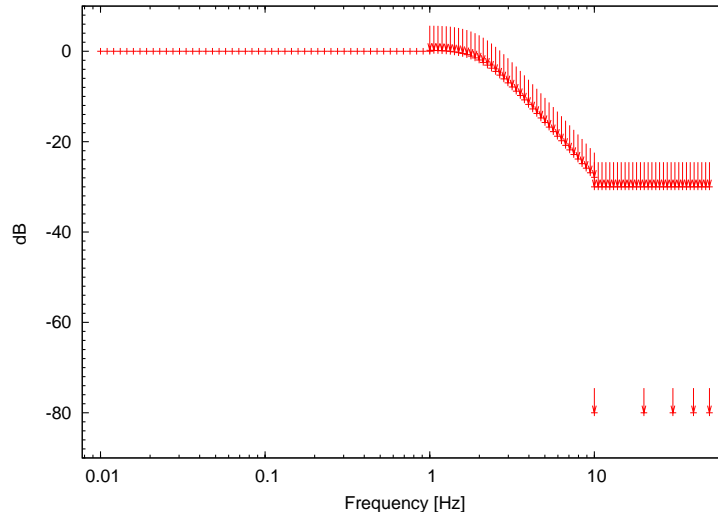
Figure 8.8: Target amplitude of the anti-aliasing filter transfer function.

1. The amplitude of the transfer function.

2. The phase of the transfer function.

3. The group delay of the signal.

4. The response to a step in the input time series.

5. The noise contribution in the measurement band assuming the same spectral amplitude at all frequencies between DC and 50 Hz if the signal at f and signals at frequencies aliasing to f are coherent.

6. The noise contribution in the measurement band assuming the same spectral amplitude at all frequencies between DC and 50 Hz if the signal at f and frequencies aliasing to f are incoherent.

For these requirements, the best IIR filter found is shown in Figure 8.9. This filter has the following coefficients:

$$
\begin{aligned}
a_0 &= +0.0051 \\
a_1 &= -0.0041 \\
a_2 &= +0.0045 \\
b_0 &= +1.0000 \\
b_1 &= -1.9267 \\
b_2 &= +0.9327.
\end{aligned}
$$

It matches the target function quite well. It is almost flat up to 1 Hz and has an f^{-4} slope above the cut-off frequency. Frequencies above 10 Hz are suppressed by
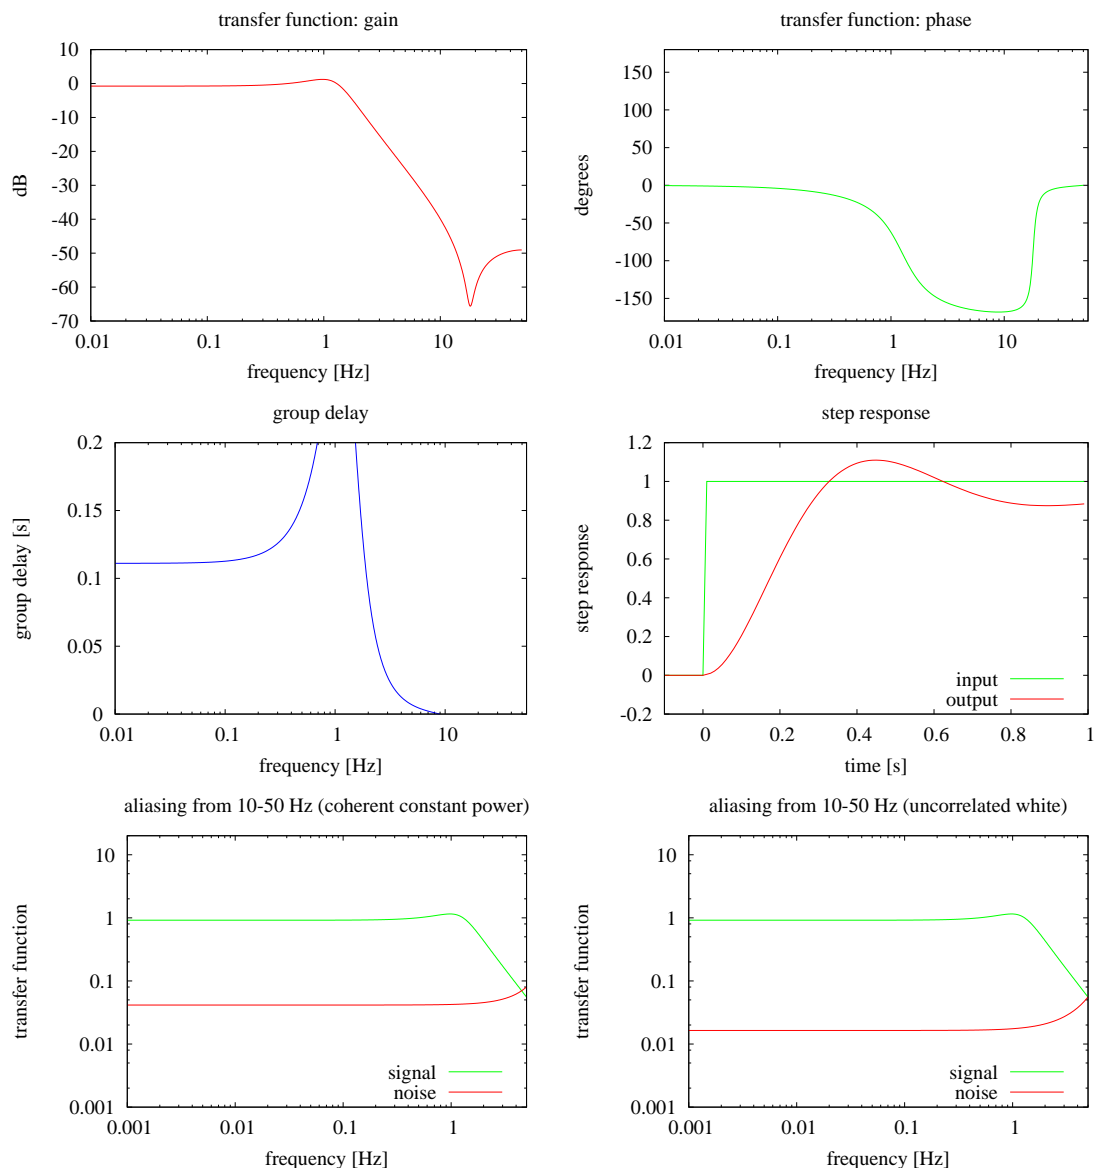
Figure 8.9: Transfer function of an IIR filter with three a and three b coefficients.
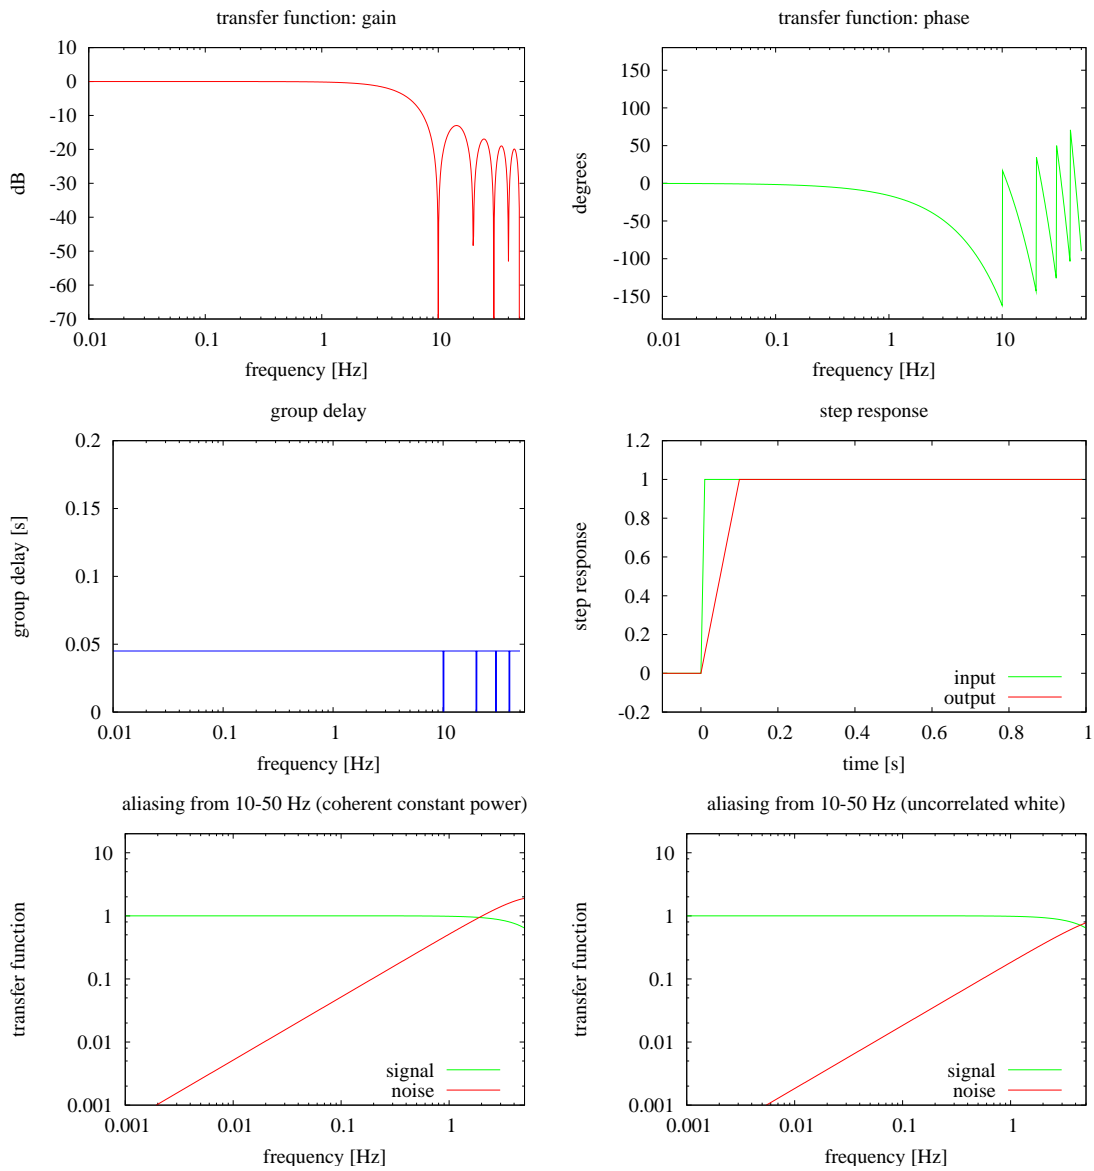
Figure 8.10: Transfer function of an FIR filter with 10 coefficients (moving average).

more than 40 dB. The group delay is acceptable. In case of noise aliasing down to the measurement band of LISA Pathfinder, the filter behaves well. In case of a coherent noise with the same amplitude at high frequencies as at low frequencies, the contribution is more than one order of magnitude less than for the 'real' low frequency signal. Due to the filter internal feedback, the output of IIR-filters tends to oscillate when a step occurs in the input signal. This so-called ringing of the filter is also acceptable: The filter responds to a step function with an overshoot of about 10 %.

In case of an FIR filter, it turned out that the best choice is the straight forward average of the last ten samples as shown in Figure 8.10. Here the transfer function is also nearly flat up to about 1 Hz with a f^{-1} roll-off of the envelope for high frequencies, which is worse than for the IIR filter shown before. However, it has additional notches at multiple frequencies of 10 Hz. The aliasing into the measurement band is not flat and much smaller at low frequencies (below 10 mHz), but increases to high frequencies with an f-slope such that slightly above 1 Hz, it has the same strength as the signal itself. The group delay is small compared to the IIR filter, and of course, it shows no ringing in the time domain due to absence of filter internal feedback.

The resulting filter characteristics are summarised in Table 8.1. In contrast to first expectations, the behaviour of FIR filters with a quite small amount of coefficients is very efficient in the case of DMU internal anti-alias filtering, but this is also helpful in a more general environment. In all cases of digital downsampling, averaging is a suitable choice of low-pass filtering as long as the frequency range of interest is well below the sampling frequency.

Table 8.1: Characteristics of filters under investigation.

property/filter	FIR	IIR
slope	f^{-1}	f^{-4}
suppression above 10 Hz	10 dB	40 dB
suppression @ multiples of 10 Hz	inf	> 40 dB
group delay (in band)	50 ms	110 ms
ringing overshoot	no	10 %
aliasing in measurement band	<2 %	5 %

8.3 OMS debugging

The DMU does the entire data processing of the Optical Metrology System (OMS) (see Section 5). The On-board Computer is connected to the DMU, but has no direct access to most of the DMU internal data. Diagnostic channels are included

inside the DMU, which are passed to the OBC in the so-called 'housekeeping' (H/K) data at 1 Hz. These schemes are adequate and approved on system level and used for checks of environmental conditions like supply voltages or protocol errors, but do not cover interferometer errors or anomalies.

Only the main science data, plus a small set of flags, are usually transferred to the OBC, but in addition sometimes intermediate results are needed. This is necessary for debugging or for characterisation of the interferometer performance, like frequency stability of the laser or thermal characterisation [66].

For this purpose, a set of programmable data channels that feed a cyclic RAM is forseen. A flow diagram is shown in Figure 8.11.
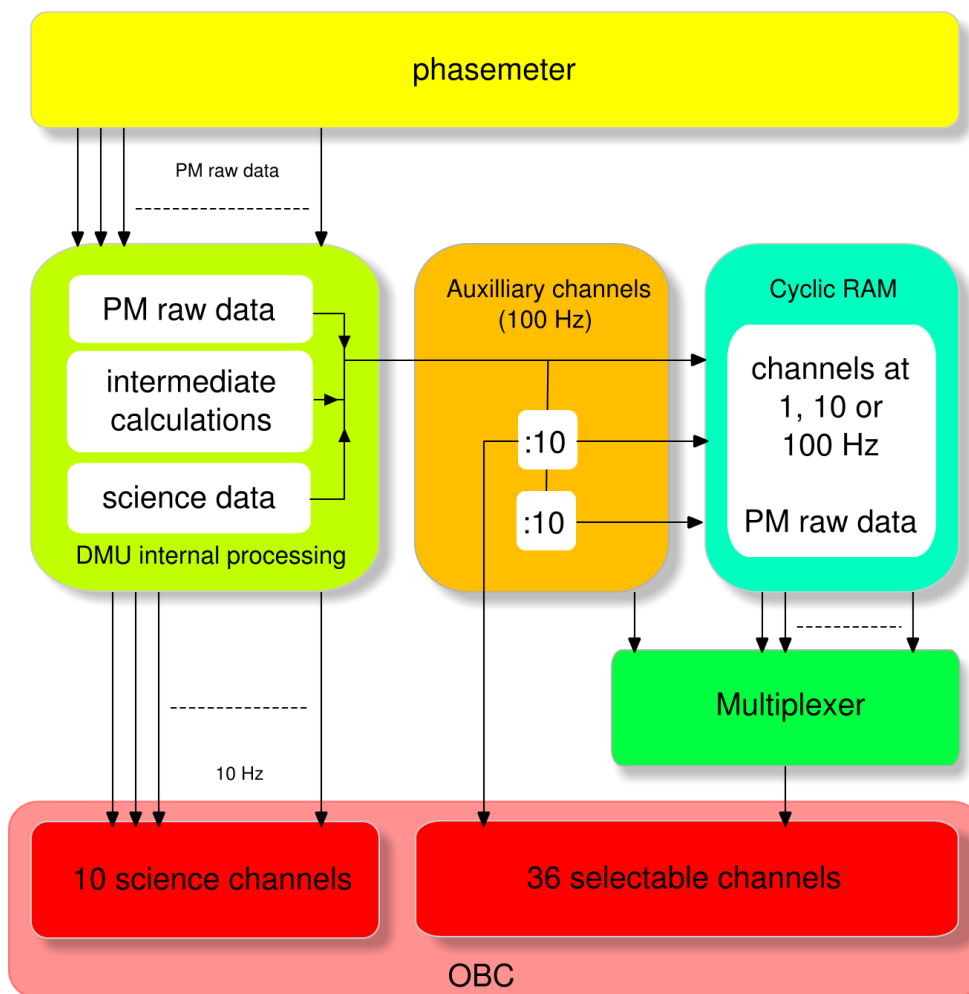
Figure 8.11: DMU flow diagram with selectable channels and cyclic RAM.

