

Laser Frequency Stabilisation and Interferometer Path Length Differences during the LISA Pathfinder Satellite Mission

Von der Fakultät für Mathematik und Physik
der Gottfried Wilhelm Leibniz Universität Hannover

zur Erlangung des akademischen Grades
Doktorin der Naturwissenschaften
Dr. rer. nat

genehmigte Dissertation von

MSc Sarah Katharina Paczkowski

2021

Referent: Apl. Prof. Dr. Gerhard Heinzl
(Institut für Gravitationsphysik, Leibniz Universität Hannover)

Korreferent: Apl. Prof. Dr. Benno Willke
(Institut für Gravitationsphysik, Leibniz Universität Hannover)

Korreferent: Prof. William Joseph Weber
(Dipartimento di Fisica, Università di Trento)

Tag der Promotion: 28.08.2020

In der vorliegenden Arbeit wurden gegenüber der am 18.06.2020 bei der Fakultät für Mathematik und Physik der Leibniz Universität Hannover eingereichten Version einige Tippfehler korrigiert. Vereinzelt wurden Textstellen und Abbildungen im Sinne eines besseren Verständnisses leicht überarbeitet.

Abstract

This thesis focusses on laser frequency stabilisation and interferometer path length differences on LISA Pathfinder (LPF). LPF was a satellite mission, in operation from December 2015 until July 2017, to demonstrate key technologies for the future spaceborne gravitational wave observatory, the Laser Interferometer Space Antenna (LISA). It successfully showed that the undesired disturbances, the so-called residual acceleration noise, of a pair of free-falling test masses (TMs) could be limited to less than $2 \text{ fm/s}^2/\sqrt{\text{Hz}}$ at mHz Fourier frequencies and thus paved the way for LISA.

This level of residual acceleration noise could only be reached by the interaction of several key subsystems. One of these, the Optical Metrology System (OMS), provided the LPF science measurement: a heterodyne interferometry readout of the relative positions of the free-floating TMs. The OMS showed excellent performance over the mission and is studied here in detail to make the best use of this unique measurement data and to learn as much as possible for future interferometer development.

The OMS is subject to several noise sources. This thesis focusses on one of these: laser frequency fluctuations. In the design phase of the OMS, a laser frequency stabilisation technique via a dedicated interferometer measurement and a nested control loop was developed. In this thesis, we show the in-flight results of the planned loop characterisation experiments and noise measurements. We prove that the stabilisation worked as expected from ground tests and was reliable over the mission duration. We also identified periods of slightly increased laser frequency fluctuations whose origin could not yet be identified. This analysis was restricted by the limited number of laser telemetry channels and their low sampling frequency.

The coupling of laser frequency noise depends on the optical path length difference between the measurement and the reference beam. Two experiments optimised for determining this quantity have been designed and executed on LPF and are analysed here. The evidence collected in the offset experiment allows us to associate a change in measured path length mismatch to a commanded offset at a 3σ uncertainty level, both in direction and amplitude. In addition, we confirmed that the estimated path length mismatch is independent of the laser frequency modulation amplitude and frequency. In general, we find the path length mismatch is only a few hundred μm , which is a sign of excellent integration. We also provide an example where these experiments were used as a means to measure absolute distances on LPF.

During these experiments, we also observed short term path length mismatch variations which are not believed to be caused by a true TM motion, as well as spurious signals in the angular measurements of the OMS. Within the scope of this work, a definite reason could not be found for either of these two observations but different cross-coupling hypotheses to explain the signals in the angular measurements could also not be fully rejected. However, we have found no reason to believe that these angular signals could be an indication of a mechanism that impacts the longitudinal measurements and thus adds a systematic error to the path length mismatch numbers reported.

Keywords: LISA Pathfinder, space interferometry, data analysis

Contents

List of Figures	5
List of Acronyms	8
List of Terms	11
1 Interferometric gravitational wave detection in space	13
1.1 Gravitational waves and their detection	13
1.1.1 Definition and properties of gravitational waves	13
1.1.2 Detection of gravitational waves with laser interferometry	14
1.1.3 Sources of gravitational waves	14
1.1.4 LISA concept	16
1.2 The LISA Pathfinder mission	17
1.2.1 LPF subsystems	17
1.2.2 Development and testing	20
1.2.3 LPF mission operations	20
1.2.4 LPF residual acceleration measurement results	21
1.3 This thesis	21
2 Selected aspects of the optical metrology system	23
2.1 Overview	23
2.1.1 Hardware	23
2.1.2 Data processing	27
2.1.3 Noise sources in the OMS	31
2.2 Coupling of laser frequency noise into the phase measurement	31
2.3 The path length difference	32
2.4 The laser frequency stabilisation control loop	34
2.5 Laser frequency noise requirement	35
2.6 The in-flight laser frequency modulation experiments	36
3 Control loop characterisation	39
3.1 Design of loop characterisation measurements	39
3.1.1 Overview of the measurements performed	40
3.2 Analysis procedure for the frequency control loop characterisation measurements	41
3.2.1 Timing correction procedure	42
3.2.2 Transfer function estimation and their errors	43
3.2.3 Controller transfer function measurement procedure	46

3.2.4	Actuator characterisation procedure	46
3.2.5	OLTF procedure	47
3.3	Results of control loop characterisation measurements	47
3.3.1	Controller transfer function measurements	47
3.3.2	Actuator gains and delays	49
3.3.3	OLTF estimate	52
4	Laser frequency noise characterisation	57
4.1	Impact of the reference interferometer subtraction	57
4.2	Laser frequency fluctuations over the course of the mission	59
4.2.1	Times of non-stationary behaviour at lower frequencies	64
4.3	Investigation of changing noise levels	66
4.3.1	Hypothesis: aliasing	67
4.3.2	Hypothesis: correlation to laser settings	68
4.3.3	Hypothesis: overall satellite temperature influence	70
4.3.4	Hypothesis: back scatter into the laser at certain test mass angles	71
4.4	Stabilised and free-running laser frequency noise measurements	73
5	Interferometer path length difference measurements	77
5.1	Design of path length difference experiments	77
5.1.1	First path length mismatch experiment: offset	77
5.1.2	Second path length mismatch experiment: varying modulation amplitudes	79
5.2	Analysis procedure for the path length difference experiments	80
5.3	Path length difference estimates	83
5.3.1	Results from offset experiment	84
5.3.2	Results from amplitude experiment	86
5.3.3	A comment on the path length difference in the reference interferometer	86
5.3.4	Results over the mission	87
5.3.5	Explaining the changes in path length difference	89
5.4	Variations of path length mismatch estimates during individual laser frequency modulations	93
5.4.1	Description of the observation	93
5.4.2	First checks	94
6	Investigation of DWS signals during frequency modulations	103
6.1	Description, relevance and first checks	103
6.1.1	Observation	103
6.1.2	Relevance	105
6.1.3	The structure of the chapter	106
6.1.4	First checks	108
6.2	Hypothesis: a true phase difference between the quadrants	110
6.3	Hypothesis: related to beam tilts	112
6.3.1	In-flight investigations	112
6.3.2	Follow up of beam tilt in the laboratory	114
6.3.3	Misalignments - laboratory	116
6.4	Hypothesis: power related	121
6.4.1	The observed signals	123
6.4.2	Polarisation as a possible explanation for the observed power modulation	124
6.4.3	Checking the coupling from power to DWS	127

6.4.4	Investigation of coupling of laser frequency modulation to power in the laboratory	128
6.5	Hypothesis: related to angular offsets	130
6.5.1	Checks of in-flight data	131
6.5.2	Checks in laboratory	133
6.6	Hypothesis: related to cross-coupling in the phasemeter	134
6.6.1	Possible signal leakage from longitudinal to DWS	134
6.6.2	Electrical/PD cross-talk	135
6.6.3	Cross-coupling related to the heterodyne frequency	136
6.6.4	Investigation of cross-coupling in the laboratory	136
6.7	Further ideas for investigations	138
6.8	Conclusions	139
7	Summary and outlook	141
7.1	Summary of the thesis	141
7.2	Contribution of laser frequency noise to the total LPF OMS noise	143
7.3	Discussion: Limits of the laser frequency stabilisation scheme	145
7.4	Outlook: Applications of the laser frequency stabilisation scheme	147
	Bibliography	148
	Appendices	154
A	Details on parameter estimation	155
A.1	A brief comment on the Nelder-Mead simplex method	155
A.2	A brief summary of Markov Chain Monte-Carlo (MCMC) uncertainty estimation	156
A.2.1	From model to posterior probability density function	156
A.2.2	Definition of the term MCMC	157
A.2.3	Important technical aspects	158
A.3	MCMC uncertainty estimation results and residuals for the actuator parameters	158
A.3.1	FM test campaign experiment	159
A.3.2	Loop characterisation experiment 1	162
A.3.3	Loop characterisation experiment 2	164
A.3.4	Loop characterisation experiment 3	166
A.3.5	Loop characterisation experiment 4	168
A.3.6	Loop characterisation experiment 5	170
B	List of OMS telemetry channels	173
C	Free-Fall Experiment	183
C.1	Concept of the experiment	183
C.2	Preparations for operations	185
C.2.1	The local simulation environment	186
C.3	The experiment during operations	188
	Acknowledgements	189
	List of Publications	191
	Curriculum Vitae	203

List of Figures

1.1	Polarisation states of gravitational waves.	14
1.2	Principle of gravitational wave detection using interferometry.	15
1.3	Sketch of the LISA orbit.	16
1.4	Core of LISA Pathfinder.	18
1.5	Sketch of the DFACS control strategy.	19
2.1	LPF OMS overview.	24
2.2	Four interferometers on the optical bench.	25
2.3	LPF coordinate system.	26
2.4	Diagram of OPD control loop.	26
2.5	Sketch of a quadrant photodiode.	26
2.6	Sketch of the DWS working principle.	29
2.7	Sketch of the path length mismatch.	34
2.8	Example: coupling of laser frequency modulation to o12.	34
2.9	Diagram of laser frequency control loop.	35
3.1	ASD of Ψ_F during experiment 3.	45
3.2	Comparison of OLTF at the precise modulation frequencies and the best bins.	46
3.3	Fast frequency controller transfer function results.	48
3.4	Slow frequency controller transfer function results.	48
3.5	Laser output power and the pump current for different temperatures.	50
3.6	Fast and slow actuator response.	52
3.7	Measured OLTF of experiment 1.	53
3.8	Measured OLTF of experiment 3.	53
3.9	Measured OLTF of experiment 4.	54
3.10	Measured OLTF of experiment 5.	54
3.11	Comparison of the OLTF from all experiments and ground tests.	55
4.1	Time series data on reference interferometer subtraction.	58
4.2	Spectra on reference interferometer subtraction.	58
4.3	Zoom into Ψ_F noise over mission.	60
4.4	Histogram of Ψ_F noise over mission.	61
4.5	Zoom into the laser frequency noise average time series.	62
4.6	Zoom into the laser frequency noise time series.	62
4.7	Spectrogram during a transition from the lower level to the upper level.	63
4.8	Spectrogram during a transition from the upper level to the lower level.	63
4.9	Histogram of the duration of Ψ_F in the upper noise level.	64

4.10	Low-passed laser frequency fluctuations.	65
4.11	Histogram of the time between two glitches.	66
4.12	Laser frequency noise recorded at 10 Hz and 100 Hz.	67
4.13	Laser frequency noise and the RLU pump current.	68
4.14	Ψ_F noise level versus the mean pump current.	69
4.15	Ψ_F noise level versus the mean RLU output power.	70
4.16	Ψ_F noise level versus the mean RLU baseplate temperature.	70
4.17	Laser frequency noise compared to the OB temperature during cool-down.	71
4.18	Laser frequency noise at angular positions.	72
4.19	Open-loop vs. closed-loop laser frequency noise.	74
4.20	Laser frequency noise and noise suppression function.	75
5.1	Concept of first dedicated path length mismatch experiment.	78
5.2	Flow chart of the path length mismatch procedure.	81
5.3	o12 measurement during the path length mismatch offset experiment.	84
5.4	Sketch of the path length mismatch offset experiment.	84
5.5	Results of the first path length mismatch experiment.	85
5.6	Results of the second path length mismatch experiment.	87
5.7	Hypothesis investigated to explain the change in path length mismatch.	90
5.8	Estimated true test mass position from June 2016 to January 2017.	92
5.9	OMS and GRS measurement at the time of a commanded offset.	93
5.10	Illustration of splitting the modulation segments into shorter segments.	94
5.11	Estimated path length mismatch from experiment 4 using shorter segments.	95
5.12	Estimated path length mismatch from experiment 6 using shorter segments.	95
5.13	GRS measurements of both TMs during experiment 4.	96
5.14	Estimated path length mismatch for shorter segments of experiment 4.	97
5.15	Amplitude spectrum of the residuals of the linear fit to the o12 measurement.	97
5.16	Estimated path length mismatch for different data segment durations.	98
5.17	Estimated path length mismatch for different data segment overlap.	99
5.18	ASD of the o12 measurement during experiment 4.	99
5.19	ASD of the o12 measurement with and without the band-reject filter.	100
5.20	Estimated path length mismatch with and without the band-reject filter.	100
6.1	Comparison of the o12 and the DWS signals during exp. 1, 3 and 5.	104
6.2	Diagram of the hypotheses investigated to explain the DWS signals.	107
6.3	ϕ_1 angle during a laser frequency modulation with caged TM.	108
6.4	Example of the signal in DWS during a laser frequency modulation at 1.123 Hz.	109
6.5	DWS channel in-flight and in the laboratory.	109
6.6	Sketch of a possible hypothesis for the measured signal in DWS.	111
6.7	DC angles in both the reference and the frequency interferometer.	113
6.8	Power measurements during a laser frequency modulation of experiment 3.	113
6.9	DC angle example during OPD modulation.	114
6.10	AS of DC angles while either of the beams was blocked.	115
6.11	Σ_A and Σ_B measurements during one power and one frequency modulation.	116
6.12	DC angle signal during laser frequency and power modulations: η and ϕ	117
6.13	AS of the DC angle ϕ_{1B} channel before adjustments and after.	118
6.14	AS of the DWS ϕ_{1B} channel before the alignment was improved and after.	118
6.15	AS of the DWS ϕ_{1B} channel during a power modulation on different days.	119
6.16	ϕ_{1B} DWS channel before the improvement of the alignment and after.	119

LIST OF FIGURES

6.17	Estimated coupling from Ψ_F to the respective DWS channel in the laboratory. . .	120
6.18	Sketch of the hypothesis: indirect coupling via associated power modulations. . .	121
6.19	Corresponding AS of RIN in the XR IFO.	124
6.20	AS of the Σ channels of the four IFOs during experiment number 5.	125
6.21	Comparison of Σ_F and DC angles during a power and a frequency modulation. . .	126
6.22	Comparison of Σ_R and DC angles during a power and a frequency modulation. . .	126
6.23	Coupling from power measurement to η_1	127
6.24	η_{1B} angle during power modulations in the laboratory	128
6.25	Position of the PD before the light reaches the OB.	128
6.26	AS of RIN during the laser frequency modulations at 4.5 Hz.	129
6.27	Example of the measured η_2 value in DWS during experiment in flight.	131
6.28	Coupling from Ψ_F to DWS ϕ_1 depending on the mean ϕ_1 position.	132
6.29	Zoom into the data of Figure 6.28.	132
6.30	Estimated coupling to DWS for three different rotations of TM1.	133
6.31	Check for cross-talk during a dedicated modulation.	135
6.32	Coupling coefficient from Ψ_F to ϕ_1 for different modulation frequencies.	136
6.33	AS of DWS signals along η_1 and ϕ_1 while removing cables.	137
6.34	AS of DWS signals along η_{1A} and ϕ_{1A} while using different ports.	138
7.1	Contribution of laser frequency noise to the o12 noise in June 2016.	143
7.2	Subtraction of frequency noise removes periods of increased noise in o12.	144
7.3	Contribution of laser frequency noise to the o12 noise in January 2017.	145
7.4	OLTF of loop model from experiment 3 and with increased gain.	147
A.1	MCMC chains ground	160
A.2	Covariances ground	160
A.3	Relative error of the fit to ground test measurements.	161
A.4	MCMC chains experiment 1	162
A.5	Covariances experiment 1.	162
A.6	Relative error of the fit to experiment 1	163
A.7	MCMC chains experiment 2	164
A.8	Covariances experiment 2.	164
A.9	Relative error of the fit to experiment 2.	165
A.10	MCMC chains experiment 3	166
A.11	Covariances experiment 3.	166
A.12	Relative error of the fit to experiment 3.	167
A.13	MCMC chains experiment 4	168
A.14	Covariances experiment 4.	168
A.15	Relative error of the fit to experiment 4.	169
A.16	MCMC chains experiment 5	170
A.17	Covariances experiment 5.	170
A.18	Relative error of the fit to experiment 5.	171
C.1	LPF acceleration noise budget current best estimate before launch.	184
C.2	Example of the simulated force impulses and free-flights.	185
C.3	Concept of the state-space model free-flight simulation.	186
C.4	Fit residuals of the fit to the simulated free-fall data.	188

List of Acronyms

ADC	Analogue to Digital Converter
AEI	Albert Einstein Institute
aLIGO	advanced Laser Interferometer Gravitational-Wave Observatory
AOM	Acousto-Optic Modulator
AS	Amplitude Spectrum
ASD	Amplitude Spectral Density
CDS	Control and Design System
DFACS	Drag-Free and Attitude Control System
DFT	Discrete Fourier Transform
DMU	Data Management Unit
DOY	Day Of Year
DPS	Differential Power Sensing
DRS	Disturbance Reduction System
DWS	Differential Wavefront Sensing
EM	Engineering Model
EMRI	Extreme Mass Ratio Inspiral
ESA	European Space Agency
FI	Faraday Isolator
FIOS	Fibre Injector Optical Sub-Assembly
FM	Flight Model
FT	Fourier Transform
GOCE	Gravity field and steady-state Ocean Circulation Explorer
GRS	Gravitational Reference Sensor

LIST OF ACRONYMS

GPRM	Grabbing, Positioning and Release Mechanism
GW	Gravitational Wave
iABG	Industrieanlagen-Betriebsgesellschaft mBH
IDL	Interferometer Data Log
IFO	Interferometer
IIR	Infinite Impulse Response
IS	Inertial Sensor
LISA	Laser Interferometer Space Antenna
LMU	Laser Modulator Unit
LPF	LISA Pathfinder
LPSD	Logarithmic Power Spectral Density estimate
LTP	LISA Technology Package
LTPDA	LISA Technology Package Data Analysis
MCMC	Markov Chain Monte-Carlo
NPRO	Non-Planar Ring Oscillator
OB	Optical Bench
OBC	Onboard Computer
OLTF	Open-Loop Transfer Function
OMS	Optical Metrology System
OPD	Optical Pathlength Difference
OSTT	On-Station Thermal Tests
PD	Photodiode
PM	Phasemeter
PSD	Power Spectral Density
PT	Phase Tracking
PZT	Piezo-electric Transducer
QPD	Quadrant Photodiode
RF	Radio Frequency
RIN	Relative Intensity Noise
RLU	Reference Laser Unit
SBDFT	Single-Bin Discrete Fourier Transform
SNR	Signal to Noise Ratio
SOVT	Science Operations Verification Test

ST7	Space Technology 7
TC	Telecommand
TDI	Time-Delay Interferometry
TFE	Transfer Function Estimate
TIA	Trans-Impedance Amplifier
TM	Test Mass

List of terms

X12	The interferometer to measure the changes in relative distance in between the two TMs, see Figure 2.2. Its measured phase is denoted φ_{12} .
X1	The interferometer to measure the changes in relative distance of TM1 with respect to the OB, see Figure 2.2. Its measured phase is denoted φ_1 .
XR	The interferometer to measure common mode fluctuations of the optical path length, see Figure 2.2. Its measured phase is denoted φ_R .
XF	The interferometer to measure laser frequency fluctuations, see Figure 2.2. Its measured phase is denoted φ_F .
Ψ_R	The reference interferometer signal after application of the phase tracking. It is used by the OPD control loop, compare Figure 2.4. This quantity is available from LPF satellite telemetry, see Appendix B.
Ψ_F	The frequency interferometer signal after subtraction of the reference interferometer measurement and application of the phase tracking. It is used by the frequency control loop, compare Figure 2.9, and thus represents an in-loop measurement. This quantity is available from LPF satellite telemetry, see Appendix B.
o1	The interferometric measurement of the changes in relative distance of TM1 with respect to the OB after it has been processed to a measurement in metres. This quantity is available from LPF satellite telemetry, see Appendix B.
o12	The interferometric measurement of the changes in relative distance of the two TMs after it has been processed to a measurement in metres. This quantity is available from LPF satellite telemetry, see Appendix B.

LIST OF TERMS

η	The angle of TM rotation around the y -axis, compare Figure 2.3. It can be measured interferometrically using the DWS or DC angles. These quantities are available from LPF satellite telemetry, see Appendix B.
ϕ	The angle of TM rotation around the vertical axis, compare Figure 2.3. It can be measured interferometrically using the DWS or DC angles. These quantities are available from LPF satellite telemetry, see Appendix B.
x1 control degree of freedom	The control degree of freedom of the satellite with respect to the TM1 x position.

Chapter 1

Interferometric gravitational wave detection in space

1.1 Gravitational waves and their detection

1.1.1 Definition and properties of gravitational waves

A central concept of general relativity is, simply speaking, that matter affects spacetime and that in turn, spacetime determines the paths of free-falling test particles, the geodesics, as summarised for example in [Giu05]. Assuming that spacetime is nearly flat, it can be shown that perturbations of spacetime propagate as waves, the so-called gravitational waves (GWs) [Hen09].

GWs can be described by several key features. They are transversal waves which propagate at the speed of light [SS09]. For a binary neutron star system, the GW signal depends, for example, on the orbital period, the distance to the binary source and the masses. In addition, two independent polarisation states of GWs exist. These can be understood by considering the impact of a passing GW on a ring of free-falling proof masses in general relativity, see Figure 1.1. Here, it is important to note that the relative distance of these proof masses is changed and that this change is direction dependent. For example, in the left panel of Figure 1.1, the relative distance in the vertical direction is first continuously prolonged and the relative distance in horizontal direction is shrunk in the same manner while preserving the area of the ring. Then, the ellipsoid is smoothly deformed back into a circle and then into the other ellipsoid and again back into a ring. This is happening during a period of a gravitational wave. The impact of GWs on relative distances provides one way to detect them. A precise measurement of such changes in relative distances can be implemented using a Michelson-like laser interferometer, as described in Section 1.1.2. As the speed of light is constant, it can faithfully measure the changes in relative distance in between the proof masses. However, these changes caused by a passing GW are exaggerated in Figure 1.1. In reality, they are extremely small as the coupling from a gravitational wave to spacetime is weak. As much as this makes the measurement of GWs a technological challenge, it allows GWs to travel a long way through spacetime while remaining quasi unaffected by their environment [SS09]. Consequently, they provide a mean to obtain information about their sources, even though these sources might be far away and/or not directly accessible by electromagnetic radiation.




Figure 1.1: The polarisation states of GWs illustrated by their impact on a ring of free-falling proof masses in the plane of the paper. A GW in plus (left) and cross (right) polarisation is propagating orthogonal to this plane. The ring is smoothly distorted into one of the ellipsoid states, then back to a ring and into the other ellipsoid states of that polarisation. Reprint of [SS09].

1.1.2 Detection of gravitational waves with laser interferometry

As already mentioned, a key property of gravitational waves is the change in relative distance of two proof masses depending on the polarisation state which can be measured by a Michelson interferometer. This measurement principle is illustrated in Figure 1.2. Here, the ideal case of a GW in $+$ polarisation propagating orthogonally to the plane of the paper and the detector is shown. Originally, both arms have the same length L as shown in the left panel. During the passage of a GW, a maximal distortion in both directions, as shown in the centre and right panel, can be measured, since the differential arm length changes result in an interferometric signal at the output port.

This is the very basic scheme of the gravitational wave detection on Earth as applied in the advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) detectors, the Virgo and the GEO600 detector for example. In reality, the set-up is significantly more complicated. Key aspects of the aLIGO detectors include additional mirrors placed in the two interferometer arms to enhance the effect of a GW signal and two further semi-transmissive mirrors to increase the laser power in the interferometer and to optimise the signal detection. Another aspect is the vibration isolation of the mirrors by a pendulum suspension system [A+16b].

1.1.3 Sources of gravitational waves

The first indirect detection of gravitational waves was from a binary system consisting of a pulsar and a neutron star denoted PSR 1913 + 16. Observations from 1974 to 78 showed a change in cumulative shift of periastron time of approximately 1s in agreement with the predictions of




Figure 1.2: The principle of gravitational wave detection using interferometry. Light of wavelength λ enters the interferometer at point A and travels along the arms of length L before it is detected at the output port B. A GW changes the length in the two interferometer arms differently such that a phase change can be detected. Reprint of [Hen09].

general relativity for the expected loss of energy due to the emission of gravitational waves. This was the first indirect proof of the existence of GWs [TFM79].

On September 14th 2015, a GW signal was directly measured for the first time in the two aLIGO detectors in the United States [A+16b]. A change in length per interferometer arm length of up to $1 \cdot 10^{-21}$, corresponding to a peak displacement measured in the interferometer arms of only ± 0.002 fm was detected. This signal was associated to two individual masses of $36^{+5}_{-4} M_{\odot}$ and $29^{+4}_{-4} M_{\odot}$ merging into an object of $62^{+4}_{-4} M_{\odot}$ with the mass difference being radiated as GWs. From the inferred masses, the orbital frequency and the decay of the waveform, the only known sources to fit the picture are two black holes merging into one. This was the first direct measurement of GWs. In the first two observation runs 9 more black hole mergers were measured [A+19a]. In addition, the aLIGO detectors and the advanced Virgo detector observed GW signals which can most likely be attributed to merging neutron stars [A+17a]. This event could also be associated to a γ -ray burst and detectable electromagnetic radiation over longer time scales, showing the existence of matter [A+17a]. Studying these astrophysical objects in synergy with electromagnetic radiation and GWs is called multi-messenger astronomy. Additionally, a compact binary coalescence, whose origin could so far not be unambiguously determined but was found to be consistent with an unusually heavy binary neutron star system, was measured in 2019 [A+20a]. In the same year, the coalescence of a black hole binary with unequal masses was observed, too [A+20b].

Despite the immense successes of GW detectors over the past years, their detection capacities on Earth are limited towards lower frequencies. This is because seismic noise, suspension thermal noise and fluctuations of the local gravitational gradient become dominant noise sources in GW detectors on Earth below 10 Hz [A+16a]. Additionally, it is difficult to further lengthen the interferometer arm length. These obstacles can be tackled by a space-borne gravitational wave detector. Simply speaking, by moving to space, the Earth seismic noise is avoided, the interferometer arm length can be made longer much easier and the local gravity gradients on the satellites can be much better controlled. However, GW observation in space brings along new technological challenges. A well established concept to observe the universe through GWs is called LISA and will be explained in Section 1.1.4.

Space-based GW observation is a promising effort because at frequencies complementary to those accessible by gravitational wave detectors on Earth, rich sources of gravitational waves are

expected. LISA aims to measure GWs with frequencies in between $20 \mu\text{Hz}$ and 1 Hz [AS⁺17]. Although a complete list of the sources is beyond the scope of this work, some examples shall be given.

Compact binaries in the Milky Way, consisting of white dwarfs, neutron stars or stellar mass black holes, are a frequent source of GWs in the LISA band [AS⁺17].

Some of these are studied in more detail by electromagnetic astronomy, such that estimates of the mass, the distance and the inclination can be used to predict their GW signal. If the predicted GW signal is expected to reach a Signal to Noise Ratio (SNR) of 5 or above within 4 years of observation with LISA, they are called ‘verification binaries’ and will in turn allow scientists to test the functionality of LISA [K⁺18]. Moreover, up to about 10% of the stellar mass black hole binary systems observable with LISA are radiating GWs with increasing frequency as they inspiral such that their time of coalescence is within 10 years of the start of LISA operations and the emitted GWs pass the aLIGO frequency band. LISA is expected to predict the time to coalescence to within approximately 10s and the sky location to within one square degree [Ses17].

Another important source of gravitational waves in the LISA band are supermassive black hole binaries and their coalescences. From such observations, a better understanding of the evolution of black holes through the inspiral of binaries, accretions and mergers is expected [AS⁺17].

A third category of sources of GWs in the LISA band are Extreme Mass Ratio Inspirals (EMRIs). In broad terms, EMRIs can be defined as a binary system in which a compact object, as for example a stellar origin black hole, is in a highly relativistic orbit around a much heavier massive black hole. Thus, GWs from EMRIs provide a chance to study spacetime in the strong-field regime provided an SNR above 50 [AS⁺17].

After the identification of a number of sources, a stochastic background of GWs is expected which can be thought of as ‘the gravitational wave equivalent of the cosmic microwave background’ [Wan19].

Finally, the most interesting sources may perhaps be those which are nowadays unforeseen.

1.1.4 LISA concept

LISA is a space mission concept to measure GWs which relies on the same principle of gravitational wave detection as the aLIGO detectors. First proposals to the European Space Agency (ESA) date back to the 1990s [Tea11]. To cut a long story short, in 2017, ESA’s Science Programme Committee selected LISA as the third L-class mission in the ESA programme, denoted L3, with an expected launch in the 2030s [ESA19]. At the time of writing, LISA is under detailed study to prepare for the mission adoption in the next years [ESA19].

Simply put, the mission concept aims at constructing a virtual Michelson-like interferometer in space.

To allow for large interferometer arm lengths, the light is interchanged between three satellites




Figure 1.3: Sketch of the LISA orbit. Reprint of [AS⁺17].

which form a triangle with three nearly equally long arms. The current mission concept aims at a mean distance of 2.5 million km from one satellite to the next [AS⁺17]. As shown in Figure 1.3, the satellites shall be put on a heliocentric orbit which enables the formation of a triangle. It follows the Earth at a distance of 50 to 65 million km corresponding to a time varying angle around 20 deg [AS⁺17]. The stability of this triangle with a suitable inter-satellite distance and a small enough distance to Earth for data transfer, is given by orbits which include a so-called ‘cartwheel motion’. Again, the GWs will be measured by the changes in relative distance they induce between the endpoints of the interferometer arms. In space, these end points are not mirrors but TMs which are freely floating inside the satellites. From these thoughts it follows that the TMs must be isolated from other influences to a certain level to be able to attribute changes in relative distance to GWs. The minimal amplitude of GW signals that can still be detected is, however, a trade-off between many factors, but the arm length and the residual acceleration level of the TMs caused by all other impact factors apart from gravitational waves are two major ones. In more precise terms, the current mission concept requires that the Amplitude Spectral Density (ASD) of the residual acceleration of a single test mass, g , is below

$$S^{\frac{1}{2}}_{g,LISA} \leq 3 \cdot 10^{-15} \text{m s}^{-2} \sqrt{\text{Hz}}^{-1} \sqrt{1 + \left(\frac{0.4 \text{ mHz}}{f}\right)^2} \sqrt{1 + \left(\frac{f}{8 \text{ mHz}}\right)^4} \quad (1.1)$$

at least in the frequency range from 100 μHz to 0.1 Hz but the goal is to achieve this from 20 μHz to 1 Hz [AS⁺17]. This is a very small level of residual acceleration and especially in the early days of the LISA mission design, this was about three orders of magnitude below the planned acceleration levels of the Gravity field and steady-state Ocean Circulation Explorer (GOCE) accelerometers [RYS11]. So, to prove the feasibility of the LISA concept, a technology test was deemed necessary. Even though a torsion pendulum test facility approaches the required accuracy, mechanical thermal noise is limiting the performance at lower frequencies and readout noise towards the higher end of the spectrum [A⁺19c][HCD⁺02]. Tests under free-fall conditions can also be performed in a drop tower. A measurement of residual acceleration in the mHz regime, however, requires a measurement duration of ≈ 17 min which is too long to be feasible in a drop tower on Earth. That is why a test mission in space was deemed necessary already in 1998 [M. 15][Aud14].

1.2 The LISA Pathfinder mission

The key idea of LPF can be thought of as testing one of the three arms in LISA and shrinking it from 2.5 million km to less than 1 m such that it fits into a single satellite [M. 15]. For historical reasons, this LPF satellite contains two technology contributions: the ESA LISA Technology Package (LTP) and the NASA payload Space Technology 7 (ST7)-Disturbance Reduction System (DRS) [M. 15]. In the following, we will mainly focus on the LTP.

1.2.1 LPF subsystems

The objective of LPF was to show that LISA is feasible and that the required tiny level of residual acceleration can be reached. To verify this experimentally, it was required that the residual *differential* acceleration noise, Δg , was below [A⁺16c]

$$S^{\frac{1}{2}}_{\Delta g,LPF} \leq 30 \cdot 10^{-15} \text{m s}^{-2} \sqrt{\text{Hz}}^{-1} \sqrt{1 + \left(\frac{f}{3 \text{ mHz}}\right)^4} \quad (1.2)$$

for frequencies from 1 mHz to 30 mHz. Even though this level is relaxed by one order of magnitude in comparison to the LISA requirement, it is a much more stringent requirement than the planned acceleration levels of the GOCE accelerometers [RYS11].




Figure 1.4: The core of LISA Pathfinder. The two golden cubes are the test masses inside their respective electrode housings inside a vacuum tank. The optical bench is discernible in the centre. Image: ESA/ATG medialab.

To prove this level of residual acceleration, LPF was equipped with two free-falling TMs. This means they were not in mechanical contact with any other satellite components during nominal operations. These were the two golden quasi cubes in Figure 1.4. They consisted of a gold-platinum alloy which had a low magnetic susceptibility and a high density [F. 11]. As each side was (46.000 ± 0.005) mm long, this yielded a mass of (1.928 ± 0.001) kg per TM which aided in minimising undesired gravitational interaction with the surroundings [A⁺16c]. The TMs were inside the Gravitational Reference Sensor (GRS). Each of the two GRSs consisted of 18 electrodes mounted on the electrode housing and the corresponding front-end electronics. This system allowed for electrostatic actuation and sensing of the TM along all six degrees of freedom. The GRS with the respective TM was in one vacuum tank. The ‘pink flashes’ around the two

TMs are illustrating the discharging mechanism via the photoelectric effect by the use of UV light [Lc17].

Simply speaking, the residual differential acceleration Δg was then obtained as the second derivative of changes in the relative distance between the two TMs along their line of sight. These changes in relative distance were measured by the OMS using heterodyne laser interferometry. This concept required a laser unit which is not shown in Figure 1.4. The light entered the Optical Bench (OB) via the fibres, whose ends are shown in green in the centre left, and Fibre Injector Optical Sub-Assemblies (FIOSs). The OMS required in addition a number of beam splitters, mirrors and photodiodes (PDs) shown in the centre of Figure 1.4. The laser beam which travelled to the two test masses is called the measurement beam and is shown in red. It could reach the TMs because the vacuum tanks were equipped with an optical window. The reference beam is sketched in blue and the interference of the two in pink. To ensure stability throughout launch and operations, the components of the interferometer were bonded onto the OB made out of Zerodur[®] because of its ultra-low thermal expansion. The OMS will be explained in more detail in Chapter 2.

The required residual acceleration level was achieved by the interaction of several key subsystems. These were not only the GRS and OMS already explained but also the DFACS. It allowed for several operational modes which ranged from safety modes to the nominal science mode and the so-called free-fall or drift mode (compare [Sch14]).

In the nominal science mode the control scheme along the measurement axis had two major constituents: the drag free and the suspension loop. This is also illustrated in Figure 1.5. As explained also in [M. 15], the drag-free control loop used the interferometric readout of the relative position of one of the two TMs, often called TM1, with respect to the satellite to determine the necessary force that needed to be applied on the satellite such that it followed the free-falling TM1. The corresponding readout is denoted o_1 in Figure 1.5. The suspension loop calculated the necessary force which needed to be applied via the electrostatic actuation system on

TM2 such that it followed TM1. To this end, the suspension loop used the interferometric readout of the relative distance in between the two TMs, named o_{12} in Figure 1.5. In the measurement band, the suspension control loop had close to negligible gain, given its unity gain frequency was around 1 mHz [A⁺16c]. This was different for the drag-free controller whose bandwidth was around 200 mHz for the x_1 degree of freedom [Sch12]. To apply the necessary commanded forces and torques onto the satellite with minimal acceleration noise, LPF had been equipped with μ Newton cold gas thrusters (see [A⁺19d]) to successfully apply the forces and torques. In addition, as part of the ST7-DRS payload, colloidal μ N-thruster have also been tested [A⁺18a]. Additionally, to achieve the required level of residual acceleration, the temperature on board the LPF satellite should remain stable to avoid thermal gradients which lead to acceleration noise in the GRS due to the radiometer effect, radiation pressure and temperature dependent outgassing [D⁺03]. The temperatures on LPF and dedicated experiments to characterise the




Figure 1.5: A sketch of the DFACS control strategy used along the sensitive axis. Reprint from [M. 15].

thermal coupling are reported in [A⁺19f]. Moreover, LPF contained magnetometers which allowed to estimate a possible contribution of magnetic fields to the residual acceleration [A⁺16c].

1.2.2 Development and testing

The development of the technology demonstrator mission LPF was accompanied by extensive testing of all the subsystems. For the OMS, the Engineering Model (EM) test campaign was performed between October 2010 and January 2011 [Aud14]. Early test results of the EM can also be found in [H⁺05]. In addition, the OMS pre-flight model phase 1 (often denoted Flight Model (FM)) test campaign took place from July to September of 2010 [Aud14]. The test campaign data for the laser frequency modulations used for comparison in Chapter 3 has been taken during this period. In 2011 the On-Station Thermal Tests (OSTT) campaign took place at Industrieanlagen-Betriebsgesellschaft mBH (iABG) where a number of flight hardware components had already been assembled into the satellite. These measurements could then be performed at different temperatures [Aud14]. The stabilised and free-running in-loop measurements used for comparison in Chapter 4 have been recorded during this test. For more details and the differences between the test campaigns see [Aud14].

Before LPF was launched, a dedicated data analysis MATLAB toolbox was developed. It is called LISA Technology Package Data Analysis (LTPDA) and wraps the data into a data object called `ao` which has the special feature of automatically tracing the history of the data through all methods applied to it [H⁺09]. In addition, it includes a user-friendly implementation of signal processing algorithms and data analysis pipelines for planned experiments on board LPF to allow for an efficient analysis of LPF data. The analysis presented in this thesis has been performed using the LTPDA toolbox and extending it where needed. The development of the toolbox is closely linked to the preparations for mission operations (compare [D⁺18]). Shortly before launch the Science Operations Verification Test (SOVT) was successfully passed. This is a common formal step for ESA missions but the preparation of this test was also used to practise operations with the scientists involved.

1.2.3 LPF mission operations

LPF was launched on December 3rd 2015 at 04:04 UTC from Kourou, French Guiana on board an ESA-Vega rocket by Arianespace. After several apogee-raising manoeuvres, LPF left its low-earth orbit and headed towards a 800,000 km x 500,000 km Lissajous orbit around the Earth-Sun Lagrange point L1 at approximately 1.5 Gm from Earth [Lc17]. In simple terms, this region is characterised by the same angular velocity as the Earth such that the relative position of comparably small objects with respect to Earth and Sun is maintained. This is ideal for a measurement of tiny residual acceleration levels as it is the aim of LPF.

In January 2016, the OMS (see Chapter 2) was switched on for the first time during the commissioning of LPF. In this period, one unit after the next was switched on. A major milestone was also the handover of the two TMs (see details in Section 1.2.1) from the launch lock fingers to the Grabbing, Positioning and Release Mechanism (GPRM) and the subsequent release into free-fall. The successful commissioning was followed by the beginning of the nominal operations phase of the LTP on March 1st 2016. The planned operations phase was extended such that LPF remained operational until the shutdown command was sent on July, 18th 2017.

1.2.4 LPF residual acceleration measurement results

Already in June 2016, the first results of LPF were published [A⁺16c]. Between 0.7 and 20 mHz, a residual differential acceleration of only $(5.2 \pm 0.1) \text{ fm s}^{-2} \sqrt{\text{Hz}}^{-1}$ was measured. This is more than a factor 5 below the requirement (see Equation 1.2) and thus a huge success. The interferometric sensing noise was determined to be only $(34.8 \pm 0.3) \text{ fm} \sqrt{\text{Hz}}^{-1}$, hence also significantly below the requirement (explained in Chapter 2). These already excellent results could be further improved in the course of the mission due to three factors. First, the amount of residual gas particles surrounding the TM was decreasing in the course of the mission due to continued venting to space which leads to a decrease in Brownian force noise. Second, improved understanding of the applied electrostatic actuation force enabled a better estimate of the residual low frequency force noise. Third, an inertial force due to satellite rotation was accounted for in the calculation of the residual acceleration measurement. These factor resulted in a residual acceleration of only $(1.74 \pm 0.05) \text{ fm s}^{-2} \sqrt{\text{Hz}}^{-1}$ between 2 mHz and 7 mHz [A⁺18b]. This result could be confirmed in a DFACS actuation mode which replaces the continuous control of TM2 along the x-direction with an intermittent control scheme [A⁺19e]. This mode is called drift mode (see Appendix C).

LPF did not only exceed the required residual acceleration level but it also provided an excellent opportunity to characterise all subsystems as well as the satellite and space environment. This was a unique opportunity to investigate LISA technology in space in as much detail as possible.

1.3 This thesis

The work presented in this thesis focusses on the LPF OMS. Therefore, it will be explained in more detail in Chapter 2. Its performance in terms of the sensing noise level along the sensitive axis of LPF was expected to have several main contributions as listed in [Aud14]. One of these are frequency fluctuations of the laser which resulted in differential readout noise on the level of some tens of $\text{fm} \sqrt{\text{Hz}}^{-1}$ in the frequency range from roughly 0.1 to 1 Hz. This noise level was a result of two factors. First, the laser frequency stabilisation worked as expected, as proven by the analysis in Chapter 3. The resulting noise performance in the course of the mission is the subject of Chapter 4. Second, the coupling factor, the path length difference Δs , is small due to excellent OB construction and satellite integration. This is the subject of Chapter 5. During several laser frequency modulation experiments performed on LPF, an unexpected modulation of other OMS channels was observed. In Chapter 6, the investigations aimed to find the origin of these modulation signals are summarised. A summary and the conclusions of this work are given in Chapter 7.

Each of the following chapters will begin with a small introduction on its own and conclude with a corresponding summary.

Chapter 2

Selected aspects of the optical metrology system

In this chapter, the LISA Pathfinder Optical Metrology System (OMS) is briefly introduced and we summarise the path from the laser to the downloaded telemetry. This telemetry will be used for the analyses in the following chapters. Next, it is shown how laser frequency noise couples into the phase measurement of a heterodyne interferometer. In addition, the required precision of the OMS and the required level of frequency stability are explained. On LPF, the laser frequency fluctuations are suppressed by a dedicated control loop, which is introduced here. At the end of this chapter, we summarise all laser frequency modulation experiments performed during LPF operations.

2.1 Overview

2.1.1 Hardware

On LPF, a diode-pumped Non-Planar Ring Oscillator (NPRO) using a Nd:YAG crystal produces a nominal wavelength of 1064 nm. The nominal output power is approximately $32.5 \text{ mW} \pm 30\%$. Measurements of the in-flight output power can be found for example in Section 3.3.2. The laser is contained inside the so-called Reference Laser Unit (RLU), shown on the right in Figure 2.1. [Aud14] From there the light is transmitted through a fibre to the Laser Modulator Unit (LMU). Inside the LMU, the beam is split into two parts with equal power [Weg14, Section 3.6]. Each of these beams is frequency shifted by an Acousto-Optic Modulator (AOM) such that the resulting frequency difference is 1 kHz nominally [H+03]. The LMU also contains the actuators of the Optical Pathlength Difference (OPD) control loop (compare Figure 2.4) [Aud14]. Single mode polarisation maintaining fibres are used to bring the light from the LMU to the OB [Aud14]. The light is injected onto the OB via Fibre Injector Optical Sub-Assemblies (FIOSs). Simply speaking, in the FIOS, the fibre end is glued into a mounting block and the light that comes out of the fibre is then collimated by a lens and its polarisation is cleaned [Fit10]. On LPF, some features of the RLU and the subsequent units are monitored most of the time. A list of relevant parameters for this thesis can be found in Appendix B. The OB is the base plate onto

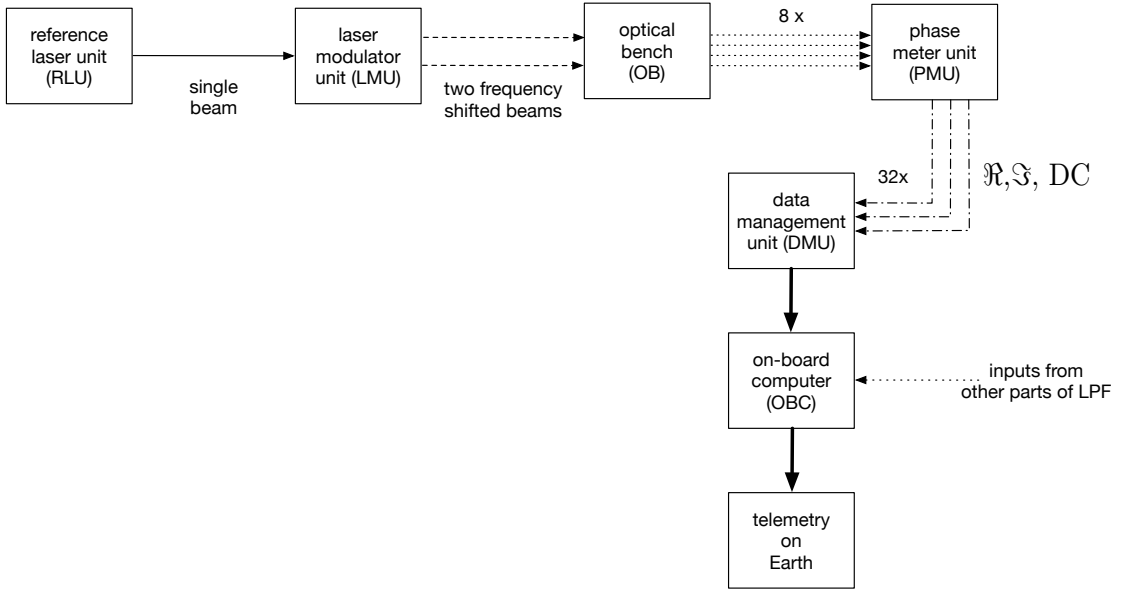


Figure 2.1: A simplified diagram of the different units on-board LPF that are involved in producing the data that will be analysed in the following sections.

which the components for the four interferometers are bonded as shown in Figure 2.2. The following description of the four interferometers can be found in many places, for example also in [Wan10]. The X12 interferometer measures the changes in relative distance between the two test masses. In addition to this measurement, it provides the difference in orientation of the two test masses. The X12 interferometer can be easily recognised because its measurement beam hits both test masses. In the X1 interferometer, the change in relative distance and orientation of TM1 with respect to the OB is sensed. As the OB is rigidly mounted to the sideslabs and thus the satellite, the X1 interferometer measures the motion of TM1 with respect to the satellite. The measurement beam of this interferometer is reflected by TM1 but not by TM2. In combination with the measurement of X12, this allows us to estimate the η and ϕ angles, see Figure 2.3, of the two test masses. The angular measurements are obtained via the Differential Wavefront Sensing (DWS) scheme which will be explained in Section 2.1.2.

On the OB, two more auxiliary interferometers can be found. In these interferometers both beams stay on the optical bench.

One of these two interferometers is the reference interferometer, XR. It is sensitive to common-mode path length fluctuations which can occur before the light reaches the ultra-stable OB due to changes in temperature or mechanical vibrations. Its measurement is denoted φ_R . To minimise the common-mode path length noise in the other interferometric phase measurements, φ_R is subtracted. During the LTP interferometer development, it was found that this subtraction alone was not sufficient as cross-talk of the AOM frequencies combined with fluctuations of the differential path length between the measurement and the reference beam lead to an additional noise source [Mar07]. Therefore, the path length difference between the measurement and the reference beam was actively stabilised by the OPD loop. As shown in Figure 2.4, the XR measurement, after the application of the Phase Tracking (PT) [U. 09], is the error signal of the OPD loop, as described in [Aud14]. It is used by the digital controller which is part of the Data

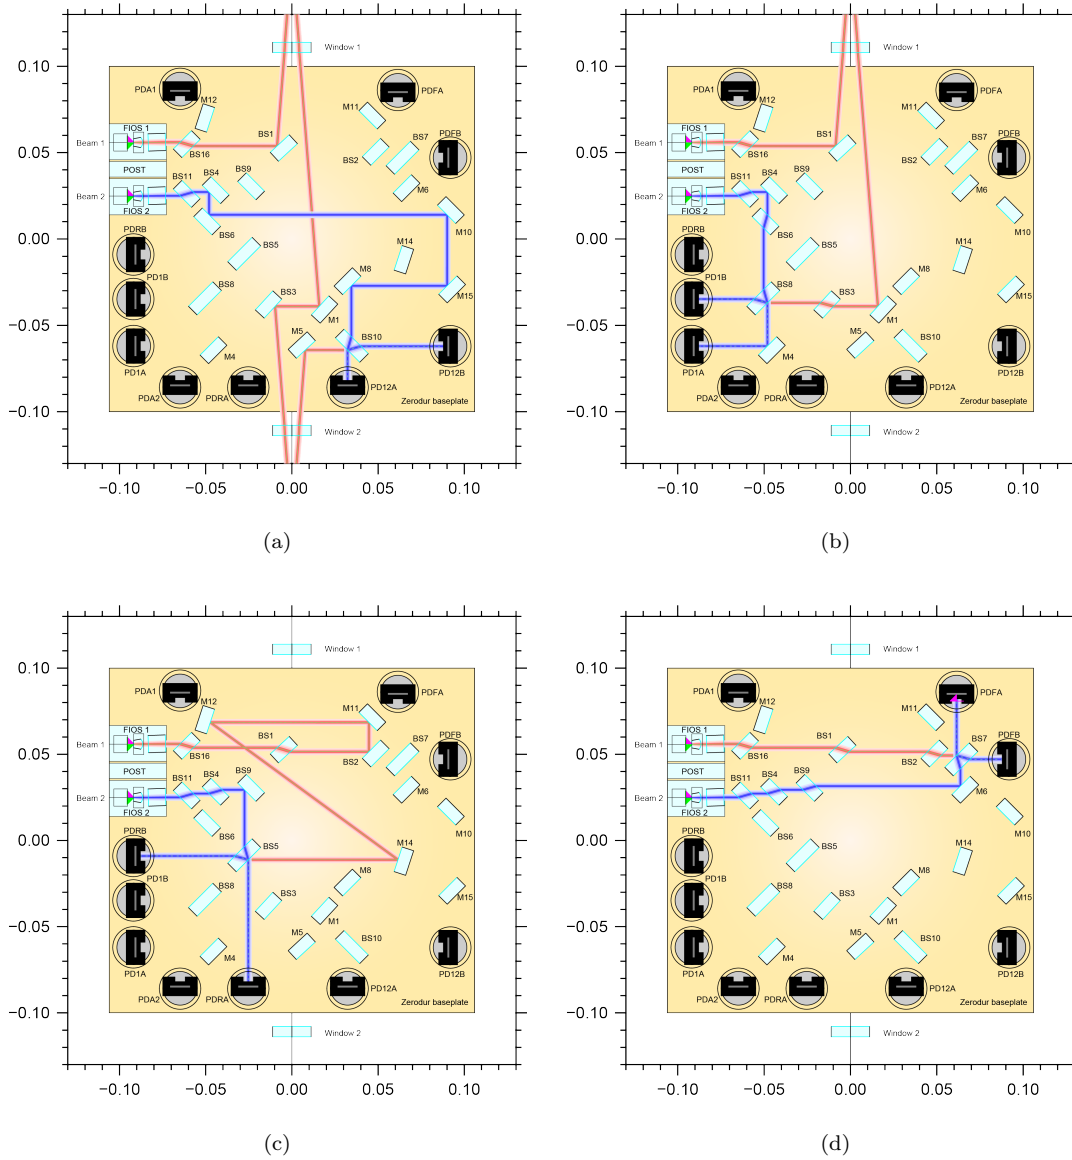


Figure 2.2: The four interferometers on the LISA Pathfinder optical bench. (a) X12 interferometer. (b) X1 interferometer. (c) Reference interferometer XR. (d) Frequency interferometer XF. Diagrams produced with OptoCAD by E D Fitzsimons.

Management Unit (DMU). The thus determined voltage is commanded to the two actuators which are operated in a push-pull configuration. The two actuators consist of a half wave plate and a triple prism mounted on a Piezo-electric Transducer (PZT) actuator located inside the LMU [Aud14]. Before launch, the unity gain frequency of the OPD actuators was estimated to be 1.65 Hz for the FM test campaign [Aud14].




Figure 2.3: The LRF coordinate system as used by DFACS. Picture taken from [Sch12].




Figure 2.4: The OPD loop to suppress common mode path length noise. C_{OPD} denotes the OPD controller and A_{OPD} the OPD actuators. The control loop operates using the φ_R measurement processed into Ψ_R as defined in Equation 2.5. The channels marked with the dashed lines are available as telemetry.

The frequency interferometer XF has a deliberate path length mismatch of 0.382 m [Rob13, Page 22] which occurs in the fibres before the OB [Wan10]. This means that the optical path of the measurement and reference beam are the same on the OB but the reference beam travels longer in the fibres [Aud14]. This path length mismatch amplifies the laser frequency noise, as explained in Section 2.2, that is measured by the frequency interferometer XF. This measurement is used as the input to the laser frequency stabilisation control loop, as outlined in Section 2.4.

In Figure 2.2, we also see that each interferometer output is recorded by two photodiodes. This is not only because redundancy is favourable in a space mission but mostly because they allow us to apply the so-called balanced detection scheme in the data-




Figure 2.5: A sketch of a QPD seen from the front with the labelling of the quadrants as used in the text. These enable DWS.

processing [CM68]. This allowed to suppress a noise source in the OMS, see Section 2.1.3. To be more precise, the PDs of the four interferometers described so far are all InGaAs QPDs as shown in Figure 2.5 [Aud14]. In addition, there are two additional QPDs, of which only a single quadrant is being used. They serve as power monitors and are the input to the fast amplitude control loop that stabilises the laser amplitude. For information about the laser amplitude stabilisation, we refer the interested reader to [Aud14].

Each quadrant of each PD is then transmitted to the Phasemeter (PM), see Figure 2.1. In this unit, the current is converted to a voltage, via a Trans-Impedance Amplifier (TIA), which is low-pass filtered and then digitised by an Analogue to Digital Converter (ADC).

2.1.2 Data processing

The measured signals on-board LPF This paragraph summarises Sections 1.4.3 and 1.4.4 of [Aud14], compare also [H⁺04]. The PM applies a Single-Bin Discrete Fourier Transform (SBDFT) to the data. This results in the real and imaginary part of the complex amplitude at the heterodyne frequency. In addition a DC part, corresponding to the mean of the data, is calculated. These three quantities per quadrant are passed on to the DMU. For the purpose of this thesis, the following data streams, applying to the case of nominal performance, are relevant:

- In a first step, the DC values of the A and B PD are averaged into one value per quadrant, $DC_{i,j}$, and the real and imaginary part of the A and B PD are combined into a complex value, usually denoted $F_{i,j}$. With these data streams, i denotes the interferometer $i = 1, 12, R, F$ and j indicates the quadrant $j = A, B, C, D$.
- From the DC signals, a total of 5 intermediate signals per interferometer i are calculated:
 1. the sum of all 4 $DC_{i,j}$ values of the 4 quadrants in one interferometer i . This result is named Σ_i and measures the power in an interferometer up to a scaling factor (see Section 6.4.1);
 2. the sum of the A and B quadrant, denoted as Σ_{up} ;
 3. the sum of the C and D quadrant, denoted as Σ_{down} ;
 4. the sum of the A and C quadrant, denoted as Σ_{left} ;
 5. the sum of the B and D quadrant, denoted as Σ_{right} .
- The same intermediate signals are obtained from the $F_{i,j}$ per interferometer i :
 1. the sum of all 4 $F_{i,j}$ values of the 4 quadrants in one interferometer i . This result is named F_i and will be further processed to measure the longitudinal phase in an interferometer;
 2. the sum of the A and B quadrant F values, denoted as F_{up} ;
 3. the sum of the C and D quadrant F values, denoted as F_{down} ;
 4. the sum of the A and C quadrant F values, denoted as F_{left} ;
 5. the sum of the B and D quadrant F values, denoted as F_{right} .

- From the F_i values, a phase is obtained via the application of the `atan2` function to the real and the imaginary part, also denoted `arg`:

$$\varphi_i = \arg(F_i), \quad (2.1)$$

which results in φ_1 , φ_{12} , φ_F and φ_R in the interval $(-\pi, +\pi)$, respectively. The reference interferometer phase, φ_R , is then subtracted from the other three phase signals. In addition, the phase tracking algorithm (PT) is applied. This procedure enables the correct measurements of changes in phase of values outside the interval of $(-\pi, +\pi)$. In simple terms, a counter of n is stored and multiplied with a factor of 2π to access values outside of this regime. The results of this step are frequently used in the subsequent chapters, thus let us define (see also [Wis17])

$$\Psi_R := \text{PT}(\varphi_R) \quad (2.2)$$

$$\Psi_F := \text{PT}(\varphi_F - \varphi_R) \quad (2.3)$$

$$\Psi_{12} := \text{PT}(\varphi_{12} - \varphi_R) \quad (2.4)$$

$$\Psi_1 := \text{PT}(\varphi_1 - \varphi_R) \quad (2.5)$$

The phases Ψ_1 and Ψ_{12} are also called the electrical phases in Chapter 6. The impact of the subtraction of the reference interferometer measurement on the laser frequency measurements is analysed in Section 4.1.

- The phase measurements Ψ_{12} and Ψ_1 are then scaled to the corresponding relative TM positions by

$$o1 = \frac{1}{2} \frac{\lambda}{2\pi} \frac{1}{\cos(\alpha)} \cdot \Psi_1 \quad (2.6)$$

$$o12 = \frac{1}{2} \frac{\lambda}{2\pi} \frac{1}{\cos(\alpha)} \cdot \Psi_{12}. \quad (2.7)$$

Here, we have a factor $\frac{1}{2}$ to account for the fact that the light travels to the TM and back. λ denotes the laser wavelength and accordingly, the factor $\frac{\lambda}{2\pi}$ converts our phase measurement in radian to a distance measurement in metres. The division by $\cos(\alpha)$ is the result of the laser beam hitting the TMs under an angle of $\alpha = 4.5^\circ$. From the LPF satellite telemetry, we obtain Ψ_R and Ψ_F as well as $o1$ and $o12$ but none of the intermediate products.

- For the angular measurement via the DWS technique, the other four F measurements in an interferometer i are used. The raw or electrical DWS phase means:

$$\text{DWS}^\phi = \arg\left(\frac{F_{\text{left}}}{F_{\text{right}}}\right) \quad (2.8)$$

$$\text{DWS}^\eta = \arg\left(\frac{F_{\text{up}}}{F_{\text{down}}}\right), \quad (2.9)$$

as illustrated in Figure 2.6. To obtain the orientation of TM1, the DWS measurements of the X1 interferometer are scaled using experimentally obtained scaling coefficients, the g-coefficients. The measurements are described in [RF12] and [R+13] and the values used

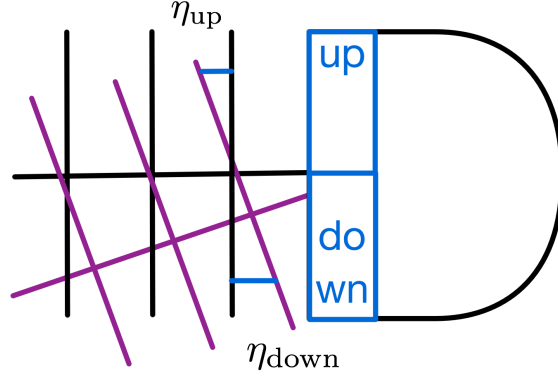


Figure 2.6: The working principle of DWS as seen from the side. The reference beam (sketched in black) hits the PD centre. The measurement beam and its wavefronts (sketched in purple) are tilted with respect to the reference beam and hit the PD slightly off-centre. On the upper two quadrants, a phase signal η_{up} is measured which is different from the phase signal measured by the lower two quadrants, η_{down} , due to the tilt. From this difference, a TM tilt can be deduced.

in Chapter 6 to scale from the TM rotations back to the measured phase are taken from [Wis17]. To measure the pitch, corresponding to η , and yaw, corresponding to ϕ , of TM2, the angular measurements of the X1 and X12 interferometer have to be scaled and combined. This process results in DWS measurements of ϕ_1, ϕ_2, η_1 and η_2 which are available from the LPF satellite telemetry.

- Similar angular measurements can also be obtained using the DC measurements. Then, the horizontal position of the beam on the QPD of an interferometer i (compare also [Wan10, p.201]) is determined from

$$\text{DC}^\phi = \frac{\Sigma_{\text{left}} - \Sigma_{\text{right}}}{\Sigma} \quad (2.10)$$

and the vertical position of the beam is calculated as

$$\text{DC}^\eta = \frac{\Sigma_{\text{up}} - \Sigma_{\text{down}}}{\Sigma}. \quad (2.11)$$

These values can be determined for all four interferometers on LPF. As for the DWS angles, the measurements of the X1 and X12 interferometers can be scaled and combined to a corresponding TM orientation using the respective g-coefficients. The results are called $\eta_1^{\text{DC}}, \eta_2^{\text{DC}}, \phi_1^{\text{DC}}$ and ϕ_2^{DC} and are available from LPF telemetry. They are also known as Differential Power Sensing (DPS) in the literature. The DC angle measurements of the reference and frequency interferometer, however, have no unit.

- A combination of DC and F values is the measurement of the interferometer contrast c_i which is given by

$$c_i = \frac{|F_i|}{\Sigma_i}. \quad (2.12)$$

In addition, the DMU contains many error flags and error handling strategies. The DMU also runs the OMS control loops as explained in Section 2.4. Other telemetry channels used in this

thesis, such as the laser temperature and the measurements which monitor the laser frequency stabilisation for example, are subject to a much simpler processing. Often, this time series data is only downsampled prior to the analysis on Earth. Therefore, this telemetry is listed in Appendix B and not further discussed here. Apart from the telemetry related to the OMS, there is a lot more data available from LPF. In principle, all the subsystems introduced in Section 1.2.1 are accessible via the respective data channels. In this work, we will also use temperature measurements of the OB (compare Section 3.3.2) and the measurements of the relative distance of the two TMs to their respective housings along the x-direction as measured by the GRS and denoted as GRS_x1 and GRS_x2 (compare Chapter 5).

From the LPF satellite to repositories on Earth We have so far discussed the measurements which have been taken by the LPF OMS and how they are processed in the PM and the DMU. As can be seen from Figure 2.1, there are two more steps until the data is available for the analysis of the subsequent chapters.

The measurement data streams are downsampled in the DMU and combined into predefined packets which are transferred from the DMU to the Onboard Computer (OBC) [U. 09].

This process depends on the measurement channel. Some measurements, such as the laser frequency control loop signals, are downsampled from 100 Hz to 1 Hz using a moving average filter and then transferred to the OBC as housekeeping packets. Other signals, such as the o12 measurement, are downsampled from 100 Hz to 10 Hz using a moving average filter and then transferred to the OBC as science data packets. The DMU has its own clock that is synchronised to the OBC by a 1PPS [D⁺18, Section 6.10].

An exception to the standard procedures in the DMU is the use of the Interferometer Data Log (IDL). This option allowed us to record OMS channels at 100 Hz for a limited duration. This duration is determined by the number of channels recorded for the total data storage of the IDL cannot exceed 256 kbytes [U. 09]. This feature has been used for some laser frequency control loop characterisation experiments. However, using this data recording feature requires a more complicated commanding which resulted in one parameter not being properly recorded during one of the laser frequency loop characterisation experiments (see Section 3.1.1).

Among other tasks, the OBC implements the data transfer to ground. In contrast to laboratory experiments on Earth, the data rate is very limited. Consequently, during nominal operations, only carefully selected data packets with sampling frequencies from approximately 0.03 Hz to 10 Hz are transferred to ground.

The AEI LPF laboratory The Albert Einstein Institute (AEI) LPF laboratory hosts the EM of the LPF OB. Therefore, the OMS measurement channels discussed so far are also available for measurements taken in the lab. There are several slight differences in the whole set-up from the laser to the data processing (see [Wit14]). Nonetheless, the longitudinal and DWS measurements are comparable with the advantages that a sampling rate of 256 Hz is usually available and that the data from the A and B diode can be obtained individually if the X1 interferometer measurement is recorded but not the X12 interferometer measurement as there is a limited number of phasemeter channels available. Data from several AEI laboratory measurements have been used for the investigations of Chapter 6.

2.1.3 Noise sources in the OMS

To achieve the required levels of residual acceleration on LPF (see Chapter 1), a high-precision readout of the relative TM positions is required. The detailed requirement will be given in Equation 2.30. Here, the focus is on some of the many noise sources which can impact such a high sensitivity measurement (see [Aud14] and [U. 09] for a more complete list). Several of these have already been mitigated by the design and the construction of the LPF OB. One of these is the common-mode path length noise which is mitigated by the subtraction of the reference interferometer measurement and the path length stabilisation using the OPD control loop (see Figure 2.4). Another possible noise source is temperature fluctuations. Their impact is minimised by the use of ultra-low thermal expansion material for the OB itself and the optical components. The optical components couple twofold to temperature: via thermal expansion of the material resulting in path length changes and changes in refractive index. In addition, also the PDs may be subject to thermal influences. However, these couplings only impact the phase measurements if they are different for the two beams [R⁺05]. In addition, laser amplitude fluctuations deteriorate the measurement performance. First, slow fluctuations of this amplitude cause changes in radiation pressure on the TMs, which results in undesired acceleration noise [A⁺16c]. Second, fluctuations of the laser power relative to the mean power, denoted as Relative Intensity Noise (RIN), around the heterodyne frequency contribute to the noise in the phase measurement. On LPF, the coupling of RIN at the heterodyne frequency is suppressed by the balanced detection scheme. However, a spurious coupling of RIN at twice the heterodyne frequency has been investigated for both the longitudinal and angular interferometric phase measurements [W⁺17][Wis17]. The processing of the PD measurements may also be affected by shot noise at the first stage in the photoreceivers as well as by electronic noise of the processing circuitry and noise which is caused by the signal digitalisation by the ADCs [Aud14].

Another noise contribution is due to laser frequency fluctuations, which are studied in this thesis. Therefore, they will be explained in more detail in the following Sections.

2.2 Coupling of laser frequency noise into the phase measurement

The path length difference between the measurement and the reference beam in an interferometer, Δs , is related to the measured phase, $\Delta\varphi$, by

$$\Delta s = \frac{\lambda}{2\pi} \Delta\varphi \quad (2.13)$$

with λ being the wavelength of the laser light. Using $c = \lambda f$, with c as the speed of light, we can write Equation 2.13 as

$$\Delta\varphi = \frac{2\pi}{c} \Delta s f. \quad (2.14)$$

Hence, laser frequency fluctuations δf cause fluctuations in the measured phase $\Delta\tilde{\varphi}$ [Hei02]

$$\Delta\tilde{\varphi} = 2\pi \frac{\Delta s}{c} \delta f. \quad (2.15)$$

These fluctuations are by principle indistinguishable from a true TM motion, which corresponds to a change of the path length of the measurement beam while the path length of the reference

interferometer remains unaffected. From Equation 2.15, it can be seen that by minimising the static path length mismatch during the integration of the satellite, the impact of the frequency fluctuations decreases proportionally. In addition, the laser itself can be frequency stabilised to allow for a more precise phase measurement. In the early development phases of the LPF mission, it was also investigated to use a free-running laser without active frequency stabilisation, to then record the measured frequency fluctuations and subtract them from the science measurements but this technique was found to be limited in performance [Mar07].

On LPF, these two means of direct minimisation have been combined. The laser frequency stabilisation method is explained in Section 2.4. For the performance of this scheme, we refer to Chapters 3 and 4. The path length difference will be explained in more detail in Section 2.3 and the analysis of the path length mismatches achieved on LPF will be presented in Chapter 5.

2.3 The path length difference

The path length difference relevant to the coupling of laser frequency noise is the difference between the optical path of the measurement and the reference beam in an interferometer.

We have already seen that the telemetry we obtain for the o12 measurement, the o1 measurement and the frequency interferometer measurement has, by default, the reference interferometer measurement already subtracted and that the distance measurements are already scaled to metres.

Hence, let us look at the coupling of laser frequency noise into these combined measurements step by step. Equation 2.15 reads in the X12, XR and XF interferometer individually as

$$\Delta\tilde{\varphi}_{X12} = \frac{2\pi}{c} \Delta S_{X12} \delta f, \quad (2.16)$$

$$\Delta\tilde{\varphi}_{XR} = \frac{2\pi}{c} \Delta S_{XR} \delta f, \quad (2.17)$$

$$\Delta\tilde{\varphi}_{XF} = \frac{2\pi}{c} \Delta S_{XF} \delta f. \quad (2.18)$$

Each of these path length mismatches is the sum of the common path length mismatch occurring in the fibres before the light enters the optical bench and after. This allows us to write:

$$\Delta\tilde{\varphi}_{X12} = \frac{2\pi}{c} (\Delta S_{\text{fibre}} + \Delta S_{X12 \text{ on OB}}) \delta f \quad (2.19)$$

$$\Delta\tilde{\varphi}_{XR} = \frac{2\pi}{c} (\Delta S_{\text{fibre}} + \Delta S_{XR \text{ on OB}}) \delta f \quad (2.20)$$

$$\Delta\tilde{\varphi}_{XF} = \frac{2\pi}{c} (\Delta S_{\text{fibre}} + \Delta S_{XF \text{ on OB}}) \delta f \quad (2.21)$$

Now the o12 measurement as from the LPF telemetry has not only the XR measurement subtracted but is also scaled to TM position. This results in a laser frequency noise contribution of

$$\tilde{o}12 = \frac{\lambda}{4\pi \cos(\alpha)} \left(\frac{2\pi}{c} (\Delta S_{\text{fibre}} + \Delta S_{X12 \text{ on OB}}) - \frac{2\pi}{c} (\Delta S_{\text{fibre}} + \Delta S_{XR \text{ on OB}}) \right) \delta f \quad (2.22)$$

which simplifies to

$$\tilde{o}12 = \frac{\lambda}{4\pi \cos(\alpha)} \frac{2\pi}{c} (\Delta S_{X12 \text{ on OB}} - \Delta S_{XR \text{ on OB}}) \delta f, \quad (2.23)$$

and similar for the frequency interferometer telemetry

$$\tilde{\Psi}_F = \frac{2\pi}{c} (\Delta_{S_{XF \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}}) \delta f. \quad (2.24)$$

So the coupling from laser frequency fluctuations to the science measurement is given by

$$\tilde{o}12 = \frac{\lambda}{4\pi \cos(\alpha)} \frac{\Delta_{S_{X12 \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}}}{\Delta_{S_{XF \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}}} \tilde{\Psi}_F. \quad (2.25)$$

We define the numerator terms as

$$\Delta_{S_{X12 \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}} =: \Delta_{S_{o12}}. \quad (2.26)$$

Analogously, we define

$$\Delta_{S_{X1 \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}} =: \Delta_{S_{o1}} \quad (2.27)$$

for the o1 measurement. Both $\Delta_{S_{o12}}$ and $\Delta_{S_{o1}}$ are required to be below 1 mm [R+13]. In Equation 2.25 the denominator $\Delta_{S_{XF \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}}$ is dominated by the path length mismatch in the XR interferometer on the OB. As mentioned in Section 2.1, the path length mismatch in the frequency interferometer occurs in the fibres before the OB and on the OB itself the path lengths are matched to within construction limits. This is not the case in the reference interferometer where the path of the beams on the OB differ to compensate for the difference in path length occurring in the fibres. In the following, we therefore write

$$\Delta_{S_{XF \text{ on OB}}} - \Delta_{S_{XR \text{ on OB}}} =: \Delta L = 0.382 \text{ m}. \quad (2.28)$$

With these definitions, we can rewrite Equation 2.25 as

$$\tilde{o}12 = \frac{\lambda}{4\pi \cos(\alpha)} \frac{\Delta_{S_{o12}}}{\Delta L} \tilde{\Psi}_F, \quad (2.29)$$

and analogously for o1. The path length mismatch is illustrated in Figure 2.7. Note that this measurement does not allow us to determine where the path length mismatch occurs. As sketched in Figure 2.7, the path length mismatch occurs on the OB. In reality, the path lengths on the OB are very well matched [R+13]. In addition, the path length mismatch in o12 and o1 is subject to the TM position with respect to the OB because a change in position of TM1, for example, also alters the path of the measurement beam in o1 but the path of the reference beam remains unaffected. As a consequence, the path length mismatch in these interferometers is affected by the assembly of the OB and the vacuum tanks with the GRS inside as well as the absolute position of the TM with respect to the electrodes at a certain time. Consequently, the path length mismatch provides a means to measure absolute distances with the interferometry set-up on LPF.