The DMU obtains phasemeter raw data at 100 Hz and processes them at the same frequency. The data transfer to the OBC is performed at 10 Hz. The 10 main science channels are hard-wired and always transferred to the OBC. In addition

to those 10 channels, 36 further channels are forseen, which can be filled with selectable variables from:

- DMU intermediate results,

- phasemeter raw data and

- flags.

DMU intermediate results are all results coming from operations performed on the phasemeter raw data, which are variables of type *double*. For debugging purposes, also the phasemeter raw data (*integer* variables) can be transferred. The status bits play an important role in the error handling inside the DMU and will be discussed later (see Section 8.4). The intermediate results can be downsampled using the same anti-aliasing filters as for the science channels, before they are written into the additional channels. The flags are filtered using the flowing rule: A flag is set to '1' if an error occurs. Therefore, a filtered flag is set if the original 100 Hz flag was set at least once during the last 10 samples. No filters are forseen for the phasemeter raw data.

Some stabilisation loops implemented in the DMU software as the laser frequency and the amplitude stabilisation run at 100 Hz. Therefore, the high data rate signals are also of interset for debugging. In order to get a glimpse of those signals a cyclic RAM is implemented inside the DMU. All selected channels can be stored in this RAM directly, but also be downsampled to 10 Hz or 1 Hz. The size of the RAM is 256 kbytes.

Cyclic RAM means that the data is written until the RAM is full. The next data points overwrite the oldest data in the RAM. When the writing is stopped, a time-stamp is added to the last sample in the RAM. Start and stop of writing and reading the RAM are commanded by the OBC.

8.4 OMS error handling

A very important part of the DMU tasks is an adequate error handling for the science channels, which are 2 longitudinal measurements of the test masses and 2×2 angles either derived from DWS or DC calculations. Possible errors may occur already inside the phasemeter, which does not include any error detection on the calculated data itself. Therefore, the DMU has to:

1. detect all possible errors,

2. exclude all affected channels from further data processing and

3. inform the OBC about the quality of the transferred science data in real-time.

First, the error must be detected, and affected channels have to be excluded from the data processing. This is an autonomous DMU task, since only the DMU has the 100 Hz real-time information. The information about erroneous channels is transferred to the OBC via event reports, which do not come in real-time. This is a dangerous situation for the LISA Pathfinder mission, since the DMU output data is used for test mass and satellite control. If erroneous data was transferred to the OBC, the test masses might be commanded in the wrong direction and might reach a point from where they cannot return. Therefore, the multitude of possible errors during one processing step have to be reduced to a small amount of states, which can then be transferred to the OBC in real-time and easily be interpreted by the OBC without knowledge of DMU internals. An overview of the DMU error detection scheme is shown in Figure 8.12 and will be explained in more detail in this section.

Figure 8.12: Overview of the Optical Metrology System error handling.

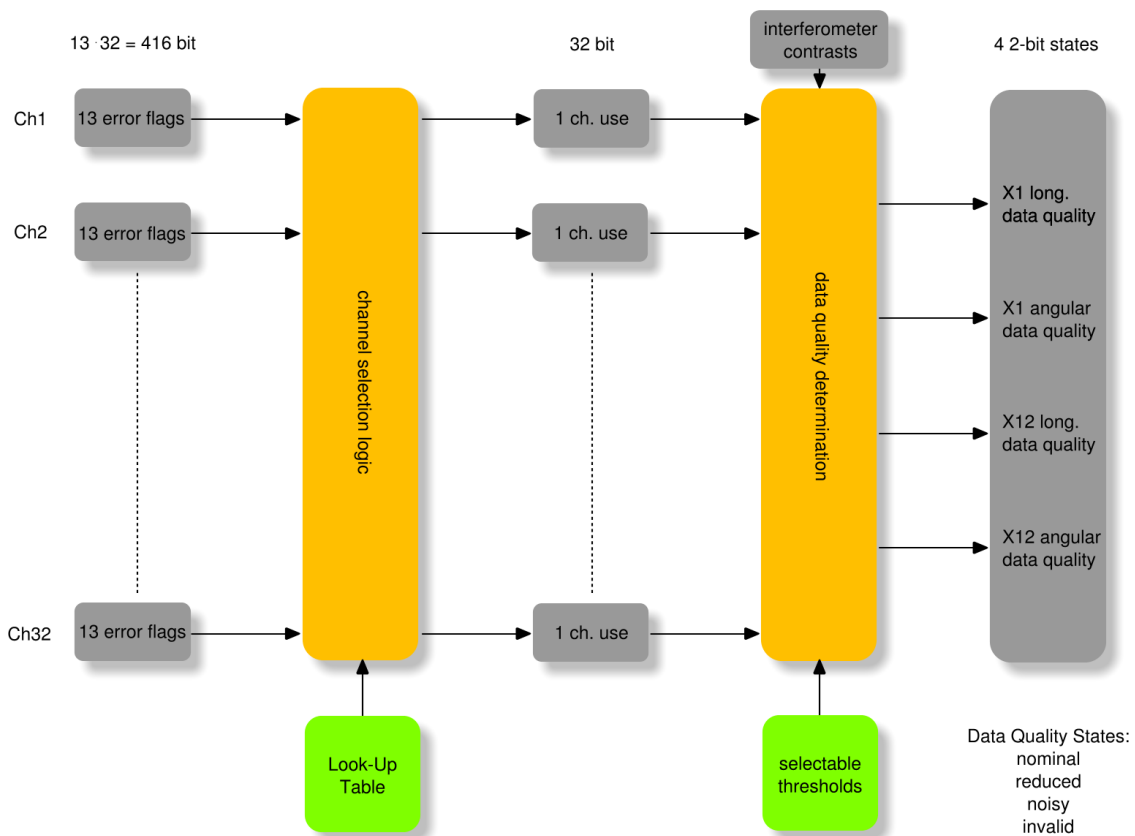13 possible errors may occur in each channel during the data processing. If one of those errors occurs, that channel is excluded from further data processing. However, in some cases one may ignore some errors for certain channels. For this purpose, a Look-Up Table (LUT) is forseen, which allows to enable or disable error

checks. This is implemented in the *channel selection* which produces a *channel use flag* for each channel. Those and the interferometer contrasts are used to define the *data quality* of the science data, which is different for each interferometer and for the angular and longitudinal output data. The data quality is a 2-bit state and can be used by the OBC in real-time in order to avoid usage of corrupted data.

In LISA, a similar error detection scheme can also be applied. Since more interferometers are involved and not only the test mass orientation, but also the relative spacecraft alignment is measured, the decision logic for the data quality classification must be modified, but can be adapted in the same manner as described here.

8.4.1 Channel selection

The possible anomalies in the data processing are listed in Table 8.2. Each of them is indicated by a flag which is raised if the respective error occurs. The indices i, j and k describe the channel of the flags. $i = 1, 12, \mathrm{F}, \mathrm{R}$ is the interferometer, $j = \mathrm{A}, \mathrm{B}, \mathrm{C}, \mathrm{D}$ represents the quadrant designation after channel reassignment according to the description in Section 8.1.2 and $k = \mathrm{A}, \mathrm{B}$ is the phasemeter index. If an index is left out, the flag is set for all channels irrespective of that index. The error checks shown in the table are in detail described in Appendix B.1.

As already mentioned in the description of the possible errors above, the system needs to behave differently for different operational modes. If, for example, LTP is in a highly sensitive measurement of differential test mass acceleration, it is expected that all channels are well utilised, meaning they are not saturated and have an adequate signal strength. However, during the alignment procedure of the test masses, it is expected that the measurement beam is far off the centre of the photodiode, which leads to the absence of a heterodyne signal on some of the quadrants. This causes an underrange condition, which is expected and does not indicate a failure in the interferometer data. In the first case, however, this underrange conditions indicates a significant error.

In order to be able to switch checks on or off depending on the measurement mode running, a selection logic as shown in Figure 8.13 is implemented. Overall 13 error checks are performed on each of the 32 data channels. For all 416 checks, a binary entry in a Look-Up table (LUT) can be chosen in order to enable or disable a check. Following the flow diagram in Figure 8.13 and assuming an overrun occurred in quadrant A of photodiode A in the X1 interferometer ($o_{1,\mathrm{A},\mathrm{A}} = 1$), this flag is connected with a LUT entry via a logical AND operation. If the LUT entry $\mathrm{LUT}_{1,\mathrm{A},\mathrm{A},\mathrm{o}}$ is 1, the output is 1. In case the LUT entry is 0, the output is 0. This operation is performed for every check for this channel with individual Look-Up table entries.

The outputs of the AND operations are then connected via a logical NOR, whose output is 1 only if all outputs of the AND operations before were 0. Those operations are performed for each individual channel. This means if the *channel use flag*

Table 8.2: Flags indicating errors during the data processing. f_u is the update rate.

No.	FLAG	f_u (Hz)	name	description
1	$o_{i,j,k}$	100	overrange	Maximum photo current exceeds threshold (e.g. 99 %)
2	$u_{i,j,k}$	100	underrange	Minimum photo current does not reach threshold (e.g. 1 %)
3	$crc_{i,,k}$	100	Cyclic redundancy check	packet validity
4	$\delta_{i,j,k}$	100	δ-check	Difference between current and previous value exceeds threshold
5	$L_{i,,k}$	100	Latch-Up	ADC disabled, most probably a temporary condition caused by high-energy cosmic radiation
6	$T_{i,,k}$	100	Transmission	Transmission error (Mil-BUS)
7	$O_{i,,k}$	100	Overrun	Overrun during data transfer (Mil-BUS)
8	$M_{i,,k}$	100	Mil-Bus error	Error in the Mil-Bus
9	$PMC_{,,k}$	100	PM configuration	Failure in the PM conf. register
10	$PML_{,,k}$	100	PM LUT checksum	Failure in the PM LUT checksum
11	$PMH_{,,k}$	1	PM State of Health	Error in the PM housekeeping data
12	$\varepsilon_{i,j,}$	100	ε-check	Difference between redundant and nominally equivalent channel exceeds threshold
13	$G_{i,j,k}$	100	Ground com. disable	Manual switch for disabling channel from ground
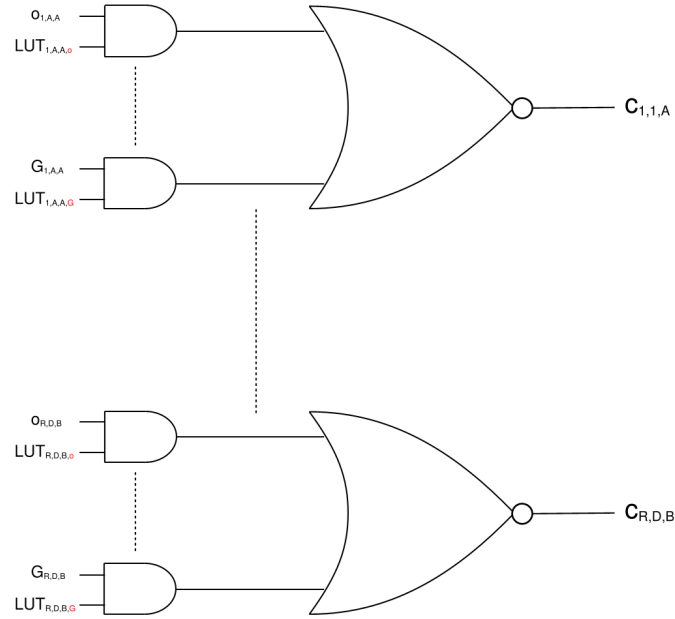
Figure 8.13: Logic flow of the channel selection.

is raised

$$c_{i,j,k} = 1 = \text{`no error'}, \tag{8.56}$$

the channel is used for further data processing. In case it is zero,

$$c_{i,j,k} = 0 = \text{`error'}, \tag{8.57}$$

the channel is ignored. If the LUT entry for a certain check is 1, the check will be performed. If it is 0, the flag is ignored. Attention: For the individual errors like an ε-flag, a 1 means that the check failed, and a 0 means that no failure occurred, which is contrary to the *channel use flag*.

In the actual software implementation the realisation of this logical flow might be different, but behaves as described above. However, for further understanding of the data quality classification the output of this flow diagram is very useful.

8.4.2 Signal calculation

This section deals with the modification of the calculation rules in case of temporary or permanent errors in the data processing. In the following channels will sometimes be called redundant and sometimes independent. *Redundant* channel means the channel which nominally carries the same information as the channel in the context. *Independent* channel means any other channel that is not redundant.

In case of errors occurring during the data processing, one part of the calculation rules is already given in Equations (8.24) and (8.25). If a channel use flag is 0,

the respective channel is excluded from further data processing. If the redundant channel is still available, this is not a big problem, since the science output channels can be calculated as before. If both a channel and its redundant counterpart have a failure, both channels are excluded. What does this mean for the calculations of longitudinal and angular signals?

The longitudinal signals are proportional to the phase of the sum of the complex vectors. If less than 4, but at least 1 independent quadrant is available, the erroneous data channel(s) are omitted from the averaging in Equations (8.26) and (8.27). For the angular signals, the calculation rules change: In case quadrant A fails, Equations (8.34) and (8.37) change to

$$DC_i^\phi = \frac{2 \cdot DC_{i,C} - \Sigma_i^{\mathrm{right}}}{\Sigma_i^*} \tag{8.58}$$

$$DC_i^\eta = \frac{2 \cdot DC_{i,B} - \Sigma_i^{\mathrm{down}}}{\Sigma_i^*} \tag{8.59}$$

$$DWS_i^\phi = \arg\left(\frac{2 \cdot F_{i,C}}{F_i^{\mathrm{right}}}\right) \tag{8.60}$$

$$DWS_i^\eta = \arg\left(\frac{2 \cdot F_{i,B}}{F_i^{\mathrm{down}}}\right). \tag{8.61}$$

Σ_i^* means here the DC sum without the invalid $DC_{i,A}$. In the same manner, the calculation rules change if the other quadrants are erroneous. In case more than one quadrant and its redundant counterpart are missing, the angular signals are not calculated, since no meaningful results can be obtained.

8.4.3 Data quality

The output of the error handling is a set of status bits which is transferred to the OBC as part of the science data stream at 10 Hz. This set is shown in Table 8.3.

Table 8.3: Bits transferred as science data to the OBC.

variable	recommended alignment signals	data quality
x_1	–	2 bits
φ_1, η_1, φ_1^{DWS}, η_1^{DWS}	1 bit	2 bits
x_{12}	–	2 bits
φ_2, η_2, φ_2^{DWS}, η_2^{DWS}	1 bit	2 bits

The longitudinal calculations and the angular calculations are classified in a 2-bit state for each interferometer individually. Additionally, the DMU recommends one particular set of alignment signals – either from DC or DWS calculations –

to be used as error signals for the angular alignment servos in the DFACS. The recommended alignment depends on the contrast in the interferometer. The data quality is distinguished between the states shown in Table 8.4, which are *nominal*, *reduced*, *noisy* and *invalid*. If, for example, four independent channels for a measurement interferometer are available, but a certain amount of the redundant channels is missing, the data quality of the angular measurement is set to *reduced* (4+3...4+0 quadrants), whereas the longitudinal measurement is still *nominal* (4+x quadrants). The reference interferometer data is handled differently. A loss of a quadrant in there is less significant than in a measurement interferometer, since both beams are, to first order, not moving on the photodiode, and for the angular alignment signals, the reference signal is not even needed.