During LPF operations, we obtain the path length mismatches, Δs , with the highest precision during laser frequency modulation experiments. During these experiments, the laser frequency is modulated sinusoidally. Then we estimate the path length mismatch by comparing the signal, for example in o12, to the signal in Ψ_F . This leads to a comparison of peak amplitudes. One example segment is shown in Figure 2.8 which illustrates that the signal is well discernible in both channels. For path length mismatches of the order of several hundred μm , the coupling coefficient is on the order of $10^{-11} \text{ m rad}^{-1}$. This analysis is explained in detail in Chapter 5.




Figure 2.7: A sketch of the reference laser unit, the laser modulation unit and the OB with a path length mismatch in the X12 interferometer placed at an arbitrary position in the beam path and not drawn to scale.

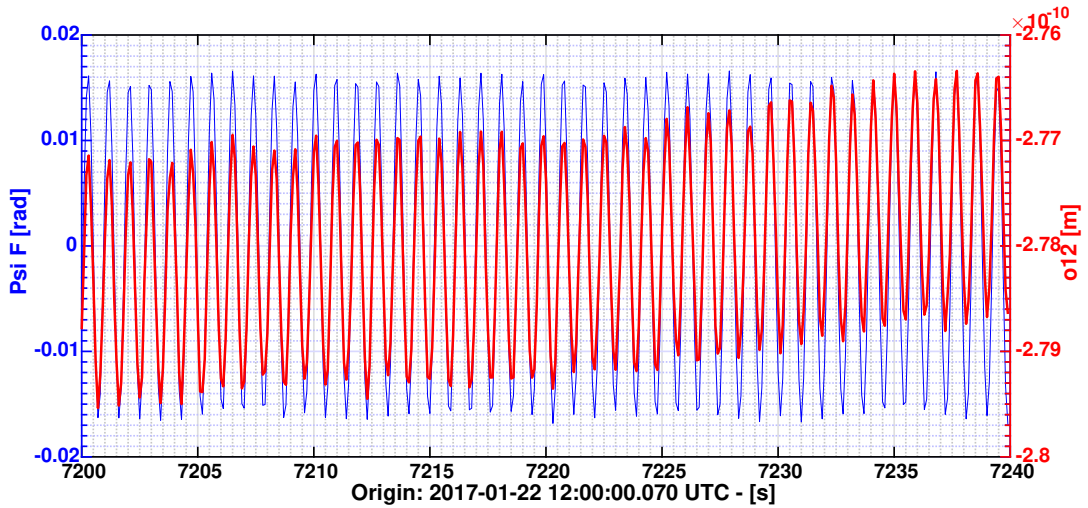


Figure 2.8: A stretch of measurement data of Ψ_F (blue trace) and o_{12} (red trace) during a laser frequency modulation experiment. We can clearly distinguish the modulation in both channels. The path length mismatch $\Delta s_{o_{12}}$ is then inferred from the comparison of the peak amplitudes.

2.4 The laser frequency stabilisation control loop

From Equation 2.15, it can be seen that laser frequency fluctuations lead to phase noise. Even though the laser source used has a comparatively low intrinsic frequency noise level, temperature changes in the laser crystal affect the resonator length and thus the frequency. These may arise due to changes in the ambient temperature and due to changes in the power of the pump light [Aud14]. Thus, to achieve the required laser frequency stability (see Equation 2.36) the laser on LPF is frequency stabilised by a dedicated nested control loop as shown in Figure 2.9. The measurement of the XF interferometer with the reference interferometer measurement already




Figure 2.9: The nested control loop for the laser frequency stabilisation. The components of this loop are explained in the text. The channels marked with the dashed lines are available as telemetry, see Appendix B.

subtracted is fed into the fast frequency controller C_F [U. 09]. It is also possible to inject signals into the control loop at the modulation input. This was used for the experiments explained in Chapter 3.1 and 5.1. However, in contrast to all the other channels marked in Figure 2.9 these signal input channels cannot be downloaded as telemetry. Both controllers are Infinite Impulse Response (IIR) filters running at 100 Hz in the DMU. The coefficients can be found in [P+17] and [Aud14]. The fast frequency controller output is passed to the fast frequency actuator, A_F . This is a PZT acting on the laser crystal located in the RLU. The output of the fast frequency controller is also fed to the slow frequency controller C_S which calculates the necessary voltage to be applied in the slow actuator, A_S , which is a heater. Thus, the long-term fluctuations of the laser frequency are suppressed and the remaining frequency fluctuations stay within the limits of the fast frequency actuator. The current laser frequency noise is then measured in the XF interferometer, the reference phase is subtracted and this signal is then again used as the input to the fast frequency controller. This description of the control loop is based on [Ker11], [P+17] and [Aud14]. The unity gain frequency of this loop was estimated to be 0.78 Hz in the FM test campaign and the slow frequency controller was estimated to dominate up to approximately 0.24 Hz [Aud14].

2.5 Laser frequency noise requirement

The requirement for the readout precision of the whole OMS system is [N. 15]:

$$S_{\delta x}^{1/2} \leq 9 \frac{\text{pm}}{\sqrt{\text{Hz}}} \sqrt{1 + \left(\frac{3 \text{ mHz}}{f}\right)^4}, \quad (2.30)$$

in terms of actual TM displacement sensing noise and *not* in terms of optical path length [N. 15, p. 111] and assuming uncorrelated noise sources (compare [Wan10]). In the early phase of the

development, it was decided that the laser frequency noise contribution to the total OMS noise should not exceed [Ker11][Aud14]

$$S^{1/2}_{\text{frequency}} \leq 2 \frac{\text{pm}}{\sqrt{\text{Hz}}} \sqrt{1 + \left(\frac{3 \text{ mHz}}{f}\right)^4}. \quad (2.31)$$

In addition, it was required that the path length mismatches have to be smaller than [Ker11]

$$\Delta s_{\text{Req}} \leq 1 \text{ cm}. \quad (2.32)$$

From Equation 2.31 and Equation 2.32, we can then derive the laser frequency noise requirement. Frequency fluctuations translate into phase noise with a delay corresponding to the prolonged path of a beam:

$$\delta\varphi = \tau\delta\omega = \frac{\Delta s_{\text{Req}}}{c} 2\pi\delta f. \quad (2.33)$$

Using Equation 2.7, we can rewrite the left hand side to obtain

$$\frac{2\pi}{\lambda} \frac{2\cos(\alpha)}{1} \delta s = \frac{\Delta s_{\text{Req}}}{c} 2\pi\delta f. \quad (2.34)$$

Solving for the frequency fluctuations, we find:

$$\delta f \leq \frac{2\cos\alpha\delta s c}{\lambda\Delta s_{\text{Req}}} \quad (2.35)$$

As a consequence, if we insert the maximum allowed contributions from laser frequency noise to the total OMS noise, $S^{1/2}_{\text{frequency}}$, for δs here, we obtain

$$\delta f \leq 112 \text{ kHz}/\sqrt{\text{Hz}} \sqrt{1 + \left(\frac{3 \text{ mHz}}{f}\right)^4} \quad (2.36)$$

which can also be found in [Aud14] and [Ker11].

2.6 The in-flight laser frequency modulation experiments

During the LPF mission, seven dedicated laser frequency modulation experiments were performed. Five of these aimed at characterising the laser frequency control loop and will be studied in Chapter 3. Two laser frequency modulation experiments were optimised for path length difference estimation. They will be explained and analysed in Chapter 5. An overview of all laser frequency modulation experiments is given in Table 2.1.

experiment	loop	path length	date
chronological	characterisation	difference	
1	1	-	DOY 027: 27-01-2016
2	2	-	DOY 153: 01-06-2016
3	3	-	DOY 164: 13-06-2016
4	-	1	DOY 166: 14-06-2016
5	4	-	DOY 022: 22-01-2017
6	-	2	DOY 022: 22-01-2017
7	5	-	DOY 096: 06-04-2017

Table 2.1: An overview of all laser frequency modulation experiments performed during the LPF mission. The first column gives the chronological experiment number. We have five experiments which aim at loop characterisation. They are chronologically numbered in the second column. The third column represents the numbering of the two experiments which are optimised for path length mismatch estimation. The DOYs are starting at 08 UTC in the convention used here which accounts for the difference in DOY and date from experiment 3 to 4.

In this chapter, we introduced in some detail the LPF OMS and listed the components of this subsystem as well as the measurements taken on board LPF. Furthermore, we showed how these measurements are processed into the telemetry channels which were used for the analyses in the subsequent chapters. The coupling of laser frequency noise into the measurement of changes in the relative position of the two TMs was found to depend on the path length mismatch Δs between the measurement and reference beam. Consequently, both the path length mismatch and the laser frequency fluctuations should be minimised. Therefore, the nested digital control loop, which served to stabilise the laser frequency on LPF, was summarised in this chapter. The performance of this control loop will be the subject of the following chapter and the achieved path length mismatches will be analysed and discussed in Chapter 5.

Chapter 3

Control loop characterisation

The following chapter presents the laser frequency control loop characterisation experiments. We used periodic excitation signals to characterise the transfer function of the loop and its components. We begin with a description of the test signals and the experiment settings. After a small intermezzo on the error estimation for these transfer functions, we will then inspect the results over the course of the mission compared to the measurements from a test campaign performed on ground.

3.1 Design of loop characterisation measurements

The aim of each of these loop characterisation experiments is to ensure that the laser frequency stabilisation works as expected from the measurements and design on ground. By repeating the measurements over the mission, we want to check for possible degradations of the actuators over time. In each of the experiments, the laser is frequency modulated with a sinusoidal test signal. It is applied at the modulation input before the fast frequency controller C_F , see Figure 2.9. We used frequencies, amplitudes and durations as listed in Table 3.1. These parameters

number	frequency [Hz]	amplitude [rad]	duration [N half cycles]
1	0.011	0.05	40
2	0.0396	0.05	40
3	0.0852	0.05	40
4	0.1421	0.05	40
5	0.237	0.05	40
6	0.3953	0.05	40
7	0.6594	0.05	40
8	1.123	0.05	40

Table 3.1: The frequencies, durations and amplitudes for the fast frequency loop injections for loop performance assessment.

follow the recommendations from the test campaigns [Aud14, p. 165] with the exception of the

highest modulation frequency. This has been shifted from 1 Hz to 1.123 Hz due the presence of a known peak in the o12 spectrum, similar to those reported for the angular measurements in [Wis17, p.30]. For the frequency loop characterisation itself, the line in o12 is irrelevant but as some of the loop characterisation experiments have been conducted with free-falling test masses, they can in principle also be used to estimate path length mismatches. The frequencies have been selected for three reasons [Aud14, p. 168]. First, they are not harmonics of each other and second, they are distributed across the LPF measurement band and above. The third reason to select these was to avoid using the IDL (see Section 2.1). This turned out to be impossible in flight because the in-loop quantities were recorded at 1 Hz only during the mission. However, no systematics using the IDL have been observed. The modulation amplitudes have been selected with respect to the actuator ranges and the duration was a result of the desired fractional error with margin added [Aud14, p. 167/168].

3.1.1 Overview of the measurements performed

Over the course of the mission, this experiment was repeated five times. However, the number of modulation frequencies that we applied or the data recording changed with each experiment. This is summarised in Table 3.2. Note that the Days Of Year (DOYs) are following the convention that a new DOY starts at 8 UTC. For all the experiments in which the IDL was not used, the modulation frequencies seven and eight were not applied. In the second experiment, only the first modulation was applied successfully. Due to a commanding error, the others were not applied. In experiment four, the IDL was not configured properly, resulting in the slow frequency controller output not being recorded. Hence, in this experiment we can determine the fast frequency controller transfer function and the Open-Loop Transfer Function (OLTF) also for frequencies seven and eight but not the slow controller transfer function. Note that we can use the frequency loop characterisation experiments to estimate the path length mismatch, see Chapter 5, but we cannot use the arm length mismatch investigations to investigate the control loop performance because the loop parameters are recorded at 1 Hz unless the IDL is used whereas the modulations in the path length mismatch experiments have been performed with 1.123 Hz and 2.879 Hz, respectively.

3.2. ANALYSIS PROCEDURE FOR THE FREQUENCY CONTROL LOOP CHARACTERISATION MEASUREMENTS

exp number	date	satellite configuration	IDL use	comment
1	DOY 027: 27-01-2016 11:15:00 UTC	TMs grabbed	x	highest frequency of 0.3953 Hz
2	DOY 153: 01-06-2016 12:05:10 UTC	TMs drag-free	(✓)	only first injection successful
3	DOY 164: 13-06-2016 06:05:10 UTC	TMs drag-free	✓	
4	DOY 022: 22-01-2017 10:35:10 UTC	TMs drag-free	(✓)	IDL not configured properly
5	DOY 096: 06-04-2017 22:30:08 UTC	TMs grabbed	✓	adapted to also record the injection at 0.3953 Hz with IDL

Table 3.2: An overview of the laser frequency modulations for loop performance assessment. The experiment number is the loop characterisation experiment number (see Table 2.1). The time indicates the start of the experiment.

3.2 Analysis procedure for the frequency control loop characterisation measurements

To characterise the laser frequency control loop, we estimate, for each experiment (compare Figure 2.9):

1. the fast frequency controller transfer function, T_f , from the fast frequency error signal to the fast frequency controller output;
2. the slow frequency controller transfer function, T_s , from the slow frequency error signal to the slow frequency controller output;
3. the response in terms of the gain and delay of the fast and slow actuator, A_f and A_s ;
4. and the open-loop transfer function, T_{OL} , from the fast frequency error signal to the frequency interferometer measurement Ψ_F .

Each analysis is slightly different and outlined in Sections 3.2.3 to 3.2.5. Let us begin with explaining the timing correction procedure and the estimation of transfer functions and their errors. These steps are common to all four analyses.

3.2.1 Timing correction procedure

To understand the timing correction procedure, some insight into the data processing on LPF and on ground is necessary. A complete list of the relevant telemetry can be found in Appendix B. The DMU obtains 100 Hz data from the phasemeter and operates itself at 100 Hz. For the science telemetry, as for example the Ψ_F measurement, the 100 Hz data is converted to 10 Hz data by the use of a moving average filter. This means that ideally, 10 samples are averaged into a single one. For the Ψ_F measurement, this description applies only to selected times of the mission, mostly when OMS experiments had been performed. For a large part of the mission, Ψ_F was downloaded at 1 Hz (see Section 4.2). For the housekeeping telemetry, as for example the fast frequency controller error channel, the 100 Hz data is directly converted to 1 Hz data (see [U. 09, p. 4-31]), again via a moving average filter. Up to here, the data is taken with respect to the DMU time. Then, the science data channels at 10 Hz are timestamped according to OBC time. This science data and the unsynchronised 1 Hz housekeeping data are then passed on to the OBC. From the OBC the requested data packages are telemetered to ground. In general, the data gets corrected for the difference between on-board computer time and UTC.

Two problems may arise due to the data processing in the DMU:

- As the housekeeping telemetry is a low-priority task in the DMU, the times when the science and housekeeping data recording processes are undertaken may be offset to each other. This offset may also change with time, depending possibly on the total workload in the DMU.
- The moving average filter is not an ideal low-pass filter and thus may also cause aliasing, see also [Wis17].

With the exception of aliasing, these problems are specific to the analysis of the OPD and frequency loop characterisation experiments, for only in these experiments science and housekeeping telemetry are compared to each other.

The idea is now to find the timing offset between the science and the housekeeping data and then to correct the times of science data channel to match those of the housekeeping. Here are the steps of the timing correction procedure¹:

- If possible, that is if the OBC to UTC correlation parameter is available, the science data is converted back to the OBC time.
- The time differences between the data samples of the thus converted science data is then inspected for the duration of the experiment. It is assumed to differ only from the nominal sampling time of 0.1 s to either 0.09 s or 0.11 s which corresponds to one sample too many or too few of the 100 Hz data being taken. Possible outliers due to other anomalies would be detected at this step. Thus, the science data channel is set on an unevenly sampled grid with only these time steps.
- The science data channel is interpolated to an evenly sampled 10 Hz grid. This means, we assume that the DMU clock originally was running correctly and that the difference to the nominal sampling time of 0.1 s is due to the difference between DMU clock and OBC clock which run independently and hence at different rates. So, in a way of speaking, we resample this data to DMU time.

¹as developed by M. Born

3.2. ANALYSIS PROCEDURE FOR THE FREQUENCY CONTROL LOOP CHARACTERISATION MEASUREMENTS

- To mimic the behaviour of the DMU and to obtain comparable channels, the science data is filtered with a moving average filter.
- The housekeeping data is assumed to have been taken without timing problems at DMU time. Therefore, the time of the data points is set to a fixed sampling time of 1 s without changing the data values themselves.
- The 10 Hz data is downsampled to 1 Hz without any additional filters being applied but with different offsets. This can be understood as selecting for example the samples number 1, 11, 21, then with some offset the samples 3, 13, 23 and so on. Note that this may lead to resulting data that is one or two samples shorter than the original data. This can be understood by the following example: Assume our data length of the 10 Hz data is 30.2 s. If we then arrive at an offset of two, that is we select the samples 3, 13, 23, . . . , the last one will not be available any more which leads to one sample missing.
- This downsampled data is corrected for the offsets introduced in the selection process.
- The resulting data is compared to the HK data for each of the offsets.
- By inspection, the offset corresponding to the maximum in correlation is taken.
- To verify the result, the transfer function from one housekeeping to a science channel is checked for several offsets. The correct offset should have a flat phase and amplitude up to the Nyquist frequency.
- The HK data is split to match the length of the now possibly one or two samples shorter science channel.

The correlation and transfer function checks are performed at the beginning and end of the experiment. If then the offset determined remained fixed, we can safely assume that it hasn't drifted during this experiment and then use a single offset to correct the data. This is a cross check because a drift of this offset and especially drifts in opposite directions would require anomalies in the DMU itself which would have been detected by the first check of the time difference between the data samples already. Additional information can be found in [D⁺18].

This is one example where the use of different clocks on-board the LPF satellite causes difficulties for the analysis. In this case, we can correct for the timing differences between two channels at the cost of significant data analysis efforts but of course we cannot correct for amplitude errors that result from not taking the right number of samples in an average. However, this procedure is not immediately transferable to cases where we have to compare data from the DMU and the Inertial Sensor (IS) which is again synchronised with the DMU [SW11, p.33]. Thus, we would like to emphasize that a single clock per satellite is important for future missions of this kind.

3.2.2 Transfer function estimation and their errors

There are several possibilities to estimate the transfer function T from a set of sinusoidal test signals applied at known frequencies. One of them is to compute the Fourier transform of both input I and output signal O at the known modulation frequencies f_{mod} and then take the ratio of these:

$$T(f_{\text{mod}}) = \frac{\mathcal{F}(O)(f_{\text{mod}})}{\mathcal{F}(I)(f_{\text{mod}})}. \quad (3.1)$$

One way to estimate the random error of $\mathcal{F}(O)$ is to estimate the noise level. This can be done by computing the Fourier transform slightly above and below the signal frequency. This is implemented in the LTPDA method `ao/dft`. Provided the noise is flat or of a simple shape and there are no disturbing lines, this yields a reasonable estimate. However, this is a single error for amplitude and phase error combined. If this is propagated through the division of the complex quantities, \mathcal{F} , under the assumption of Gaussian noise, the result is a complex quantity and thus not a standard deviation of a Gaussian-distributed quantity.

Alternatively, we chose to estimate the uncertainty on the transfer function, T , via the coherence. This approach also works for the case where the excitation signal is noise and not a periodic signal. It has also been used in [N+13] following [BP80]. We apply

$$\sigma_T \simeq \frac{[1 - \gamma_{OI}^2(f)]^{\frac{1}{2}}}{|\gamma_{OI}(f)|\sqrt{2n_d}} \quad (3.2)$$

with n_d being the number of averages and the coherence function estimate γ_{OI}^2 being calculated as

$$\gamma_{OI}^2(f) = \frac{|G_{OI}(f)|^2}{G_{OO}(f)G_{II}(f)}. \quad (3.3)$$

Here, G_{OI} is the cross-power spectral density of the output and input. G_{OO} and G_{II} are the Power Spectral Density (PSD) of the output and input signal, respectively. This error estimate has been implemented in `ao/tfe`. We will use this approach because it takes into account the signal coherence intrinsically and because it is the most general.

Apart from random errors, transfer function estimates can be biased. One mechanism is if the frequency under study doesn't fall into the exact centre of a bin, signal power will leak into adjacent bins. This is another expression for the signal not being periodic in the given observation time which means that discontinuities are present at the beginning and end of the available data which leads to a bias in the spectral estimate across many frequencies, a phenomenon called spectral leakage. This can be counteracted by the use of a window function [HRS02][Har78].

For the transfer function estimates of this chapter, a Hanning window was chosen. Even though its amplitude error is larger than those of flat top windows, compare [HRS02], it would cancel anyway in the transfer function estimation. In addition, we choose the number of data points in each Fourier transform such that the non-integer modulation frequencies fall, as closely as possible, into the centre of the bin. This minimises the spectral leakage and amplitude error anyway, independent of the window used.

During the laser frequency control loop characterisation experiments, we are in the beneficial situation of estimating the amplitude of a clearly discernible peak in the presence of noise without adjacent lines which would affect the choice of the window function. To illustrate this claim, we show the ASD of the Ψ_F measurements during the third laser frequency modulation experiment in Figure 3.1. Here we plot the ASD using a single average per modulation segment and a Blackman-Harris window. Even though the peak amplitudes are biased in the ASD, this plot illustrates that we have good SNR in this experiment and no unwanted spectral lines. We have calculated the SNR according to Equation 16 in [LC09] as

$$\text{SNR} = \frac{\text{AS}(f_{\text{mod}}) \cdot (\sqrt{2})\sqrt{T}}{\sqrt{\text{S}(f_{\text{mod}})}}. \quad (3.4)$$

Here, AS denotes the amplitude spectrum of Ψ_F evaluated at the modulation frequencies f_{mod} . This result is in rms units and, since we are dealing with sine waves, we can convert this result

3.2. ANALYSIS PROCEDURE FOR THE FREQUENCY CONTROL LOOP CHARACTERISATION MEASUREMENTS




Figure 3.1: The ASD of Ψ_F during the laser frequency modulations of experiment 3. Even though the peak amplitudes are biased in the ASD, this plot illustrates that we have good SNR in this experiment and no unwanted spectral lines.

to peak unit by a multiplication with a factor of $\sqrt{2}$. T is the duration of the modulation in seconds. S is the PSD of the Ψ_F noise floor, estimated from the plot at the respective modulation frequency. For this experiment it was found to be above 300 for all modulations. This is expected to be comparable for the other modulation experiments because the commanding is the same and the noise floor changes only slightly (see Chapter 4) and thus has not been further investigated.

Indeed, for a given record of data, the bins of a power spectral density estimate or the transfer function estimate here are located at integer multiples ($n = 1, \dots, \frac{1}{2}f_s$) of

$$\Delta f = \frac{f_s}{N} \quad (3.5)$$

with f_s being the sampling frequency of the data and N the number of data samples [HRS02]. Now for a given maximum data length, N , and a given sampling frequency f_s , we can chop N for each modulation segment in such a way that a bin falls at or very close to the respective modulation frequency. Requiring a minimum of two averages, we could find the optimum by comparing all possible values of N . Indeed, for each modulation frequency, all possible data lengths and the resulting Δf have been calculated. For each N , we found the bin n_{opt} such that $n_{\text{opt}}\Delta f$ becomes closest to the respective modulation frequency. The minimum of these provided the best choice of N for a given modulation frequency and sampling frequency. Finally, for each modulation frequency and sampling frequency, we only need to get the value and the error from the transfer function, T , at the bin n_{opt} for the best choice of data points in each average N . Note that we can be sure that the commanded modulation frequency was indeed applied due to the signal demodulation which has been done for path length mismatch estimates, see Section 5.2. The locations of the frequencies are compared in Figure 3.2. We find a good agreement between the frequencies so we can proceed using this method and the uncertainties provided.




Figure 3.2: An example of the open-loop transfer function computed using the DFT at the precise modulation frequencies and the best bins. The agreement is sufficient.

3.2.3 Controller transfer function measurement procedure

The procedure is the same for the fast and the slow frequency controller and consists only of the two previously explained steps:

1. Apply the timing correction to the respective controller error and output signal data;
2. Estimate the transfer function values as described in the previous Section 3.2.2.

3.2.4 Actuator characterisation procedure

The actuator characterisation is not so obvious because we do not measure the two actuator outputs separately, we only measure their combined impact as can be seen in Figure 2.9. The fast frequency controller output signal is available on ground and the frequency interferometer measurement Ψ_F . However, from the control loop diagram, we see this transfer function is given in the frequency domain by

$$T_{\text{ff}_{\text{out}} \rightarrow \Psi_F} = T_{A_{\text{fast}}} + C_s \cdot T_{A_{\text{slow}}} . \quad (3.6)$$

$T_{\text{ff}_{\text{out}} \rightarrow \Psi_F}$ can be obtained by transfer function estimation as previously described. We obtain C_s from the respective analysis explained in Section 3.2.3. $T_{A_{\text{fast}}}$ and $T_{A_{\text{slow}}}$ are the transfer functions of the two actuators. These transfer functions are modelled as

$$T_{A_{\text{fast}}} = 3.88 \text{ MHz V}^{-1} \quad (3.7)$$

and

$$T_{A_{\text{slow}}} = \frac{5.42 \cdot 10^8}{s + 2\pi^2} \text{ Hz V}^{-1} . \quad (3.8)$$

Here, s is a complex frequency in units of rad s^{-1} . The slow frequency actuator transfer function has a pole associated with the frequency of π Hz. The models are based on [Aud14] and [Ker11, Figure 6.4-1] with an adaptation of the pole location of the slow frequency actuator and the gains. The gains are now modelled as the results from the FM campaign [Aud14]. However, this model fits the data recorded during the FM test campaign and in-flight best while providing reasonable positive delays. We characterise them on behalf of two parameters, a gain G_{A_f}/G_{A_s} and a delay d_{A_f}/d_{A_s} . Thus we fit

$$T_{\text{ff}_{\text{out}} \rightarrow \Psi_F} = T_{A_{\text{fast}}}(G_{A_f}, d_{A_f}) + C_s \cdot T_{A_{\text{slow}}}(G_{A_s}, d_{A_s}) \quad (3.9)$$

to measurements of $T_{\text{ff}_{\text{out}} \rightarrow \Psi_F}$ at the modulation frequencies. We write the actuator transfer functions as

$$T_{A_{\text{fast}}} = G_{A_f} e^{-s d_{A_f}} \quad (3.10)$$

and

$$T_{A_{\text{slow}}} = \frac{G_{A_s} e^{-s d_{A_s}}}{s + 2\pi^2} \text{Hz V}^{-1}. \quad (3.11)$$

This is a non linear minimisation problem in frequency domain which we solve by combining a Nelder-Mead simplex and a MCMC approach. More details can be found in Appendix A.

3.2.5 OLTF procedure

The OLTF can be measured by computing the transfer function from ff_{err} to Ψ_F as

$$\text{OLTF} = T_{\text{ff}_{\text{err}} \rightarrow \Psi_F} = \frac{\mathcal{F}(\Psi_F)(f_{\text{mod}})}{\mathcal{F}(\text{ff}_{\text{err}})(f_{\text{mod}})} \quad (3.12)$$

and previously described. Alternatively, we can use the previously determined actuator and controller transfer functions and calculate the thus expected OLTF as

$$\text{OLTF} = T_{\text{ff}_{\text{err}} \rightarrow \Psi_F} = (A_s C_s C_f + A_f C_f) \text{ff}_{\text{err}}. \quad (3.13)$$

Both methods have to agree.

3.3 Results of control loop characterisation measurements

3.3.1 Controller transfer function measurements

The loop characterisation measurements allow us to study the individual controllers and actuators of the nested laser frequency stabilisation loop (see Figure 2.9). Let us begin with the fast controller transfer function. The measured results are shown as dots in Figure 3.3. The black trace is the controller model. It is based on the filter coefficients as given in the references in Section 2.4. We note that the measurements agree to each other within the errors. They also agree well to the controller model. This means that the fast frequency controller performance has not changed from ground to space and that it remained constant from June 2016 to April 2017. This agrees to expectations given that the fast frequency controller is implemented digitally in the DMU.




Figure 3.3: Estimate of fast frequency controller transfer function. Here we show all results of the five loop characterisation measurements compared to the controller model. The controller behaves as expected and does not change in the course of the mission.

The slow frequency controller results are shown in Figure 3.4 and compared to the controller model in black. As for the fast frequency controller, the measurements agree well to each other and the model. This allows us to confirm that the slow frequency controller has remained unchanged since the ground measurements as well.




Figure 3.4: Estimate of slow frequency controller transfer function. Here we show all results of the five loop characterisation measurements compared to the controller model. The controller behaves as expected and does not change in the course of the mission.

3.3.2 Actuator gains and delays

A model has been fitted to the timing-corrected measurements of the transfer function $T_{\text{ff.out} \rightarrow \Psi_F}$ as described in Section 3.2.4. The results are shown in Table 3.3. The chains of the MCMC algorithm, the parameter covariance matrices, the estimated parameter distributions and the fit residuals are shown in Appendix A for each experiment. The measurements taken during the

experiment	G_{A_f} [MHz V ⁻¹]	d_{A_f} [s]	G_{A_s} [100 MHz V ⁻¹]	d_{A_s} [s]
ground	$4.51 \pm 7 \cdot 10^{-3}$	$0.04 \pm 0.4 \cdot 10^{-3}$	$5.36 \pm 9 \cdot 10^{-3}$	$0.24 \pm 4 \cdot 10^{-3}$
1	$4.88 \pm 28 \cdot 10^{-3}$	$0.08 \pm 2 \cdot 10^{-3}$	$5.33 \pm 9 \cdot 10^{-3}$	$0.41 \pm 12 \cdot 10^{-3}$
2	nA	nA	$5.37 \pm 8 \cdot 10^{-3}$	nA
3	$4.38 \pm 13 \cdot 10^{-3}$	$0.04 \pm 0.6 \cdot 10^{-3}$	$5.28 \pm 6.5 \cdot 10^{-3}$	$0.2 \pm 6.3 \cdot 10^{-3}$
4	$4.32 \pm 11 \cdot 10^{-3}$	$0.04 \pm 0.4 \cdot 10^{-3}$	$5.34 \pm 6.9 \cdot 10^{-3}$	$0.18 \pm 5 \cdot 10^{-3}$
5	$4.36 \pm 10 \cdot 10^{-3}$	$0.04 \pm 0.4 \cdot 10^{-3}$	$5.4 \pm 7 \cdot 10^{-3}$	$0.19 \pm 5 \cdot 10^{-3}$

Table 3.3: The results for the actuator parameters from the FM ground test and each of the loop characterisation experiments. We mostly find only slight parameter changes over the experiments.

FM test campaign, have been fitted again using the exact same model and minimisation function as has been used for the in-flight data. In experiment 2, only the slow frequency loop gain value is reported because in this experiment only the modulation at 11 mHz was executed successfully. We applied the same data processing and fit to this experiment but the fast frequency loop parameters cannot be reasonably estimated from this modulation. The slow actuator delay is correlated to the two fast parameters and is therefore also not reported here. However, we want to include all data available.

The values reported are the mean values of the MCMC chains and the uncertainties are the calculated standard deviations. This is a reasonable approximation given the symmetric probability density functions for the individual parameters as shown in Appendix A. Our parameter estimation is limited here by fitting 4 correlated parameters to six to eight data points only. Therefore, it was decided to neglect the correlation between the parameters for the uncertainty estimation. This experiment could be improved by additional laser frequency modulations using the IDL at above 1 Hz to further disentangle the actuators, provide more data points for the fit and to be able to estimate the gain margin, too.

On DOY 13 2016, a single triangular input signal has been applied to the input of the fast and the slow frequency controller. By comparing the changes in the controller output signals to the changes in the in-loop measurement of the laser frequency, the fast actuator gain has been estimated to be $(5.6 \pm 0.4) \text{ MHz V}^{-1}$ and the slow frequency actuator gain $0.53 \pm 0.01 \text{ GHz V}^{-1}$. This fast actuator result is in agreement to three of the results reported in Table 3.3 at the

3σ level whereas the agreement of the slow frequency controller estimate from the triangular test signal is at 1σ with all six results obtained from the sinusoidal modulation experiments. Consequently, this is a confirmation obtained with a different but less precise experimental and data analysis method.

The estimated gain of the fast frequency actuator changes by up to approximately 12%. The estimate of the delay of the fast frequency actuator is the same for all experiments except for experiment 1 in flight. The slow frequency actuator gain estimate changes are below 5% throughout the experiments. The measurement result on ground and from experiment 2 as well as from experiments 1 and 4 are in agreement to each other. The delays of the slow frequency actuator are around 0.2s for experiments 3 to 5 but approximately twice this number for experiment 1. This difference in both actuator delays of experiment 1 compared to the other experiments is believed to be caused by the lack of modulation data above 0.3953 Hz in experiment 1. Thus, the fast frequency actuator gain estimate may also be more difficult in this experiment. This gain estimate, however, is correlated to both delays (compare Figure A.5). Such a correlation was also observed by repeating the parameter estimation while either both delays or both gains remained fixed. Fixing for example the actuator gains to the values of the other experiments yielded delays very close to the results of the other experiments and vice versa.

Even though no changes in the actuator parameters were expected, the slight variations observed may perhaps be caused by different thermal environments. Indeed, already in the commissioning phase it was found that the laser output power and consequently the laser behaviour depends strongly on the RLU temperature and is only stable for certain environments, as shown in Figure 3.5. This led to the idea that also the actuator parameters may depend on the thermal environment. Table 3.4 lists averaged selected on-board temperatures for the frequency loop characterisation experiments.

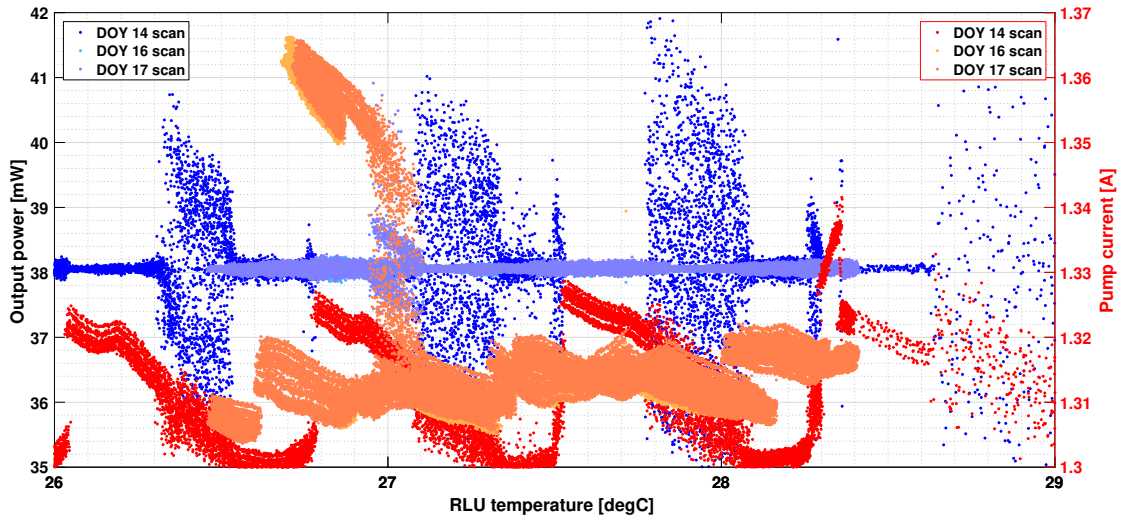


Figure 3.5: The laser output power and the pump current for different temperatures of the RLU as measured during the commissioning phase in January 2016. We note that the laser pump current and output power are stable for certain temperatures but not for others. Hence, the RLU temperature had to be carefully adjusted at the beginning of the mission.

exp number	1	2	3	4	5
RLU T [°C]	28.4	28.2	28.2	23.3	22
AOM T [°C]	25.3	22.6	22.6	23.5	20.1
TS 13 [°C]	nA	21.1	21.4	21.9	12
TS 14 [°C]	nA	21.3	21.4	21.8	11.9
TS 15 [°C]	nA	21.3	21.4	21.8	11.9
TS 16 [°C]	nA	21.1	21.2	21.6	11.8

Table 3.4: The averaged temperatures at different positions on-board the satellite during the laser frequency loop characterisation experiments. During experiment 1, several temperature measurements have not been available (nA).

The RLU T is the temperature measured in the RLU unit, the AOM T is the temperature measured in the LMU in between the two AOMs and the temperatures TS 13-16 are measured on the OB as shown in the top panel of Figure 1 in [A⁺19f]. The selected OB temperature measurements reported here are in agreement to the measurements shown in Figure 2 of [A⁺19f] which also explains the changes in temperature in the course of the mission.

The RLU temperature changes the most from experiment 3 to 4. The largest change in the AOM temperature can be found in between experiments 1 and 2 as well as four to five. The temperature sensors TS 13-16 on the OB are more or less constant, except for the 5th experiment where the temperature is significantly lower. The temperatures are comparable for experiments 2 and 3 but in experiment 2, many parameters are not available. So the comparison is limited to the slow actuator gain of these two experiments which are close to each other but not quite in agreement on the 3σ level. Thus, we conclude that the temperature may be a cause for the actuator parameter fluctuations but it is likely there may also be other reasons which could not be identified yet.

In Figure 3.6, we show the response of the fast actuator and the combined response of the slow controller and actuator for experiment 3 as an example. We see that the slow frequency controller and actuator are dominating up to 0.1 Hz. At frequencies above this, the fast frequency controller transfer function is dominating. This frequency can be read off to be approximately 0.23 Hz in [Aud14, p.153]. This slightly different value may be caused by the use of a different slow frequency actuator pole in the model. However, as shown in Appendix A, the model used here shows very small residuals. These are not known for the other pole location, so we will proceed with this model.




Figure 3.6: A plot of the response of the fast frequency actuator (blue trace) and the slow frequency controller and actuator combined (red trace) as well as the combined response of the two (magenta trace) and the measured data (black crosses). The uncertainty estimates are not discernible on this scale. This is an example using the data from experiment 3. We note that the slow frequency controller and actuator are dominating up to 0.1 Hz.

3.3.3 OETF estimate

Figures 3.7 to 3.10 show the OETF estimate for the nested laser frequency control loop. For each experiment, we show the estimated OETF from the measured signals following Equation 3.12 and from the control loop parameter estimation results according to Equation 3.13. From the single modulation of experiment 2, no transfer function can be shown.




Figure 3.7: The measured OLTF of experiment 1 compared to the model of the loop evaluated with the fit parameters for the actuator values of this experiment.




Figure 3.8: The measured OLTF of experiment 3 compared to the model of the loop evaluated with the fit parameters for the actuator values of this experiment.




Figure 3.9: The measured OLTF of experiment 4 compared to the model of the loop evaluated with the fit parameters for the actuator values of this experiment.




Figure 3.10: The measured OLTF of experiment 5 compared to the model of the loop evaluated with the fit parameters for the actuator values of this experiment.

Throughout all four experiments, the two ways of estimating the OLF agree well to each other. Figure 3.11 shows all OLF estimates compared to each other and the results from the latest tests on ground. As expected, they agree to each other within the errors. From Figure 3.8, we




Figure 3.11: A comparison of the OLFs obtained via the TFE of all five experiments. As expected, they agree well to each other. The result of the latest tests on Earth is shown in grey. The measured transfer functions in-flight agree well to this result, too.

find a unity gain frequency of 0.8 Hz and the phase margin is 77° . As all OLFs of the different experiments are comparable, these values are representative. To determine the gain margin, the laser would have to be modulated at higher frequencies, too. However, extending the OLF model to higher frequencies allows us to estimate the gain margin to be approximately 19dB. Given the reading-off accuracy, the unity gain frequency and the phase margin agree well to the results from the FM test campaign where a unity gain frequency of 0.78 Hz and a phase margin of 74.5° had been found [Aud14, p. 157].

Thus, we conclude that the laser frequency control loop operates as expected from the FM ground test campaign. Fluctuations in the actuator parameters and the OLF over the mission are small and no degradation could be identified. As already suggested, this type of experiment could be further improved by increasing the number of modulation frequencies above 1 Hz.

On the previous pages, we analysed the data of five laser frequency control loop characterisation experiments performed in the course of the LPF mission. The controller transfer functions, the OLFs and the actuator parameters were comparable to the results from the test campaign data analysed in the same way and over the course of the mission. Slight changes in the actuator parameters may be related to changes in the temperatures of the different units on board but there may also be other reasons not identified yet. In summary, we can say that the laser frequency stabilisation worked as expected.

Chapter 4

Laser frequency noise characterisation

In the previous chapter, we showed that the nested laser frequency control loop works as expected from the test campaigns. Now, we will study the resulting laser frequency noise performance achieved with this stabilisation technique. We will explore measurements of the residual laser frequency fluctuations over the course of the LPF mission. In addition, we will compare the laser frequency noise when the stabilisation was active to periods where the stabilisation was switched off as well as to test campaign measurements.

4.1 Impact of the reference interferometer subtraction

The laser frequency measurement on LPF as obtained from the satellite telemetry has the reference interferometer output already subtracted: $\Psi_F = \text{PT}(\varphi_F - \varphi_R)$ (see Section 2.1). Here, the impact of this subtraction is investigated. Therefore, we will show an example of the laser frequency fluctuations for a selected timespan where the data was recorded at a sampling frequency of 10 Hz. At that time, the laser frequency stabilisation control loop and the OPD control loop were active. The data segment has been chosen such that no glitches are discernible in the low-passed data.

Note that we can only approximate the φ_F signal because the subtraction is computed at 100 Hz inside the DMU whereas there is only the downsampled 1 Hz or 10 Hz data available during most of the mission. In addition, φ_R is not available from telemetry so that we use Ψ_R instead.

In Figure 4.1 the Ψ_F data where the reference interferometer output has already been subtracted is shown in red and an estimate of the measurement prior to this subtraction φ_F is shown in blue. We find that the approximate φ_F measurement, where Ψ_R has been added again, has a larger fluctuation amplitude than the frequency interferometer measurement Ψ_F , as expected. Obviously, the subtraction seems reasonable.

In Figure 4.2 the amplitude spectral density of the time series data of Figure 4.1 is shown. A Blackman-Harris window was used and 20 averages have been taken, resulting in uncertainty estimates which are not discernible at this scale. In Figure 4.2 we observe that the common




Figure 4.1: Example time series to illustrate the effect of the subtraction of the reference interferometer channel from the frequency interferometer measurement. The blue trace is an approximation of the phase of the frequency interferometer measurement directly converted to frequency fluctuations. In red, the same measurement with the reference interferometer subtracted and then converted to frequency fluctuations is shown. During this measurement, the laser frequency is stabilised. We note an increase of the laser frequency noise towards the end of the measurement period shown. Such increases will be analysed in Section 4.2.




Figure 4.2: Amplitude spectral density of the time series data of Figure 4.1. The common mode noise dominates at the lower and higher end of the frequency spectrum if the reference interferometer output is not subtracted.

mode noise, as measured by the reference interferometer, dominates at the lower and higher end of the frequency spectrum if the reference interferometer output is not subtracted (blue trace). At frequencies above approximately 1 Hz, the common mode path length noise is increasing while the true frequency fluctuations are decreasing and hence the subtraction is significant. There are two possible reasons for this increase: one reason could be that with the 10 Hz data, we cannot access the spectrum above 5 Hz and thus we do not note that we are affected by a servo-bump around the expected unity gain frequency of 3.1 Hz [G. 11]. A second possible reason is that aliasing could also affect the Ψ_R channel. However, the Ψ_R spectrum above 1 Hz is still under investigation.

Since the data shown in Figures 4.1 and 4.2 is only a single example, it may be possible that the impact of the subtraction is time-dependent. To assess this possibility, the Ψ_F and φ_F spectra, obtained using a Blackman-Harris window, at seven different times during the mission, when the Ψ_F and Ψ_R channels were available at a sampling rate of 10 Hz, have been compared. A difference between Ψ_F and φ_F above 1 Hz was discernible in all of the spectra examined but the difference between the two spectra ranged from approximately a factor three to eight in power at around 2 Hz. At frequencies below 10 mHz, the subtraction made a difference in some cases, in others not. Possibly, this could be due to different levels of fluctuations in the reference interferometer but this analysis is beyond the scope of this thesis.

Consequently, the subtraction of the reference interferometer measurement is necessary to recover the laser frequency fluctuations for the whole frequency range of interest.

4.2 Laser frequency fluctuations over the course of the mission

In this paragraph, the level of residual laser frequency fluctuations over the course of the LPF mission is studied. For this investigation, data recorded during periods of quiet noise measurements on LPF, the so-called ‘noise runs’, has been used. This means no satellite manoeuvres or experiments of any kind have taken place. A single unusually noisy period from 2016-03-22 07:00 to 2016-03-22 14:30:00, which corresponds to less than 1% of the total noise run time, has been observed. Given that unusual spikes occurred in some of the RLU telemetry during this period, it is possible that this behaviour originated in the RLU itself. Since only a very small fraction of the data used is concerned, this issue has not been investigated further. With this exception all of this data could be included in the analysis and no segments had to be excluded for glitches or transients.

To estimate the power of the remaining laser frequency fluctuations, all Ψ_F data is split into segments of 15 min duration. For each of these segments the PSD is computed using a Blackman-Harris window and 36 averages with 50% overlap. Then, every 4th bin, as recommended by [V+14], is selected. This could be improved by using all the data available while taking the correlation introduced by the window and the overlap into account. With this number of averages, the data in every fourth bin can be considered to closely resemble a Gaussian random variable with corresponding variance. Next, the data in a certain frequency range is averaged to obtain a single noise estimate for a given time. Hence, the uncertainty estimate is obtained by the propagation of the uncertainty in each bin. To be more precise, the uncertainty in each bin is given by the spectral value at the respective frequency divided by the square root of the number of averages. This uncertainty is then propagated through the mean following the rules for Gaussian

uncertainty propagation with vanishing covariance. More details on the method can be found in [Wis17].

As will be discussed in more detail in Section 4.3, the measurement of the laser frequency noise with the stabilisation activated is frequency dependent. Therefore, we average the noise power over the range from 0.2 Hz to 0.5 Hz here because in this range, the closed-loop spectrum changes the least and the resulting noise power uncertainty estimate is small. The optimal duration of each short data segment is a consequence of the frequency range chosen. The data segment should not be very long because possible fast changing features in the noise power might become indiscernible this way. For the frequency range chosen, the duration of each segment is close to the minimum because with 36 averages, we result in 25 sec of data in each Fourier transform. Given that with a Blackman-Harris window, we have to ignore the first three bins, 25 sec duration is close to the shortest duration possible. The square root of the estimated average noise power in a 15 min segment is shown in Figure 4.3.

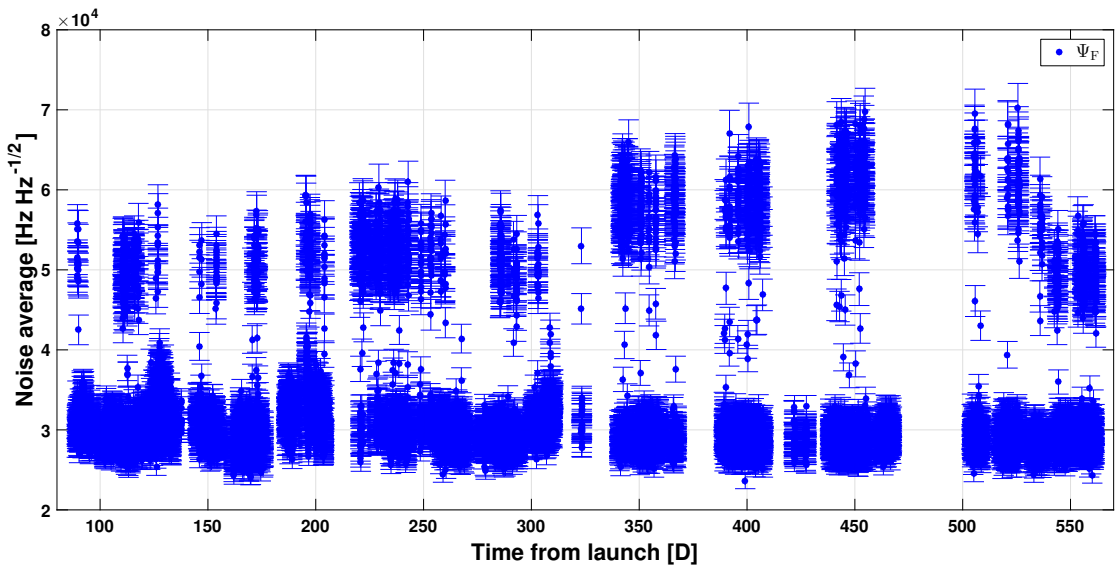


Figure 4.3: A zoom into the average noise in the band from 0.2 Hz to 0.5 Hz of Ψ_F in the course of the mission. Each data point corresponds to 15 min of data.

In Figure 4.3, it is evident that the average laser frequency noise level in the band investigated is not the same over the whole mission. On the contrary, a bi-modal behaviour seems present as approximately 86 % of the averaged laser frequency noise estimates are in the range from 23 to 38 kHz/ $\sqrt{\text{Hz}}$ and 13 % in the range from 46 to 66 kHz/ $\sqrt{\text{Hz}}$. This can be clearly seen in the histogram in Figure 4.4. These two ranges are also described as a lower and upper level of laser frequency fluctuations. We note that the upper level is a factor 1.6 - 2 above the lower level. We note a change in the upper level around November 2016 which occurs at the same time as a decrease in RLU temperature by approximately 4 °C. In Section 4.3, it will be further investigated whether the laser temperature causes this bi-modal behaviour or whether the occurrence at the same time is a mere coincidence. Since the lower level stays quasi constant over the mission and the upper level changes, the distance between the two levels is increased for the period from November 2016 until May 2017.

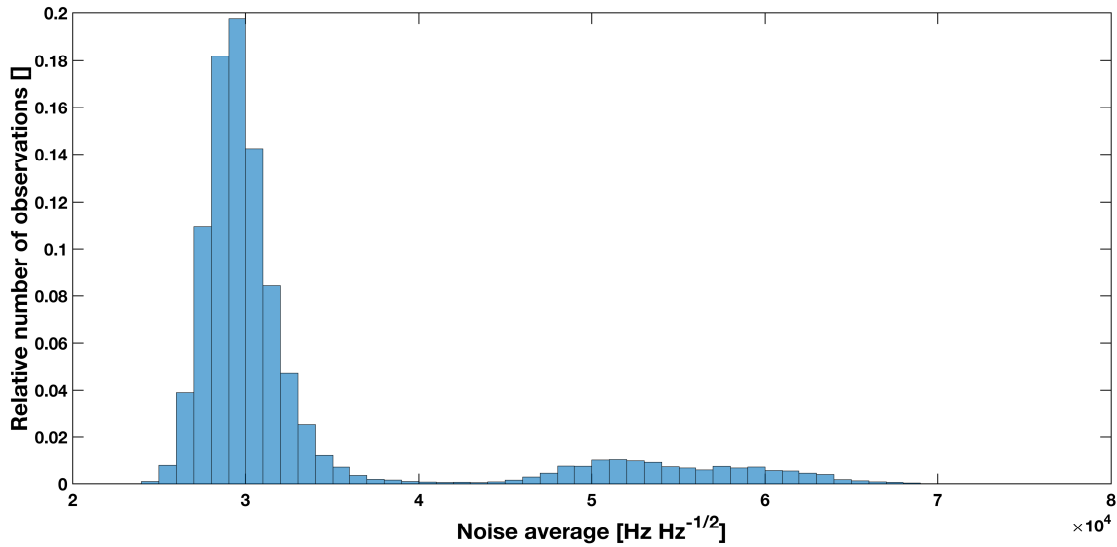


Figure 4.4: A histogram of the average laser frequency noise in the frequency band from 0.2 – 0.5 Hz.

The histogram in Figure 4.4 consists of the data of Figure 4.3 and is normalised such that the height of each bar represents the relative number of observations in this range with respect to the total number of observations. As in Figure 4.3, the two ranges of laser frequency fluctuations are easily discernible. The upper level is broader than the lower level which is caused by the change of the upper level in the course of the mission, see Figure 4.3.

To further describe the time-dependent laser frequency noise level, Figure 4.5 shows the same analysis for a single noise run. This figure illustrates that both the lower and the upper level can occur during the same noise run. However, this does not have to be the case for all of the data under investigation. By comparing to the time series data shown in Figure 4.6, it becomes evident that the times of the higher and the lower laser frequency noise level are discernible by eye in the time series and that the noise estimate tool detects these periods correctly in this example.

Figures 4.7 and 4.8 are spectrograms of Ψ_F data segments recorded at 1 Hz during the example period of quiet noise measurement shown in Figures 4.5 and 4.6. The data segments have been adjusted such that the transition from laser frequency noise at the lower level to the upper level and vice versa are shown. The spectrograms have been obtained using a Blackman-Harris window and 1 min of data in each average of the PSD with 50% overlap. The logarithm of the result is shown. Five of these transitions have been checked via the spectrogram and all showed a quick change of less than a minute of the predominant normalised amplitude range. In the two examples in Figures 4.7 and 4.8 these transitions occur around $9.074 \cdot 10^4$ s and $9.861 \cdot 10^4$ s, respectively, when the dominant colour changes from blue to yellow. The lower frequencies may be affected by the DC values but above 0.2 Hz, the change is even more clear. These spectrograms allow to conclude that there is no drift from one level of laser frequency noise to the other but a quick change.




Figure 4.5: A zoom into Figure 4.3 for a single period of quiet noise measurement.




Figure 4.6: The time series data corresponding to Figure 4.5.




Figure 4.7: A spectrogram of the laser frequency noise measurement that shows the switch around $9.074 \cdot 10^4$ s from a period of laser frequency noise in the lower level to the upper level. The change is immediate.




Figure 4.8: A spectrogram of the laser frequency noise measurement that shows the switch around $9.861 \cdot 10^4$ s from a period of laser frequency noise in the upper level to the lower level. The change is immediate.

From the analysis of the data during the periods of noise measurement on LPF, 80 periods of increased laser frequency noise have been identified. They can last for less than an hour up to a day with a wide spread distribution, as indicated by the histogram of Figure 4.9. However, a 70% majority of these periods lasts up to 10 hours. The maximum duration found here is not




Figure 4.9: A histogram of the duration of Ψ_F in the upper noise level. The periods of increased laser frequency noise last up to one day with a widespread distribution. However, 70% of these periods last up to 10 hours.

limited by the duration of the periods of noise measurement which lasted up to more than two weeks, as for example the period shown in Figure 4.6.

The time in between two periods of increased laser frequency ranges from 2.5 up to 97 hours in the 45 examples measured. Approximately 27 % of these periods are shorter than 10 hours and 69 % last less than 30 hours. Here, only the time in between two periods of increased noise in one noise run is considered. Thus, the values obtained are not expected to be biased by the start and stop times of quiet noise measurement periods on LPF.

In the EM test campaign, the laser has shown frequency jumps which do not look similar to this observed bi-modal behaviour, compare Figure 33 in [M. 06]. A similar behaviour of the laser frequency noise was also observed in the FM and OSTT test campaign but no reason could be identified at that time [A⁺12b][MH13]. Section 4.3 summarises the analysis that has been undertaken to find the reason for the two observed noise levels using the in-flight data. Despite the periods of laser frequency noise at a higher level, the in-loop measurements of the laser frequency fluctuations with the stabilisation being active, as shown in Figure 4.19, are well in agreement in the respective frequency range to the lower fluctuation level, as of Figure 4.3. From both figures, the laser frequency noise is around $30 \text{ kHz}/\sqrt{\text{Hz}}$ from 0.2 to 0.5 Hz. As expected, this lower level also remains more or less constant over the mission. In this sense, this analysis confirms the choice of closed-loop measurements in Figure 4.19.

4.2.1 Times of non-stationary behaviour at lower frequencies

To complete the laser frequency noise analysis, we are also interested in the behaviour of Ψ_F at lower frequencies. However, in the case of the closed-loop measurements, the noise at higher frequencies is dominating the time series. Therefore, noise run data of Ψ_F has been low-pass filtered using a cutoff frequency of 0.01 Hz. Figure 4.10 shows the result for a noise run at the

beginning of LPF operations. We note several spikes revealing a non-stationary behaviour, which




Figure 4.10: Low-passed in-loop measurement of laser frequency fluctuations during a noise run on LPF. The spikes reveal a non-stationary noise behaviour of the laser frequency noise.

we also call ‘glitch’ here. This data segment is considered a worst case example considering the frequency of the glitches. The glitch amplitude is larger in other noise measurement data but was never observed to be above roughly $300 \mu\text{rad}$.

For this exemplary data stretch, none of the observed glitches occurs at the same time as the single glitch observed in the residual differential acceleration measurement, Δg , in the same noise measurement period. This was checked for four more noise runs and none of the glitches discernible in the low-passed laser frequency data coincided with a glitch in Δg . The glitches of our example segment shown in Figure 4.10 also do not occur at the same time as glitches in the laser parameters as listed in Appendix B. During the first three periods of quiet noise measurements on LPF, most of the glitches observed in the low passed Ψ_F data coincided with anomalies discernible in the low-passed reference interferometer data. For the other periods of noise measurement, this is often not the case. This is sufficient to conclude that not all of the glitches in Ψ_F are common to all OMS channels and that they are not systematically caused by the subtraction of the reference interferometer measurement. Given that the coupling from $\Psi_F \rightarrow \text{o12}$ is on the order of $10^{-11} \text{m rad}^{-1}$ for path length mismatches of hundreds of μm (see Section 2.3), frequency glitches on the level of mrad would be necessary to cause a notable glitch in the o12 time series at a noise level of tens of fm [A⁺16c]. This is in accordance with the fact that we have a well discernible o12 signal for the laser frequency modulations at 1.123 Hz with a commanded modulation amplitude of 0.025 rad. The glitches found in the low passed data are approximately two orders of magnitude below the mrad level.

A smaller test subset of the data recorded during the noise runs on LPF has been searched for glitches. The search was done in the following way: For each noise run, the Ψ_F data has been low-pass filtered with a cutoff frequency of 0.01 Hz. In addition, the first derivative of the original Ψ_F data and the low-passed data has been taken. These quantities have then been split into segments of 5 min duration. If then data points of these quantities exceed 5 times the standard

deviation of this segment, the 5 min segment is considered to contain a glitch and its start time and end time are stored. Next, the PSD of each of the data segments in between two glitches is visually inspected and data segments with discernible anomalies have been excluded from the further analysis. With this procedure, approximately 76% of the data remain. An estimate of the amount of false positives would require a more detailed definition of the term ‘glitches’, which is beyond the scope of this thesis. The time which passed in between two data segments with glitches is shown in the histogram in Figure 4.11. Here, we note that a majority of the glitches




Figure 4.11: A histogram of the time between two glitches identified in the Ψ_F data during noise measurement periods on LPF. A significant fraction of the glitches occurs less than five hours after the previous glitch.

occurs less than five hours after the previous glitch. Again, the maximum of the x-axis is not limited by the duration of the noise measurements on LPF.

4.3 Investigation of changing noise levels

In Section 4.2, periods of increased laser frequency noise have been observed and their characteristics have been identified. In theory, they could be caused by a malfunctioning of the control loop, increased laser frequency noise at the input of the loop or a combination of both. As we have shown that the laser frequency control loop works as expected over the course of the mission (see Chapter 3), we conclude that the laser frequency control loop always suppresses the noise by the same factor. Thus, increased laser frequency noise at the input of the loop seems more likely. In this section, the hypotheses investigated to identify the cause of this behaviour on behalf of the in-flight data are summarised.

4.3.1 Hypothesis: aliasing

At first, it was investigated whether the two levels of laser frequency noise are due to an aliasing artefact of the downsampling process of the DMU or whether it is a feature across a larger frequency band. This investigation uses laser frequency noise data with sampling rates of 10 Hz and 100 Hz. The Ψ_F data at 100 Hz was recorded on 2016-06-12, 2017-01-22 and on several days in February 2017. It is compared to the 10 Hz data recorded during a period of quiet noise measurement in March 2016. On each of these days, the data has been split into segments in the lower and upper level. One 100 Hz data segment recorded in June 2016 could not be associated to either level and is therefore discarded.

For spectral estimation, a different version of Welch's periodogram with a Blackman-Harris window is used here compared to other parts of this work. This method is called Logarithmic Power Spectral Density estimate (LPSD) [TH06] where the number of averages is proportional to the frequency that is computed. Consequently, the number of samples used in each average is smaller for higher frequencies. Thus, a spectrum with low noise over a large frequency range is possible.




Figure 4.12: The ASD of laser frequency noise recorded at 10 Hz (red traces) and at 100 Hz (all other traces) at different times throughout the mission. The flat spectrum above 20 Hz and a similar shape of the spectrum recorded with different sampling rates make aliasing an unlikely reason for the two levels of laser frequency noise observed.

The result is shown in Figure 4.12. The 100 Hz data recorded matches quite well the 10 Hz data and is constant above 20 Hz. With these observations, the hypothesis of aliasing in the DMU becomes unlikely. However, this analysis does not allow us to comment on aliasing which may possibly occur when processing the analogue PD signals to the 100 Hz DMU data. Given the short durations of the 100 Hz data segments, this data is not sufficient to find out whether the two levels are a broad band feature. From the 10 Hz data spectrum below 0.1 Hz, it looks rather like the difference in between the two levels decreases towards lower frequencies. However, a definite answer would require a close analysis for possible ‘glitches’.

4.3.2 Hypothesis: correlation to laser settings

None of the laser parameters available from telemetry exceed their defined nominal ranges or threw warnings consistent with the appearance of the different laser frequency fluctuation levels. Of course, this statement assumes that the mechanism that leads to these levels is always the same. In addition, even during the noise runs, a restricted number of commands is applied to the satellite. A certain fraction of those is related to the OMS. It was checked for the first seven noise runs that none of these commands coincide in time with the beginning of a noisier period of the laser frequency. In flight, it is impossible to access all laser data and the data available is mostly sampled at a rate of 0.2 Hz. At first, the time series of the laser parameters is compared to the time series of the measured laser frequency fluctuations for the long noise measurement from 2016-12-26 to 2017-01-13. The laser frequency fluctuations seem typical in this measurement. We show a comparison to the laser pump current in Figure 4.13. When comparing the time




Figure 4.13: The laser frequency noise measurement compared to the RLU pump current. At the times of changes in the laser frequency noise level, no unusual behaviour is discernible in the pump current measurement.

series data, no features or glitches in the laser telemetry have been found at the same time as the change from a period of low laser frequency noise to a period of high laser frequency noise. A slight decrease in the pump current observed before a period of increased laser frequency noise cannot be the main reason for the two levels because it does not fit to a period of increased noise

starting around 400.8. In other noise measurements, there are examples during which the pump current is quasi constant while the laser frequency noise level increases or when the measured pump current decreases but no change in the laser frequency noise level can be observed. All parameters listed in Appendix B have been compared to the laser frequency data in this manner and no correlation could be identified.

In addition, it was investigated if the observed changes in laser frequency noise are depending on the laser temperature, the pump current or the laser output power. The mean of these RLU quantities in a 15 minutes segment is therefore compared to the respective laser frequency noise average at that time. Figures 4.14 to 4.16 show the results for the respective quantities, discarding the outliers above $70 \text{ kHz}/\sqrt{\text{Hz}}$, corresponding to the single glitch also previously discarded in the Ψ_F noise average estimate. This comparison indicates that the noise levels




Figure 4.14: The Ψ_F noise level versus the mean pump current over segments of 15 minutes at the same time.

may be closer to each other for RLU pump currents between 1.34 and 1.36 A and for higher laser temperatures above 26°C . This observation is in agreement to the observed increase in difference of the two levels in Figure 4.3 which coincided with a decrease in RLU temperature. However, from the given data, no temperature or pump current range without the two levels of laser frequency noise can be identified. The RLU output power seems to be the same for the two levels of laser frequency noise. At certain times during the LPF commissioning phase, the three laser parameters of Figures 4.14 to 4.16 have been telemetered at 1 Hz but also with this data, no time correlation to the bi-modal behaviour could be found.

In summary, the analysis presented does not allow us to identify a laser parameter or a combination of laser settings causing the observed bi-modal behaviour. However, given that the data access is limited in terms of the number of parameters and their sampling frequency, we cannot rule out the hypothesis under investigation, even though it seems unlikely with the analysis of important parameters summarised here.




Figure 4.15: The Ψ_F noise level versus the mean RLU output power over segments of 15 minutes at the same time.




Figure 4.16: The Ψ_F noise level versus the mean RLU baseplate temperature over segments of 15 minutes at the same time.

4.3.3 Hypothesis: overall satellite temperature influence

This paragraph aims at understanding whether the OB temperature has an impact on the bi-modal behaviour of the laser frequency noise. Therefore, the laser frequency noise measurements are compared to the temperature of the LTP as measured by the temperature sensors on the OB. A period before a cool-down of the LPF satellite has been chosen because it exhibits a comparatively large change in temperature. However, during this period, the interferometer data

was interrupted by periods where one or both beams have been turned off. The Ψ_F measurement has been set to 0 during these times. There are four temperature sensors on the OB. As they measure the same trend, only one measurement is shown alongside the laser frequency noise in Figure 4.17. From this time series, a linear dependence of the bi-modal laser frequency noise




Figure 4.17: The laser frequency noise compared to the temperature on the OB during the cool-down period. The measurement is set to zero during the gaps. These occur either due to the beams being switched off for spot-position measurements or due to glitches.

behaviour on the overall temperature seems unlikely. More sophisticated dependencies, however, may not be excluded.

4.3.4 Hypothesis: back scatter into the laser at certain test mass angles

The hypothesis investigated in this paragraph is that for certain TM angles, a fraction of the light on the optical bench is back-scattered into the laser and thus increases the laser frequency noise fluctuation level. This hypothesis is plausible as the LPF optical set-up is reciprocal, meaning that in theory, the beams can not only travel as explained in Section 2.1 but also in the opposite direction, up to the Faraday Isolator (FI).

To investigate this hypothesis, the laser frequency noise level is compared to the TM angular measurements. The noise run data has been used alongside the measurements recorded during the second version of the so-called DWS step experiment. During this experiment, both TMs have been tilted to reach large angular offsets. With this experiment, a significant fraction of the possible test mass orientations could be included. More details on this experiment can be found in [Wis17].

For the noise runs, the laser frequency noise level estimate as shown in Figure 4.3 has been used. For the data of the DWS step experiment, the same processing has been applied. The angular measurements are the mean values over 15 min segments of data.

The comparison of orientation and laser frequency noise level is shown in Figure 4.18. It was not




Figure 4.18: The laser frequency noise level for different angular offsets for both TMs: (a) TM1 (b) Zoom into TM1 orientation (c) TM2 (d) Zoom into TM2 orientation. Periods of increased laser frequency noise can be found at several orientations of either TM. This may also be consistent with measuring increased laser frequency noise at orientations most frequently used during the noise runs.

possible to identify an orientation of either TM with increased laser frequency noise. Rather, the data may also be consistent with measuring increased laser frequency noise more often at certain orientations because these were the most frequent ones.

Another possible reason for the bi-modal behaviour that we cannot exclude with the telemetry available are different modes in the pump diode. In addition, for this investigation the current AEI laboratory set-up is not very useful as the flight-spare laser is not integrated in the set-up at present.

4.4 Stabilised and free-running laser frequency noise measurements

In this section, we compare the in-loop measurements of the laser frequency fluctuations with active stabilisation to the measurements of the free-running laser frequency fluctuations.

For the measurements of the stabilised laser frequency noise, appropriate measurement times have been selected. One selection criterion was the availability of 10 Hz data for the Ψ_F channel. The second criterion was the absence of non-stationarities at lower frequencies, as discussed in Section 4.2.1. The third criterion was that during the selected measurement periods, the average laser frequency noise in the band from 0.2 to 0.5 Hz is in the lower level from 23 to 38 kHz/ $\sqrt{\text{Hz}}$. As seen in Section 4.2, this level is present during most of the noise measurement periods and thus this criterion ensures representativeness. The free-running laser frequency noise measurements could only be recorded during certain periods of LPF operations. All of this data is used. For comparison, data from the latest ground measurements, the OSTT campaign, has been included. The data quality of the ground measurements has been improved as outlined in [Aud14].

For spectral estimation, the LPSD method has been used. In addition, the ASD values and uncertainties have been estimated according to the supplemental material of [A⁺18b]. As for the LPSD method, the segments are shorter for the ASD estimation at higher frequencies. But in addition, the frequencies are chosen such that a possible correlation between the different bins becomes negligible. It is also possible to calculate the correlation for each bin and account for it but this has not been used in this thesis. The uncertainty estimate of this method is calculated from the cumulative distribution function of the spectral estimate in each bin.

In Figure 4.19 the ASDs of the free-running laser frequency fluctuations are compared to the ASDs of the laser frequency fluctuations while the stabilisation is active. The results of the LPSD method are given by the lines and the squares represent the result following the analysis method as given in the supplemental material of [A⁺18b]. Both methods have been applied to each data set which is indicated by the colour of the trace and the square. The uncertainty is negligible except for the three lowest frequencies estimated following the supplemental material of [A⁺18b]. The estimates obtained with the two methods agree well. Moreover, the requirement is shown as the grey dashed trace.

From Figure 4.19, a good agreement between the test campaign and the flight data can be found. This is very useful for future interferometer development because it provides one example where the ground tests are representative of the in-flight behaviour. In addition, the results of measurements taken at different times of the LPF mission are very similar to each other, even though there are up to several months in between them. This is another example which is consistent with the laser frequency control being stable over the mission as already found when studying the transfer function of the control loop in Chapter 3. It is apparent from Figure 4.19 that the stabilised laser frequency fluctuations are below the requirement at all investigated times.




Figure 4.19: Amplitude spectral density of the laser frequency fluctuations in-flight with and without active stabilisation compared to the requirement (grey dashed trace) and the latest ground measurements. Each trace represents the LPSD result and each square the result following the analysis method as given in the supplemental material of [A⁺18b] for the respective measurement. The OSTT campaign measurements show increased noise in between 4 and 5 Hz which has not been further investigated in this work.

4.4. STABILISED AND FREE-RUNNING LASER FREQUENCY NOISE MEASUREMENTS




Figure 4.20: Amplitude spectral density of the laser frequency fluctuations in-flight with and without active stabilisation on DOY 21 and 22 2017 compared to the expected stabilised frequency fluctuations using the control loop analysis results (black trace). A good agreement between expectation and measurement is found.

Concerning the verification of the results shown in Figure 4.19, there is no independent sensor for the laser frequency fluctuations on LPF to compare the measurements against. The order of magnitude of the free-running measurements, however, is confirmed up to 0.5 Hz by the measurements of the two controller outputs multiplied by the gains reported in Table 3.3 and summed up. This check does not take into account the presence of the feedback and therefore confirms the free-running and not the stabilised laser frequency noise. However, from the OLF analysis presented in Chapter 3, the amount of suppression of laser frequency noise by the control loop is known. Thus, to check the stabilised laser frequency measurements, the free-running noise measurement of DOY 22 2017 is multiplied with the absolute value of the noise suppression function and compared to the measured frequency fluctuations with the laser stabilisation being active, as recorded on DOY 21 2017. This comparison is shown in Figure 4.20. The expected fluctuations agree well to the measured closed-loop noise across all frequencies.

From the analysis presented in this chapter, we conclude that the laser frequency stabilisation worked reliably over the course of the LPF mission. Also, measurements of the free-running laser frequency noise and the in-loop measurements of the stabilised laser frequency fluctuations agreed well to the results from the OSTT ground test campaign. This is a helpful conclusion for future interferometer development because it provides one example where the ground tests are representative of the in-flight behaviour. As in the ground test campaigns, periods of higher laser frequency noise have been observed but could not be understood yet. Moreover, low frequency glitches in the frequency interferometer measurements have been identified. Fortunately, they do not deteriorate the LPF performance. However, from these two observations, laser frequency noise measurements of ≥ 10 h duration are recommended for future test campaigns of laser frequency stabilisation for satellite missions where similar increased noise periods or glitches are of interest.

Chapter 5

Interferometer path length difference measurements

In this chapter the path length differences between the measurement and the reference beam in the LPF interferometers will be studied using frequency modulations of the laser. We will begin with the explanation of two dedicated experiments for this analysis and the verification opportunities they provide. We will then trace the evolution of the path length mismatch over several months of the LPF mission and during individual experiments.

5.1 Design of path length difference experiments

Table 2.1 lists all laser frequency modulation experiments performed during the LPF mission. In principle, all of these can be used to estimate the path length mismatch according to Equation 2.15 for the XR, the X1 and the X12 interferometers. We will focus on the X12 path length mismatch and just comment briefly on the path length mismatch in X1 and XR in the following sections. This path length mismatch is the most interesting one because it determines the coupling of laser frequency noise into the o12 measurement - the primary science signal.