Table 8.4: Data quality states. The examples are for loss of quadrants in the X1- and X12 interferometer. Loss in the reference or frequency interferometer have to be handled differently.

state	alignment example	longitudinal example
nominal	2×4 quadrants	4+x quadrants, $c > 70\,\%$
reduced	4+3...4+0 quadrants	3+x quadrants, $50 < c < 70\,\%$
noisy	3+3...3+0 quadrants	2+x, 1+x quadrants, $2 < c < 50\,\%$
invalid	< 3 quadrants	0 quadrants, $c < 2\,\%$

In Appendix B.2, the intention of this classification is described in more detail and a decision logic is introduced, which also takes the differentiated handling of the interferometers and the contrast into account.

Since the angular data for test mass 2 that is transferred to the OBC is calculated from both the X12 and X1 interferometer data (Equations (8.39) – (8.46)), the X12 angular data quality can have at maximum the X1 angular data quality state. A logical flow diagram of the angular data quality classification is shown in Figure 8.14.

The same is shown in Figure 8.15 for the longitudinal data quality. It is much more complex, since three different interferometers have to be taken into account. First all interferometers are handled individually and get an intermediate data quality state. The frequency interferometer cannot make the data quality *invalid*, since in that case, the longitudinal measurements are indeed more noisy but still reliable. In the end, the lowest state of the intermediate data quality classification is taken as the final output. Additionally, the data is defined as *invalid* if one of the DMU internal stabilisation loops (OPD or laser frequency) fails. The invalidity lasts until the loop states are again well defined. This can also mean that they are turned off.

Bit operations for angular signals
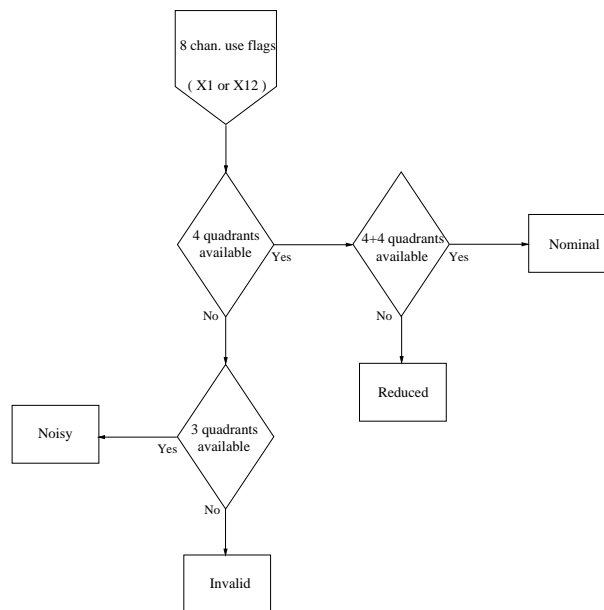
Figure 8.14: Logical flow of the angular data quality classification.

Recommended alignment signals

There are two types of alignment signals, DC and DWS. The former has a wider dynamic range and is more robust to abnormal conditions, while the latter has the best noise performance under near-nominal conditions. The DMU always computes both and transmits both of these signals to the OBC. In addition, it computes and transmits one bit per interferometer that summarises the state of the individual channels as a recommendation to the OBC, which of the two types of signals are more reliable and recommended to be used. This bit is called *recommended alignment signal*.

The recommended alignment signals depend on the contrast in the interferometers. The DMU transmits, in addition to the data quality flags, a *recommended alignment* bit to the OBC with the science data stream.
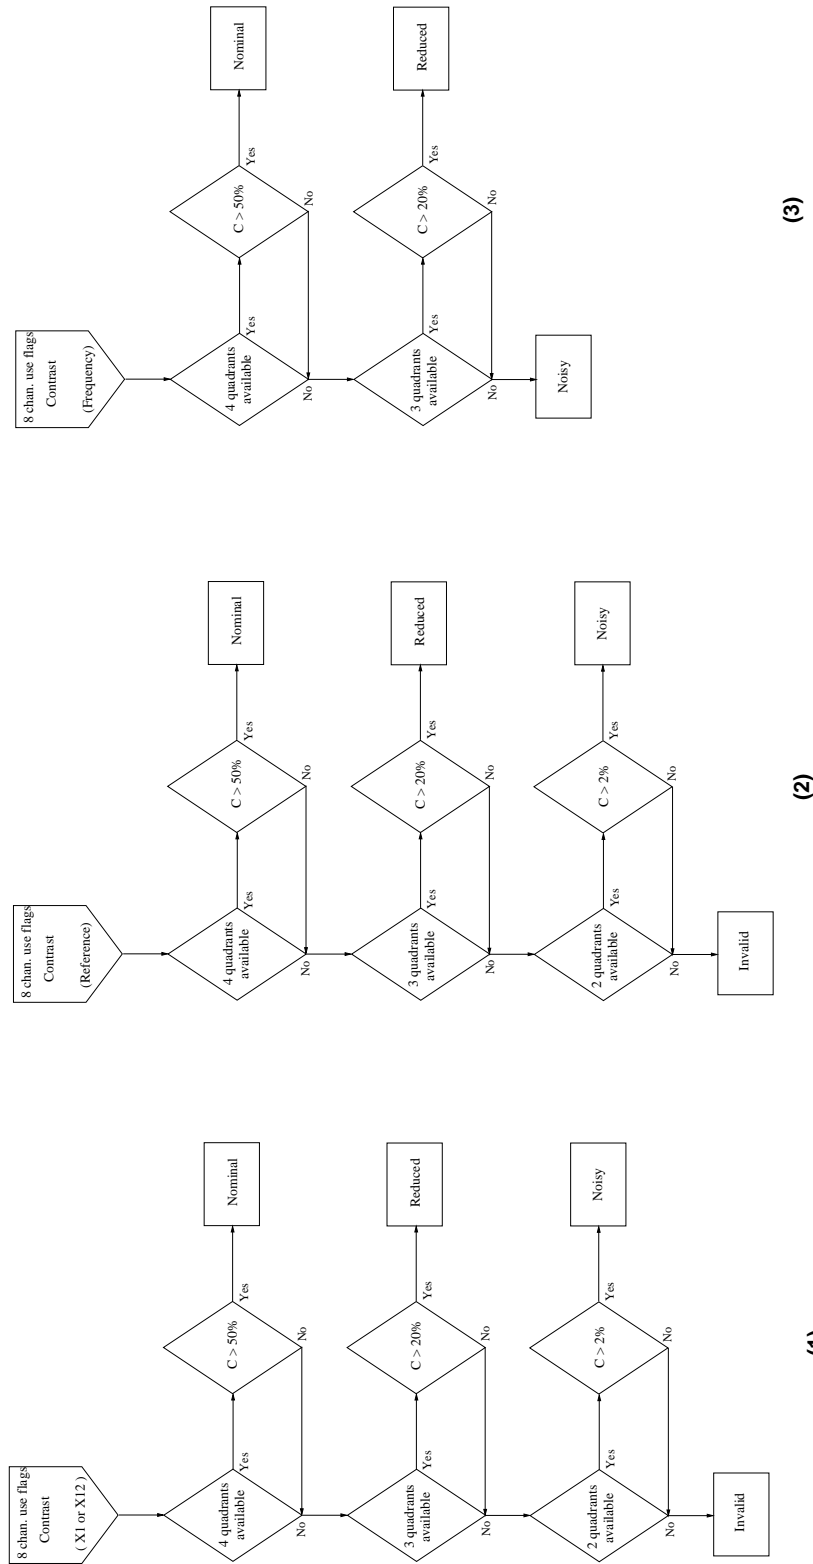
Figure 8.15: Logical flow of the longitudinal data quality classification.

Testing of the on-board data processing

The software as defined in Chapter 8 was thoroughly tested at AEI as part of this work, which was carried out together with Felipe Guzmán Cervantes, Antonio García Marín, Vinzenz Wand and Gerhard Heinzel. A software written in C including all aspects of the DMU software concerning the interferometric data processing was used. Special aspects like the OMS debbuging (Section 8.3) and the error handling (Section 8.4) were not included, since those are very specific for the real flight hardware and could therefore not be implemented in the testbed.

The experimental results shown here demonstrate the performance with averaged signals that are downsampled and asynchronously transferred to an OBC simulating software. The on-orbit test mass alignment is demonstrated with asynchronously transferred data and wide range PZT actuated mirrors. The effect of channel failures was investigated. Additionally, the cross-talk of the interferometric output channels was reduced by including an orthogonalisation scheme in soft- and hardware.

The tests presented here can be consulted for the definition of procedures and the data analysis for the on-ground testing of flight hardware and the mission. Especially the software orthogonalisation and cross-talk are an essential part of the LISA Pathfinder experimental master plan and relevant also for LISA.

9.1 Channel failures

During the data processing inside the phasemeter and/or the DMU, temporary channel failures are expected. Probably most common will be a Latch-Up event in an ADC of the phasemeter. Even if just one channel is affected, all channels of one quadrant photodiode will be switched off for a short time due to the phasemeter

hardware design. Since the channels of redundant quadrants are averaged and there is a slight mismatch between the redundant measurements, this effect causes small jumps in the DFACS data. Mismatch can be caused by:

- contamination of photodiodes or the recombination beamsplitter,

- spatially unbalanced response of the photodiodes,

- different photodiode amplifier gain between the nominal and redundant channels,

- photodiode alignment and

- differing beam parameters of the interfering laser beams in combination with different distances to the photodiodes.

This mismatch was investigated on the Engineering Model optical bench. The effect on the Flight Model optical bench may be different due to different occurrence of the effects listed above, but similar. The results shown here are very important in order to see characteristic behaviour and implement precautions that can adequately handle these failures in the software.

Figure 9.1 and 9.2 show a time series of the the DFACS channels as transferred to the OBC. These are the downsampled longitudinal and angular signals calculated by the DMU. They are used by the DFACS to control the satellite and the test masses. The data is already anti-alias filtered with a moving average filter of 10 samples. In the first graph of Figure 9.1, a Latch-Up or similar effect is simulated, which switches a whole quadrant photodiode off. In time segment 3 the averaged information of both QPDs (nominal and redundant) is used. In time segment 4 PD12A is switched off. In time segment 5 PD12B is switched on again and in segment 6 PD12B is switched off.

In the X12 channel there is a change of 1 mrad if a photodiode is switched off. This corresponds to about 80 pm test mass displacement. If DFACS is using the OMS data for test mass control, the test mass will be moved to a different position. This is not critical for the mission, but in the data analysis these changes have to be taken into account.

The DC angles have huge jumps of the order of $100 \, \mu$rad. This is not very important, since usually the DWS data is used for test mass control, but also in these channels the effect is non-negligible: If a photodiode is switched off, the change in the angle is of the order of $1 \, \mu$rad.

In Figure 9.2 the DFACS channels are shown under the same conditions as before, but here just individual quadrants of a QPD were switched off. It is indicated which quadrant is switched off above the upper x-axis: E.g. BC means quadrant C of photodiode B is switched off. The data processing starts while all channels are active. In time segment 4 quadrant A of PD12A is switched off and in segment 5 switched on again. In segment 6 the same quadrant of the redundant quadrant
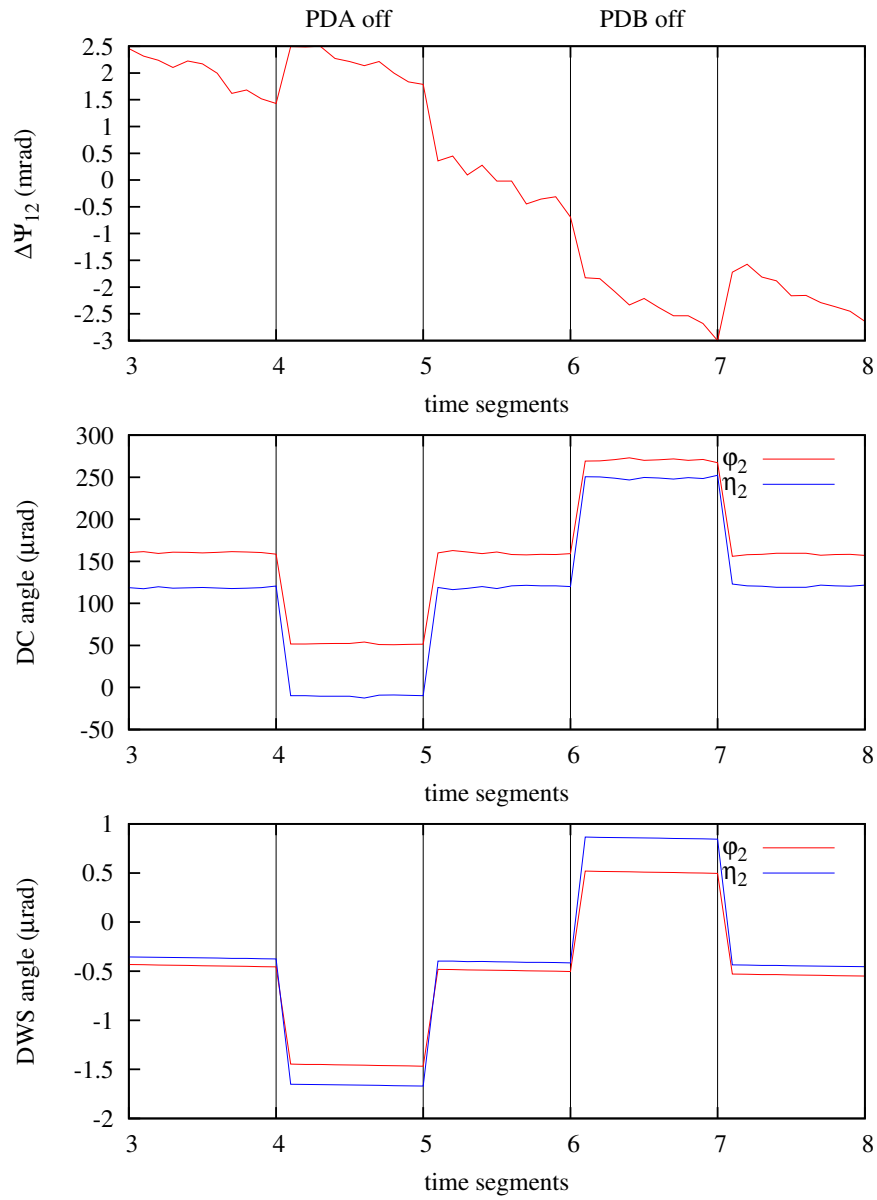
Figure 9.1: DFACS channel response when one of the redundant photodiodes is switched off.
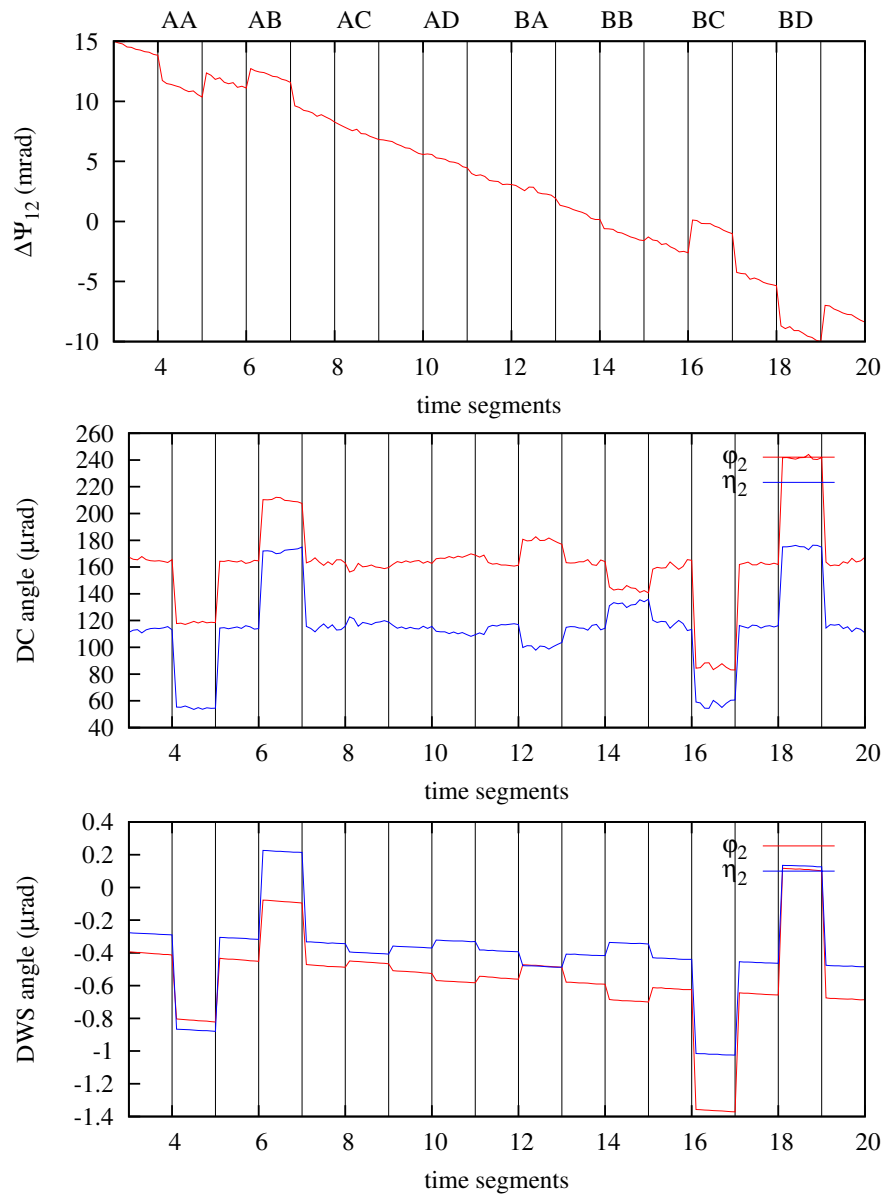
Figure 9.2: DFACS channel response, when quadrants are switched on and off separately.