We can, and will, use the data from the frequency loop characterisation experiments with free-falling test masses for path length mismatch estimates. They are not optimal because we do not record many cycles at frequencies above 1 Hz where the o12 measurement reaches the lowest noise levels and thus allows for a high accuracy path length mismatch estimation. Therefore, two laser frequency modulation experiments aiming directly at path length mismatch estimation have been designed and executed. These are experiments number 4 and 6 in Table 2.1. They will be further explained in the next paragraphs.

5.1.1 First path length mismatch experiment: offset

The first dedicated path length mismatch experiment is number 4 in Table 2.1. This version of the experiment consists of two injections into the fast frequency loop which together last

approximately 1 hour. They are repeated after TM2 has been commanded to an offset position of nine times the laser wave length corresponding to $9.576 \mu\text{m}$. This is shown in Figure 5.1 where the longitudinal position of the two TMs, as measured by the GRS, is shown. The times of the laser frequency modulation at 1.123 Hz and 2.879 Hz are shown in black and grey, too. In the design of this experiment integer multiples of half the laser wavelength have been chosen to




Figure 5.1: The time series data of the x positions of TM1 and TM2 as measured by the GRS during the first path length mismatch experiment. The times of the laser frequency modulations are marked in black and grey, depending on the modulation frequency. This figure does not only show the concept of the experiment but the GRS measurements also confirm the commanded offset was executed correctly.

minimise the impact of RIN. For details see [W⁺17]. A large offset is desired to produce a well measurable change in path length mismatch. For operational safety reasons, we chose an offset slightly below approximately $10.5 \mu\text{m}$ since the latter had already been achieved successfully in the commissioning phase of LPF. The measurement at the offset position is designed as a sanity check: the commanded offset and its direction are known and the measured path length change should agree to within the errors.

Accordingly, the frequency modulation amplitudes and durations have to be chosen in such a way that the error on the resulting path length difference measurements at the two positions is smaller than the deliberate change of nine times the laser wavelength or $9.576 \mu\text{m}$. Thus, the expected uncertainty of a path length difference estimate has to be determined. It results from the expected noise level of the path length mismatch, $\Delta\bar{L}$. This noise level is given by the phase fluctuations, $\tilde{\phi}$, as

$$\frac{\widetilde{\Delta\bar{L}}}{\overline{\Delta\bar{L}}} = \tilde{\phi}, \quad (5.1)$$

with $\widetilde{\Delta\bar{L}}$ denoting the fluctuations of the path length mismatch over its mean value $\overline{\Delta\bar{L}}$. This equality holds true because we compare two phase measurements to obtain the path length

mismatch. Thus, we can estimate the phase noise as [Sch18]

$$\tilde{\phi} = \frac{\tilde{n}}{\tilde{s}}, \quad (5.2)$$

where \tilde{n} is the ASD of the noise, which is assumed to be additive at the beat-note frequency, and \tilde{s} denotes the rms signal amplitude. Then, $\tilde{\phi}$ is the inverse of the Carrier-to-Noise density C_0/N , where C_0 indicates the carrier and N the noise, denoted in units of Hz:

$$\tilde{\phi} = \frac{1}{\frac{C_0}{N}}. \quad (5.3)$$

In the present case, the Carrier-to-Noise-density is proportional to the SNR, see Chapter 3, with a factor of \sqrt{T} . Here, T is the duration in seconds. This yields the expected relative uncertainty as being proportional to

$$\tilde{\phi} \propto \frac{N}{S}. \quad (5.4)$$

The expected signal was calculated according to

$$S = \frac{A}{\sqrt{2}} \cdot \frac{|G|}{1 + |G|} \cdot \frac{\lambda \Delta s_{o12}}{4\pi \cos \alpha \Delta L} \quad (5.5)$$

for an estimated path length mismatch, Δs_{o12} , in o12. A is the commanded peak amplitude of the laser frequency modulation, which is converted to the rms amplitude by multiplication with $1/\sqrt{2}$. G denotes the estimated gain of the laser frequency control loop at the respective modulation frequency. The coupling of this signal to o12 is then also influenced by the laser wavelength λ , the angle of incidence on the TM, α , and the intentional path length mismatch, ΔL . The noise was estimated following:

$$N = \sqrt{\text{PSD}(o12@f)} \cdot \frac{1}{\sqrt{T}}. \quad (5.6)$$

The modulation amplitude at 1.123 Hz was kept at 0.05 rad, which was known to work well from previous laser frequency loop characterisation experiments. It was decreased to 0.02 rad for the modulation at 2.879 Hz to remain well within actuator limits. This is because, simply speaking, the same modulation amplitude at a higher frequency requires more energy and at some point a physical limit will be reached. All LPF experiments have been designed with the idea of minimising any potential risk to the equipment and thus a large safety margin was chosen.

The expected relative uncertainty $\tilde{\phi}$ can thus be calculated by inserting the expressions for S and N from Equations 5.5 and 5.6 into the expression 5.4 which then only depends on the duration T . This duration has been chosen to yield values of $\tilde{\phi}$ around or below 0.01 for each modulation and to account for the unknown error on the assumed path length mismatch and o12 noise floor level, the durations have been doubled. This lead to 4042 half-cycles at 1.123 Hz and 10364 at 2.879 Hz. For operational reasons, the durations have been slightly shortened again to the values in Table 5.1. This experiment has also been described in [P+17].

5.1.2 Second path length mismatch experiment: varying modulation amplitudes

In the first dedicated path length mismatch experiment, during modulations 1 and 3, we find the third harmonic of the modulation at 3.369 Hz well discernible. As such harmonics could be

frequency	amplitude	duration
[Hz]	[rad]	[N half cycles]
1.123	0.05	3642
2.879	0.02	9364

Table 5.1: The injection parameters for the first path length mismatch experiment. These two laser frequency modulations have been executed at the nominal and at the offset test mass position. Accordingly, this experiment consists of 4 laser frequency modulations.

frequency	amplitude	duration
[Hz]	[rad]	[N half cycles]
1.123	0.5	1886
1.123	0.05	4716
1.123	0.025	11858
2.879	0.05	12090

Table 5.2: The injection parameters for the second path length mismatch experiment.

a sign of non-linearities, the second path length mismatch experiment was designed. These laser frequency modulations aim at systematically ensuring that the measured path length mismatch does not depend on the amplitude of the fast frequency loop injection. Therefore we have a set of three injections at the same frequency but with varying amplitudes as can be seen from Table 5.2. To reach comparable uncertainty levels, the frequency modulations with smaller amplitude have longer durations. The calculation of the necessary durations is the same as for the first path length mismatch experiment, see Section 5.1.1. As a sanity check, this experiment also contains a single injection at the same amplitude but with a different frequency. As a side effect, the high amplitude injection also allows for a more precise determination of the path length difference, especially in the noisier X1 interferometer.

5.2 Analysis procedure for the path length difference experiments

The path length mismatch estimates are obtained following the procedure as sketched in Figure 5.2. The analysis has to start with a check of the operational state of the laser system. This is done by a check of the parameters listed in Appendix B. The next step is to split the data into a segment per modulation frequency. In addition, the data is checked for the coherence in

5.2. ANALYSIS PROCEDURE FOR THE PATH LENGTH DIFFERENCE EXPERIMENTS

between the two channels under investigation. Then a path length mismatch for each of these is obtained. This means that the remainder of the procedure is iterated until all segments have been treated. The key idea of estimating the path length mismatch is to compare the amplitude




Figure 5.2: A flow chart of the path length mismatch procedure. There are four different ways to estimate the path length mismatch which are explained in more detail in the text. We iterate over the modulation segments in this part.

of a sinusoidal laser frequency modulation in the X12 or X1 interferometer to the amplitude of the modulation data recorded in the frequency interferometer measurement Ψ_F , in other words, the ratio of two amplitudes of a sine wave.

Using the o12, o1, and Ψ_F telemetry implies that the reference interferometer measurement has already been subtracted, as explained in Section 2.2. We also refer to this section for the definition of the path length mismatches. The path length mismatch estimates in the reference interferometer will be discussed separately, see Section 5.3.3.

Several laser frequency modulation experiments took place while the TMs on board of LPF were free-falling and the DFACS controller used the OMS measurements as a sensor. This is the case for the nominal science mode, as explained in Section 1.2.1. As the laser frequency modulations couple to o12 and o1, and this coupling is indistinguishable from true TM motion, the control system aimed to counteract the false TM motion. Consequently, we can observe modulations of the commanded forces, too, if the corresponding data channels have been recorded with a sufficient sampling rate. This is relevant because the determination of the path length mismatch should only take into account the signal amplitude in o12 due to frequency modulation and should not be biased by a fraction of the signal originating from true TM motion. However, the unity gain frequency of the suspension control loop is around 1 mHz [A⁺16c] and thus below our laser frequency modulation frequencies, so the expected impact is small. Nonetheless, it was checked that the impact on the path length mismatch in o12 is negligible within the errors via an LPF state-space model simulation.

The same is true for the drag-free loop and the estimated path length mismatch in $o1$ because we can only detect a signal in $o1$ during our long modulations at 1.123 Hz and 2.879 Hz, frequencies at which the noise in $o1$ is relatively small and which is significantly above the drag-free controller bandwidth around 250 mHz for the $x1$ control degree of freedom [Sch12, Table 4-8]. Again, via a LPF state-space model simulation, the correction was estimated to be below 2% for the modulation at 1.123 Hz and below 1% for the modulation at 2.879 Hz. However, because the accuracy of the model has not been quantified and because the uncertainty of the correction factors is not known, the correction was deemed negligible given that it is on the order of magnitude of the uncertainty. This scenario is different for the path length mismatch estimation in the reference interferometer, as explained in Section 5.3.3.

To estimate the amplitude of a sine wave in noisy data, several data processing methods have been applied in this thesis:

- Discrete Fourier Transform (DFT),
- Transfer Function Estimate (TFE),
- demodulation,
- linear fit.

Each of these options can also be found in Figure 5.2. Using the **DFT** method, the data of the frequency interferometer measurement and the interferometer under study is Fourier transformed at the known modulation frequency using a Blackman-Harris window. The uncertainty is estimated from the amplitude of the noise floor at frequencies around the modulation frequency. This implies that a certain minimal duration of the data is required. Only then useful data is available for the uncertainty estimate in the bins around the modulation peak. As both data channels under comparison are sampled in the DMU no significant delay in between the channels is expected. This has been checked for some example times. Thus, phase properties are neglected and the amplitude estimate can be obtained by dividing the absolute values of the two Fourier transforms:

$$\text{DFT}(f_{\text{mod}}) = \frac{|\mathcal{F}(o_{12})(f_{\text{mod}})|}{|\mathcal{F}(\Psi_{\text{F}})(f_{\text{mod}})|}. \quad (5.7)$$

In this context, the **TFE** method aims to estimate the transfer function from Ψ_{F} to $o12$ using the `ao/tfe` method. It computes the ratio of the cross-power spectral density between the input and the output to the PSD of the input [N+13][BP80]. Similar to the transfer function estimation aiming at the loop characterisation, the number of data points in each average has been chosen such that the modulation frequency is located at a bin frequency, see Section 3.2.2. The path length mismatch estimate is then obtained from the absolute value of the function at the modulation frequency using the same Blackman-Harris window. Its uncertainty estimate is a result of the coherence of the signals, as given by Equation 3.2. This is the main difference to the DFT method for cases of high coherence between the input and output signal. For both the DFT and the TFE method, the choice of the Blackman-Harris window over a flat-top window can be justified by considering that the bins match the modulation frequencies well.

All of the methods presented here can be understood as a way of demodulation. However, the results denoted ‘**demodulation**’ in this chapter are from multiplication of the data with a sine or cosine wave at the modulation frequency with a phase offset to adjust for the time of the beginning of the data. This is implemented in `ao/heterodyne`. The modulation frequency and the offset have been chosen such that all of the signal is in the sine quadrature and the cosine

quadrature is as small as possible. The multiplication results have then been filtered to remove the signal at twice the modulation frequency. There, a 5th order low pass filter with a cutoff frequency of 0.2 Hz has been applied forwards and backwards to the data which is subsequently downsampled to 0.05 Hz. The resulting time series data is then chopped at the beginning and the end to discard filter transients. A path length mismatch is then obtained from dividing the sine component estimate of o12 by the sine component estimate of the frequency interferometer measurement.

The fourth option is the **linear fit**. The sinusoidal modulation in both the o12 and frequency interferometer measurement channels can also be fitted in the time domain. In general, this would be a non-linear model

$$s = A \cdot \sin(\omega t + \phi) , \quad (5.8)$$

with the amplitude A , the frequency ω and the phase ϕ . However, the frequency and the phase are correlated here. Thus, it was assumed that the laser frequency has been modulated at the commanded modulation frequency. This has also been verified using the demodulation strategy. Then

$$s = A \cdot \sin(\omega_{\text{mod}}t) + B \cdot \cos(\omega_{\text{mod}}t) \quad (5.9)$$

can be used to obtain the in-phase amplitude, A , and a phase estimate from B . When this method was used to estimate the path length mismatch for a small part of a modulation segment, an offset phase accounting for the start time of the data to be fitted was added to the argument of the sine and cosine term. In addition, for all o12 data, a 4th degree polynomial has been fitted to the data to account for small drifts over time. To obtain the path length mismatch, the ratio of the two in phase-amplitudes is taken.

The last step of the procedure outlined in Figure 5.2 is then again common to all methods. The ratio of the amplitudes needs to be scaled to the corresponding optical path length mismatch.

In Section 2.2, the coupling of laser frequency noise into the o12 and o1 measurements has been explained. Solving Equation 2.29 for the path length mismatch and replacing the fluctuations with the previously determined amplitudes yields

$$\Delta s = \frac{A_o}{A_F} 4\pi \cos(\alpha) \cdot \frac{\Delta L}{\lambda} . \quad (5.10)$$

$\frac{A_o}{A_F}$ is the ratio of the signal in the interferometer o, which is o12 or o1, to the signal in the frequency interferometer. The scaling factor is a result of the conversions explained previously. α is the angle under which the light hits the TM. ΔL is the intentional path length mismatch and λ the wavelength of the laser light. For the estimation of Δs_{o12} and Δs_{o1} a path length mismatch of $\Delta L = (38.2 \pm 0.1)$ cm is used throughout [Rob13]¹. This corresponds to a relative error of less than 1%. This path length mismatch could not be assessed independently in flight because the only measurement of laser frequency noise is through the XF interferometer and there is no out-of loop measurement possible.

5.3 Path length difference estimates

In this section, the results of the two dedicated path length mismatch experiments are explained. In addition, the results from all laser frequency modulation experiments are combined into a

¹uncertainty estimate is from private communication with D. Robertson

table such that the path length mismatch estimates in the course of the LPF mission can be studied.

5.3.1 Results from offset experiment




Figure 5.3: The o12 measurement during the four modulations of the path length mismatch offset experiment. The commanded shift of TM2 to the offset position is clearly discernible. The inset figure is a zoom into the o12 measurement during the first laser frequency modulation. A sinusoidal signal on the pm level is visible.

At first, we need to verify the experiment was executed as designed. The shift of TM2 to the offset position can be clearly seen in the o12 relative distance measurement, compare Figure 5.3. The times of the laser frequency modulations are marked by the four blocks. At the beginning of the second modulation a jump in timing has occurred. As a result, the data starts approximately 4 minutes later. As planned, the first two laser frequency modulations took place at the nominal position and the second two at the offset position.

Note that o12 is defined as $x_2 - x_1$ and as such increases if the TM2 moves towards the centre of the satellite (corresponding to an increase in x_2) and TM1 remains at the same position (that is x_1 remains constant), as was commanded for this experiment. This is illustrated in Figure 5.4.

The expected change in TM position is measured in the GRS, as can be seen from Figure 5.1. In the DFACS operational mode chosen during the laser frequency modulations of this experiment, the GRS is an out-of-loop sensor and thus provides an independent position measurement. The laser has been frequency modulated as expected. This can be seen

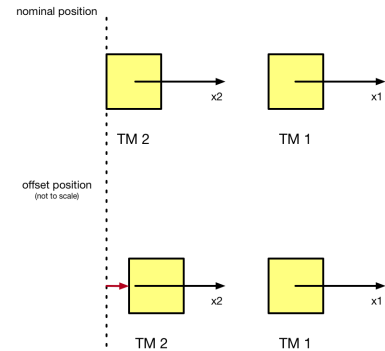


Figure 5.4: A sketch of the offset experiment to determine the path length mismatch. The definition of the axes follows [Sch12], as already shown in Figure 2.3.

from the larger viewgraph of Figure 5.10 where the four in-loop laser frequency modulation segments have been shown.

In Figure 5.5, we show the results of estimating the path length mismatch for each of the four modulations of this experiment, along with a model of the expected behaviour.

We estimated the path length mismatch applying the DFT and TFE method to each of the modulation segments, as explained in Section 5.2. The model used is quite simple here: we assume the path length mismatch is constant to the μm level at both the nominal and the offset position. These two positions are fixed to differ by the commanded offset of $9.576\ \mu\text{m}$, approximately translated into optical path length via a factor two. As the path length mismatch at neither the nominal nor the offset position are known from the experiment design, we fit this model linearly to the TFE results. This is a one-parameter fit where the difference between the two positions stays fixed and only the path length mismatch at the nominal position is affected by the fit.

We see the two methods do agree with each other at roughly the 3σ level. The estimated uncertainty levels of the DFT at the offset position are larger than the commanded offset. This is due to the increase in the noise floor at the offset position. The residual amplitude spectrum at the modulation frequency is at approximately 8% of the o12 peak amplitude spectrum or less for each of the four modulations. However, the same slight variations of the path length mismatch estimate increase the residuals from the modulations at the nominal position more than those at the offset position. This indicates the data is not as sensitive to the path length mismatch at this position and thus larger uncertainty estimates seem reasonable there. However, at the same




Figure 5.5: The results of the first dedicated path length mismatch experiment using a single path length mismatch estimate for each of the four modulation segments, alongside a model of the expected change in path length mismatch. We fit the model to the measurements as the path length mismatch at none of the two positions is known in advance. Solely the expected change is known, as it is part of the experimental design and has been verified by the GRS measurements. The result of this fit is shown in green. It is described above and should not be confused with the fit explained in Section 5.2. The caveats are discussed in Section 5.4.

3σ uncertainty level, the model agrees to the measured TFE results. Thus, we can associate the measured path length mismatches to the commanded offset at a 3σ uncertainty level, both in direction and amplitude. Unfortunately, the evidence does not seem as compelling as aimed for in the design of the experiment. However, we have only a few data points here. Dividing the four modulations into shorter periods and estimating the path length mismatch for each of those, as explained in more detail in Section 5.4, has not provided a solution yet. A comparatively large slew as required for this experiment takes about half an hour per direction. Therefore, it was not possible to repeat this experiment many times during LPF operations to obtain a statistical uncertainty estimate.

From this experiment, we can suspect that the measurement beam is longer than the reference beam. This is because if we command TM2 towards the centre of the satellite, the optical path of the measurement beam becomes shorter while the path of the reference beam is not affected. Our experimental results suggest that the path length mismatch is likely to become smaller with the commanded offset and, hence, the path of the measurement beam must be longer. This agrees to the change in path length mismatch due to a commanded offset on TM1, as will be shown in Figure 5.9. Here, TM1 was actuated in negative x1 direction thus towards the centre of the satellite, as can be seen in Figure 5.4, and the path length mismatch decreased, see Table 5.3. This change in path length mismatch over time is a more convincing evidence that the path length mismatch estimation represents changes correctly in amplitude and direction.

5.3.2 Results from amplitude experiment

The second path length mismatch experiment was executed as planned on January 22nd 2017. During this experiment, jumps in the timing were present which look like frequency jumps. With the same procedure as for the frequency loop characterisation experiments (see Section 3.2.1), the data of this experiment was timing corrected as well. The only difference to previous timing corrections is that in this experiment, the step down to 1 Hz data was not necessary. For each of these modulation segments, the path length mismatch was estimated using the DFT method. The result is shown in Figure 5.6 where the commanded modulation amplitude is shown on the x-axis. We find the estimated path length mismatch is independent of the laser frequency modulation amplitude and frequency, as expected. However, one needs to keep in mind that this approach is valid but so is the idea to split the data of each of the modulations into smaller segments which revealed short term path length mismatch fluctuations, see Section 5.4. The consolidation of these two approaches and their results requires more investigation. For the sake of completeness, note that also during this experiment and especially during the largest amplitude modulation, harmonics in the o12 spectrum have been discernible. In addition, the time series data of Ψ_F of this modulation shows periodic changes in amplitude at a frequency below the modulation frequency whose origin could not yet be identified. However, even though this could be an interesting feature, this is not relevant for the path length mismatches because they are obtained from the data at the modulation frequency only.

5.3.3 A comment on the path length difference in the reference interferometer

To estimate the path length mismatch in the reference interferometer, we need to approximate the laser frequency fluctuations measured prior to the subtraction of the reference interferometer




Figure 5.6: The results of the second dedicated path length mismatch experiment using a single path length mismatch estimate for each of the modulation segments. The results are consistent with the estimated path length mismatch being independent of the laser frequency modulation amplitude. The caveats are discussed in Section 5.4.

measurement by adding the XR data again:

$$\phi_F = \Psi_F + \Psi_R. \quad (5.11)$$

This is only an approximation because the subtraction has originally been performed on the 100 Hz data in the DMU. However, it is not sufficient to use ϕ_F and Ψ_R for the path length mismatch estimation yet. This is because the OPD control loop is effective at the modulation frequencies and Ψ_R is an in-loop measurement. Hence, the fluctuation level that would have been present if the OPD control loop were inactive needs to be estimated. Consequently, the precision of this estimate relies on the accuracy of the knowledge of the OPD control loop. Nonetheless, an upper limit of the order of a few cm could be obtained. This order of magnitude is plausible given that this path length mismatch arises at some point in between the first beam splitter in the LMU and the PD on the OB. In the LMU, the optical path length from the beam splitter element to each of the output fibre ends has a required accuracy of ± 1 cm [Bod07]. In the worst case, a path length mismatch of up to 2 cm could arise there. This is much smaller than the intentional path length mismatch $\Delta L = (38.2 \pm 0.1)$ cm. In addition, it was shown in Section 2.3 that a path length mismatch before the OB cancels in the telemetry data due to the subtraction of the reference interferometer measurement.

5.3.4 Results over the mission

The laser frequency modulation experiments aiming at characterising the laser frequency control loop and the path length mismatch are distributed over the LPF mission duration. Hence, it is possible to estimate the path length mismatch on each of these days. The preliminary results of this analysis are shown in Table 5.3. Here, we focus on the experiments performed while the TMs have been free-floating. More precisely, no path length mismatch was estimated from the first and

date	experiment	#	Δs_{o12} [μm]	Δs_{o1} [μm]
DOY 153: 01-06-2016	frequency loop characterisation	2	368 ± 5	-
DOY 164: 12-06-2016	frequency loop characterisation	3	351.4 ± 0.8	-
DOY 166: 14-06-2016	path length difference offset	4	352 ± 1	148 ± 4
DOY 022: 22-01-2017	path length difference amplitudes	6	329.6 ± 0.8	130 ± 1
DOY 022: 22-01-2017	frequency loop characterisation	5	330.7 ± 0.8	-

Table 5.3: Summary of the measured path length differences in o12 and o1 with free-floating TMs. The experiment numbers refer to the chronological numbers of Table 2.1. These results are obtained by applying a single DFT per modulation frequency and then estimating the weighted mean and its error. On DOY 166, only the first two modulations at the nominal position are included. Note also that there is no signal in the o1 channel during the loop characterisation experiments such that no path length mismatch estimate can be provided. We note comparable changes in Δs_{o12} and Δs_{o1} in the course of the mission due to a commanded offset on TM 1.

last loop characterisation experiment because those would be affected by the positions at which the two TMs get grabbed and thus are not representative for the integration and alignment quality or the coupling of laser frequency noise to the total sensing noise of the LPF OMS. In principle, the path length mismatch can also be determined during a period of quiet noise measurement on LPF. At these times, a coupling from the remaining in-loop laser frequency fluctuations to o12 in a frequency range from approximately 0.7 to 2 Hz can also be determined but not to the accuracy levels reported here. This is also true for several large slews which took place towards the very end of the LPF mission. Therefore, such values are not reported here.

The values reported in Table 5.3 have been obtained by applying the DFT method to the data recorded during the laser frequency modulations. On DOY 153 2016, only a single modulation of the frequency loop experiment was successful. Hence, the value reported here is the DFT result and its error. On the other days, we have averaged the path length mismatches obtained from each of the modulation frequencies. This was possible since no frequency dependency of these estimates was found, as expected. Here, the weighted mean with the corresponding uncertainty estimates is reported. However, these uncertainty estimates may be biased due to the small number of laser frequency modulations that we can include in the average, where 8 is the maximum. On DOY 166 2016, only the first two modulations at the nominal TM position have been used to estimate a path length mismatch in o12. Of course, the individual modulations can also be split into smaller segments and those could also be overlapped with each other to obtain better statistics. However, as long as the short term fluctuations as described in Section 5.4 are not fully understood, the uncertainty estimates reported are potentially not free of systematics, regardless of the number of segments included in the average. The data of the frequency loop characterisation experiments on DOY 164 2016 and 22 2017 has also been demodulated to estimate the path length mismatch. However, the modulation segments of all frequencies above 0.011 Hz are too short to apply the demodulation and the filtering, as described in Section 5.2. Therefore, a different demodulation scheme, as implemented in `ao/stability_demod`, was used. Here, data segments of a specified number of cycles at the

modulation frequency are demodulated. In this process, a mean and a drift can be subtracted and a notch filter is implemented prior to demodulation in such a way that no data has to be discarded. In addition, with the successful demodulation, the modulation frequencies have been verified in these two experiments. Given the loop characterisation modulation frequencies have been the same for all 5 runs of this experiment, one can safely assume that all of them have been applied as commanded.

In general, the coupling of the laser frequency modulations to o1 is only observable during the modulations of the dedicated path length mismatch experiments. During an example measurement period of quiet free-fall data from December 26 2016 until January 13 2017, we can see that the o1 noise floor is frequency dependent and is on the $\text{nm}/\sqrt{\text{Hz}}$ level from approximately 10-100 mHz and decreases to the level of $\text{pm}/\sqrt{\text{Hz}}$ around 1 Hz. From Table 5.3, we find that the path length mismatch in o12 estimated on DOY 153, 164 and 166 is consistent with each other to within 3σ . However, we find a change in Δs_{o12} of approximately $20 \mu\text{m}$ from the measurement on DOY 166 to DOY 22. Accounting for this, the two measurements on DOY 22 are again consistent with each other. The estimated path length mismatches in o1 are in accordance with a $20 \mu\text{m}$ change at the 1σ level.

Order of magnitude of the path length mismatch estimates A path length mismatch of the order of several hundred μm is an indicator of an excellent overall satellite integration as all in-flight results are significantly below the requirement of 1 cm (see Equation 2.32). However, it is impossible to determine where exactly the path length mismatch occurs. The first reason is that the absolute position of either TM is a result of the alignment of the OB with respect to the sideslabs, the adjustment of the sideslabs with respect to the vacuum housing and from there to the electrode housing of the GRS and then finally to the TMs. So the whole integration and assembly of the LTP core assembly is involved. The second reason is that the path length mismatch could also be caused on the OB itself but also there, no individual components can be identified. These two possible contributions are likely to be both present.

The comparison to path length mismatches on ground is limited in the sense that the TMs had to be replaced with mirrors positioned to the nominal TM locations. However, on ground, the path length mismatch has been experimentally confirmed to be below $400 \mu\text{m}$ [R⁺13]. The path length mismatches in the X1 and X12 interferometer should agree to the path length mismatch in XR to better than $100 \mu\text{m}$ [R⁺13]. This estimate was obtained from calculations involving measured component positions and a model of the OB. However, as the uncertainty of these numbers is not well known, the expected LPF integration inaccuracy is unknown as well. Hence, there is little to gain from further detailed analysis of the integration documentation.

5.3.5 Explaining the changes in path length difference

The results for the path length mismatch in o12 obtained in June 2016 differ from the results obtained in January 2017 by approximately $20 \mu\text{m}$ in optical path length, as shown in Table 5.3. Naturally, we would like to understand the origin of this difference in path length mismatch. We make the obvious assumption that the components on the OB stayed fixed and that the distance change occurred on the way to the TMs and back. This is confirmed by the fact that the position and orientation of the two test masses did not change by such an amount according to both OMS and GRS. Also changes in alignment in the course of the mission have been deemed unlikely to account for this level of change. Similarly, major changes in temperature have been commanded mostly in the period after January 22nd 2017, compare [A⁺19f].

Hypothesis under investigation Hence, a combination of hardware drift and mode transitions as sketched in Figure 5.7, was suspected. During a noise run, the satellite is in Sci 1.2 DFACS mode. This implies that the reference TM position is controlled according to the optical readout of the relative distance $x_2 - x_1$. If during such a period, the GRS would drift, possibly due to a drift of the hardware, this would not affect the o12 readout. When the operating mode changes to one using the GRS as a position sensor, the reference TM would be moved such that the GRS readout shows 0. This 0 may also be offset with regard to the OMS 0, as illustrated in Figure 5.7. Before the next transition to an optically controlled mode, the phase tracking algorithm is reset. In case the phase tracking counter was at 1 or more before the reset, this results in the OMS showing the same measurement as the last time the satellite was in an optically controlled mode but the absolute position of the reference TM has changed since that time. This change of the counter means that a difference in OMS readout of more than half a wavelength before and after the phase tracking reset is present. An accumulation of this effect over the months in between the two path length mismatch experiments could then maybe explain the change in path length mismatch. Note that such a hardware drift would be very interesting




Figure 5.7: A sketch to illustrate the hypothesis that hardware drifts lead to a change in estimated path length mismatch.

for the further development of the LISA hardware because possible drifts over long time have to be accounted for in the mounting process of the OB and the GRS.

Description of procedure to identify changes in absolute positions As a first step, all o12 data from in between the two measurements at a low sampling frequency of approximately 0.033 Hz was obtained and processed to an evenly sampled grid. From searching the Telecommand (TC) history, the times of the phase tracking resets in between the two measurements have been identified and 30 min of o12 data before and after that have been downloaded. If the absolute value of the mean position before and after the phase tracking reset differed by more than half a wavelength, the change in phase tracking reset integer is calculated and stored. The change in phase tracking index is then summed up over the relevant resets. This allows us in principle to reconstruct the absolute position of the reference TM. We also store the times of the phase tracking resets with the large steps found.

When studying the resulting σ_{12} in between two of those resets, one would expect a smooth curve. However, several steps that coincide in time with the spot position experiments can be identified. As a part of these experiments, the measurement beam, the reference beam or both beams have been turned off. Four more periods with beams switched off outside these experiments have been found. For these periods of invalid OMS data, it was checked whether the difference from the start of this period to the end is larger than 300 nm. If that was the case, a position estimate for this time with GRS data had to be used.


This replacement was done in the following way: From the DFACS mode parameter the last time the satellite was on OMS control before this period and the first time after was identified. At these times, the GRS was acting as an out-of-loop sensor and thus the value it reports is not controlled to 0. Next, the data was checked for an offset between GRS and OMS measurement and the GRS data was adjusted such that there was no offset present.

Afterwards, the search for steps at phase tracking resets was again applied to all the OMS data in between the two experiments but now, if there were no valid OMS measurements, the replaced GRS data was used. With this procedure, all of the data except for one phase tracking reset could be used. This had to be excluded because it provided unreasonable data after a short period of LTP emergency shut down. Again, the change in phase tracking index for each of these resets was traced.

Finally, the whole σ_{12} data, which has in part been replaced with offset-corrected GRS data, in between the two path length mismatch experiments has been split into segments. Each segment starts with a phase tracking reset with a large step and ends at the next. Each of these segments was corrected for the accumulated phase to obtain a plot of the estimated true test mass position over time. The result is shown in Figure 5.8.

We note that the estimated true test mass position changes in both directions. It remains within a small range of the starting position until January 18th 2017 (day 412 since launch) where we accumulate most of the change. This period is highlighted in orange. It looks like on that day, the reference test mass was actually moved. A detailed view of this period is given in Figure 5.9. During this transition from DFACS mode NOM1 to NOM2, it looks like according to the GRS TM1 moved immediately by more than 10 μm . Then, while the DFACS was still in NOM2 mode, TM1 was moved to the 0 of the GRS x_1 measurement. This can be easily understood, given that in this DFACS mode, there is no drag-free control but each of the TMs is controlled to be in the centre of the respective housing, which means that the respective GRS sensor measures 0. In this mode, the OMS is the out-of-loop sensor and measures the motion of TM1. Then, in science mode 1.1, the reset of the phase tracking is applied and the σ_{12} measurement is set to 0. The observed behaviour was due to an error in the commanded offset at DFACS mode transitions from NOM1 to NOM2. An offset of 11.9 μm was set where it should have been 0.2491 μm ². Accordingly, the evidence presented in Figure 5.9 shows that TM1 moved by 11.7 μm even though both sensor measurements are around 0 at the beginning and the end of this period. This motion was to compensate the offset, so in $-x_1$ direction towards the centre of the satellite (see Figure 5.4).

²private communication with L. Mendes



92

Figure 5.8: The estimated true test mass position in between the laser frequency modulation experiments of June 2016 and January 2017. The data is noisy because it contains all kinds of experiments and station keeping times. While for most of the time the estimated true test mass position is within a range of $4 \mu\text{m}$ of the starting position, there was a large reset in January 2017. It is marked in orange.

5.4. VARIATIONS OF PATH LENGTH MISMATCH ESTIMATES DURING INDIVIDUAL LASER FREQUENCY MODULATIONS

Reason for observed changes The TM motion thus identified corresponds to a change in optical path length of the measurement beam in the X12 interferometer while the path of the reference beam remains unchanged. The amplitude of this change in path length difference is approximately twice the TM motion, $23.4\ \mu\text{m}$. From the offset experiment, it was deduced that the path of the measurement beam is longer and hence a decrease in path length mismatch is expected from the motion of TM1 towards the centre. So the amplitude and the direction of the changes in path length mismatch introduced by the TM1 motion are consistent on the 2σ level with the changes in path length difference in o12 in Table 5.3 that we set out to understand. In addition, this motion explains comparable changes in Δs_{o12} and Δs_{o1} .

To conclude, this TM motion is the main reason for the measured changes in optical path length mismatch. Thus, one can also argue that hardware drifts may be present but on a much smaller level, compare Figure 5.8. In addition, this TM offset is confirmed by a stray electrostatic field measurement campaign³. Both of these experiments are useful cross-checks in such a novel and experimental mission. Even though the impact of the absolute differential TM position is limited, it is at least relevant for the coupling of laser frequency noise to the o12 noise budget.




Figure 5.9: The o12 measurement of the OMS and the GRS x1 measurement at the time of a mode transition in January 2017. Due to a wrong offset in the GRS, TM1 moved by $11.7\ \mu\text{m}$.

5.4 Variations of path length mismatch estimates during individual laser frequency modulations

5.4.1 Description of the observation

In the previous section, we have estimated the path length mismatch using a single DFT estimate per modulation. However, unless either of the test masses moves over time, the estimated path

³private communication with V. Ferroni

length mismatch should remain the same to within the estimated uncertainty and should be scattered randomly if we divide this one modulation segment into smaller segments, as illustrated in Figure 5.10. However, the estimated path length mismatch seems to vary during most of




Figure 5.10: An illustration of how the modulation segments of the offset path length mismatch experiment are split into shorter segments. The inset figure shows a zoom into the second modulation of this experiment.

the individual modulations of the dedicated path length mismatch experiments. In Figures 5.11 and 5.12, we show the estimated path length mismatches of non-overlapping segments of different short durations. Since the laser is frequency modulated at $f_{\text{mod}} = 1.123 \text{ Hz}$ or $f_{\text{mod}} = 2.879 \text{ Hz}$, the selected durations corresponds to 12 to 32 cycles nonetheless and should therefore provide a reasonable estimate, as outlined in Section 5.2. In Figures 5.11 and 5.12, a Blackman-Harris window was used. The path length mismatch estimation on the shorter segments with other methods is shown in Figure 5.14. We note that with the exception of the first modulation of the experiment with varying amplitudes, all path length mismatch estimates fluctuate above the estimated uncertainty. The modulation during which the estimated path length mismatch does not fluctuate is at the same time the one with the largest commanded modulation amplitude of 0.5 rad. This amount of fluctuation could not be observed during the path length mismatch estimates using the data recorded during the laser frequency modulations applied for loop characterisation. It is unclear whether the same effect would be observed if the modulations for the loop characterisation would have lasted for a comparable number of cycles.

5.4.2 First checks

At first, we need to ensure that neither of the test masses is moving on the μm level and thus causing the observed fluctuations in path length mismatch. As an example, we show the GRS measurement recorded at 1 Hz during the second modulation of experiment four in Figure 5.13. We note no changes in position on the μm level. In this DFACS mode, the GRS is an independent out-of-loop sensor. Consequently, the variations which we note must be due to a sensing problem, as expected given their order of magnitude.

5.4. VARIATIONS OF PATH LENGTH MISMATCH ESTIMATES DURING INDIVIDUAL LASER FREQUENCY MODULATIONS




Figure 5.11: The estimated path length mismatch for shorter segments of the four laser frequency modulations of the dedicated path length mismatch experiment including the offset, experiment number 4.




Figure 5.12: The estimated path length mismatch for shorter segments of the four laser frequency modulations of the dedicated path length mismatch experiment with varying amplitudes, experiment number 6. With the exception of the first modulation with the largest commanded modulation amplitude (from approximately 300s until 1116s), the path length mismatch estimates fluctuate above the estimated uncertainty.

Another first check is made on the operational state of the laser system, exemplary for experiment 4. No unusual measurement of the laser data streams, as listed in Appendix B, could be found. The same is true for the laser frequency control loop parameters and the OB temperature




Figure 5.13: The GRS measurements of both test masses along the x direction during the second modulation of experiment 4. No change in position on the μm level can be observed.

sensors. Even though it would not be expected that changes in the reference interferometer could cause the observed short term fluctuations of the estimated path length mismatch, as they are subtracted from both measurements prior to the path length mismatch estimate, these measurements have been checked nonetheless. No unexpected features were discernible in the data stream. The contrast in o12 is above 97% during all of the time of the path length mismatch experiment including the offset. Calculating the DFT and fitting the sinusoidal modulations relies on knowing the correct modulation frequency in advance. This is the case here as we command the modulation and, to check that we modulated indeed at the commanded frequency, demodulated the data, as explained in Section 5.2. In addition, when fitting a short segment of a laser frequency modulation, jumps in the frequency can be compensated by a different phase. For the second modulation segment of experiment 4, we noted a small remaining signal at the modulation frequency in the fit residuals. This was due to a jump in the timing. However, this single jump does by far not explain the drift in path length mismatch during this modulation which can also be observed in the timing corrected data.

Several steps to exclude a data processing artefact have been undertaken. One of these is to apply different methods to the same measured data during the second modulation of experiment 4. The DFT, the demodulation and the linear fit method, as explained in Section 5.2, have all been applied. Note that due to the bandpass-filter applied, the beginning and the end of the demodulation results has been split such that they are shorter. All of the estimated path length mismatches shown in Figure 5.14 are starting later than the respective modulation because they are split for the timing jump. The results shown in Figure 5.14 show that all methods reveal a drift in path length mismatch. Figure 5.15 shows the amplitude spectrum of the the residuals of the linear fit to the o12 measurement during the second modulation of the path length mismatch experiment, experiment 4. Each trace corresponds to the data which was used to generate one path length mismatch data point estimate. The legend has been omitted for clarity. We see that the modulation has been reasonably well fitted in the sense that no peak around 2.879 Hz is visible here. In addition, the linear fit was verified on simulated data where the modulation

5.4. VARIATIONS OF PATH LENGTH MISMATCH ESTIMATES DURING INDIVIDUAL LASER FREQUENCY MODULATIONS




Figure 5.14: The estimated path length mismatch for shorter segments of the second modulation of the experiment 4.




Figure 5.15: The amplitude spectrum of the residuals of the linear fit to the o12 measurement during the second modulation of the path length mismatch experiment, experiment 4. Each trace corresponds to the data which was used to generate one path length mismatch data point estimate. The legend has been omitted for clarity. We see that the modulation has been reasonably well fitted.

amplitude was correctly estimated. The splitting of modulation data into smaller segments and the amplitude estimation via the DFT has also been validated by a simulation where the known modulation amplitude could be recovered as expected.

Another step to ensure the data processing is performed correctly is to vary the segment length and the overlap between the shorter segments of a single modulation. In case that the trend observed is truly present, it should not be affected by changes in either of these two parameters. Figure 5.16 shows the estimated path length mismatch for different durations of the data con-




Figure 5.16: The estimated path length mismatch during the second modulation of experiment 4 for different numbers of cycles used to perform the DFT. The drift in path length mismatch seems independent of the number of cycles chosen.

sidered in each DFT. The duration of each short segment varies between 32 and 384 cycles at 2.879 Hz and the short segments have a maximum overlap of two samples as a result of aiming at full modulation cycles. However, the trend in the path length mismatch estimate seems to be consistent among the different choices of data length. A similar check has been performed for the overlap in between the modulation segments. The results, as shown in Figure 5.17, indicate that the observed drift in estimated path length mismatch is independent of this choice.

Another possible reason for the observed changes in path length mismatch could be a time-dependent mixing of the lines in the o12 measurement at 1 and 3 Hz into the coupling from the laser frequency modulations at the modulation frequencies of 1.123 Hz and 2.879 Hz. This problem is illustrated in Figure 5.18. There we see the o12 spectrum during the second and fourth modulation of the path length mismatch offset experiment. The data has been split into 110 segments with the default overlap of 66.01% and a Blackman-Harris window has been used. We observe the peak at 2.879 Hz resulting from the coupling of the laser frequency modulation. During the fourth modulation, which took place at the offset position, the peaks at 1, 2, 3 and 4 Hz are well discernible and so is the proximity of the modulation peak to the 3 Hz line. In addition, the noise floor is higher at that time and not flat. Therefore, we have applied a notch-filter to the laser frequency modulation experiment data. The resulting amplitude spectral density for the fourth modulation of this experiment is shown in Figure 5.19. However, this has no impact to within the uncertainty levels on the estimated path length mismatch, as shown for the example of the fourth modulation of experiment 4 in Figure 5.20. This is different to Figures 5.16 and 5.17 where data of modulation 2 of this experiment has been used as an example. This

5.4. VARIATIONS OF PATH LENGTH MISMATCH ESTIMATES DURING INDIVIDUAL LASER FREQUENCY MODULATIONS




Figure 5.17: The estimated path length mismatch during the second modulation of experiment 4 for different overlap between the segments used to perform the DFT. The drift in path length mismatch does not seem to be affected by this.




Figure 5.18: The ASD of the o12 measurement during the laser frequency modulations at 2.879 Hz of the dedicated path length mismatch experiment including the offset (experiment 4). We note the coupling from the laser frequency noise is well discernible and so are the lines at 1 Hz and its multiples for the fourth modulation. In addition, the noise floor has changed.

check allows us to conclude that a mixing of the spectral lines at 1 Hz and its multiples is not the cause of the observed short-term path length mismatch fluctuations.

In addition, using the amplitude and the phase estimated via the linear time-domain fit, it was checked whether the estimated path length mismatches remain constant over time if the complex




Figure 5.19: The amplitude spectral density of the o12 measurement during the fourth laser frequency modulation at 2.879 Hz without and with the band-reject filter applied.




Figure 5.20: The estimated path length mismatch using the measurements as they are and data filtered as shown in Figure 5.19. No difference can be found to within the errors.

amplitudes

$$A_{12} = \sqrt{A_{12\sin}^2 + A_{12\cos}^2} \quad (5.12)$$

and analogous for Ψ_F are used for the estimation of the path length mismatch. However, this is not the case. The change in path length mismatch could not be attributed to changes in the phase estimated as

$$\phi_{12} = \text{atan2}(A_{12}^{\sin}, A_{12}^{\cos}) \quad (5.13)$$

5.4. VARIATIONS OF PATH LENGTH MISMATCH ESTIMATES DURING INDIVIDUAL LASER FREQUENCY MODULATIONS

or correspondingly for Ψ_F , either. To investigate unexpected DWS signals during the laser frequency modulation experiment, laser frequency modulation experiments have also been undertaken in the AEI LPF laboratory, as discussed in more detail in Chapter 6. During one of these experiments, the laser was frequency modulated with a modulation amplitude at the same order of magnitude as during the in-flight investigations. However, even though no short-term variations could be observed in this case, their presence cannot be ruled out, since the noise level of the longitudinal measurement in the laboratory on this day was around 5 to 10 pm/ $\sqrt{\text{Hz}}$ and thus significantly above the noise level in-flight. Unfortunately, no explanation for the observed short-term variations of the path length mismatch has been found yet. Given that we do not observe this fluctuation at the highest laser frequency modulation amplitude but that we observe it strongly during the modulations with low amplitude, it could be suspected to arise from noise in the o12 interferometer. It varies over time, perhaps similar to the time-dependency of the DWS noise floor due to aliasing, as discussed in Section 4.5.3 of [Wis17]. Also, in the time series of the demodulated o12 results stronger fluctuations can be observed than in the same time series of the in-loop laser frequency fluctuations, but further investigations might be required here.

These short term fluctuations make us consider the results reported in Section 5.3 as preliminary. However, we have found no reason to believe that the observation of spurious signals in DWS, as described in Chapter 6, could be an indication of a mechanism that impacts the longitudinal measurements and thus adds a systematic error to the path length mismatch numbers reported.

We reported on the results of two dedicated path length mismatch experiments. The evidence collected in the offset experiment allows us to associate the measured path length mismatches to the commanded offset at a 3σ uncertainty level, both in direction and amplitude. In addition, we could confirm that the estimated path length mismatch is independent of the laser frequency modulation amplitude and frequency, as expected. In general, we found the path length mismatch is only a few hundred μm which is a sign of excellent integration. We found that the path length difference in the X12 and X1 interferometers changed in the course of the mission due to a commanded offset. A detailed investigation of the path length mismatch experiments also revealed short-term variations of the path length difference. Even though those are not fully understood yet they are not believed to be caused by real motion of either TM, as this would have been measured by the GRS, or by any path length changes on the OB. Interestingly, these variations could not be found in the case of the highest amplitude modulation.

Chapter 6

Investigation of DWS signals during frequency modulations

During several laser frequency modulation experiments we noticed a signal in the DWS measurements at the same time. In this chapter we will describe the observation and the hypotheses that we investigated. Each of these has been investigated in-flight and in the AEI LISA Pathfinder laboratory. Unfortunately, we have not found the cause of this unexpected observation which makes it difficult to estimate the relevance of this finding for other scenarios.

6.1 Description, relevance and first checks

6.1.1 Observation

In several but not all laser frequency modulation experiments we note approximately 10 % of the signal amplitude in the longitudinal raw phase also in the DWS raw phase measurements when the test masses are free-falling. Figure 6.1 provides three examples. In these examples a sinusoidal signal at the frequency at which we modulate the laser is sometimes observable in the DWS channels. The raw or electrical phases (see Chapter 2) are used here for comparison between the signal in the longitudinal and DWS channels. In flight, we have to recombine the data from η_1 and η_2 to create η_{12} and similarly for ϕ_{12} . Consequently, a discernible spike in η_{12} or ϕ_{12} does not necessarily imply a discernible spike at the modulation frequency in the two TM orientation measurements. Figure 6.1 (a) shows data collected during experiment 1 (see Table 2.1). Note that this experiment took place before TM release. We assume poor contrast but during this experiment no contrast telemetry is available. Part (b) and (c) of this figure show examples collected during experiments performed with free-falling test masses.




Figure 6.1: A comparison of the o12 and the DWS signals converted to electrical phase at selected modulations of different laser frequency modulation experiments (see Table 2.1). (a) Experiment 1. (b) Experiment 3. (c) Experiment 5. In contrast to the modulations shown in (a) and (b), no impact of the laser frequency modulation onto the DWS channels is discernible in (c).

Such signals are not expected and thus are the subject of this chapter. Since these coincident signals in DWS have only been discovered after the end of the LISA Pathfinder mission operations, no dedicated experiments could be performed on the satellite. In the course of the investigations in-flight and in the AEI laboratory we found the following characteristics:

1. The coupling is spurious in the sense that it is not discernible in all four DWS channels during every laser frequency modulation in flight. These modulations are grouped into the experiments listed in Table 2.1. In the laboratory we also note changes of the coupling (see Section 6.3 for more details);
2. It seems to be different for all DWS channels;
3. It seems to depend linearly onto the laser frequency modulation amplitude.

The first two characteristics can already be observed in Figure 6.1.

6.1.2 Relevance

The effect described is interesting in itself but becomes even more so if we assume that the modulation signal we observe in DWS is the result of some kind of error which could be present in other measurement channels, such as for example o12, too. However, it is difficult to derive the amplitude of a possible error in a longitudinal measurement from the DWS measurement. On ground, we obtain two DWS measurements per TM. These are derived from the phase measurements of the X1 and X12 interferometers. Each of these interferometers uses two PDs with four quadrants each. As a result, we have four measured DWS signals obtained from 16 PD quadrants. Consequently, without additional telemetry such as a single quadrant for example, we cannot derive the error of a quadrant because the system is underdetermined. In addition, we measure only a difference between the two sides (left minus right and up minus down) of a PD in DWS, so possible common errors are already suppressed in the DWS.

Therefore, let us assume the signal we observe in DWS corresponds to an error of the raw longitudinal phase before scaling to TM position. This may be too pessimistic or optimistic but it is not possible to tell. For the remainder of this paragraph, we will focus on one observation at 0.1421 Hz during experiment 3 shown in Figure 6.1(b) because this experiment was performed with free-falling TMs and thus we assume a better contrast than during the fixed TMs case of Figure 6.1(a). Thus, we read off the mean signal amplitude in η_{12} and ϕ_{12} from the case depicted in Figure 6.1(b) to be $\approx 3 \mu\text{rad}$ [rms] and take this difference to be the phase error. This assumes implicitly that one side of the QPD causes the error or that the error occurs with opposite sign in the two signals which are subtracted in the DWS processing. One might as well assume that each side contributes half to this error. Consequently, the estimate would be reduced by 50% but here we stay with the assumption that the signal we observe in DWS corresponds to an error of the raw longitudinal phase before scaling to TM position. Hence, scaling this raw phase error to a displacement signal following Equation 2.7 yields

$$\Delta x_{12\text{err}} \approx \frac{\lambda}{4\pi\cos(\alpha)} \cdot 3 \mu\text{rad} [\text{rms}] \approx 0.25 \text{ pm} [\text{rms}] , \quad (6.1)$$

neglecting the subtraction of the reference interferometer signal and the phase tracking algorithm. Given the uncertainty of our assumption, this estimate has to be considered with care. However, it shows us that it could possibly reach relevant orders of magnitude. As also seen in Figure 6.1(b), this would correspond roughly to a 10 % amplitude error for the longitudinal measurement. The

maximum of the coupling on days of free-falling test masses was found to be to $(2.6 \pm 0.2) \cdot 10^{-8} \frac{\text{rad}}{\text{rad}}$. The units correspond to the TM orientation in rad per modulation amplitude in the phase of the frequency interferometer signals which are also measured in rad. The rms signal amplitude is in this case around 0.6 nrad in radian scaled to TM orientation. This is a stronger coupling than the one shown in Figure 6.1(b).

6.1.3 The structure of the chapter

This chapter summarises the investigation of the observed DWS signals. For clarity, all hypotheses are summarised in Figure 6.2. We will begin on the left of this diagram after a first check of an example time series.




Figure 6.2: A diagram of the hypotheses investigated in this chapter to explain the observed DWS signals. We set out at the left and already saw why a true TM rotation is unlikely.

6.1.4 First checks

As a first check, we show a time series measurement of ϕ_1 during a laser frequency modulation of experiment 1 in Figure 6.3. A modulation is visible in ϕ_1 , too.




Figure 6.3: An example of the measured ϕ_1 angle during the slowest laser frequency modulation at 11 mHz during experiment 1 with caged TM. We can see a modulation at this frequency for most of the time. This is very unlikely to be a true rotation of the satellite or TM1 given that the latter is still held by the caging mechanism at this time of the mission.

Of course, in nominal mission operations, the DWS measures the TM or satellite rotation. In the cases we investigate in this chapter, this is not a rotation measurement but a cross-coupling. One reason is that we see this effect also with grabbed test masses, as during experiment one. Another example of the DWS signals during a frequency modulation at 1.123 Hz is given in Figure 6.4. Unfortunately, for the timespan of the modulation shown in Figure 6.4, the GRS measurements are only sampled at 1 Hz which does not allow to cross check the rotation with the GRS in this case. In addition, the modulation in Figure 6.4 is below the sensing noise of the GRS of $83 \text{ nrad}/\sqrt{\text{Hz}}$ to $170 \text{ nrad}/\sqrt{\text{Hz}}$ [A⁺17b]. Even though the quoted numbers are estimated from data recorded at 1 Hz, they are also valid at modulation frequencies of up to 2.879 Hz because the sensing noise model is flat towards higher frequencies. However, even though we cannot double check with the GRS measurements here, a true TM rotation is very unlikely at this frequency given the bandwidth of the control loops for these angular degrees of freedom is around 2 mHz [Sch12, Table 4-9].

In flight, the time for the experiment is very limited and so is the data access. Thus, we decided to additionally study the coupling in the laboratory to see if we could observe the same effect¹. This seems to be the case as shown in Figure 6.5. This observation provides another argument against a true rotation since we also observe DWS modulations in the AEI laboratory (see Figure 6.5) where the TMs have been replaced by fixed mirrors.

¹all AEI laboratory measurements have been planned by A. Wittchen and S. Paczkowski. A. Wittchen has performed the measurements and adjusted the set-up where necessary. S. Paczkowski has analysed the data.




Figure 6.4: An example of the signal in DWS during a laser frequency modulation at 1.123 Hz. Note that the signals at 1 Hz (and its multiples) are almost always present. It is suspected that they are caused by electrical cross-coupling from a pulse per second timing signal on LPF but the details are still under investigation. As seen before, the coupling seems to be different for each degree of freedom.




Figure 6.5: The laboratory data is the η_{1A} channel recorded on 2018-11-07 during a laser frequency modulation with an rms amplitude of ≈ 75 MHz at 1 Hz. A third order polynomial has been subtracted from the data. The flight data is the η_1 measurement during the laser frequency modulation on 2017-01-22 with an rms amplitude of 31 MHz at 1.123 Hz. The same trend correction has been applied as for the lab data. We note that we can observe a modulation in both data sets. The signal amplitude is larger in the lab than in flight but we can observe the same effect.

Additionally, we would like to ensure the laser during the LPF mission is operating as expected during these frequency modulations and is not the cause of the spurious DWS signals. In some of the laser frequency modulation experiments, we noted a drift in the RLU temperature during the modulations. But this doesn't seem to be a necessary condition to see the coupling to DWS because even with a stable temperature in the course of the experiment, the coupling was discernible. One example is experiment 4. The mean RLU or AOM temperature is also unlikely to be the cause of the observed behaviour because we can observe the spurious coupling in experiment 3 with an RLU temperature of 28.2 °C and in experiment 7 with an RLU temperature of 22 °C. The same is true for the AOM temperature. For the Radio Frequency (RF) amplitudes 1 and 2 we also note a drift in experiments 1, 6 and 7 but again, this cannot be a necessary condition because it is not found in experiment 4. The same seems to hold true for the power stabilisation feedback 1 signal. All other relevant laser telemetry, as listed in Appendix B, did not show any features during the modulation experiments in the time or frequency domain.

6.2 Hypothesis: a true phase difference between the quadrants

In this section, we investigate two ideas how a truly different signal on the PD sides during a laser frequency modulation could arise which would then lead to a DWS signal. First, let us assume the difference in frequency between measurement and reference beam is not equal to the heterodyne frequency all the time because one beam travels a slightly different path than the other. Hence, in the worst case, one beam reaches the PD while its frequency is modulated to be the nominal beam frequency (including the respective frequency shift introduced by the AOM) plus the maximum laser frequency modulation amplitude while the other beam is at the nominal beam frequency (again including the respective frequency shift introduced by the AOM) minus the laser frequency modulation amplitude when it reaches the PD, as would be the case at the time T in Figure 6.6. This hypothesis could theoretically not only account for a longitudinal signal but also for a DWS signal, if we assume that the rotation causes a difference in path length and thus in the deviation from the nominal beam frequency. However, with a realistically small rotation both beam paths are still ≈ 38 cm and thus the light reaches the PD after ≈ 1.2 ns, which is orders of magnitude smaller than a quarter of the modulation period at our highest modulation frequency of 2.879 Hz.

Another hypothesis is to assume the same coupling mechanism as for the longitudinal signal and to associate a rotation of the measurement beam with respect to the reference beam to a path length mismatch. However, this approach implies that a laser frequency modulation affects the part of the beam incident on the left side of the QPD differently to the part that is incident on the right side and, correspondingly, for the upper and lower part. But so far, we have found no mechanism or component that could have caused this asymmetric beam modulation. Even though we do not expect this observation to be caused by the shape of the wavefronts, clipping on the PD or effects of the beams hitting the PD not perfectly centred, this has been cross checked by a dedicated IfoCAD² simulation. In this simulation, the path of the measurement beam (length L_1) and the reference beam (length L_2) has a mechanical difference in length. Hence, the maximum impact of the frequency modulation is also different. To see this, we can

²A laser interferometer simulation tool as described for example in [K+12] and available under gitlab.aei.uni-hannover.de/ifocad

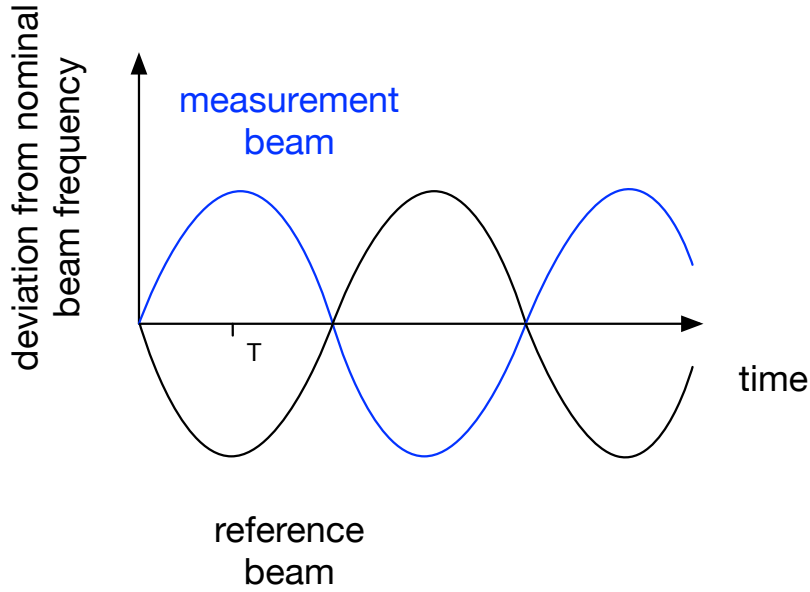


Figure 6.6: A sketch of a possible hypothesis for the measured signal in DWS. The idea here is that the two path lengths of the measurement and reference beam differ such that at a time T , the difference in frequency between the measurement and the reference beam is not the heterodyne frequency but, in the worst case, additionally twice the amplitude of the frequency modulation. This hypothesis can be discarded.

combine the phase expressed as

$$\varphi = \tau \cdot \Delta\omega \quad (6.2)$$

and the optical path length, written as

$$s = \frac{\lambda}{2\pi} \varphi, \quad (6.3)$$

into the maximum change in optical path length for a beam i , $L_{i\max}$, resulting in

$$L_{i\max} = \frac{\lambda}{2\pi} \frac{L_i}{c} \Delta\omega. \quad (6.4)$$

In the next step we have adjusted the length of the beams to account for the impact of laser frequency modulation. This is how we simulate the wavefronts of frequency modulated light that reach the PD. From this simulation we could see that this effect is significantly too small.

In addition, these two systematic explanations seem unlikely given that we cannot observe the same coupling in each degree of freedom, on each day and for each frequency in the LPF satellite data or in the laboratory data (see Section 6.3).

6.3 Hypothesis: related to beam tilts

After we saw that the signal in DWS during the laser frequency modulations is unlikely to be caused by a true test mass or satellite rotation, let us investigate if it can be due to a tilt of the beam that occurs during the modulation.

One possible cause could be that the OPD actuators do not only increase or decrease the path length but also tilt the beam slightly. This idea came up because the path length mismatch in the reference interferometer causes the frequency modulations to couple into Ψ_R and, as the OPD loop is active during the in-flight laser frequency modulations, it tries to counteract the perceived changes in optical path length.

A second mechanism that could lead to beam tilt is related to the laser crystal. At frequencies above approximately 0.1 Hz, the response of the fast actuator dominates the response of the slow controller and actuator, see Section 3.3.2. Since the fast frequency actuator is a piezo acting on the laser crystal, this may also cause mechanical stress on the crystal and thus affect the beam pointing [Heu04, p. 50]. The loop is effective up to the unity gain frequency at approximately 0.8 Hz, see Section 3.3.3. Hence there is a small frequency range where this mechanism may occur. However, in both of the above cases the laser light is afterwards transmitted to the OB via fibres and FIOSs, as explained in Section 2.1. Nominally, we expect the pointing of the beam that comes out of the FIOS onto the optical bench should stay the same independent of the pointing of the light that enters the fibre for changes in beam direction which are small enough such that the light is still well coupled into the fibre. However, this may not be true for short fibres, which is not likely to be the case on LPF, or fibres that are not sufficiently rolled. In these cases, also in a single mode fibre, cladding modes may appear which may affect the output. Therefore, let us see nevertheless what we can find out about beam tilt from the in-flight data.

6.3.1 In-flight investigations

In the frequency and reference interferometer there are no moving components so we should be able to discern a beam tilt here. However, a DWS signal of the frequency and reference interferometer is not available from telemetry. That is why we have to use the DC angles which have been explained in Section 2.1. For the frequency and reference interferometer, these measurements have not been calibrated to a rotation and therefore are shown without units. They should also show the effect of a possible beam tilt where the point of rotation lies at the fibre output. For such a rotation, the DC angles should be the same in both the frequency and the reference interferometer if the different path lengths from the fibre output to the respective PDs are taken into account as lever arms. From the design of the OB, we find that for the reference beam, the mean of the distance from the fibre output to the A and B diode of the frequency interferometer is a factor 1.7 larger than the mean of the distance to the diodes of the reference interferometer. Similarly, a factor 0.3 was found for the measurement beam.

One example for the AS of the DC angles during a laser frequency modulation experiment is shown in Figure 6.7. Indeed, the two η measurements and the two ϕ measurements of the two interferometers show a signal at the modulation frequency of the laser. This is also true for the other modulations of the third laser frequency modulation experiment which was recorded with 1 Hz data. Therefore, further checks of these angles during the other laser frequency loop characterisation experiments have not been considered necessary. Despite the observation of a




Figure 6.7: An example for the DC angles during a laser frequency modulation of Exp 3 (June 2016) in both the reference and the frequency interferometer. The detected power modulations agree quite well (see Figure 6.8) but the measured angles don't.




Figure 6.8: An example for the power measurements during a laser frequency modulation of Exp 3 (June 2016) in both the reference and the frequency interferometer. The detected power modulations agree quite well. As the fast power loop is active at the same time, this is an in-loop measurement and the true power fluctuations could be even larger. The power fluctuations are further investigated in Section 6.4.

modulation, the ratios of the signal amplitudes do not agree to the expectations of a rotation as described above. Consequently, Figure 6.7 is contradicting the beam tilt hypothesis.

However, these DC angles may also be affected by changes in the total power of the beam on the optical bench, see Section 6.4. But if we compare the power changes measured in these two interferometers, they agree quite well, as can be seen in Figure 6.8. This agreement is expected. The observed power modulation will also be further discussed in Section 6.4 but it cannot explain the observed ratios of the signal amplitudes in Figure 6.7.

This check is not in favour of the beam tilt hypothesis, which was unlikely in the first place. Especially the OPD actuator hypothesis is further contradicted because it implies larger changes in *all* four DC angle measurements during *all* of the OPD modulation experiments than during the laser frequency modulation experiments which was not found to be the case. This is not the case as shown in Figure 6.9.




Figure 6.9: An example where the DC angle modulations in the fixed interferometers during the OPD modulation experiment are not above those recorded during the frequency loop modulation experiment. This is contradicting the OPD actuator hypothesis.

Another implication of the OPD actuator hypothesis is that the DC angles should change with the OPD loop being open or closed which was again not the case. However, we may also observe a slight DWS modulation during at least one modulation of one OPD loop characterisation experiment.

But as these two do not necessarily have the same cause, we will not investigate this observation further in this thesis. In addition, it may also be possible that the DC angles studied in this paragraph are subject to significant cross-talk. This is further investigated in Section 6.6.1.

6.3.2 Follow up of beam tilt in the laboratory

A follow-up of the beam tilt hypothesis in the laboratory was undertaken, as the reason for the modulation of the DC angles is not clear. We have repeated a similar frequency modulation experiment in the laboratory but here either the measurement beam or the reference beam have

been blocked such that only a single beam hits the PDs while the laser was frequency-modulated. We then record the measured DC angles as it is not possible to calculate the DWS signal with a single beam. The observed DC angle measurement during the laser frequency modulation depends on three factors. These are the location of the beam spot on the PD, the tilt of the beam and changes in the beam power. Now if none of these factors is changed during the laser frequency modulation, no signal in the Amplitude Spectrum (AS) of the DC angle measurements is expected. The DC angle measurements from this experiment are shown in Figure 6.10. Here




Figure 6.10: The AS of the DC angle measurements on the two diodes of the X1 interferometer while either of the beams was blocked. We note the DC angle signal at the modulation frequency which has the same order of magnitude except for η_{1B} while blocking the reference beam. For comparison, measurements of the respective channels, while the reference beam is blocked but no modulation is applied are shown as straight lines. This comparison ensures that the peak at 8 Hz is indeed caused by the modulation, contrary to the other peaks in the AS.

we show the AS of the DC angle measurements on the two diodes of the X1 interferometer. The laser has been frequency modulated at 8 Hz and either one of the beams was blocked. We note the DC angle signal at the modulation frequency which has the same order of magnitude except for η_{1B} while blocking the reference beam. The DC angle readout signal is a unitless quantity because the difference in power between left and right or up and down PD quadrants is divided by the calculated total power, denoted Σ . Also note that these DC angle values are not calibrated to test mass rotation. Before we can attribute the signal here to beam tilt, we need to exclude that it is due to changes in power of a slightly misaligned beam. With perfect alignment, a change in the total power would not affect the DC angle signal but during our experiment, we have a slight misalignment as listed in Section 6.3.3.

That is why we have performed an additional experiment to determine the coupling coefficient from power modulations to DC angle for the measurement beam while the reference beam was blocked. The power modulations have been applied by changing the AOM output power of the measurement beam accordingly. Figure 6.11 shows the power measurements during one of the power modulations and the frequency modulation. We find the deliberate power modulations match the order of magnitude of those we observe during a laser frequency modulation, even




Figure 6.11: The Σ_A and Σ_B measurements during one power modulation with a small modulation amplitude and one laser frequency modulation. The data has been filtered with a pass-band from 7.7 Hz – 8.3 Hz and the modulation was applied at 8 Hz. The power modulation amplitude on the B diode is larger than on the A diode. The Σ values are so far uncalibrated and in arbitrary units.

though the power fluctuations during the laser frequency modulation are very non-stationary in amplitude. We calculate the coupling from power to DC angle, assuming it is the same despite the small difference in power and signal shape. As the coupling from power to both beams seems to be comparable, except for η_{1B} , compare Figure 6.10, it was not deemed necessary to repeat the measurement with a power modulation of the reference beam only.

The next step to find out how much of the DC angle signal is due to change in power is to multiply the measured power signal during the laser frequency modulation with the coupling coefficient determined from the power modulations. The result is compared to the measured DC angle signal. This is shown in Figure 6.12. Since we can explain three out of four signals to a significant fraction with the coupling from power to DC angle, the beam tilt hypothesis becomes unlikely. However, as mentioned before, the amount of coupling from changes in power to DC angle depends on the alignment of the beam. This is studied separately in the next subsection.

6.3.3 Misalignments - laboratory

In general, the alignment of the two beams may affect the DWS measurements and thus the observation of spurious signals therein. In flight the contrast in both TM interferometers was above 95% for most of the time of the nominal mission operation period, compare Figure 3.2 in [Wis17]. An exception are certain experiments or the times of grabbed test masses. Consequently, the beams must have been well aligned. Nonetheless, this is a possible impact factor that we investigated in the laboratory.




Figure 6.12: A comparison of the DC angle signal during the laser frequency modulation to the expected DC angle signal from power modulation for η and ϕ of both diodes. (a) η_{1A} (b) ϕ_{1A} (c) η_{1B} (d) ϕ_{1B} . We can explain the majority of the DC angle signal with the result from power modulation except for the ϕ_{1B} quadrant. Note that also here the DC angle measurement has not been calibrated to a test mass rotation and thus is a measurement of how well centred the beams are on the PD.

The DC angle measurements are affected by the location of the beams on the PD. This is because if the centre of either of the beams does not hit the exact centre of the PD, a larger fraction of the light power reaches for example the top two quadrants than the two bottom quadrants (in case of a misalignment in η) and a non-vanishing DC angle signal is measured. If then the power of a single beam is modulated, the total power on the top and bottom quadrants changes differently and we see a modulation in the respective DC angle measurement.

This is the first check that has been done in the laboratory. Next we checked whether this is the same for DWS. However, as long as it is not certain that the coupling from frequency modulation to DWS is via the changes in power, a third check is necessary: we also have to compare the DWS signal during a laser frequency modulation before and after the alignment. Therefore, we compare the DC angle and DWS measurements during a power modulation and during the same power modulation after alignments in the laboratory. This is being done for the ϕ_{1B} channel in which the alignment could be improved best: from a static DC angle value of $-6.14 \cdot 10^{-4}$ rad on December, 12 2018 to $6.7 \cdot 10^{-5}$ rad on December, 17 2018. Note that these mean DC angle levels are in rad because the measurements have been scaled to test mass rotation. The DC angle measurement before and after the alignment is shown in Figure 6.13. As expected, the improved alignment reduces the coupling from the power modulation into the DC angle channel. Next we check whether this is the same for DWS. The corresponding DWS channel is shown in Figure 6.14. This data suggests that the impact of the power modulation onto the ϕ_{1B} DWS signal does not improve with alignment. In fact, we also have to consider that the coupling from power modulation to DWS ϕ_{1B} varies with time as indicated by Figure 6.15.




Figure 6.13: The amplitude spectrum of the DC angle ϕ_{1B} channel during the same power modulation before the adjustments on December, 17 2018 (dotted black trace) and after (blue trace). As expected, the improved alignment reduces the coupling from the power modulation into the DC angle channel.




Figure 6.14: The amplitude spectrum of the DWS ϕ_{1B} channel during the same power modulation before the alignment was improved on December, 17 2018 and after. Despite the improved alignment, a coupling from the power modulation into the DWS channel is discernible in this case.

For our example channel we have seen that the power modulations are less discernible in the DC angles if we improve the alignment, as expected. However, even with improved alignment, the coupling from power to DWS does not seem to vanish. This coupling seems to fluctuate by a factor three between different runs of the experiment on different days, compare Figure 6.15. In




Figure 6.15: The amplitude spectrum of the DWS ϕ_{1B} channel during the same power modulation on different days. The alignment of the beams has not been adjusted in between. The power has been modulated with 50 counts modulation amplitude to mimic the small power modulation which we observe during a frequency modulation.

the other DWS channels, the signal amplitude during the power modulations is fluctuating as well. However, as long as we are not sure that the coupling from laser frequency modulation to DWS is due to the fact that the frequency modulation also affects the power which then couples to DWS, we also have to compare the DWS signal from a frequency modulation before and after the improvement of the alignment. The results of this comparison are shown in Figure 6.16. In




Figure 6.16: The ϕ_{1B} DWS channel during a laser frequency modulation before the improvement of the alignment (black trace) and after (blue trace). The OPD is deactivated in both cases.

CHAPTER 6. INVESTIGATION OF DWS SIGNALS DURING FREQUENCY MODULATIONS

both experiments the laser was frequency modulated at 8 Hz with an amplitude of approximately 0.048 rad rms and the OPD loop was off. With improved alignment, the peak amplitude of the DWS ϕ_{1B} channel decreased by roughly a factor three. To claim that the alignment is a driving




Figure 6.17: The estimated coupling from Ψ_F to the respective DWS channel during selected laser frequency modulation experiments in the laboratory. We note that the coupling to ϕ_{1B} changes by approximately one order of magnitude.

factor for the observed coupling, the change in Figure 6.16 should be larger than the changes we observe over time.