photodiode is switched off and so on. Here, changes of the same order of magnitude can be seen. The deviation of the signals, when quadrants are switched off, are summarised in Table 9.1.

Table 9.1: Deviations in the science data channels in case of channel failures.

PD/chan.	$\Delta\Psi_{12}$ (mrad)	φ_2^{DWS} [μrad]	η_2^{DWS} [μrad]	φ_2^{DC} [μrad]	η_2^{DC} [μrad]
PD12A	1.0	-1.0	-1.3	-108	-131
PD12B	-1.5	1.5	1.3	108	131
A/A	-1.55	-0.37	-0.55	-45	-58
B/A	1.55	0.37	0.55	45	58
A/B	-0.3..0.3	0.05	-0.04	-5..5	-5..5
B/B	-0.3..0.3	-0.05	0.04	-5..5	-5..5
A/C	0.4	0.10	-0.07	19	-16
B/C	-0.4	-0.10	0.07	-19	16
A/D	2.45	-0.75	-0.58	-78	-58
B/D	-2.45	0.75	0.58	78	68

This knowledge is important especially for the Flight Model optical bench, since the information of temporary failures in some channels come with event reports from the DMU to the OBC and can be transferred to ground. Therefore, it might be possible to correct for the deviations in long-term measurements in the data post-processing on ground, since the affected channels can be identified in the data analysis.

9.2 Downsampling and asynchronous data transfer

In order to verify the performance of the LTP including the asynchronous data link between the DMU and the OBC described in Section 8.2, two investigations were done: A real-time simulation and a post-processing analysis. The real-time simulation was performed with two deviating clocks driving the OBC and DMU processing. The frequency difference of the clocks was of the order of 1 %. The data link was implemented using a shared memory. In this real-time simulation the clock deviation between the two processes could not be controlled. Therefore, a more analytic investigation was performed in addition that characterises the data link using post-processing methods. Both results are presented here.

9.2.1 Real-time simulation

In the laboratory implementation of the LTP interferometer, results are usually taken at full rate from the raw data without any of the downsampling described in

Section 8.2. One main goal of the measurements described here is the verification of an implementation similar to the real LTP experiment.

The phase data is calculated with about 32 Hz instead of 100 Hz in LISA Pathfinder in order to reduce the data rate and to avoid noise introduced by the 50 Hz harmonics of the power supply. The reason for 100 Hz sampling on-board LISA Pathfinder is a sufficient update rate for the laser feedback signals, which are in the present implementation of the laboratory experiment produced by an analogue circuit that bypasses the digital phasemeter. The interferometer performance is reached also with lower sampling rate as can be seen in Figure 9.3. The different curves are the raw DMU channel, and both the either synchronously or asynchronously filtered and downsampled data. The anti-aliasing filter used is an IIR low-pass filter for downsampling by a factor of 10, which has a sufficient noise suppression in the measurement band, as described in Section 8.2.1. Above about 200 mHz aliasing shows up in the X1 interferometer data. However, it is above the measurement band. The asynchronism can also be seen in this figure: The characteristic peaks around 1 Hz in the spectrum are slightly shifted between the synchronous downsampled data and the asynchronous data.
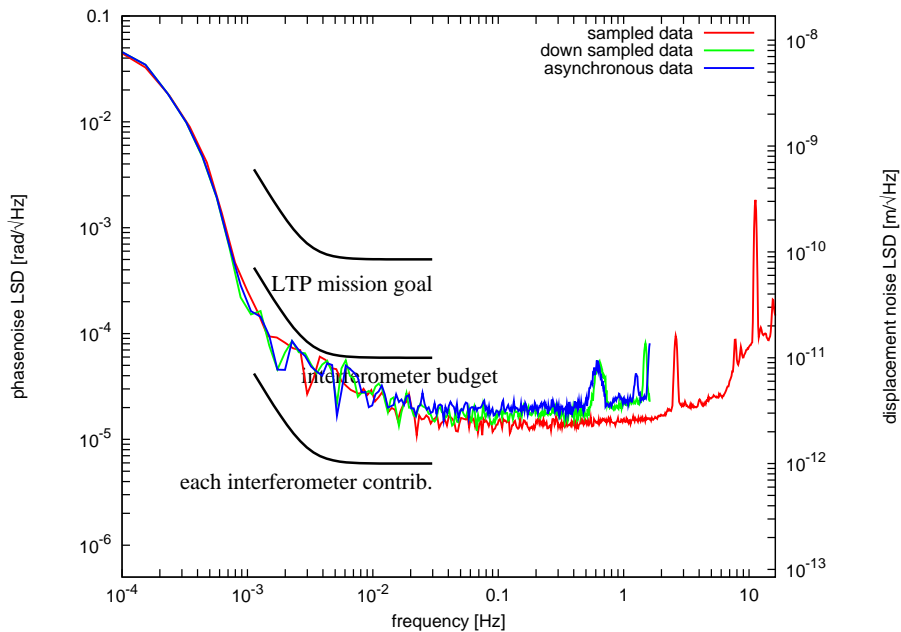
Figure 9.3: Interferometer performance with asynchronous data transfer.

This result was the main purpose of the software tests. The data processing includes all critical parts in the software which have not been tested before. During writing this thesis, tests of Engineering and Flight Model units of the OMS including the Data Management Unit are in progress. The presented results are representative also for these units and can, therefore, be used for comparison.

9.2.2 Data link analysis

For the analysis of the asynchronous data link between the DMU and the OBC, artificially generated data was used [61]. Spectra of the expected interferometer outputs on-board LISA Pathfinder were taken from simulation results from EADS Astrium [67]. A time series was generated with the expected spectral behaviour using the noise generator of LISO [65, 68]. A plot of this time series is shown in Figure 9.4.
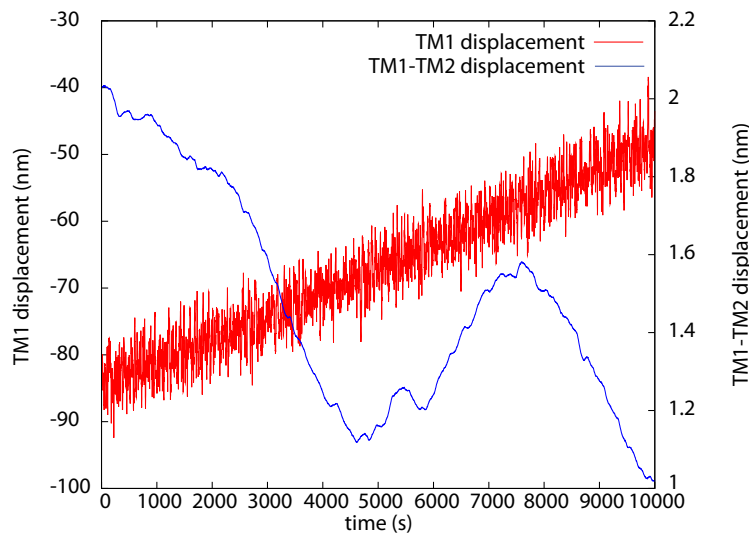
Figure 9.4: Generated time series of expected test mass motion.

In order to simulate a realistic timing jitter between OBC and DMU, the phase noise spectrum of a commercially available 48 MHz quartz oscillator was measured. This oscillator is not particularly well specified in absolute frequency and phase noise behaviour and therefore representative for the oscillator characteristics on LISA Pathfinder. As for the test mass motion, a time series was also generated for the oscillator with the LISO noise generator. It is shown in Figure 9.5(a).

The deviation of the oscillator frequency in comparison to a stable reference clock is about 800 Hz. On-board LISA Pathfinder the DMU is not synchronised to the OBC, which might either have a faster or slower clock. In the simulations presented here, the OBC has a faster clock. On-board LISA Pathfinder, this would cause a loss of one sample every 600 s, which can be seen in Figure 9.5(b). The time stamps produced by the DMU deviate from the OBC time stamps.

However, in the LISA Pathfinder experiment the steps in the time stamps are not very important, since a quasi-synchronisation is implemented. A mechanism is forseen, inside the DMU that always averages the latest 10 data points, before the averaged result is transferred to the OBC (Section 8.2). This feature of the LTP system was simulated with a C-program by post-processing the high data rate output of the DMU-simulator. The resulting spectra of test mass motion with and

(a) Oscillator deviation

(b) Step function in the time stamps
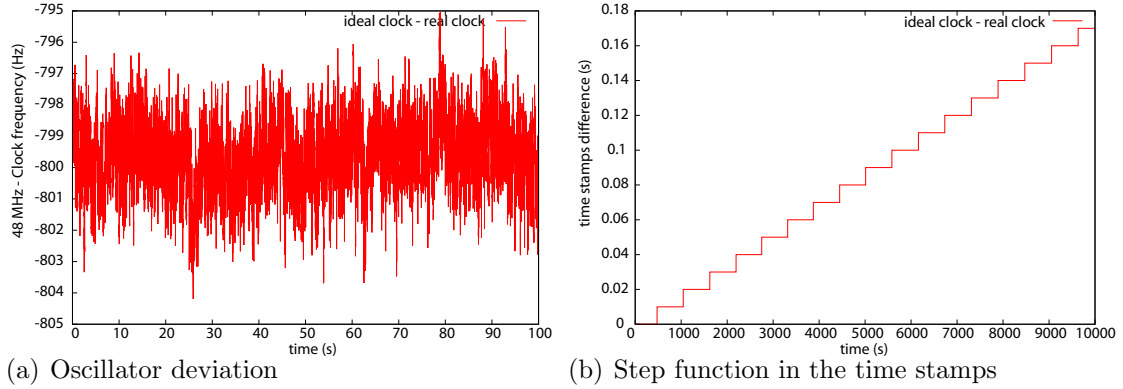
Figure 9.5: Difference of clocks in the OBC and DMU, and step function in the time stamps due to clock difference.
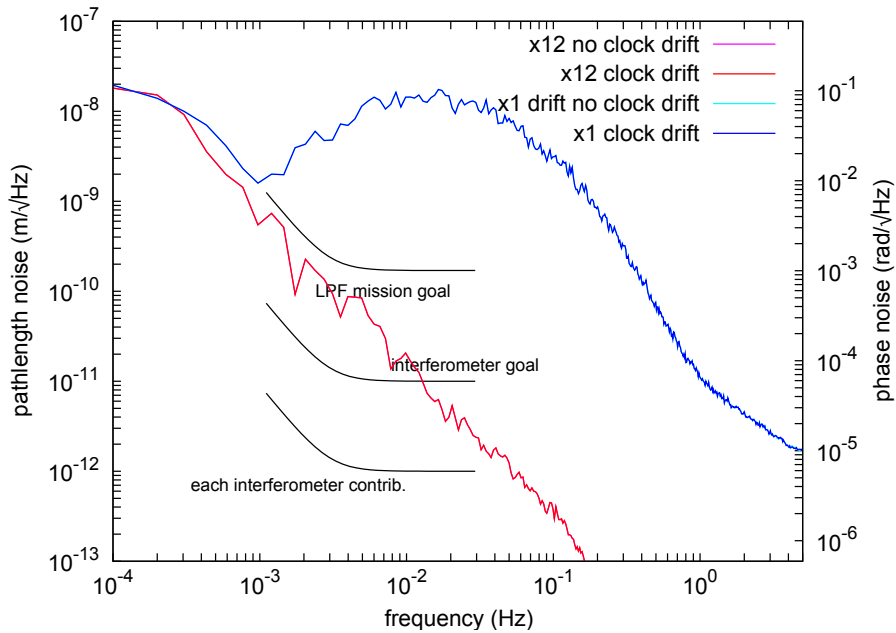
Figure 9.6: Comparison of test mass motion with and without asynchronous data link. The effect of clock drift is invisibly small.

without asynchronously sampled data is shown in Figure 9.6. It shows no visible effect of asynchronous sampling.

The remaining uncertainty in this simulation is the influence of the dynamic behaviour of the system. Such a simulation was not done, since in that case a dynamic model of the satellite has to be included, which exceeds the scope of this thesis.

9.3 DFACS Loops

On-board LISA Pathfinder the test masses are controlled via an electrostatic actuator inside the Inertial Reference Sensor (IRS). This sensor is also able to measure the test mass positions. The sensitive x-axis and two angles, φ and η, are controlled in science mode using the interferometer output signals. All other degrees of freedom are controlled using the more noisy IRS signals. In the Engineering Model optical bench wide range 3-axes PZT actuated mirrors were inserted in order to test not just the pure interferometric readout, but also the LTP dynamics.