This is not the case for ϕ_{1B} , as can be seen from Figure 6.17. Here, we show the coupling from Ψ_F to each of the respective DWS channels over several laser frequency modulations applied in the laboratory. The DWS measurements have been scaled to TM orientation radian prior to taking the ratio with the Ψ_F measurement. Thus, the coupling is in rad rad^{-1} and as such without units. On the x-axis, we denote the number of the modulation. Several modulations have been part of one experiment. Here, we only used data from frequency modulations recorded with the nominal set-up in terms of PM cables plugged in. We did include data recorded at different TM angles because we saw, as described in Section 6.5, that this does not have a significant impact on the coupling. We have taken the ratio of the Fourier transforms at the modulation frequency of 1 Hz, 2 Hz or 8 Hz, respectively. The uncertainty estimates attributed to the data points in Figure 6.17 are a result of the noise floor estimate. Given modulation durations of at least 50 sec for these modulations at 1 Hz and above, this is considered reasonable.

The data in Figure 6.17 suggests that the coupling to η_{1B} is the smallest. Most of the time it increases for ϕ_{1A} and η_{1A} . The largest coupling has been found for ϕ_{1B} . The changes in the estimated coupling also depend on the DWS measurement. Indeed, in ϕ_{1B} , the coupling ranged from $(14 \pm 2) \text{ nrad rad}^{-1}$ to $(158 \pm 15) \text{ nrad rad}^{-1}$ whereas the change in the coupling to η_{1B} is the smallest.

Now, to state that the alignment is a significant contribution to the observed coupling to DWS, it would have to be larger than the fluctuations we observe over time. This is not the case and

hence we have to conclude that the alignment is not the main reason for the observed coupling from laser frequency modulations to DWS.

6.4 Hypothesis: power related

During our laser frequency modulations we sometimes observed a modulation of the powers on the optical bench at the same time. The hypothesis investigated in this section is the following: the laser frequency modulation causes a power modulation which in turn couples into DWS, as illustrated in Figure 6.18. With this analysis and the laboratory experiments, we want to find out if one observation is the cause of the other.

This idea seems valid given the order of magnitude of the couplings, as listed in Table 6.1. This table compares estimates of the coupling coefficients from laser frequency modulation to




Figure 6.18: A sketch of the idea that the coupling from laser frequency to DWS is an indirect one from laser frequency modulation (Ψ_F) via a power modulation (Σ) to DWS (ϕ, η).

power and from power to DWS to the ratio of the signal measured in DWS per Ψ_F modulation amplitude for experiment 1. The data of modulations 1 to 4 has been used where the impact of the modulation has been discernible by eye. The coefficients have been determined in the frequency domain using a Hanning window. The values reported are averaged for each channel over the four modulations and the error on the mean is reported. We find that the product of the coupling from $\Psi_F \rightarrow \Sigma_1$ and $\Sigma_1 \rightarrow$ DWS matches the measured coupling from $\Psi_F \rightarrow$ DWS. However, these values have to be considered with care given that during this experiment no contrast data is available and the coupling from laser frequency modulation to DWS is fluctuating from one experiment to the next.

Table 6.2 shows the same comparison for a frequency modulation experiment performed in the AEI laboratory on 2018-11-15. Here we note that the calculated product of the two couplings matches very well the measured coupling. The uncertainty estimates are from the estimates of the noise floor at adjacent frequencies of the modulation. For this case we perform a single Fourier transform using a Hanning window for at least 240 sec of modulation data at 8 Hz. Thus, we obtain a reasonable error estimate here.

	$\Psi_F \rightarrow \Sigma_1$ [$\mu\text{W rad}^{-1}$]	$\Sigma_1 \rightarrow \text{DWS}$ [rad W^{-1}]	calculated [rad rad^{-1}]	$\Psi_F \rightarrow \text{DWS}$ [rad rad^{-1}]
η_1	0.22 ± 0.08	$(3.3 \pm 1.5) \cdot 10^3$	$(0.75 \pm 0.45) \cdot 10^{-3}$	$(0.47 \pm 0.1) \cdot 10^{-3}$
η_{12}	0.22 ± 0.08	$(15 \pm 7) \cdot 10^6$	3.5 ± 2.1	2 ± 0.5
ϕ_1	0.22 ± 0.08	$(82 \pm 49) \cdot 10^3$	$(18 \pm 13) \cdot 10^{-3}$	$(8 \pm 0.3) \cdot 10^{-3}$
ϕ_{12}	0.22 ± 0.08	$(378 \pm 227) \cdot 10^6$	84 ± 62	39 ± 1

Table 6.1: Estimates of the in-flight coupling coefficients from laser frequency modulation to power and to DWS for experiment 1. We find that the product of the coupling from $\Psi_F \rightarrow \Sigma_1$ and $\Sigma_1 \rightarrow \text{DWS}$ is compatible with the coupling measured from $\Psi_F \rightarrow \text{DWS}$.

	$\Psi_F \rightarrow \Sigma_{A/B}$ [arb. rad^{-1}]	$\Sigma_{A/B} \rightarrow \text{DWS}$ [rad arb.^{-1}]	calculated [rad rad^{-1}]	$\Psi_F \rightarrow \text{DWS}$ [rad rad^{-1}]
η_A	$(5.3 \pm 0.8) \cdot 10^{-4}$	$(1.6 \pm 0.3) \cdot 10^{-4}$	$(8.9 \pm 2) \cdot 10^{-8}$	$(8 \pm 0.9) \cdot 10^{-8}$
η_B	$(7 \pm 0.6) \cdot 10^{-4}$	$(3.2 \pm 0.9) \cdot 10^{-5}$	$(2.2 \pm 0.6) \cdot 10^{-8}$	$(2.2 \pm 0.7) \cdot 10^{-8}$
ϕ_A	$(5.3 \pm 0.8) \cdot 10^{-4}$	$(1.1 \pm 0.3) \cdot 10^{-4}$	$(6 \pm 1.9) \cdot 10^{-8}$	$(5 \pm 1) \cdot 10^{-8}$
ϕ_B	$(7 \pm 0.6) \cdot 10^{-4}$	$(1.8 \pm 0.2) \cdot 10^{-4}$	$(1.3 \pm 0.19) \cdot 10^{-7}$	$(1.3 \pm 0.05) \cdot 10^{-7}$

Table 6.2: Estimates of the coupling coefficients, measured in the AEI laboratory, from laser frequency modulation to power and to DWS on 2018-11-15. Again, we find that the calculated twofold coupling matches quite well the measured signal amplitude.

Nonetheless, this table is in favour of a further investigation of the twofold coupling as shown in Figure 6.18. We find a comparable match between the product of the twofold coupling and the direct measurement in the laboratory measurements, as shown in Table 6.2. Note that the laboratory power measurements are in uncalibrated arbitrary units. The DWS values have been scaled to TM angles. Therefore the entries in the first two columns of Tables 6.1 and 6.2 are different. However, this is a useful check because we have good contrast in the lab and we can average over modulations at 8 Hz, each of them recorded for at least 1920 cycles in contrast to the flight data where each modulation has only been recorded for 20 cycles.

However, this twofold coupling hypothesis may seem exotic because we need to explain why the power is modulated simultaneously with the laser frequency. One hypothesis investigated in Section 6.4.2 is that the frequency modulations have an impact on the laser polarisation which in turn affects the measured power. In addition, we investigated in the laboratory whether the power fluctuations occur already before we couple the laser light into the fibre and then transmit it onto the optical bench, as explained in Section 6.4.4.

The next step in this hypothesis is to explain the coupling from power to DWS. Note that this is a very unusual hypothesis in the first place. This is because the DWS phase is computed as the argument of a complex vector, compare Equation 2.8. In the phasor picture, the argument of the complex number is the phase of the vector in the complex plane which is independent of the vector's length unless other spurious signals are mixing in. Moreover, a coupling from a power modulation to a phase modulation at the same frequency is unusual, in contrast to RIN noise, for example which couples at the heterodyne frequency and twice the heterodyne frequency [Wis17]. Depending on the order of magnitude, this could be due to a computational inaccuracy, compare [Kau21].

In this section, we summarise our experiments concerning a possible coupling from laser frequency modulation to the measured power on the optical bench and from there to the spurious DWS signal. Let us begin with a description of the observed power modulation.

6.4.1 The observed signals

During the laser frequency modulation experiments we sometimes observe power modulations. Figure 6.19 shows the AS of the RIN in the reference interferometer during several laser frequency modulation experiments. During the two dedicated path length mismatch experiments the laser frequency modulations took place above the data recording frequency of the Σ channels. So Figure 6.19 only shows the data recorded during the modulations for loop characterisation. Therefore, we also cannot say whether we always see a power modulation during the laser frequency modulation. What we can say is that there is, for example, a power modulation discernible in Σ_R in the second frequency loop characterisation experiment where there is no DWS modulation present. We show one amplitude spectrum for all modulation frequencies. Note that the conversion from the Σ measurement to power has been done using an estimated coefficient of 1.024 mW. This coefficient was obtained by first calibrating the measurements of the power monitor diodes PDA1 and PDA2 (see Figure 2.2) back to the power of the incident light using the known ADC range, TIA resistance and an assumed PD efficiency of 0.8, as explained in [Aud14]. These measurements have then been propagated to the X12 PDs taking into account the optical components along the beam paths. The coefficient used in this analysis is an estimate of the constant of proportionality between the thus estimated power and the Σ measurement. This calibration coefficient neglects the slight differences between the interferometers. A more detailed and precise calibration based on the measurements of changes in radiation pressure noise




Figure 6.19: The amplitude spectrum of RIN in the reference interferometer for all laser frequency loop characterisation experiments. With the exception of experiments 1 and 7, a RIN signal is discernible at the lowest modulation frequency of 11 mHz.

can be found in [Kau21]. However, for this purpose and the argumentation in Section 6.4.3, the simple estimate of the calibration coefficient is sufficient.

In Figure 6.19 at the lowest modulation frequency, the measured peak in power seems much smaller in the first (blue trace) and last experiment (yellow trace) even though the commanded amplitude at that frequency was constant over the experiments. However, from the in-flight data we cannot estimate how the commanded modulation amplitude affects the power measured in the Σ channels. Even though we have a dedicated path length mismatch experiment with varying modulation amplitude, we modulate the laser frequency in this experiment at 1.123 Hz but the Σ values of the four interferometers are only available at 1 Hz.

In addition, the laser amplitude control loop telemetry has been checked for changes in power. The amplitude control loop error signal, its feedback signal and the measured power on two dedicated PDs are the parameters available. These PDs, called PDA1 and PDA2, are not part of any of the four interferometers, see Figure 2.2. They are an in-loop measurement of the fast amplitude control loop. However, no effect of the frequency modulation was seen in these three channels. The available power measurements from telemetry are too noisy to detect what we see in the Σ values. In addition, many laser frequency modulations have been above the Nyquist frequency of 0.1 Hz in the available channels.

6.4.2 Polarisation as a possible explanation for the observed power modulation

To understand the twofold coupling from laser frequency modulations via power modulations to DWS, we also need an explanation for the observed power modulations. One possible reason we want to check here is the polarisation. A polarisation-related effect should always be smaller in the X12 interferometer than in the frequency interferometer. Also the ratio between the effect in

the X12 and frequency interferometer should always remain the same. These two characteristics have been identified and justified in [Kau21] and can be checked in Figure 6.20. Here we show




Figure 6.20: An example for the AS of the Σ channels of the four interferometers during two laser frequency modulations of the frequency loop characterisation experiment number 5. We note that at 11 mHz the signal is the strongest in Σ_R and at approximately 39 mHz the signal in Σ_F is the strongest. Furthermore, note that the ratio between the signal amplitudes in Σ_F and Σ_{12} does not stay constant even in this example.

an example for the AS of the Σ channels of the four interferometers during two laser frequency modulations of the frequency loop characterisation experiment number 5. We note that at 11 mHz the signal is the strongest in Σ_R and at approximately 39 mHz the signal in Σ_F is the strongest. Furthermore, note that the ratio between the signal amplitudes in Σ_F and Σ_{12} does not stay constant even in this example. Consequently, we can discard this explanation but this is not sufficient to exclude a coupling from frequency modulation to power via a different mechanism.

At this point, we would also like to test the hypothesis that the power modulations explain the signal in the DC angles and not the beam tilt. If that was true, the same effect with a larger amplitude would also have been seen in the dedicated power modulation experiments. This means, the ratio between the signal in the DC angles in the fixed interferometers and the respective Σ values should be the same in the laser frequency and laser power modulation experiments. This is not the case as can be seen from Figures 6.21 and 6.22. To be more explicit: the ratio between the solid blue curve and the solid red and green trace should be the same and between the dashed lines of the same colour. Here, we show an amplitude spectrum using a Blackman-Harris window and we take a single average of the data.

We have also checked the RF amplitude signals during experiment 3 for modulations at the laser modulation frequency but they were not discernible. Since the sampling frequency of the RF amplitude is 0.2 Hz, it was not possible to check for all modulations of this experiment. The data settings have been different for experiment 1, but also there, no modulations of the RF amplitude have been seen. During the modulation shown in Figures 6.21 and 6.22, the laser frequency was modulated with an rms amplitude of 4.4 MHz. Note that even though the commanded modulation amplitudes are 0.05 rad for all frequencies of the loop characterisation

CHAPTER 6. INVESTIGATION OF DWS SIGNALS DURING FREQUENCY MODULATIONS




Figure 6.21: Comparison of Σ_F and the DC angles in the frequency interferometer during the dedicated power modulation experiment and during the laser frequency modulation experiment (Exp 3). Note that we never modulated the laser frequency with a frequency below 0.011 Hz so we cannot compare at the same frequency.




Figure 6.22: Comparison of Σ_R and the DC angles in the reference interferometer during the dedicated power modulation experiment and during the laser frequency modulation experiment (Exp 3). Note that we never modulated the laser frequency with a frequency below 0.011 Hz so we cannot compare at the same frequency.

experiments, the respective Ψ_F modulation amplitudes measured depend on the frequency of interest due to the control loop being active.

6.4.3 Checking the coupling from power to DWS

In the course of LPF data analysis, a coupling from power measurements to DWS of up to $3.157 \text{ mrad W}^{-1}$ was observed in a jump in the DWS ϕ_1 time series that coincides with a step in the laser power. This value is an upper limit determined from the coupling to ϕ_1 with approximately $50 \mu\text{rad}$ offset, scaled to TM radian, and may also be depending on the power on the PD. It is suspected that this coupling depends on the absolute phase of the respective DWS angle which is proportional to the TM tilt [Kau21].

To check if this effect could also explain the coupling from power to DWS during the laser frequency modulation experiments, the orders of magnitude should be compared. Therefore, the ratio of the amplitudes in η_1 and laser power at the modulation frequency is computed. The laser power is derived from the Σ_1 measurement, as described in Section 6.4.1. η_1 is an example degree of freedom and the measurement used for this computation is scaled to TM orientation. This is different to the scaling of Table 6.1 where the DWS measurement is given as the unscaled phase measurement. Since the power measurement is only available up to frequencies of 0.5 Hz, only the data of the loop characterisation measurements could be used here. The power measurement data during the sixth modulation of experiment 7, however, seemed unusual. Therefore, this single modulation was discarded. All other modulations have been included, regardless of whether a signal in the η_1 channel was observable. The cases where no signal was discernible correspond to data points with large uncertainties.




Figure 6.23: The coupling from power measurement to η_1 for all laser frequency modulations below 0.5 Hz with nominal Σ_1 data. We note the order of magnitude is significantly above the upper limit predicted in [Kau21].

Figure 6.23 compares the thus determined coupling to the upper limit of [Kau21]. We find that the coupling during the laser frequency modulations is orders of magnitude above this upper limit for the modulations with small uncertainties. This is the case for most modulations of experiments 1 and 7. In addition, a dependency on the absolute phase of the DWS signal was not found (compare Section 6.5). For these two reasons, this already observed coupling mechanism is discarded.

6.4.4 Investigation of coupling of laser frequency modulation to power in the laboratory

The hypothesis of the twofold coupling (see Figure 6.18) was also investigated in the AEI laboratory. In a first step, a comparable coupling from power to DWS was reproduced in the laboratory. One example is shown in Figure 6.24.




Figure 6.24: (a) AS of the uncalibrated power measurement on the B diode during the three power modulations and during a noise measurement. (b) This is an example of the measured η_{1B} angle during these power modulations. A spike at the modulation frequency is discernible in the three measurements of laser power modulation but not in the noise measurement data. This is similar for the other DWS channels.

The power modulation amplitude during this experiment was chosen such that we are only a factor six above the order of magnitude of the power modulations which we record during the laser frequency modulations in the laboratory.

We would also like to understand the origin of the power fluctuations observed during the laser frequency modulations. For this reason, we tried to determine whether they originate in the laser itself or somewhere along the path to the power measurement on the optical bench.

Therefore, we measured the power during a frequency modulation at two different points in the laboratory. One measurement was taken from a reflection of the back of a coupling mirror, which is usually dumped for laser safety reasons. This reflection occurs before the light is coupled into the fibre on the modulation bench, as shown in Figure 6.25. A PD was placed into the relatively large and scattered reflection. Thus, we only detect the central part of the beam. However, at this position in the set-up, we have no reason to assume significant higher order modes which we would neglect when measuring only the central part of the beam. In addition, we assume with this set-up that the changes in power which we want to measure are independent of possible changes in polarisation which depend on the mirror on which the beam is reflected. We have found no reasons that contradict this assumption and, considering the qualitative nature of this test, this set-up is sufficient.




Figure 6.25: A sketch to illustrate the position of the PD used in the laboratory to measure the beam power before the light reaches the OB.

Hence, the PD measured the light before it entered the fibres and the OB. This is a more direct measurement of what happens in a laser, similar to the in-flight laser, itself. This measurement was recorded directly from the TIA output and not via the Control and Design System (CDS). At the same time the measured power was recorded with our standard measurement of the Σ channels on the two PDs of the X1 interferometer via the CDS and again in the X12 interferometer from the TIA output. To ensure the power measurement is working, we modulated at first the laser power directly. This was discernible in all channels for a modulation amplitude of $30 \cdot 10^3$ counts at 0.07 Hz. So are the changes in power during the laser frequency modulation at 4.5 Hz, as can be seen in Figure 6.26. The spikes above 10 Hz are unrelated to this modulation because they also appeared during other modulations. However, we cannot immediately conclude that the power changes occur in the laser itself because to be sure that the power is changing, we need to exclude that the beam is moving on any of the PDs. This is especially important for the power measurement taken before the light enters the optical bench as the reflected beam has a comparatively large area. On the other hand, we have concluded in Section 6.3.2 that we can explain our modulations of the DC angle measurements with changes in power and hence that they are not due to changes in beam pointing. This argument should also be valid here. Even though the reason for the difference in peak height at the modulation frequency between the measurement recorded in the reflection (green curve) and the others has not been fully understood yet, Figure 6.26 supports the idea that the power modulations may occur in the laser itself.




Figure 6.26: Amplitude spectrum of the measured RIN during the laser frequency modulations at 4.5 Hz. At this frequency changes in power are distinguishable in all channels.

To further assess the hypothesis of the twofold coupling, we compare the calculated coupling from the Σ measurements to DWS of Table 6.2 to the DWS signals recorded during dedicated laser power modulations in the laboratory. In Section 6.3.2, we found that a power modulation with 5 counts amplitude is comparable in order of magnitude to the power modulations measured during a laser frequency modulation. However, with this small amplitude, the resulting DWS signal is not well discernible in all four DWS channels. Therefore, we used the data with a modulation amplitude of 50 counts, assuming linearity. This observation alone already indicates that this coupling seems to be weaker in a power modulation than in a frequency modulation.

	frequency modulation	power modulation
	$\Sigma_{A/B} \rightarrow \text{DWS} [\text{rad arb.}^{-1}]$	$\Sigma_{A/B} \rightarrow \text{DWS} [\text{rad arb.}^{-1}]$
η_A	$(1.6 \pm 0.3) \cdot 10^{-4}$	$(2 \pm 0.9) \cdot 10^{-6}$
η_B	$(3.2 \pm 0.9) \cdot 10^{-5}$	$(1.3 \pm 0.6) \cdot 10^{-6}$
ϕ_A	$(1.1 \pm 0.3) \cdot 10^{-4}$	$(1.4 \pm 0.7) \cdot 10^{-6}$
ϕ_B	$(1.8 \pm 0.2) \cdot 10^{-4}$	$(0.6 \pm 0.3) \cdot 10^{-6}$

Table 6.3: The coupling from the measured power to DWS for a frequency modulation experiment and an average over three repeats of a laser power modulation experiment on different days. We note that the coupling measured during the laser power modulation experiment is smaller than that of the frequency modulation.

However, we prefer to compare the numbers given the variation from channel to channel and day to day (cf. Figure 6.17).

The couplings are compared in Table 6.3. The coupling from power to DWS during the laser frequency modulation is copied from Table 6.2 to simplify the comparison. The data for the power modulations represents the average over the result from three power modulations with the same amplitude repeated on different days. The uncertainty is estimated by the standard deviation here. It may be biased due to the small sample size. But for this comparison, a possible bias of the uncertainties is of little importance because the coupling is smaller by an order of magnitude or two in the dedicated power experiments. A possible criticism of this comparison is that the dedicated power modulation has been performed by modulating the output power of a single beam whereas the changes in power most likely occur before the OB (cf. Figure 6.26) and affect both beams.

In summary, there are two arguments supporting the power hypothesis and two against it. On the one hand, the idea of the twofold coupling is supported by the matching products in Tables 6.1 and 6.2. In addition, the changes in power may be caused already in the laser unit itself which makes the observed Σ signals more plausible. On the other hand, the coupling coefficients from changes in power to DWS measured during the dedicated power modulation experiments do not match the order of magnitude observed during the frequency modulations, as shown in Table 6.3. In addition, we have found no plausible theoretical explanation for a coupling from power to DWS at all. Consequently, we have to conclude that the given experimental evidence is not sufficient to either confirm or reject this hypothesis.

6.5 Hypothesis: related to angular offsets

In the course of the LPF mission it was found that the DWS noise level increases for larger tilts of the test masses due to an increase of RIN [Wis17]. This led to the hypothesis that a possible noise due to laser frequency noise could be proportional to the TM tilt, too. Thus we need to

check the coupling from laser frequency noise to DWS versus the rotation of the test mass at that time.

6.5.1 Checks of in-flight data

To compare the coupling coefficient over the different experiments and rotations we would like to average the angular offset during one laser frequency modulation. However, some modulations last approximately half an hour or more, so we need to ensure at first that the coupling which we do not understand yet, does not change in this time range. Some modulations of this duration are part of the dedicated path length mismatch experiment with the offset. Figure 6.27 shows for example the η_2 DWS channel during this experiment. We see small drifts in this channel throughout the experiment. For each of the small subsegments, we have computed the coupling coefficient and found that it does not change over time for the majority of data segments and DWS channels. So we can proceed with averaging over each of the modulations per experiment.




Figure 6.27: An example of the measured η_2 value in DWS during the first ALM mismatch experiment including the offset. The timespan of each of the four modulations is split into smaller segments. A dependency of the coupling amplitude on the mean position during these smaller segments could not be found.

The resulting coupling coefficients are exemplary shown for the ϕ_1 DWS channel in Figure 6.28. We have only selected the modulations where the impact of the laser frequency modulation was discernible by eye in this channel. This was not the case for experiment 2 and 3³, so there is no data from those. We note that the coupling coefficients are fairly comparable in experiments 4-7 but much higher in experiment 1. A single modulation of experiment 1 has a larger uncertainty which is in agreement to the smaller coherence in this modulation compared to the others of the same experiment. However, the data from experiments 1 and 7 has to be considered with care here due to the probably low contrast in the case of grabbed test masses. However, there is

³compare the list of experiments in Table 2.1




Figure 6.28: The coupling from Ψ_F to DWS ϕ_1 for each of the laser frequency modulations depending on the mean ϕ_1 position during the modulation. Each experiment contains several applied modulations and each of them (where the modulation was discernible) produced one of these dots. Note that experiment 1 and 7 have been performed with grabbed test masses which have generally lower contrast.




Figure 6.29: A zoom into the data of Figure 6.28 for the experiments with free-floating test masses only. Note that we take the mean of the DWS scaled to test mass tilt here. Thus, we reach offsets that are comparable to those of the DWS step experiment. However, a clear tilt dependency is not discernible.

no contrast data available for these experiments. Figure 6.29 shows a zoom into the data from

experiments with free-floating test masses only. A simple linear angular offset dependency is not discernible. The little number of data points at only three angular positions does not allow to identify more sophisticated dependencies such a sinusoidal behaviour for example. In fact these are not expected because the increase in the DWS noise level attributed to RIN is also linear [Wis17, p.52].

6.5.2 Checks in laboratory

As we only have a very restricted range of rotations at which the laser was frequency modulated in-flight, we have performed a dedicated experiment in the laboratory. TM1 has been rotated to three different angular offset positions, see Table 6.4, and the laser has been frequency modulated at 8 Hz for 5 minutes with an amplitude of 4.8 mrad rms (corresponding to 1500 counts) at each of them. Note that in this experiment the OPD was active and thus the modulation amplitude measured in Ψ_F is smaller than in other experiments with the same modulation amplitude in counts but with the OPD deactivated. In all three rotations, the contrast on both diodes was kept above 79%. To estimate the coupling the Fourier transform at the modulation frequency has been calculated using a Blackman-Harris window for each of the DWS channels and Ψ_F . This way we estimated the two amplitudes and then divided the amplitude in DWS by the amplitude in Ψ_F .




Figure 6.30: The estimated coupling from a laser frequency modulation to the different DWS channels for three different rotations of TM1, indicated via the mean angular offset in the respective degree of freedom.

The resulting coupling coefficients are shown in Figure 6.30. On the x-axis we show the mean rotation of TM1 during a laser frequency modulation as measured by η_{1A} , ϕ_{1A} , η_{1B} and ϕ_{1B} , respectively. On the y-axis the coupling from Ψ_F to each of the DWS measurements is shown. We note that the angular offset along η_{1A} has not changed much but still the coupling was different for the three modulations. Thus, this difference is an indicator of the fluctuations. We note that the change in coupling to ϕ_{1A} and ϕ_{1B} over the different rotations is comparable

	pos 1	pos 2	pos 3
DWS η_{1A} [urad]	-1.1	-0.6	-0.4
DWS ϕ_{1A} [urad]	8.6	4.4	-0.4
DWS η_{1B} [urad]	-4.1	-6.1	-8.6
DWS ϕ_{1B} [urad]	-6.8	-3.1	1.2
DC $^{\eta_{1A}}$	$-1.6 \cdot 10^{-3}$	$-1.6 \cdot 10^{-3}$	$-1.6 \cdot 10^{-3}$
DC $^{\phi_{1A}}$	$-5.5 \cdot 10^{-4}$	$-5.49 \cdot 10^{-4}$	$-5.49 \cdot 10^{-4}$
DC $^{\eta_{1B}}$	$-1 \cdot 10^{-3}$	$-1 \cdot 10^{-3}$	$-9.8 \cdot 10^{-4}$
DC $^{\phi_{1B}}$	$-1 \cdot 10^{-4}$	$8.9 \cdot 10^{-5}$	$6.8 \cdot 10^{-5}$

Table 6.4: Table of the three different angular offsets studied in the laboratory experiment to test for a dependency of the coupling on the rotation of TM1.

with the fluctuation as measured by η_{1A} . In η_{1B} we observe no change in the coupling with changing offset. From this experiment we conclude that the angular offset does not seem to be the dominant cause for the observation of modulations in DWS during the laser frequency modulations.

6.6 Hypothesis: related to cross-coupling in the phasemeter

6.6.1 Possible signal leakage from longitudinal to DWS

To investigate a possible signal leakage from a longitudinal to a DWS channel, we use the data from the second path length mismatch experiment, experiment 6. This experiment contains the modulation with the highest commanded modulation amplitude of 0.5 rad. This results in a coupling on the order of 10^{-4} $\mu\text{rad pm}^{-1}$ from o1 to ϕ_1 and from o12 to ϕ_2 . This coupling coefficient is calculated using ϕ_1 and ϕ_2 in radians corresponding to TM orientations as of the LPF telemetry. From a laser frequency modulation experiment in the laboratory with the OPD active we find the coupling from o1 to $\phi_{1A/B}$ is on the order of 10^{-5} $\mu\text{rad pm}^{-1}$, so approximately one order of magnitude smaller. The comparison of these numbers is difficult because we have to account for possible differences in the two phasemeters used. However, these two examples show an order of magnitude which can be considered as plausible for phasemeter cross talk.

This hypothesis has been cross-checked with the in-flight data for an experiment in which a guidance signal was injected into the drag-free control loop at a frequency of 0.4 Hz for 1 hour. This is clearly discernible in the o1 interferometer measurement but not in the ϕ_1 or η_1 measurement,

see Figure 6.31. As we are investigating a spurious coupling, this observation is not sufficient to completely rule out this hypothesis but makes it unlikely nonetheless.




Figure 6.31: The amplitude spectra of o_1 , ϕ_1 and η_1 during a modulation in o_1 . While the o_1 modulation is clearly discernible no cross-coupling into the angular measurements can be seen. This makes a general cross-coupling from the longitudinal to the angular measurements unlikely.

6.6.2 Electrical/PD cross-talk

Using the in-flight data it is not possible to distinguish an electrical cross-talk, occurring before the QPD signals are digitised, from cross-talk between the quadrants on the photo diodes. For LPF no requirements for these types of cross-talk have been set and thus no measurements in the test phases on ground have been undertaken. However, such signals have also not been observed during other tests and are hence a possible crosstalk between individual phasemeter channels is believed to be below -90 dB⁴. Even so, this signal suppression level cannot be immediately compared to the level of the coupling from Ψ_F to a DWS measurement because already the Ψ_F telemetry is the result of subtracting the reference interferometer measurement from the frequency interferometer measurement. The frequency interferometer measurement itself is calculated from two PD signals using four quadrants each. Consequently, additional assumptions would have to be made to compare to this noise suppression ratio.

Using the data of the second path length mismatch experiment we can assess the cross-talk hypothesis using the in-flight data. In general, such a type of cross-talk should depend on time or relative phase of two signals. Therefore, we computed a phase $\Delta\phi = \Psi_F - \phi_{o12}$ and the coupling from $\Psi_F \rightarrow$ DWS for each 5 minutes of data of this experiment. However, no change was found within the errors. This makes the electrical or PD cross-coupling hypothesis less likely.

⁴private communication with D. Hoyland

6.6.3 Cross-coupling related to the heterodyne frequency

If the heterodyne frequency was not centred in a bin or if it was drifting, then we could obtain errors from the SBDFT that is performed inside the phasemeter. This should be related in some way to the commanded laser frequency modulation frequency. Figure 6.32 shows the coupling from Ψ_F to ϕ_1 for different modulations where it has been discernible. The data is from several experiments which have been performed on different days during the mission. Given the uncertainty of the data no frequency dependency could be identified. This is also true for the other DWS channels.




Figure 6.32: The estimated coupling coefficient from Ψ_F to ϕ_1 . We obtain one value per modulation where a coupling was discernible.

6.6.4 Investigation of cross-coupling in the laboratory

A hypothesis that seems unlikely but could not be ruled out from the analysis of the in-flight data is cross-coupling. Therefore, this was investigated in the laboratory. The two phasemeters have the SBDFT algorithm in common but differ in sampling frequency and hardware used. More precisely, we have investigated cross-coupling with three experiments:

1. Physically un-plugging cables in the PM;
2. Using different channels of the PM to record the laser frequency modulations;
3. Connect only the ground of the frequency interferometer channel and then record the laser frequency modulations.

The results of systematically unplugging the cables are shown in Figure 6.33. It seems the signal in DWS on the A diode and along ϕ_{1B} becomes smaller if we remove the laser frequency cable. This does not seem to be the case for η_{1B} . Removing the other cables does not seem to have




Figure 6.33: (a) η_{1A} , (b) η_{1B} , (c) ϕ_{1A} , (d) ϕ_{1B} . The AS of DWS signals along η_1 and ϕ_1 during a laser frequency modulation at 8 Hz with a modulation amplitude of 61 MHz recorded on 2019-01-22. The angles have been converted to corresponding TM orientation. The label indicates the cables plugged into the PM.

much impact any more. However, we see a spike on the B diode if we remove all cables except for the 1B cable.

Next, the laser frequency modulation was recorded using different channels of the phasemeter. Figure 6.34 shows these measurements. With the laser frequency modulation signal recorded via different channels of the PM, the measured amplitude in the DWS channels seems to fluctuate. However, as the change is not the same for the η_{1A} and ϕ_{1A} channels, the channel used does not seem to be the major reason. With the given data, we cannot exclude that these fluctuations are due to noise. Nevertheless, we have excluded ground loops in the laboratory. This has been checked by only connecting the ground of the frequency interferometer channel during a laser frequency modulation. With ground loops present, we should have been able to observe the modulation in DWS just the same as with the cable properly connected. This turned out not to be the case.

From the analysis based on the in-flight data we conclude that a cross-talk seems unlikely because a dependency on the absolute phase of the laser frequency modulation frequency could not be identified. The conclusion of the laboratory experiments is not so clear. Even though the fluctuations between the different measurements also have in part a random character, changes in the set-up sometimes also coincide with changes in the coupling. This does not allow us to discard the cross-talk hypothesis as there could be a set-up or components causing it which have yet to be identified.




Figure 6.34: The AS of DWS signals along (a) η_{1A} and (b) ϕ_{1A} during a laser frequency modulation at 8 Hz with a modulation amplitude of 4.8 mrad rms recorded on 2019-01-22. The angles have been converted to corresponding TM orientation. Note that X is a bonus diode which copies the signal of the reference diode. Each curve shows the same η_{1A}/ϕ_{1A} measurement channel while the signal of the frequency interferometer A diode is plugged into a different PM input.

6.7 Further ideas for investigations

- To further investigate the behaviour of the laser it could be useful to place two PDs in front of the Faraday Isolator already or to use a different laser.

- In addition to changing the laser used, one could also record the DWS signal during a laser frequency modulation using a different phasemeter.
- In the experiments performed so far concerning power modulations, only one of the two beams was power-modulated and the other was left unaffected. However, if the laser output power is changed during a frequency modulation this should affect both beams the same. Thus, in a follow-up experiment, the commanding in the laboratory experiment could be changed to power modulate both beams with the same amplitude to phase relation.
- Changes in the wavefront that could possibly lead to DWS signals could also be cross-checked by first blocking a beam and then putting a laser beam profiling instrument in front of the PD during the laser frequency modulation.
- A possible up-conversion of the modulation to frequencies around the heterodyne frequency has been checked and deemed unlikely for the power modulations performed. The same could be done for the frequency modulations.

6.8 Conclusions

In this chapter we investigated possible reasons for the observation of modulations in the DWS channels during the laser frequency modulation experiments on LPF. We saw that interpreting the DWS signals as traces of a cross-coupling, which possibly also affects other OMS channels, would lead to errors of roughly 10 to 15 % in the σ_{12} measurement. In Section 6.1, it was also shown that a similar observation could be made using the AEI laboratory measurements. In the following section, we studied and discarded two possible mechanisms how the path length mismatch between the measurement and the reference beam could also result in a coupling from laser frequency noise to DWS. This was discussed in detail in Section 6.2. A very strong but simple argument we would like to emphasise here is that the fluctuations over time of the observed coupling cannot be explained with this approach. Even though a tilt of one or both beams around the fibre end is unlikely, this hypothesis was investigated using the DC angles in the frequency and reference interferometer but could not be confirmed. In addition, we frequency-modulated a laser in the laboratory and recorded the angular measurements while one beam was blocked. Most of our observed rotations there could be attributed to changes in power. Also a possible misalignment could not be attributed to the observed DWS signals because the DWS signals did not change above the fluctuation level with improved alignment. In Section 6.4, we investigated a possible twofold coupling from laser frequency modulation to changes in beam power and from those to DWS. For this hypothesis, we found evidence both in favour and against it. Thus, a final conclusion on this hypothesis is still outstanding. A simple relation of this coupling to the TM orientation could also not be proven. Possible cross-couplings in the phasemeter have been under study in Section 6.6 but a final conclusion on this hypothesis could not be reached yet.

Unfortunately, we could not identify the cause for the observed signals in DWS. However, a combination of several of the mentioned effects is also possible and further investigation is required to identify the origin of the observations. Especially the use of a different set-up should allow us to judge the importance of this finding.

Up to now, several hypothesis for the observation of spurious sinusoidal signals in the DWS measurements on board of LPF and in the AEI laboratory were investigated. None of these provided a convincing explanation. However, the theory of a twofold coupling via changes in power and the hypothesis of cross-couplings in the phasemeter could also not be fully rejected.

CHAPTER 6. INVESTIGATION OF DWS SIGNALS DURING FREQUENCY
MODULATIONS

Nevertheless, a combination of these effects or other reasons should explain the observations and further investigation is recommended.