With those PZTs it was possible to close the loop using the interferometric output for the angular test mass control. A control scheme is shown in Figure 9.7.
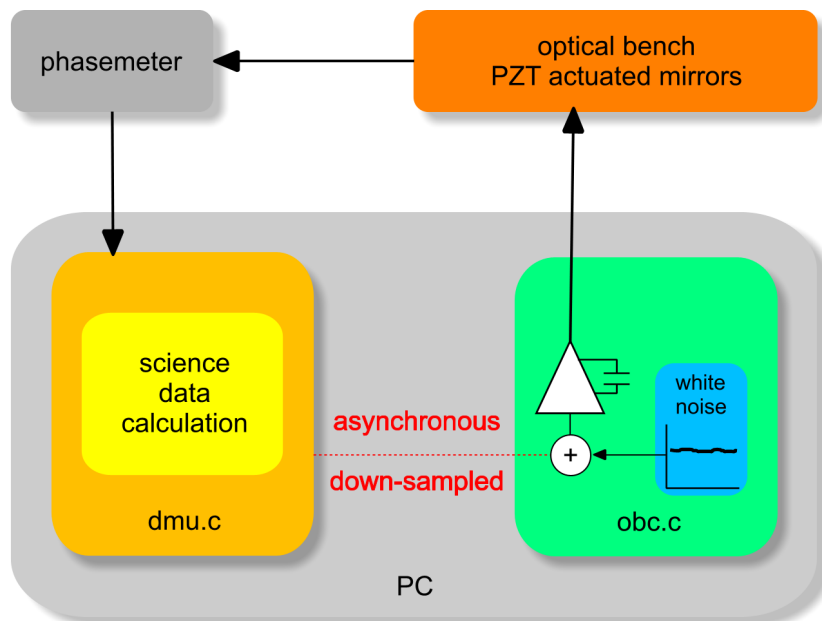
Figure 9.7: Test mass control scheme.

A control loop was implemented as already described in Section 8.2. The data processing is done in the DMU simulator, which transfers the asynchronous science data to the OBC simulator, where a servo for each test mass angle is implemented in software. The servo output data is fed through a DAC, which is connected to the

PZTs via high voltage amplifiers. In addition a noise generator is forseen, which was used to add white noise to the error signals in order to measure the open loop gain (OLG) of each servo loop. The results are shown in Figure 9.8. Each OLG has the same design: a $1/f$ part in the frequency region of unity gain and a $1/f^2$ part at lower frequencies in order to have an adequate noise suppression. The measured open loop gain transfer functions are in agreement with the servo design and have the predicted characteristics, since the frequency response of the plant is flat at low frequencies, and each angular output signal has approximately the same scaling factor.
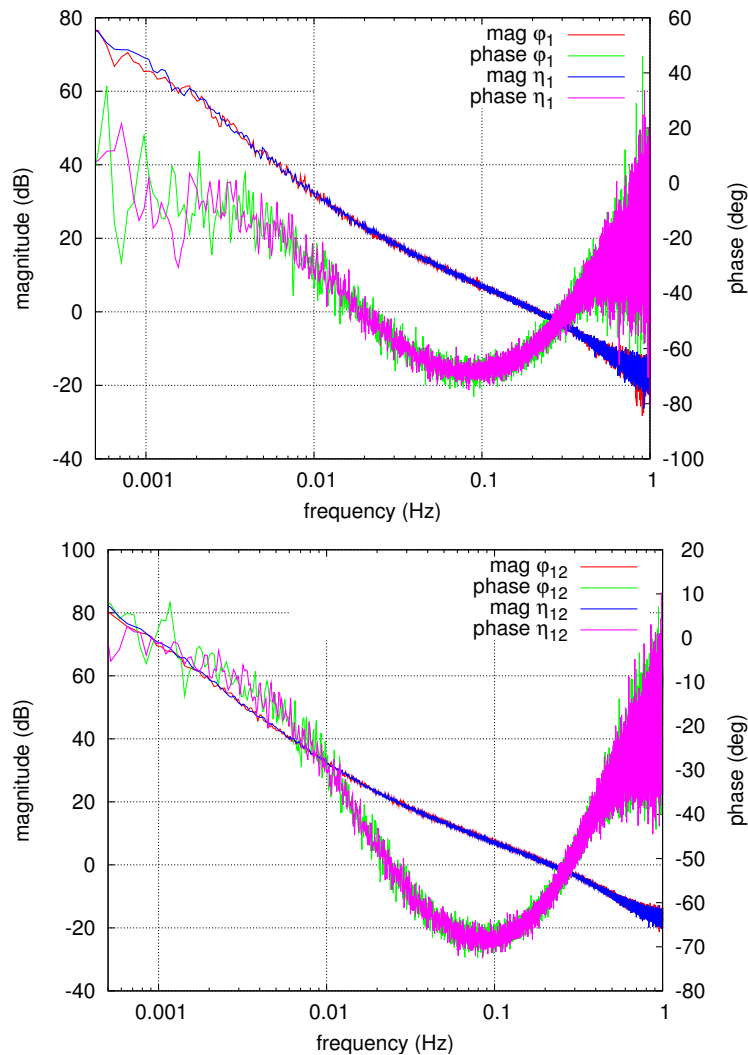
Figure 9.8: Open loop gain for each control loop of the test mass angles.

The servo can either be driven with the DC or DWS signals. The calibration of those signals is well known. Therefore, they scale the same way and both have a flat frequency response. However, the DWS signals are useful only in a range

around the nominal zero.

The angles, in which the PZTs actuate, are orthogonalised using an electronic matrix mixing circuit, which consists of a combination of amplifiers and adders for each PZT channel, which can be combined and tuned in such a way that the coupling into the wrong axes is minimised. A software implementation of this circuit was also realised and is described in the following section. The residual coupling is small, but for this measurement not optimised. However, this has almost no effect on the frequency response.

The unity gain frequency for all loops is of the order of 200 mHz. In the experiment it is limited by the sampling rate in the setup, which is 3.246 Hz on the OBC side. In LISA Pathfinder this frequency is 10 Hz. The satellite will follow the first test mass using Micro-Newton thrusters, which can move the satellite very accurately. These thrusters have a finite delay which limits the actuation bandwidth. The second test mass will be actuated using the electrostatic actuator. In order to not introduce noise in the measurement band, the unity gain frequency of this controller is in the mHz-range. Therefore, the actual implementation on-board LISA Pathfinder includes lower bandwidth in the loops compared to the ones presented here.

9.4 Software orthogonalisation

The 3-axes PZTs used for closing the test mass alignment loops are of type S-325 (Physik Instrumente). They are actuated in three axes arranged in an equilateral triangle. In order to control the PZTs in three orthogonal axes, an electronic mixing circuit was used to disentangle the different degrees of freedom. An additional orthogonalisation scheme was implemented in software.

The interferometer is sensitive to four angles: φ_1, η_1, φ_{12} and η_{12}, where the index '1' indicates the angles of TM 1 with respect to the X1 interferometer and index '12' the combined angle of TM 2 and TM 1 with respect to the X12 interferometer. φ_1 and η_1 are measurements of the TM 1 position, whereas φ_{12} and η_{12} carry the information of TM 2, but also a coupling from TM 1. By measuring the TM 1 alignment, this can be subtracted in X12.

The coordinate system for the angles is defined by the slits of the quadrant photodiodes, which are aligned differently in the X1 and X12 interferometer and even between a photodiode and its redundant counterpart (Section 8.1.3). On-board LISA Pathfinder the situation is the same: Each inertial reference sensor and each interferometric readout defines a different coordinate system. However, the misalignment can be measured.

Assuming two coordinate frames, where 'ifo' is the interferometer frame and 'ac'

the frame of the actuators, the measured angles can be written as

$$
\begin{align}
\varphi_1^{\mathrm{ifo}} &= \varphi_1^{\mathrm{ifo}}\left(\varphi_1^{\mathrm{ac}}, \eta_1^{\mathrm{ac}}\right), & (9.1)\\
\eta_1^{\mathrm{ifo}} &= \eta_1^{\mathrm{ifo}}\left(\varphi_1^{\mathrm{ac}}, \eta_1^{\mathrm{ac}}\right), & (9.2)\\
\varphi_{12}^{\mathrm{ifo}} &= \varphi_{12}^{\mathrm{ifo}}\left(\varphi_1^{\mathrm{ac}}, \eta_1^{\mathrm{ac}}, \varphi_{12}^{\mathrm{ac}}, \eta_{12}^{\mathrm{ac}}\right), & (9.3)\\
\eta_{12}^{\mathrm{ifo}} &= \eta_{12}^{\mathrm{ifo}}\left(\varphi_1^{\mathrm{ac}}, \eta_1^{\mathrm{ac}}, \varphi_{12}^{\mathrm{ac}}, \eta_{12}^{\mathrm{ac}}\right). & (9.4)
\end{align}
$$

The coordinate frames were pre-aligned using the electronic circuit by minimising the cross-talk in the graphical user interface of the interferometer. The cross-talk of the different channels is assumed to be linear. Each coupling factor was measured and the coupling was compensated in software. In the following graphs the residual coupling of test mass jitter of each angle into the orthogonal channels is shown.

φ_1 couples into η_1, φ_{12} and η_{12} as can be seen in Figure 9.9(a). φ_1 is modulated sinusoidally with an amplitude of about $25\,\mu\mathrm{rad}$. The measurement, when η_1 is modulated, is shown in Figure 9.9(b). The modulation has the same order of magnitude. For the angles of test mass 2 the results are shown in Figure 9.10.

The residual cross-talk of all channels is less than $0.2\,\%$ if the software orthogonalisation is applied. The same low level of cross-talk can also be achieved, when only the electronic mixing circuit is used for orthogonalisation. This is obvious, since even with software orthogonalisation the mixing circuit is still in the actuation chain. However, instead of manually adjusting the amplification of the input amplifiers, it is practically easier to measure the coupling numerically and use the resulting value directly as coefficient for subtraction.

The implementation in the DFACS is similar. The residual cross coupling of less than $0.2\,\%$ sets an upper limit for the interferometer internal cross coupling. The measured effect might still be introduced by the PZT, but it is not possible to measure their influence independently.

9.5 On-orbit test mass alignment

The test masses on-board LISA Pathfinder have to be initially aligned when there is no contrast in the X1 or X12 interferometer or if the measurement beam even misses a photodiode. After it is decaged, the test mass is misaligned with respect to the optical bench due to constructional constrains. This has the consequence that the beams reflected by the test masses might not hit the photodiodes, or at least do not produce a heterodyne signal.

Therefore, an autonomous alignment procedure is implemented on-board the satellite that is able to initially align the test masses without any intervention from ground. Only the OBC has all information needed to control the test masses. Therefore, the DMU calculates the OMS output data plus a set of status flags that are used by the OBC in order to decide which error signals to use for the control loops.

(a) φ_1-coupling
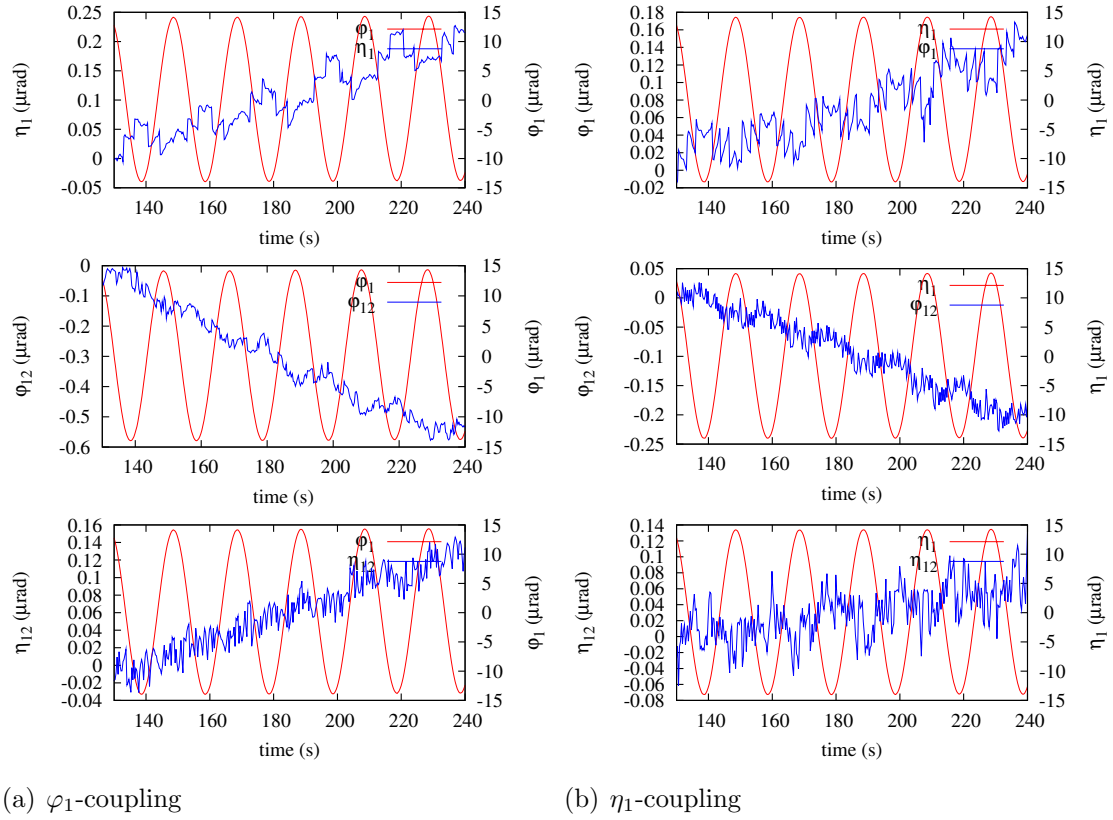
(b) η_1-coupling

Figure 9.9: Coupling of the X1 angles into the other interferometer channels.
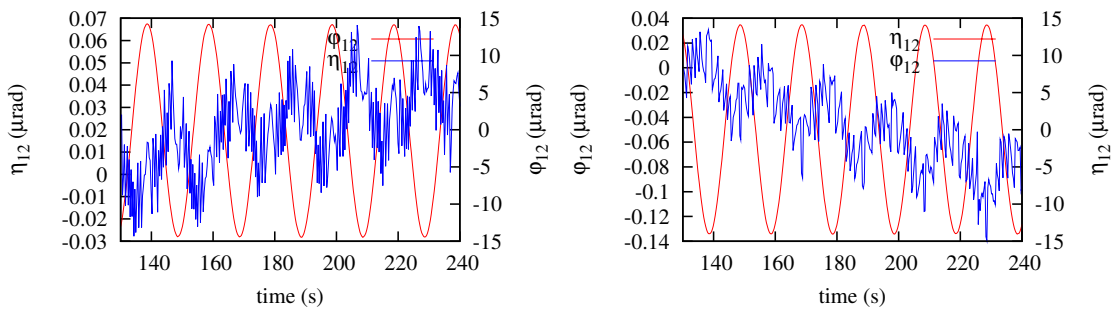
Figure 9.10: Coupling of the X12 interferometer angles into each other.

The basic idea of the alignment procedure is to spirally scan the test mass in the X1 interferometer first until light impinges on the quadrant photodiode belonging to this interferometer. If a certain threshold is reached, the DC alignment signals are used to control the test mass and, thus, force the reflected beam to be centred on the quadrant photodiode. If a certain contrast level of the heterodyne signal is reached, the DWS signals are used instead for test mass alignment. Test mass 1 is then held in this aligned orientation and the same procedure is done for test mass 2 in the X12 interferometer.

In order not to unnecessarily misalign the test masses, the alignment procedure does not unconditionally start with a spiral scan of the test mass, but first checks for contrast or DC signals on the photodiodes.

The OBC continuously receives flags for the following conditions for each test mass from the DMU. They are raised if the respective condition is fullfilled:

$$c_{1,12} > 5\%$$
$$c_{1,12} > 12\%$$
$$c_{1,12} < 8\%$$
$$c_{1,12} > 50\%$$
$$\sum_{1,12} > \sum_{1,12}(\text{RB}) + 0.14 \cdot \sum_{1,12}(\text{MB})$$
$$\sum_{1,12} < \sum_{1,12}(\text{RB}) + 0.06 \cdot \sum_{1,12}(\text{MB})$$
$$\sum_{1,12} > 0.2 \cdot \sum_{1,12}(\text{MB}).$$

$\sum(\text{RB})$ indicates the DC signal if only the reference beam impinges on the photodiode and $\sum(\text{MB})$ if only the measurement beam impinges on the photodiode. The logical flow diagram of the alignment procedure is shown in Figure 9.11.

A very similar alignment procedure was already demonstrated and described in [56, 20]. In addition to this first demonstration of the alignment procedure, a hysteresis is implemented: e.g. the DWS signals are used for TM control if the contrast is bigger than 12%. If for some reason the contrast decreases, the DC signals are used, when the contrast gets smaller than 8%. This way continuous switching between DWS and DC signals is avoided. Furthermore, the PZTs used for these measurements have a wider actuation range (linear travel: $30\,\mu$m, tilt angle 5 mrad), which leads to the possibility to test the procedure with much worse initial dummy test mass alignment than before. The alignment procedure shown here was performed using the DMU and OBC simulators with asynchronous data transfer and downsampled data. As on LISA Pathfinder, the decision which alignment step is performed is on the OBC side.

Even with the wide range PZTs used here, the measurement beam cannot be removed from the photodiode. Therefore, it is not possible to test the spiral scan

Figure 9.11: Flow diagram of the alignment procedure.

of the test mass under realistic conditions. However, the condition under which the alignment procedure goes into the scan mode was changed from

$$\Sigma_{12} > \Sigma_{12}^{\mathrm{RB}} + 0.14 \cdot \Sigma_{12}^{\mathrm{MB}} \tag{9.5}$$

to

$$c_{12} > 40\,\%. \tag{9.6}$$

Σ_{12} is the DC signal on PD12, $\Sigma_{12}^{\mathrm{RB}}$ the DC signal of the reference beam and $\Sigma_{12}^{\mathrm{MB}}$ of the measurement beam respectively. c_{12} is the contrast on the photodiode. With this new threshold it was possible to test the part of the algorithm that includes the spiral scan. Figure 9.12 shows the angular movement of TM 2 during the alignment procedure. It starts with the spiral scan until a contrast of 40 % is reached. Then the DC signals are used to move the test mass. In the end the test mass is aligned to a zero of the DWS signals. For this test the input signals for the test mass alignment were switched from DC to DWS without resetting the loop filter. This results in a much smoother change from DC to DWS signals, where only small steps appear in the angle.

The optimal DWS alignment is not the same as the zero point of the DC alignment, which can clearly be seen in Figure 9.12. Both signals differ, since the photodiodes are not optimally centred with respect to the incoming beams. This has, to first order, no influence on the DWS but on the DC measurement.

Figure 9.13 shows that aligning to the DWS zero coincides with the point, where the contrast is maximised. Here the contrast is shown as a function of the distance of the beam with respect to the DWS zero. This fact is of importance for the logical

Figure 9.12: Experimental test of the alignment procedure.

flow of the alignment procedure, where a lot of decisions are based on the contrast in the interferometer.
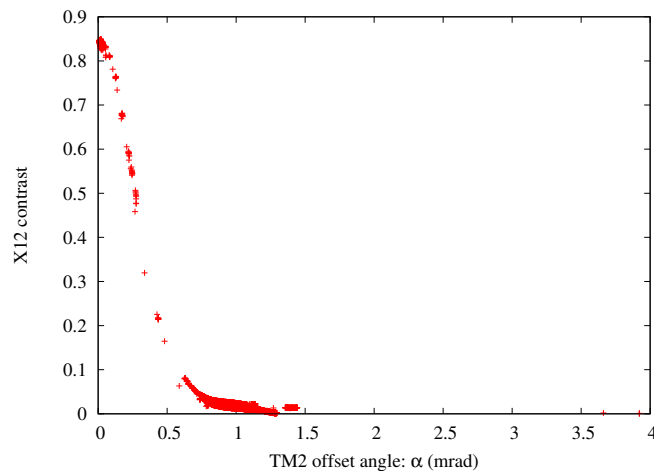
Figure 9.13: Dependency of contrast on the angular misalignment.

During the whole alignment, shown in Figure 9.12, almost all of the light of the measurement beam was impinging on the photodiode, while the worst misalignment during the scan was about 1.5 mrad. Therefore, a spiral scan of the test masses might not be necessary. However, it is implemented in the software as a fall-back option.

9.6 Conclusion

One subject of this thesis was the definition and the test of the interferometric data processing on-board LISA Pathfinder. The software as defined in Chapter 8 is part of the DMU flight software. In addition to breadboard experiments, features like suitable downsampling of the science data channels and averaging of redundant phasemeter output data were included in a hard- and software simulator.

This simulator includes separated programs simulating the DMU and OBC tasks of LISA Pathfinder. The data link between these programs is asynchronous like on-board LISA Pathfinder. In order to test also the dynamics of the test masses, 3-axes piezo actuated dummy mirrors were used instead of the test masses, and their performance was shown to be sufficient for the LISA Pathfinder sensitivity requirement.

The performance achieved with this simulator was shown to be below the required interferometer sensitivity. The alignment procedure was simulated with additional features compared to earlier experiments. It includes the task separation and interaction between the DMU and OBC and uses an asynchronous data link. Also the effect of channel failures on the science output data was investigated.

The most important outcome of the investigations presented here is the performance of the system with conditions as close as possible to those on-board LISA Pathfinder. The simulator can be used as a starting point for an end-to-end hardware model of the LTP. It includes all hardware that can practically be integrated in on-ground experiments when sufficient performance is needed. At the present status these are mainly breadboard components, but they will be replaced step by step by Engineering and Flight Models.

The results of the interferometer characterisation (Chapter 7) concerning the photodiode alignment and the simulation of channel failures (Section 9.1) can be used as reference for future Flight Model tests and are an important piece of information for the data processing and analysis of the real mission.

APPENDIX A

Electronics

For the weak light phase lock experiment, the development of suitable electronic circuits is one of the main challenges. The important schematics of the electronics used in this experiment can be found here. These are the low-pass filters using *composite amplifiers*, shown in Figure A.1 and the two photodiode pre-amplifiers under test. The discrete transistor design is shown in Figure A.2 and the design with the operational amplifier is shown in Figure A.3.
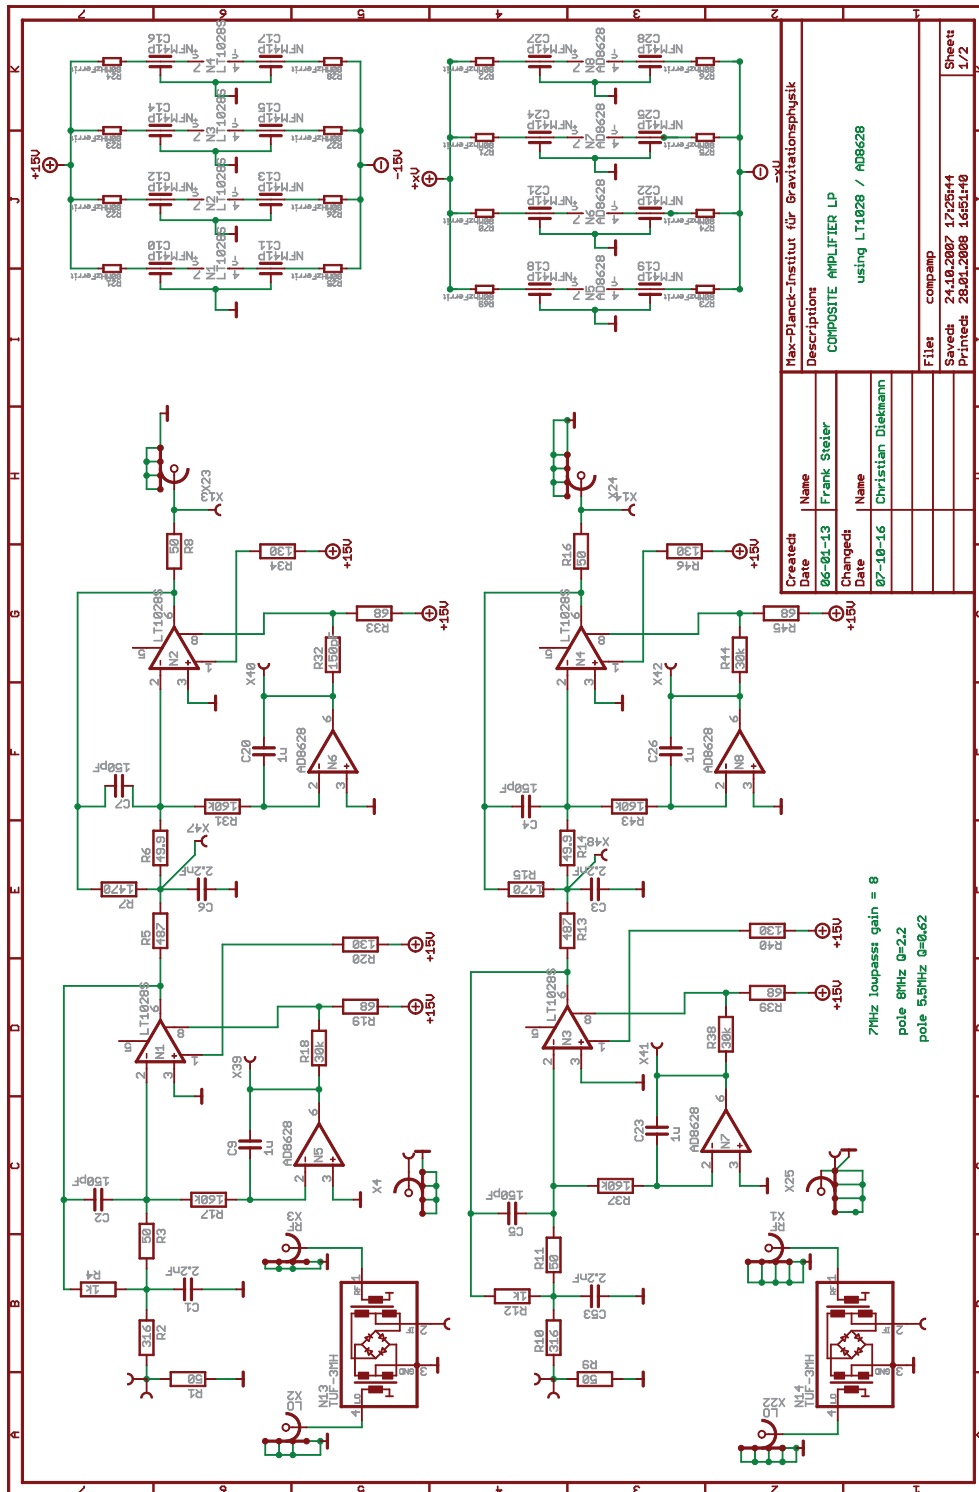
Figure A.1: Phase lock mixer electronics with composite amplifers as low-pass filter.
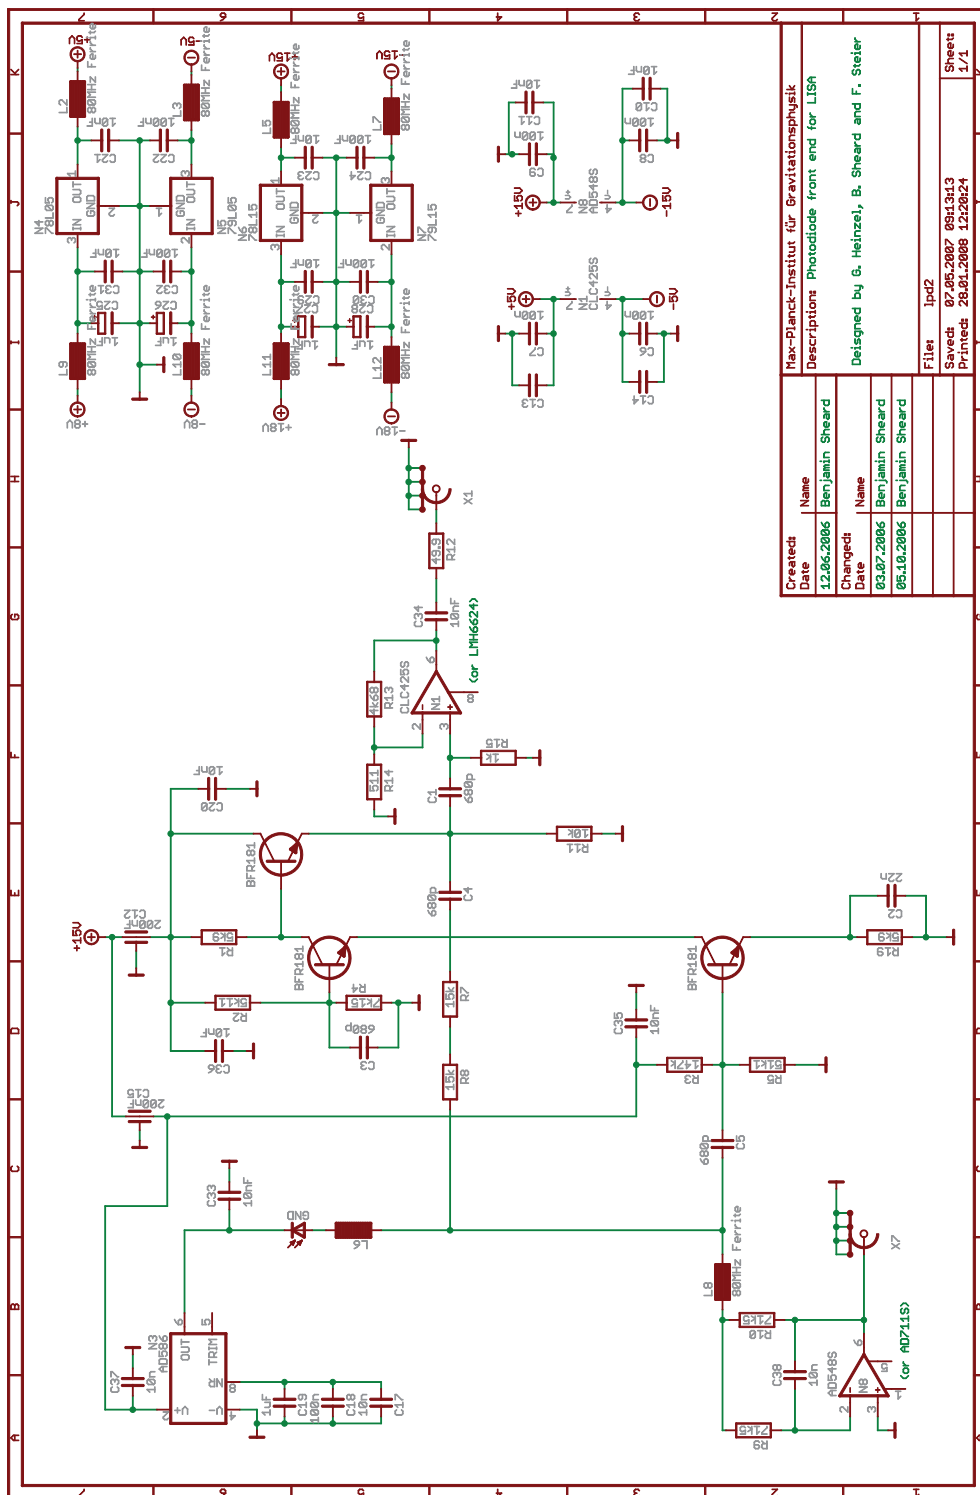
Figure A.2: Photodiode pre-amplifier using discrete transistor amplification stages.
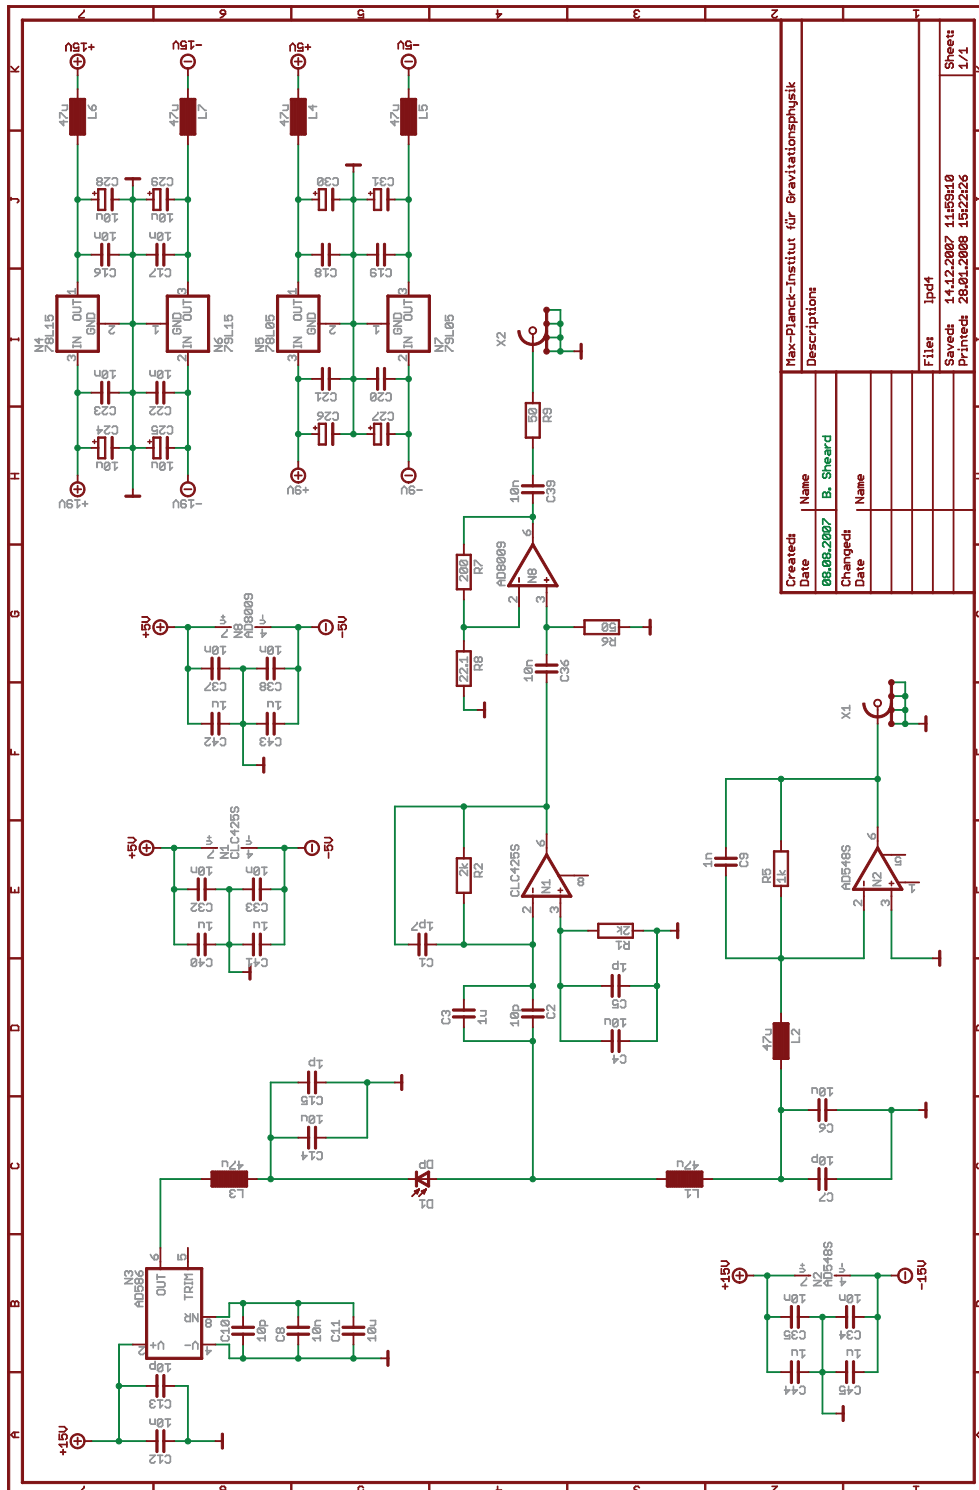
Figure A.3: Photodiode pre-amplifier using an operational amplifier.

Error flags and data quality

Here, all possible error flags that may happen during the data processing are shortly described. In addition the data quality classification is motivated and the classification criteria are summarised.

B.1 Error flags

Latch-Up ($L_{i,,k}$)

This flag is produced by the phasemeter. It is raised if an AD-converter, sampling the photodiode signal, temporarily fails. This is expected occasionally, since the ADC has to be reset, whenever high energy particles cause a latch-up condition in its CMOS circuitry. The flag affects all four channels of one quadrant photodiode and leads to temporary invalidity of the respective data.

Over- and Underrange ($o_{i,j,k}, u_{i,j,k}$)

These flags are raised if a channel is saturated or if there is almost no signal in one channel. This might happen in case of physical damage of a quadrant, but also during the on-orbit alignment of the test masses, when the beams are not centred on the quadrant photodiodes. In case of a physical damage, the data is invalid, whereas during the alignment procedure the channel might be noisy and the DC signals have to be used for the test mass alignment. In this particular case, the error check must be disabled by the OBC if the interferometer is in alignment mode, but not if it is in regular science mode.

Phasemeter errors ($\mathrm{PMC}_{,,k}$**,**$\mathrm{PML}_{,,k}$**,**$\mathrm{PMH}_{,,k}$**)**

These flags are set if the phasemeter detects an internal error. If the Look-Up table, used to calculate the sine and cosine coefficients via Fourier transform from the photodiode signals, is corrupted, $\mathrm{PML}_{,,k}$ is set. In case the configuration of the phasemeter is wrong, $\mathrm{PMC}_{,,k}$ is set. The $\mathrm{PMH}_{,,k}$ flag plays a special role, since it refers to environmental conditions inside the phasemeter like temperatures or supply voltages and is called *state of health flag*. In case of a configuration failure and a wrong Look-Up table, the phasemeter data is invalid. In case the *state of health* of the phasemeter is not correct, it might still deliver reasonable outputs.

Epsilon check ($\varepsilon_{i,j,}$**)**

The ε flag is raised if the difference between a channel and its redundant counterpart is bigger than a certain threshold. This might happen if the efficiency of one photodiode degrades slowly. The ε-flag is expected to be raised in case of a Latch-Up, Overrun, Transmission and Mil-Bus error or also for phasemeter errors. It is used just as an indicator and not to disable a channel for the data processing, since it is not possible to determine, which of the nominal or redundant channel shows an anomaly.

Delta Check ($\delta_{i,j,k}$**)**

A δ flag is raised if the difference between two consecutive samples of a signal exceeds a particular threshold. This is usually not expected, since the test mass motion is very slow, and therefore, all signals should change slowly. A raise of a δ flag is expected in case of a Latch-Up, Overrun, Transmission, Mil-Bus error and for phasemeter errors. In this case the data is usually invalid.

Transmission errors ($\mathrm{crc}_{i,,k}$**,**$\mathrm{O}_{i,,k}$**,**$\mathrm{M}_{i,,k}$**)**

Those flags are raised if a transmission error occurred during transfer of the signals from the phasemeter to the DMU or inside the phasemeter. If such a flag is raised, the effected channel is declared as invalid.

Ground commanded channel disable ($\mathrm{G}_{i,,k}$**)**

An additional flag is forseen in order to disable certain channels by manual selection. In case of performace degradation of certain channels, they can be disabled from ground and are then not used for the data processing.

The Phasemeter *state of health flag* ($\mathrm{PMH}_{,,k}$) and the ε-check ($\varepsilon_{i,j,}$) are for diagnostic purposes and will not disable any channel, while all other flags normally prevent the use of affected channels.

B.2 Data quality

Longitudinal data quality

nominal: The longitudinal data quality is defined as *nominal* if there are at least 4 independent quadrants for each of the interferometers (measurement (X1 or X12), the reference and the frequency interferometer) available. In case of temporary loss of quadrants in the measurement or the reference interferometer, the test mass position is expected to be calculated correctly, but small steps in the time series are expected if the number of used channels changes.

reduced: The data is defined as *reduced* if there are at least 3 independent quadrants in each measurement interferometer (X1 or X12), the reference and the frequency interferometer available. In this case the longitudinal data is still useful, even if data from a certain quadrant plus its redundant counterpart is missing. This is different for the alignment data. It may be even possible to get high performance longitudinal measurements from this data.

noisy: The data is marked as *noisy* if at least one, but less than 3, quadrants in the measurement (X1 or X12) and the reference interferometer is available. The data is not completely wrong, but cross-checks with the IRS data have to be performed. Even if all quadrants (4 nominal + 4 redundant) of the frequency interferometer are not available, the X1 and X12 longitudinal interferometer data is still useful and lead to correct pathlength calculations as long as the beam keeps its position on the photodiode. However, due to beam clipping and absence of frequency stabilised laser light, the data is noisy.

invalid: The data is marked as *invalid* if there is not enough information to calculate the longitudinal data. This situation appears if no quadrant from the reference or the respective measurement (X1 or X12) is available. In this case the data should never be used by the DFACS.

In case the OPD or laser frequency stabilisation loop must be reset, the data is also marked as *invalid* for a brief period, since steps are expected in the longitudinal data. If the system is well characterised, and the expected steps are small, it may be even possible to set the data quality to *noisy* in the case of a loop reset.

Angular data quality

The angular data quality does never depend on the reference and frequency interferometer measurement, since those do not contribute to the calculated signals. The quality is defined as follows:

nominal: If all quadrants (4 nominal + 4 redundant) in the measurement interferometers – X1 and X12 respectively – are used to calculate the angular alignment, the data is marked as *nominal*. This differs from the longitudinal data quality classification for practical reasons. Classifying the angular data quality leaves not much room for different states, so the distinction is done in smaller steps.

reduced: If an error occurs in one or more quadrants, but there are still 4 independent quadrants used for the calculation of the respective measurement interferometer, the data quality is *reduced*. The complete alignment information is still available, even if some quadrants are not used for the calculation. The data might be even adequate for low noise measurements. Temporary losses of certain quadrants might cause small steps in the alignment data.

noisy: If the information of one quadrant and its redundant counterpart is not available, the angular data is expected to be *noisy*. The data might be useful for the DFACS, but cross-checks with the IRS data are necessary in order to correctly control the test masses.

invalid: The data is marked as *invalid* if there is not enough information to calculate the angular data. This situation appears if less than 3 quadrants and their redundant counterparts are available. In this case the data should never be used by the DFACS.

Bibliography

[1] ESA. `http://lisa.esa.int/science-e/www/area/index.cfm?fareaid=27`. Sep. 2008.

[2] S. Dhurandhar, K. Rajesh Nayak, S. Koshti, and J.-Y. Vinet. Fundamentals of the LISA stable flight formation. *Class. Quantum Grav.*, 22:481–487, 2005.

[3] K. Danzmann and the LISA Study team. *LISA – Laser Interferometer Space Antenna: A Cornerstone Mission for the Observation of Gravitational Waves.* ESA-SCI(2000)11, 2000.

[4] W. J. Weber, A. Cavalleri, R. Dolesi, G. Fontana, M. Hueller, and S. Vitale. Position sensors for LISA drag-free control. *Class. Quantum Grav.*, 19:1751.–1756, 2002.

[5] D. Weise. LISA-ASD-RS-3300 Telescope Subsystem Specification 0.1. Technical report, EADS Astrium Deutschland, 2008.

[6] D. Weise. LISA-ASD-RS-3100 Optical Bench Subsystem Specification 0.2. Technical report, EADS Astrium Deutschland, 2008.

[7] M. Vallisneri, J. Crowder, and M. Tinto. Sensitivity and parameter-estimation precision for alternate LISA configurations. *Class. Quantum Grav.*, 25:065005 (17pp), 2008.

[8] G. Heinzel, C. Braxmaier, K. Danzmann, P. Gath, J. Hough, O. Jennrich, U. Johann, A. Rüdiger, M. Sallusti, and H. Schulte. LISA interferometry: recent developments. *Class. Quantum Grav.*, 23:119–124, 2006.

[9] E. Ellife, J. Bogenstahl, A. Deshpande, J. Hough, C. Killow, S. Reid, D. Robertson, S. Rowan, H. Ward, and G. Cagnoli. Hydroxide-catalysis bonding for stable optical systems for space. *Class. Quantum Grav.*, 22:257–267, 2005.

[10] C. Killow. *Interferometry Developments for Spaceborne Gravitational Wave Detectors.* PhD thesis, Department of Physics and Astronomy, University of Glasgow, 2006.

[11] P. McNamara. Weak light phase locking for LISA. *Class. Quantum Grav.*, 22:243–247, 2005.

[12] D. Robertson and J. Hough. Interferometry for LISA. *Class. Quantum Grav.*, 13:271–277, 1996.

[13] W. Folkner, F. Hechler, T. Sweetser, M. Vincent, and P. Bender. LISA orbit selection and stability. *Class. Quantum Grav.*, 14:1405–1410, 1997.

[14] D. Shaddock, B. Ware, P. Halverson, R. Spero, and B. Klipstein. Overview of the LISA Phasemeter. *AIP Conf. Proc.*, 873:654–660, 2006.

[15] S. Pollack and R. Stebbins. Demonstration of the Zero-Crossing Phasemeter with a LISA Test-bed Interferometer. *Class. Quantum Grav.*, 23:4189–4200, 2006.

[16] R. Cruz, J. Thorpe, A. Preston, R. Delgadillo, M. Hartman, S. Mitryk, A. Worley, G. Boothe, S. Guntaka, S. Klimenko, D. Tanner, and G. Mueller. The LISA benchtop simulator at the University of Florida. *Class. Quantum Grav.*, 23:751–760, 2006.

[17] V. Wand, F. Guzmán, G. Heinzel, and K. Danzmann. LISA phasemeter development. *AIP Conf. Proc.*, 873:689–696, 2006.

[18] G. Heinzel, V. Wand, A. García Marín, F. Steier, F. Guzmán, C. Killow, D. Robertson, H. Ward, and C. Braxmaier. S2-AEI-TN-3028 Investigation of noise sources in the LTP interferometer. Technical report, Albert-Einstein-Institut Hannover, 2005.

[19] V. Wand. *Interferometry at low Frequencies: Optical Phase Measurement for LISA and LISA Pathfinder.* PhD thesis, Institut für Gravitationsphysik, Universität Hannover, 2007.

[20] A. García Marín. *Minimisation of optical pathlength noise for the detection of gravitational waves with the spaceborne laser interferometer LISA and LISA Pathfinder.* PhD thesis, Institut für Gravitationsphysik, Universität Hannover, 2007.

[21] D. Robertson, C. Killow, H. Ward, J. Hough, G. Heinzel, A. García, V. Wand, U. Johann, and Claus Braxmaier. LTP interferometer - noise sources and performance. *Class. Quantum Grav.*, 22:155–163, 2005.

[22] B. Sheard. *Arm locking for space-based gravitational wave detectors and optomechanical effects in interferometers.* PhD thesis, The Australian National University, 2005.

[23] R. Mancini. *Op Amps For Everyone.* Texas Instruments, Aug. 2002.

[24] J. Williams. Composite Amplifiers – Application Note. Technical report, Linear Technologies, 1986.

[25] C. Diekmann. Phasenstabilisierung und -auslesung für LISA. Master's thesis, Institut für Gravitationsphysik Universität Hannover, 2008.

[26] PerkinElmer. `http://optoelectronics.perkinelmer.com/catalog/Product.aspx?ProductID=C30619G`. Sep. 2008.

[27] Hamamatsu. `http://sales.hamamatsu.com/index.php?id=131919560`. Sep. 2008.

[28] Onset. `http://www.gpd-ir.com/products.htm`. Sep. 2008.

[29] G. Heinzel. LISA related activities in Hannover. *Talk at WG2 meeting, GSFC*, June 2006.

[30] G. Heinzel, F. Steier, R. Fleddermann, B. Sheard, and K. Danzmann. Components for the LISA local interferometry. *AIP Conf. Proc.*, 873:291–296.

[31] G. Heinzel. *Advanced optical techniques for laser-interferometric gravitational-wave detectors.* PhD thesis, Max-Planck-Institut für Quantenoptik Garching, 1999.

[32] Y. Sasaki. Long-Length Low-Loss Polarisation-Maintaining Fibers. *Journal of Lightwave Tech.*, LT-5(9), Sep. 1987.

[33] M. Berger. LISA-ASD-DD-3001 Payload Preliminary Design Description 2.0. Technical report, EADS Astrium Deutschland, 2008.

[34] P. Gath. LISA-ASD-TN-5001 Requirement Breakdown 2.0. Technical report, EADS Astrium Deutschland, 2008.

[35] R. Fleddermann. Komponentencharakterisierung für LISA – Rauscharme Spannungsreferenzen und Reziprozität einer Glasfaser. Master's thesis, Institut für Gravitationsphysik Universität Hannover, 2006.

[36] R. Schilling. *OptoCad – A Fortran 95 module for tracing Gaussian beams through an optical set-up – Version 0.82d.* Garching, 2005.

[37] J. Bogenstahl. S2-UGL-PL-3011 FIOS Construction Plan 1.0. Technical report, IGR, University of Glasgow, 2007.

[38] G. Heinzel, V. Wand, A. García, O. Jennrich, C. Braxmaier, D. Robertson, K. Middleton, D. Hoyland, A. Rüdiger, R. Schilling, U. Johann, and K. Danzmann. The LTP interferometer and phasemeter. *Class. Quantum Grav.*, 21:581–587, 2004.

[39] V. Wand, J. Bogenstahl, C. Braxmaier, K. Danzmann, A. García, F. Guzmán, G. Heinzel, J. Hough, O. Jennrich, C. Killow, D. Robertson, Z. Zodnik, F. Steier, and H. Ward. Noise sources in the LTP heterodyne interferometer. *Class. Quantum Grav.*, 23:159–167, 2006.

[40] Schott. `http://www.schott.com/advanced_optics/german/our_products/zerodur/zerodur.html`. Sep. 2008.

[41] Electrolube. `www.farnell.com/datasheets/94418.pdf`. Sep. 2008.

[42] K. Middleton, M. Caldwell, R. Edeson, C. Neeson, G. Munro, C. Eley, M. Sandford, C. Braxmaier, U. Johann, C. Killow, D. Robertson, H. Ward, J. Hough, G. Heinzel, and V. Wand. Prototype optical bench instrument in the interferometer for the LISA-Pathfinder space mission. *Opt. Engineering*, 45(12), 2006.

[43] Private communication with E. Fitzsimons. *IGR, University of Glasgow*, June 2008.

[44] S. Anza et. al. The LTP experiment on the LISA Pathfinder mission. *Class. Quantum Grav.*, 22:125–138, 2005.

[45] W. Fichter, P. Gath, S. Vitale, and D. Bortoluzzi. LISA Pathfinder drag-free control and system implications. *Class. Quantum Grav.*, 22:139–148, 2005.

[46] P. McNamara. LISA Pathfinder. *Talk at 7th International LISA Symposium, Barcelona*, June 2008.

[47] R. Gerndt, W. Fichter, N. Brandt, D. Gerardi, F. Montemurro, A. Schleicher, T. Ziegler, and U. Johann. LISA technology package system design and operations. *AIP Conf. Proc.*, 873:668–674, 2006.

[48] G. Heinzel, V. Wand, and A. García Marín. S2-AEI-TN-3017 LISA Pathfinder Optical Bench Engineering Model: functional and performance tests – test report. Technical report, Albert-Einstein-Institut Hannover, 2004.

[49] E. Morrison, B. Meers, D. Robertson, and H. Ward. Experimental demonstration of an automatic alignment system for optical interferometers. *Appl. Opt.*, 33(22):5037–5041, 2004.

[50] E. Morrison, B. Meers, D. Robertson, and H. Ward. Automatic alignment of optical interferometers. *Appl. Opt.*, 33(22):5041–5049, 2004.

[51] V. Wand, J. Bogenstahl, C. Braxmaier, K. Danzmann, A. García, F. Guzmán, G. Heinzel, J. Hough, O. Jennrich, C. Killow, D. Robertson, Z. Sodnik, F. Steier, and H. Ward. Noise sources in the LTP heterodyne interferometer. *Class. Quantum Grav.*, 23:159–167, 2006.

[52] F. Guzmán Cervantes, F. Steier, G. Wanner, G. Heinzel, and K.Danzmann. Subtraction of test mass angular noise in the LISA technology package interferometer. *Appl. Phys. B*, 90:395–400, 2008.

[53] G. Heinzel, C. Braxmaier, R. Schilling, A. Rüdiger, D. Robertson, M. te Plate, V. Wand, K. Arai, U. Johann, and K. Danzmann. Interferometry for the LISA technology package (LTP) aboard SMART-2. *Class. Quantum Grav.*, 20:153–161, 2003.

[54] C. Killow, J. Bogenstahl, F. Guzmán Cervantes, M. Perreur-Lloyd, D. Robertson, F. Steier, and H. Ward. Construction of the LTP Optical Bench Interferometer. *AIP Conf. Proc.*, 873:297–303, 2006.

[55] U. Denskat. S2-ASD-RS-3018 Phasemeter Processing and Laser Control Specification 3.0. Technical report, EADS Astrium Deutschland, 2008.

[56] A. García Marín, V. Wand, F. Steier, F. Guzmán Cervantes, J. Bogenstahl, O. Jennrich, G. Heinzel, and K. Danzmann. On-orbit alignment and diagnostics for the LISA technology package. *Class. Quantum Grav.*, 23:133–140, 2006.

[57] G. Heinzel. S2-AEI-TN-3011 SMART2 LTP phasemeter. Technical report, Albert-Einstein-Institut Hannover, 2003.

[58] D. Hoyland. S2-UBI-ICD-3001 Phasemeter Electrical ICD 3.01. Technical report, University of Birmingham, 2006.

[59] V. Wand. Interferometrie und Phasenauslesung für die wissenschaftliche Weltraummission SMART2. Master's thesis, Institut für Gravitationsphysik, Universität Hannover, 2003.

[60] D. Gerardi. S2-ASD-TN-2046 Interferometer initial acquisition design and analysis 1.0. Technical report, EADS Astrium Deutschland, 2008.

[61] A. García. S2-AEI-TN-3053 Considerations on DMU Synchronization. Technical report, Albert-Einstein-Institut, 2008.

[62] M. Unser. Sampling - 50 Years After Shannon. *Proceedings of the IEEE*, 88(4), 2000.

[63] P. Vaidyanathan. Generalization of the Sampling theorem: Seven Decades after Nyquist. *IEEE transactions on Circuits and and Applications*, 48(9), 2001.

[64] G. Heinzel, F. Steier, and F. Guzmán Cervantes. S2-AEI-TN-3039 Antialiasing filters for data from DMU to OBC. Technical report, Albert-Einstein-Institut Hannover, 2007.

[65] G. Heinzel. *LISO – Program for Linear Simulation and Optimisation of analog electronic circuits.* User Manual, Max-Planck-Institut für Quantenoptik Garching, 1999.

[66] F. Steier, G. Heinzel, F. Guzmán Cervantes, and A. García Marín. S2-AEI-TN-3021 LTP extramode definition. Technical report, Albert-Einstein-Institut Hannover, 2007.

[67] F. Cirillo. S2-ASD-TN-2013 LISA Pathfinder Simulator User Manual 2.1. Technical report, EADS Astrium Deutschland, 2008.

[68] G. Heinzel. S2-AEI-TN-3034 Generation of random time series with prescribed spectra. Technical report, Albert-Einstein-Institut Hannover, 2007.

Acknowledgements

I enjoyed my work on the LISA and LISA Pathfinder project. My colleagues, the environment in the institute and my work made my time at the AEI a valuable experience. I thank Prof. Karsten Danzmann for this opportunity and his inspiration.

I learned very much from Gerhard Heinzel who always had the right hint when an experiment had to be improved. His elegant way of doing physics and his clear view were very helpful.

I want to thank all people of the LISA group: Antonio García Marín and Vinzenz Wand brought me in touch with the LISA Pathfinder experiment and let me benefit from their experience. Felipe Guzmán Cervantes and I complemented one another very well during our work on the LTP experiment. Thanks also to Roland Fleddermann, Ben Sheard and Christian Diekmann, who worked on the LISA experiments with me. I also enjoyed the work with Jens Reiche on the 'other' aspects of the LISA Pathfinder project. Boris Hage and Michael Tröbs were always open-minded for discussions even if the topic was not related to their own work. Thanks also to all other LISA and LISA Pathfinder group members and all other colleagues at the AEI for the pleasent time not only at work, but also during the leasure time.

During my time in Glasgow, I found very obliged people, who made my stay a nice residence. In particular Harry Ward, Dave Robertson and Prof. Jim Hough accommodated me very friendly at the institute. Ik Siong Heng and Christian Killow as well as Mike Perreur-Lloyd made my work and social life in Glasgow to a vitally experience. Johanna Bogenstahl was very helpful especially for bonding the interferometers, but also to keep the groups in Hannover and Glasgow together.

I thank all the members of the European Graduate College for *Interference and Quantum Applications* for the atmosphere in this college and especially Prof. Eberhard Tiemann for making my stay in Glasgow possible.

I also want to thank our colleagues at EADS Astrium Friedrichshafen for the support with technical information and nice graphs that can be found in this thesis.

For proof-reading this thesis, I thank Gerhard, Anneke Monsky, Antonio, Martin Hewitson, Oliver Jennrich, Roland, Ben and Christian. I am sure, it improved the quality a lot.

Special thanks go to my family for the unrestricted support. Its importance is invaluable. Last but not least, thank you Anneke for your crucial influence on my mind and the quality of my life and my work.

Curriculum Vitae

Frank Steier
Glockseestr. 3C
30169 Hannover

born on 1st of August in Dorsten, Germany

1985 – 1989	Grundschule & Orientierungsstufe
1991 – 1998	Abitur, Hölty-Gymnasium Wunstorf
1998 – 1999	Zivildienst
1999 – 2004	Diploma studies in physics at the Leibniz Universität Hannover
06/2004	Physik-Diplom, Leibniz Universität Hannover
2004 – to date	Scientific assistant at the Max-Planck Institut für Gravitationsphysik (Albert-Einstein-Institut) Hannover, Doctoral studies in physics at the Leibniz Universität Hannover
2004 – 2007	Member of the European Graduate College for Interference and Quantum Applications
03 – 08/2005	Doctoral studies in physics at the University of Glasgow
2007 – to date	Member of the International Max-Planck Research School

- G. Heinzel, C. Braxmaier, M. Caldwell, K. Danzmann, F. Draaisma, A. García, J. Hough, O. Jennrich, U. Johann, C. Killow, K. Middleton, M. te. Plate, D. Robertson, A. Rüdiger, R. Schilling, F. Steier, V. Wand and H. Ward, *Successful testing of the LISA Technology Package (LTP) interferometer engineering model*, Class. Quantum Grav. 22 (2005) p.149–p.154

- A. García Marín, G. Heinzel, R. Schilling, A. Rüdiger, V. Wand, F. Steier, F. Guzmán Cervantes, A. Weidner, O. Jennrich, F. Meca Meca and K. Danzmann, *Phase locking to a LISA arm: first results on a hardware model*, Class. Quantum Grav. 22 (2005) p.235–p.242

- G. Heinzel, F. Steier, R. Fleddermann, B. Sheard, and K. Danzmann, *Components for the LISA local interferometry*, AIP Conf. Proc. 873 (2006), p.291–p.296

- C. Killow, J. Bogenstahl, F. Guzmán Cervantes, M. Perreur-Lloyd, D. Robertson, F. Steier, and H. Ward, *Construction of the LTP Optical Bench Interferometer*, AIP Conf. Proc. 873 (2006), p.297–p.303

- A. García Marín, J. Bogenstahl, F. Guzmán Cervantes, F. Steier, J. Reiche, S. Skorupka, V. Wand, M. Nofrarias, J. Sanjuan, O. Jennrich, G. Heinzel, and K. Danzmann, *Interferometric characterization of the optical window for LISA Pathfinder and LISA*, AIP Conf. Proc. 873 (2006), p.344–p.348

- A. García Marín, V. Wand, F. Steier, F. Guzmán Cervantes, J. Bogenstahl, O. Jennrich, G. Heinzel and K. Danzmann, *On-orbit alignment and diagnostics for the LISA Technology Package*, Class. Quantum Grav. 23 (2006) p.133–p.140

- G. Heinzel, J. Bogenstahl, C. Braxmaier, K. Danzmann, A. García, F. Guzmán, J. Hough, D. Hoyland, O. Jennrich, C. Killow, D. Robertson, Z. Sodnik, F.

Steier, H. Ward and V. Wand, *Interferometry for the LISA technology package LTP: an update*, J. Phys.: Conf. Ser. 32 (2006) p.132–p.136

- V. Wand, J. Bogenstahl, C. Braxmaier, K. Danzmann, A. García, F. Guzmán, G. Heinzel, J. Hough, O. Jennrich, C. Killow, D. Robertson, Z. Sodnik, F. Steier and H. Ward, *Noise sources in the LTP heterodyne interferometer* Class. Quantum Grav. 23 (2006) p.159–p.167

- F. Guzmán Cervantes, G. Heinzel, A. García Marín, V. Wand, F. Steier, O. Jennrich and K. Danzmann, *Real-time phase-front detector for heterodyne interferometers*, Appl. Opt. (Vol. 46, No. 21) (2007), p.4541–p.4548

- F. Guzmán Cervantes, F. Steier, G. Wanner, G. Heinzel, K. Danzmann, *Subtraction of test mass angular noise in the LISA technology package interferometer*, Appl. Phys. B 90 (2008) p.395–p.400

- R. Fleddermann, M. Tröbs, F. Steier, G. Heinzel and K. Danzmann, *Intrinsic noise and temperature coefficients of selected voltage references*, accepted for publication in IEEExplore

- F. Steier et al, *The LISA Pathfinder interferometry: experiences for the LISA local readout*, invited for publication in the Proceedings of the 7th LISA Symposium in Class. Quantum Grav.

- R. Fleddermann et al, *Measurements of the nonreciprocal phase noise of a single mode polarization maintaining fiber*, invited for publication in the Proceedings of the 7th LISA Symposium in J. Phys. Conf. Ser.

- C. Diekmann et al, *Analog phase lock between two lasers at LISA power level*, invited for publication in the Proceedings of the 7th LISA Symposium in J. Phys. Conf. Ser.