Chapter 7

Summary and outlook

7.1 Summary of the thesis

In this thesis, detailed investigations concerning the laser frequency fluctuations and interferometer path length differences on board LISA Pathfinder (LPF) were presented. We started out with a brief review of Gravitational Wave (GW) detection on Earth and moved on to the interesting science which will be made possible by GW observation from space with the Laser Interferometer Space Antenna (LISA) mission. We have seen that such an observatory implies new technological challenges. As a consequence, the technology demonstrator mission LISA Pathfinder (LPF) was designed and operated from the launch on December 3rd 2015 until the final shutdown on July 18th 2017. After a short summary of the key components for LISA, which were tested on LPF, Chapter 1 concluded with a brief summary of the mission preparation, operation and the outstanding results.

In Chapter 2, one of the key LISA Pathfinder components, the Optical Metrology System (OMS) was introduced in more detail. In this introduction, we summarised the optical components, their measurements and the data processing. We also reviewed the coupling of laser frequency noise into the main science measurement which depends on the path length mismatch. In this chapter, the laser frequency stabilisation scheme on LPF was presented. This scheme combined a laser frequency fluctuation measurement by a dedicated interferometer with a nested digital control loop. An overview of all the laser frequency modulation experiments that were performed during the LPF mission concludes this chapter.

The performance of the laser frequency stabilisation control loop was analysed in Chapter 3. We used periodic excitation signals to characterise the transfer function of the loop and its components. Five of these laser frequency control loop characterisation experiments were performed during the mission lifetime. We found that the controller transfer functions, the OLTFs and the actuator parameters were comparable to the results from the test campaign data analysed in the same way and over the course of the mission. Slight changes in the actuator parameters may be related to changes in the temperatures of the different units on board but there may also be other reasons not identified yet. This analysis was further complicated due to the different clocks on board LISA Pathfinder and thus provides one example why a single clock would have been beneficial. In summary, we were able to say that the laser frequency stabilisation worked as expected.

In Chapter 4, we compared the resulting laser frequency noise performance achieved with this stabilisation technique to the free-running laser frequency fluctuations. These two measurements were performed at different times during the mission. The results were found to be very similar over the course of the mission and agreed well to the data from the OSTT ground test campaign. With active stabilisation, the requirement (Equation 2.36) was met with ample margin. As in the ground test campaigns, periods of higher laser frequency noise when the stabilisation was active, were observed. As we had previously convinced ourselves that the laser frequency control loop works as expected, possible reasons related to the laser were investigated. However, the cause of this behaviour was not found with the in-flight data available. This is one reason why laser frequency noise measurements of 10 hours duration are recommended for future test campaigns of laser frequency stabilisation for satellite missions where similar increased noise periods could be relevant.

Chapter 5 dealt with the interferometer path length mismatches on LPF. Two dedicated experiments were performed during the LPF mission. The first of these, the path length mismatch offset experiment, allowed us to associate the measured change in path length mismatch to the commanded offset at a 3σ uncertainty level, both in direction and amplitude. The second dedicated path length mismatch experiment confirmed that the estimated path length mismatch is independent of the laser frequency modulation amplitude. Additionally, the data of the laser frequency control loop characterisation experiments was used for path length mismatch estimation. These experiments revealed a change in path length mismatch in the course of the mission of approximately $20\mu\text{m}$. The analysis of the phase tracking reset commands has helped us to identify a commanded offset as the cause. Thereby, we have provided one example where laser frequency modulation can be used as a tool to measure absolute distances. A detailed investigation of the path length mismatch experiments also revealed short-term variations of the path length difference. Even though those were not yet fully understood they are not associated to real motion of either TM, as this would have been measured by the GRS. Path length changes on the OB itself are also unlikely to cause the observed variations because of their amplitude. All components are bonded onto the OB, hence, a physical motion on the $10\mu\text{m}$ level would imply a catastrophic failure. Moreover, the temporal evolution of the OB temperature does not correlate in time to the observed changes in path length mismatch. In addition, these variations could not be found in the case of the highest amplitude modulation.

During some of the laser frequency modulations on LPF, a modulation in one or more DWS channels was discernible. These observations are the subject of Chapter 6. A comparable effect could be reproduced in the AEI LPF laboratory. These two data sets were used to investigate several hypothesis for this observation but, unfortunately, no convincing explanation was found so far. This unknown origin makes it difficult to estimate the relevance of this finding for other heterodyne interferometer readouts. Further investigation is recommended.

In summary, the planned loop characterisation experiments and laser frequency noise measurements were performed during the LPF operations, as anticipated in [Aud14]. They were analysed successfully and their results agree to the expectations from ground test campaigns. Moreover, two laser frequency modulation experiments optimised for path length mismatch estimation with consistency checks included, have been designed, executed and analysed.

The work presented in this thesis contributes to the detailed understanding of the unprecedented performance of the LPF OMS, as described in Section 7.2. The low-noise heterodyne interferometer measurements and the detailed analysis also resulted in new observations, as discussed in Section 5.4 and Chapter 6. However interesting they are, let us emphasise here that these are very detailed questions and the performance of LPF OMS and LPF in general has been

excellent ([A⁺16c][A⁺18b]) and significantly better than the requirements or even the best-case expectations [A⁺12a].

7.2 Contribution of laser frequency noise to the total LPF OMS noise

In this section, we describe the contribution of the laser frequency noise to the total sensing noise in the o12 measurement using an example noise measurement taken on June 1st 2016 (DOY 153). To compute the contribution of laser frequency noise to the o12 relative distance measurement, the laser frequency noise is converted to the corresponding path length noise using Equation 2.29 which is reprinted here, for clarity, in the frequency domain,

$$S_{F \rightarrow o12}^{1/2} = \frac{\lambda}{4\pi \cos(\alpha)} \frac{\Delta s_{o12}}{\Delta L} S_{\Psi_F}^{1/2}. \quad (7.1)$$

For Δs_{o12} , we use the value from the same day from Table 5.3. This noise contribution is shown in blue in Figure 7.1. In this figure, the ASD of the o12 measurement is shown in magenta. We have also subtracted the time series of the laser frequency noise contribution from the o12 measurement time series. The amplitude spectral density of this result is shown for comparison in black. From Figure 7.1, we find that laser frequency noise has a very small impact in the frequency range from approximately 0.2 to 0.9 Hz.




Figure 7.1: The contribution of laser frequency noise, converted to the corresponding path length noise (blue trace), compared to the o12 measurement (magenta trace). The ASD of the subtraction of the laser frequency noise from the o12 measurement is shown in black. Laser frequency noise is a small contribution to the total noise in o12 in the frequency range from approximately 0.2 to 0.9 Hz.

However, we have seen in Chapter 4 and most clearly in Figure 4.3 that the laser frequency noise level changes in the course of the mission. For periods of increased laser frequency noise, this

noise source becomes dominant in o12 in the frequency range from approximately 0.3 to 1 Hz. This statement is supported by two pieces of evidence.

First, we observe in an example data segment that increased periods of noise in o12 in the band from 0.2 to 1.9 Hz are no longer discernible if the laser frequency noise contribution is subtracted. This check could not be repeated during other noise measurement periods, as it required the laser frequency noise to be recorded at sampling frequencies above 1 Hz while the laser frequency noise level was increased. Figure 7.2 shows this observation. The Ψ_F measurement is shown in blue, the o12 measurement in red and the o12 measurement with the frequency noise contribution subtracted in yellow. For this subtraction, the laser frequency fluctuations have been converted to a path length noise using a coupling coefficient k with an assumed path length mismatch estimate of $350\ \mu\text{m}$. All signals are bandpassed in the frequency range from 0.2 to 1.9 Hz. As anticipated, periods of increased noise in o12 occur at the same time as those in the frequency interferometer but are no longer discernible after the subtraction of the laser frequency noise contribution.




Figure 7.2: Measurements of Ψ_F , o12 and o12 with the frequency noise contribution subtracted in the frequency band from 0.2 to 1.9 Hz. We find that the periods of increased noise in o12 are no longer discernible after the subtraction of the frequency noise contribution.

The second piece of evidence is to look at the contribution from laser frequency noise to the total noise in o12 at a period of increased laser frequency noise. The data processing is identical to that of Figure 7.1 with the exception of the path length mismatch used. Here, we use data recorded on January 21st 2017 and thus use the path length mismatch of DOY 22 from Table 5.3. The resulting contribution of laser frequency noise to the noise in o12 is shown in Figure 7.3. For this period, laser frequency noise is the dominating noise source in o12 in the frequency range from approximately 0.3 to 1 Hz. Both of these arguments only consider certain data periods but so far, we have found no reason to assume a different behaviour at other times.

Before launch, laser frequency noise was a significantly smaller contribution to the total noise in the o12 measurement [Aud14]. With the excellent performance of the LPF OMS, the noise due to comparable laser frequency fluctuations as on ground has become a discernible and at certain




Figure 7.3: The contribution of laser frequency noise, converted to the corresponding path length noise (blue trace), compared to the o12 measurement (magenta trace). The ASD of the subtraction of the laser frequency noise from the o12 measurement is shown in black. Here, data from a period of increased laser frequency noise in January 2017 is shown. Laser frequency noise is the dominating noise source in o12 in the frequency range from approximately 0.3 to 1 Hz.

times a dominating noise contribution, even though this noise source itself is significantly below the required level (see Chapter 4).

In addition, the coupling from laser frequency noise to the o12 measurement is small due to the path length mismatches achieved (see Chapter 5). They were a result of minimising the path length mismatches on the OB during the bonding process and the integration of the OB into the LTP with respect to the expected TM positions which is a complex chain of alignment itself. However, the final construction tolerances of the flight hardware could only be estimated in the mission design phase. That is why a conservative laser frequency stabilisation requirement and the related control loop is important.

Figures 7.1 and 7.3 show how the laser frequency noise analysis provides a small contribution to the detailed understanding of the excellent performance of the LPF OMS. This is not only making the best use of this unique data but also helpful for future interferometer development. Perhaps this may also contribute to the understanding of why the in-flight performance is much better than in the AEI ground set-up.

7.3 Discussion: Limits of the laser frequency stabilisation scheme

There are three possible aspects which can limit the performance of the LPF laser frequency stabilisation scheme. One of them is that the XF interferometer, which is used as a laser frequency sensor for the stabilisation, depends on the intentional path length mismatch. If this mismatch

was subject to a change, the control loop would follow this change. However, with the ultra-low expansion material used and given the frequency range of interest, thermally induced changes of the path length mismatch are deemed unlikely. In other words, assuming a perfect stabilisation, one is always limited by the frequency stability of the reference.

A second aspect is the interferometer and phasemeter readout noise which may become limiting at some small level of laser frequency fluctuations. For the LPF case, let us roughly estimate this level. The o12 ASD as shown in Figure 7.3 is dominated by readout noise above approximately 0.2 Hz. In the lower part of the frequency range shown in this figure, o12 is dominated by cross-coupling from satellite jitter, as also shown in [A+16c]. The readout ASD of approximately $30 \text{ fm}/\sqrt{\text{Hz}}$ at frequencies above 0.2 Hz corresponds to a phase noise of approximately $0.35 \text{ } \mu\text{rad}/\sqrt{\text{Hz}}$ (see Equation 2.7). Assuming this phase noise level is the same for all channels and multiplying this result with a factor of $\sqrt{2}$ to account for the noise in the reference interferometer, this corresponds to a sensing noise limit of approximately $61 \text{ Hz}/\sqrt{\text{Hz}}$ (see Equation 2.15). However, the noise level above 20 Hz, as shown in Figure 4.12, is more than an order of magnitude above this limit and also shows two levels. This renders a limitation due to sensing noise unlikely. Similarly, the estimated quantisation noise is smaller than the fluctuations measured. To conclude, even though the limiting factor of the high frequency noise could not be identified yet, sensing noise seems unlikely. If an independent out-of-loop measurement of the laser frequency fluctuations on LPF was available, it could be used to double check whether sensor noise is limiting. If this was the case, the laser frequency fluctuations could also be amplified by the choice of the path length mismatch in the frequency interferometer.

The third aspect of the discussion of the limits of this laser frequency stabilisation scheme is the control loop. The phase margin of 77° and the estimated gain margin of 8.7 in magnitude (see Chapter 3) are comfortably above the recommended stability limits of $30\text{-}45^\circ$ for the phase margin and 2 for the gain margin in magnitude [Aud14][Pac14]. Thus, an increase of the loop gain would be possible if required. A unity gain frequency of 0.8 Hz has been measured in-flight. From the control loop model, we find that an increase of the fast actuator gain would also increase the unity gain frequency. The corresponding OLTF for an increase of a factor 4.3 is compared to the OLTF of experiment 3 (see Table 2.1) in Figure 7.4. With the increased fast actuator gain, we find a unity gain frequency of 3.8 Hz, along with a gain margin of 2.02 in magnitude and a phase margin of 45.7° . So this increase in fast frequency actuator gain reduces the stability margins but the loop should still be stable. For loop stability reasons a further increase of the fast actuator gain is not recommended. The delays in the MIL-Bus and DMU software latencies amount to approximately 20 ms ([Ker11]) which imposes an additional limit to the bandwidth at higher frequencies.




Figure 7.4: The OLTF model using the parameters of experiment 3 (blue) and with an increased fast frequency actuator gain of a factor 4.3 (red). We find the unity gain frequency could be increased up to 3.8Hz with a gain margin of 2.02 in magnitude and a phase margin of 45.7° .

7.4 Outlook: Applications of the laser frequency stabilisation scheme

On ground, a similar laser frequency stabilisation scheme using balanced homodyne DC readout has been designed and tested [GIM⁺17]. The authors deem this set-up beneficial for homodyne interferometric measurements and point out that this readout scheme comes along with simpler readout electronics.

As LISA Pathfinder (LPF) is the technology demonstrator mission for LISA, it is obvious to assume that also the laser and its frequency stabilisation scheme, which have been proven to fulfil the requirements and to operate reliably over the mission, will be used in LISA. However, the significance of laser frequency noise is different in LPF and in LISA.

In LPF, this noise source is a small contribution to the sensing noise of a key subsystem, whereas in the LISA data streams, laser frequency noise is the dominating noise source. We refer to [Ott16] for a detailed review on this subject. This can be easily understood, since we have seen in many places in this thesis that the coupling of laser frequency noise into the interferometer measurement is proportional to the path length mismatch. Now, in LISA, as explained in Chapter 1, the relative distance in between two TMs on two satellites 2.5 million km apart from each other is measured by heterodyne interferometry. Each of these measurements is split into three parts, known under the term split-interferometry configuration, see [Ott16]. In one of these measurements the light from a distant satellite is brought to interference with the light on another satellite. Accordingly, in this measurement, there is a path length mismatch of the order of the LISA arm length which largely amplifies laser frequency fluctuations. This implies that an effective laser frequency stabilisation is crucial. In addition, laser frequency noise has to be suppressed by a post-processing scheme, called Time-Delay Interferometry (TDI) [Ott16].

Due to the large distance in between the satellites, the laser light originating on one satellite cannot be simply reflected of the TM on board another satellite and then brought to interference again. Thus, according to the current design as described in [Hew19], there are two stabilised laser systems per satellite. One of the six lasers is operating as the so-called master laser and the other lasers are offset phase-locked to this one for frequency stabilisation. All lasers can be the master laser. The current master laser is then planned to be stabilised by the use of a cavity. Currently, the stabilised laser frequencies are expected to reach fluctuation levels below $30 \text{ Hz}/\sqrt{\text{Hz}}$ in the LISA bandwidth [Hew19, p. 15].

At earlier phases of the LISA mission concept, it was also discussed to frequency pre-stabilise the LISA lasers similarly to the LPF laser and then use the LISA arms to stabilise the laser frequency. Two advantages have been pointed out [SHD10]. One is that a Mach-Zehnder interferometer frequency stabilisation has a wide operating range which also facilitates lock acquisition and the combination with the stabilisation to the LISA arms. The second advantage is that no cavity inside a dedicated vacuum chamber would be needed which saves space and a number of components needed to implement the stabilisation. However, additional components on the OB would make its design, construction and testing even more complicated. Consequently, the laser system and the OB could also not be tested independently.

To stabilise the LISA master lasers to the required level of $\delta f = 30 \text{ Hz}/\sqrt{\text{Hz}}$ using only the frequency stabilisation approach as on LPF is difficult for the following reason: To measure this level of laser frequency fluctuations with a phasemeter readout noise of $\tilde{\varphi} \approx 6 \mu\text{rad}/\sqrt{\text{Hz}}$, as expected for LISA (see [Sch18]), the required amplification via the intentional path length mismatch in a frequency interferometer on LISA follows from Equation 2.15 to be

$$\Delta L = \frac{c}{2\pi} \frac{\Delta\tilde{\varphi}}{\delta f} \approx 10 \text{ m}. \quad (7.2)$$

This is significantly larger than on LPF and impractical to accommodate on a complex OB. Moving the path length difference into the fibres which lead to the OB, however, also implies that this path length difference needs to be compensated in the other interferometers on the OB and thus does not solve the issue but simply moves it to the other interferometers.

On the contrary, the cavity as a frequency reference does fulfil the LISA requirements as demonstrated in laboratory tests and therefore was at a technology readiness level of 8 in 2017 [AS⁺17]. In the meantime, this laser frequency stabilisation scheme has proven successful operation on the laser ranging instrument on the GRACE Follow-on mission [A⁺19b].

Bibliography

- [A⁺12a] F. Antonucci et al. The LISA Pathfinder mission. *Classical and Quantum Gravity*, 29(12):124014, 2012.
- [A⁺12b] H. Audley et al. Report on the OMS FM Phase I test campaign - S2-AEI-TN-3068. Technical report, AEI, 2012.
- [A⁺15] M. Armano et al. Free-flight experiments in LISA Pathfinder. *Journal of Physics: Conference Series*, 610:012006, 2015.
- [A⁺16a] B. P. Abbott et al. GW150914: The Advanced LIGO Detectors in the Era of First Discoveries. *Phys. Rev. Lett.*, 116:131103, 2016.
- [A⁺16b] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116:061102, 2016.
- [A⁺16c] M. Armano et al. Sub-Femto-*g* Free Fall for Space-Based Gravitational Wave Observatories: LISA Pathfinder Results. *Phys. Rev. Lett.*, 116:231101, 2016.
- [A⁺17a] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119:161101, 2017.
- [A⁺17b] M. Armano et al. Capacitive sensing of test mass motion with nanometer precision over millimeter-wide sensing gaps for space-borne gravitational reference sensors. *Phys. Rev. D*, 96:062004, 2017.
- [A⁺18a] G. Anderson et al. Experimental results from the ST7 mission on LISA Pathfinder. *Phys. Rev. D*, 98:102005, 2018.
- [A⁺18b] M. Armano et al. Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20 μ Hz. *Phys. Rev. Lett.*, 120(6), 2018.
- [A⁺18c] M. Armano et al. Calibrating the system dynamics of LISA Pathfinder. *Phys. Rev. D*, 97:122002, 2018.
- [A⁺19a] B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9:031040, 2019.
- [A⁺19b] K. Abich et al. In-Orbit Performance of the GRACE Follow-on Laser Ranging Interferometer. *Phys. Rev. Lett.*, 123:031101, 2019.
- [A⁺19c] M. Armano et al. LISA Pathfinder. *arXiv e-prints*, page arXiv:1903.08924, 2019.

BIBLIOGRAPHY

- [A⁺19d] M. Armano et al. LISA Pathfinder micronewton cold gas thrusters: In-flight characterization. *Phys. Rev. D*, 99:122003, 2019.
- [A⁺19e] M. Armano et al. LISA Pathfinder Performance Confirmed in an Open-Loop Configuration: Results from the Free-Fall Actuation Mode. *Phys. Rev. Lett.*, 123:111101, 2019.
- [A⁺19f] M. Armano et al. Temperature stability in the sub-milliHertz band with LISA Pathfinder. *Monthly Notices of the Royal Astronomical Society*, 486(3):3368–3379, 2019.
- [A⁺20a] B. P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$, 2020.
- [A⁺20b] R. Abbott et al. GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses. *arXiv e-prints*, page 2004.08342, 2020.
- [AS⁺17] P. Amaro-Seoane et al. Laser Interferometer Space Antenna. *arXiv e-prints*, page arXiv:1702.00786, 2017.
- [Aud14] H. Audley. *Preparing for LISA Pathfinder operations : characterisation of the optical metrology system*. PhD thesis, Leibniz Universität Hannover, 2014.
- [Bod07] M. Bode. OMS Requirement Specification - S2-ASD-RS3010. Technical report, EADS/Astrium, 2007.
- [BP80] J. S. Bendat and A. G. Piersol. *Engineering Applications of Correlation and Spectral Analysis*. Wiley, 1980.
- [Bra09] N. Brandt. Test Mass Actuation Algorithm for DFACS - S2-ASD-TN-2011. Technical report, Astrium, 2009.
- [CM68] H. R. Carleton and W. T. Maloney. A Balanced Optical Heterodyne Detector. *Appl. Opt.*, 7(6):1241–1243, 1968.
- [D⁺03] R. Dolesi et al. Gravitational sensor for LISA and its technology demonstration mission. *Classical and Quantum Gravity*, 20(10):S99–S108, 2003.
- [D⁺18] K. Danzmann et al. LPF final report for the German contribution to the nominal mission. Technical report, AEI, 2018.
- [ESA19] ESA. LISA Mission Summary, Sep 2019. <https://sci.esa.int/web/lisa/-/61367-mission-summary>, accessed on 2020-02-18.
- [F. 11] F. Antonucci et al. LISA Pathfinder: mission and status. *Classical and Quantum Gravity*, 28(9):094001, 2011.
- [Fit10] E. D. Fitzsimons. *Techniques for precision interferometry in space*. PhD thesis, University of Glasgow, 2010.
- [G. 11] G. Hechenblaikner et al. Digital Laser Frequency Control and Phase-Stabilization Loops in a High Precision Space-Borne Metrology System. *IEEE Journal of Quantum Electronics*, 47:651–660, 2011.
- [GFV10] A. Grynagier, W. Fichter, and S. Vitale. Parabolic Drag-Free Flight, Actuation with Kicks, Spectral Analysis with Gaps. *Space Science Reviews*, 151(1):183–196, 2010.

- [GIM⁺17] O. Gerberding, K.-S. Isleif, M. Mehmet, K. Danzmann, and G. Heinzel. Laser-Frequency Stabilization via a Quasimonolithic Mach-Zehnder Interferometer with Arms of Unequal Length and Balanced dc Readout. *Phys. Rev. Applied*, 7(2):024027, 2017.
- [Giu05] D. Giulini. Einsteins Kunstwerk Die Allgemeine Relativitätstheorie – aus mittlerer Entfernung betrachtet. *Physik Journal*, page 27–33, 2005.
- [Giu17] R. Giusteri. *Parabolic flights in pico-g for space-based gravitational wave observatories: the free-fall experiment on LISA Pathfinder*. PhD thesis, University of Trento, 2017.
- [Gry09] A. Grynagier. The Drift Mode for LISA Pathfinder. Technical report, Institute of Flight Mechanics and Flight Control, Universität Stuttgart, 2009.
- [H⁺03] G. Heinzel et al. Interferometry for the LISA technology package (LTP) aboard SMART-2. *Classical and Quantum Gravity*, 20(10):S153–S161, 2003.
- [H⁺04] G. Heinzel et al. The LTP interferometer and phasemeter. *Classical and Quantum Gravity*, 21(5):S581–S587, 2004.
- [H⁺05] G. Heinzel et al. Successful testing of the LISA Technology Package (LTP) interferometer engineering model. *Classical and Quantum Gravity*, 22(10):S149–S154, 2005.
- [H⁺09] M. Hewitson et al. Data analysis for the LISA Technology Package. *Classical and Quantum Gravity*, 26(9):094003, 2009.
- [Har78] F. J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1):51–83, 1978.
- [HBL10] D. W. Hogg, J. Bovy, and D. Lang. Data analysis recipes: Fitting a model to data. *arXiv e-prints*, page arXiv:1008.4686, 2010.
- [HCD⁺02] M. Hueller, A. Cavalleri, R. Dolesi, S. Vitale, and W. J. Weber. Torsion pendulum facility for ground testing of gravitational sensors for LISA. *Classical and Quantum Gravity*, 19(7):1757, 2002.
- [Hei02] G. Heinzel. SMART-2 interferometer - S2-AEI-TN-3010. Technical report, AEI, 2002.
- [Hen09] M. Hendry. An Introduction to General Relativity, Gravitational Waves and Detection Principles, 2009. https://www.ego-gw.it/public/events/vesf/2009/Presentations/Hendry_2.pdf, accessed on 2020-02-18.
- [Heu04] M. Heurs. *Gravitational waves in a new light: Novel stabilisation schemes for solid-state lasers*. PhD thesis, Leibniz Universität Hannover, 2004.
- [Hew19] M. Hewitson. LISA-Instrument Design Description - LISA-LCST-INST-DD-001. Technical report, LISA Consortium, 2019.
- [HFM18] D. W. Hogg and D. Foreman-Mackey. Data Analysis Recipes: Using Markov Chain Monte Carlo. *The Astrophysical Journal Supplement Series*, 236(1):11, 2018.

BIBLIOGRAPHY

- [HRS02] G. Heinzel, A. Rüdiger, and R. Schilling. Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows. *Max Plank Publication Repository*, 2002.
- [K⁺12] E. Kochkina et al. Simulating and Optimizing Laser Interferometers. *Astronomical Society of the Pacific: Conference Series*, 467:291, 2012.
- [K⁺18] T. Kupfer et al. LISA verification binaries with updated distances from Gaia Data Release 2. *Monthly Notices of the Royal Astronomical Society*, 480(1):302–309, 2018.
- [Kar14] N. Karnesis. *Bayesian Data Analysis for LISA Pathfinder*. PhD thesis, Universitat Autònoma de Barcelona, 2014.
- [Kau21] B. Kaune. *In-orbit Stability Analysis of the LISA Pathfinder Optical Metrology: Photoreceivers and Polarisation*. PhD thesis, Leibniz Universität Hannover, 2021.
- [Ker11] M. Kersten. OMS-control loop stability analysis and filter design - S2-ASD-TN-3107. Technical report, Astrium, 2011.
- [LC09] T. B. Littenberg and N. J. Cornish. Bayesian approach to the detection problem in gravitational wave astronomy. *Phys. Rev. D*, 80:063007, 2009.
- [Lc17] LISA Pathfinder collaboration. LISA Pathfinder: First steps to observing gravitational waves from space. *Journal of Physics: Conference Series*, 840:012001, 2017.
- [LRWW98] J. Lagarias, J. Reeds, M. Wright, and P. Wright. Convergence Properties of the Nelder–Mead Simplex Method in Low Dimensions. *SIAM Journal on Optimization*, 9:112–147, 1998.
- [M. 06] M. Tröbs. LTP reference laser unit engineering model characterization - S2-AEI-TN-3032. Technical report, AEI, 2006.
- [M. 15] M. Armano et al. The LISA Pathfinder Mission. *Journal of Physics: Conference Series*, 610(1):012005, 2015.
- [Mar07] A. F. García Marín. *Minimisation of optical pathlength noise for the detection of gravitational waves with the spaceborne laser interferometer LISA and LISA Pathfinder*. PhD thesis, Leibniz Universität Hannover, 2007.
- [MH13] P. McNamara and M. Hewitson. OSTT data analysis report - S2-EST-TN-3061. Technical report, ESA, 2013.
- [N. 15] N. Brandt, T. Ziegler, C. Trenkel. Experiment Performance Budget (M3 optical, Cold Gas). Technical report, Airbus, 2015.
- [N⁺13] M. Nofrarias et al. Subtraction of temperature induced phase noise in the LISA frequency band. *Phys. Rev. D*, 87:102003, 2013.
- [Ott16] M. Otto. *Time-delay interferometry simulations for the laser interferometer space antenna*. PhD thesis, Leibniz Universität Hannover, 2016.
- [P⁺17] S. Paczkowski et al. Laser Frequency Noise Stabilisation and Interferometer Path Length Differences on LISA Pathfinder. *Journal of Physics: Conference Series*, 840(1):012004, 2017.

- [Pac14] S. Paczkowski. Modelling the S/C and TM Motion for LISA. Master's thesis, Leibniz Universität Hannover, 2014.
- [Pre92] W. Press. *Numerical recipes in C : the art of scientific computing*. Cambridge University Press, 1992.
- [R⁺05] D. Robertson et al. LTP interferometer - noise sources and performance. *Classical and Quantum Gravity*, 22(10):S155–S163, 2005.
- [R⁺13] D. Robertson et al. Construction and testing of the optical bench for LISA Pathfinder. *Classical and Quantum Gravity*, 30(8):085006, 2013.
- [R⁺18] G. Russano et al. Measuring fN force variations in the presence of constant nN forces: a torsion pendulum ground test of the LISA Pathfinder free-fall mode. *Classical and Quantum Gravity*, 35(3):035017, 2018.
- [RF12] D. Robertson and E. Fitzsimons. Optical characterisation of 3OB - S2-UGL-TN-3039. Technical report, University of Glasgow, 2012.
- [Rob13] D. Robertson. 3OB As Built OptoCAD Model - S2-UGL-TN-3045. Technical report, University of Glasgow, 2013.
- [RS13] L. M. Rios and N. V. Sahinidis. Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56(3):1247–1293, 2013.
- [RYS11] R. Rummel, W. Yi, and C. Stummer. GOCE gravitational gradiometry. *Journal of Geodesy*, 85(11):777, 2011.
- [Sch12] A. Schleicher. DFACS-science mode 1 design and analysis - S2-ASD-TN-2003. Technical report, Astrium, 2012.
- [Sch14] A. Schleicher. DFACS-User Manual - S2-ASD-MA-2004. Technical report, Astrium, 2014.
- [Sch18] T. Schwarze. *Phase extraction for laser interferometry in space : phase readout schemes and optical testing*. PhD thesis, Leibniz Universität Hannover, 2018.
- [Ses17] A. Sesana. Multi-band gravitational wave astronomy: science with joint space- and ground-based observations of black hole binaries. *arXiv e-prints*, page arXiv:1702.04356, 2017.
- [SHD10] B. Sheard, G. Heinzel, and K. Danzmann. LISA long-arm interferometry: an alternative frequency pre-stabilization system. *Classical and Quantum Gravity*, 27(8):084011, 2010.
- [SS09] B. S. Sathyaprakash and Bernard F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Reviews in Relativity*, 12(1):2, 2009.
- [SW11] S. Schulze-Walewski. IS FEE Operations Manual. Technical Report S2-CSA-MA-3002, RUAG, 2011. Issue 2.8.
- [Tea11] LISA Science Team. LISA - Unveiling a hidden Universe. Technical report, 2011. Also available at <http://sci.esa.int/science-e/www/object/doc.cfm?fobjectid=48363>.

- [TFM79] J. H. Taylor, L. A. Fowler, and P. M. McCulloch. Measurements of general relativistic effects in the binary pulsar PSR1913 + 16. *Nature*, 277(5696):437–440, 1979.
- [TH06] M. Tröbs and G. Heinzel. Improved spectrum estimation from digitized time series on a logarithmic frequency axis. *Measurement*, 39:120–129, 2006.
- [U. 09] U. Denskat. Phasemeter Processing & Laser Control. Technical report, Airbus, 2009. Issue 4.0.
- [V⁺14] S. Vitale et al. Data series subtraction with unknown and unmodeled background noise. *Physical Review D*, 90(4):042003, 2014.
- [W⁺17] A. Wittchen et al. Coupling of relative intensity noise and pathlength noise to the length measurement in the optical metrology system of LISA Pathfinder. *Journal of Physics: Conference Series*, 840(1):012003, 2017.
- [Wan10] G. Wanner. *Complex optical systems in space : numerical modelling of the heterodyne interferometry of LISA Pathfinder and LISA*. PhD thesis, Leibniz Universität Hannover, 2010.
- [Wan19] G. Wanner. Space-based gravitational wave detection and how LISA Pathfinder successfully paved the way. *Nature Physics*, 15(3):200–202, 2019.
- [Weg14] R. Wegener. LTP User Manual. Technical report, Astrium, 2014.
- [Wis17] L. Wissel. In-orbit performance and behaviour of the LISA Pathfinder Optical Metrology System. Master’s thesis, Leibniz Universität Hannover, 2017.
- [Wit14] A. Wittchen. Noise Investigation on the LISA-Pathfinder Optical Bench Ground Setup. Master’s thesis, Leibniz Universität Hannover, 2014.

Appendix A

Details on parameter estimation

In this chapter, the parameter estimation procedures applied to determine the laser frequency actuator gains and delays presented in Section 3.3.2 and the free-flight quasi-parabola parameters of Appendix C will be briefly introduced.

A.1 A brief comment on the Nelder-Mead simplex method

The search for the optimal actuator parameters is started using the Nelder-Mead simplex method. It is implemented in MATLABTM following [LRWW98]. As explained in [RS13], it is an optimisation method which does not require derivatives of the objective function. As it is not possible with this algorithm to refine the search area arbitrarily, it is considered by [RS13] as a local method. The key idea is to construct a simplex in parameter space. The simplex ‘is the geometrical figure consisting, in N dimensions, of $N+1$ points (or vertices) and all their interconnecting line segments, polygonal faces, etc. In two dimensions, a simplex is a triangle’ [Pre92]. Then the function to be minimised is evaluated at each of the corner points. The worst of these points, that is to say the one where the function value is the largest, is to be replaced in the next step of the minimisation. Therefore, candidate alternative points can be calculated or all points except for the best point have to be recalculated and the whole simplex is shrunk along its axes [Pre92]. For details and a geometrical representation, we refer to [Pre92].

In this work the algorithm was started from a reasonably close starting point and was applied for three times iteratively because the simplex is quite large at the beginning of the optimisation. This way, the probability to get stuck in a local minimum can at least be reduced. The simplex minimisation result was then also used to estimate a starting covariance matrix, via the inverse of the derivative of the cost function with respect to its parameters at the simplex results, for the MCMC sampling.

A.2 A brief summary of MCMC uncertainty estimation

A.2.1 From model to posterior probability density function

To briefly summarise the Markov Chain Monte-Carlo (MCMC) formalism to estimate the uncertainty of the actuator parameters, let us start with mentioning explicitly the assumption which is often used to fit the data. We assume that our measured data \vec{y} can be explained by a model \mathcal{M} depending on parameters $\vec{\theta}$ and that there is additional noise \vec{n} [Kar14]. This means

$$\vec{y} = \mathcal{M}(\vec{\theta}) + \vec{n}. \quad (\text{A.1})$$

Here, the noise is assumed to be Gaussian, ‘uncorrelated, stationary and with zero-mean’ [Kar14]. Then we define the function to be minimised by tuning the parameters $\vec{\theta}$ as

$$\sum_{i=1}^N \left(\frac{(y_i - \mathcal{M}_i)^2}{\sigma^2} \right). \quad (\text{A.2})$$

This is the difference between data and model should be as small as possible considering the variance of the data σ^2 . This variance is the same for each measured data point i . With our assumptions and considering that the discrepancy of the measured data to the model is only due to the presence of the noise and that this contribution is independent for each data point, the probability to measure exactly a given data set is the product of the probabilities of measuring each data point. Consequently, we can speak of maximum likelihood parameter estimation if parameters are estimated using Equation A.2 [Pre92]. Under these circumstances [HBL10] show that an approach similar to Equation A.2 is indeed a good choice of the function to be minimised. In other words, we consider the parameters $\vec{\theta}$ which minimise Equation A.2 most likely the true values. A similar quantity as the one expressed in A.2, just with the variance σ depending on the data point i , leading to a term σ_i in Equation A.2, is often called χ^2 . As can be seen from A.2, a data point dependent uncertainty σ_i attributes a higher impact to data points with small uncertainties whereas Equation A.2 assigns the same influence on the result to all data points, that is they have equal weights. Other choices for a maximum likelihood estimator are also possible and for example recommended if known outliers are present [Pre92].

In the case of the frequency actuator parameter estimation, an estimate for σ_i is obtained from the transfer function estimate following Equation 3.3 which is different for each frequency. Therefore, we replace σ in Equation A.2 in this case with σ_i . In addition, we use the absolute value of the nominator in this case. Moreover, the model \mathcal{M} is the model of the control loop as given in Equation 3.9. The precise values of the optimum are naturally influenced by the choice of this function, often informally also named ‘cost function’.

Based on the previous considerations, a likelihood function has to be defined to use the MCMC approach. From Equation A.2 with our modification of known σ_i , we find that the likelihood function of the measured data \vec{y} being the result of the parameters $\vec{\theta}$, is proportional to the product of the probability of each data point to be the result of these parameters in light of the given noise with the above assumptions:

$$\mathcal{L} \propto \prod_{i=1}^N \exp \left(-\frac{1}{2} \frac{(y_i - \mathcal{M}_i^2)}{\sigma_i^2} \right) \quad (\text{A.3})$$

Often, it is easier to compute the natural logarithm of this expression, the so-called ‘log-likelihood’.

As explained in [HFM18], the term MCMC summarises methods to sample a probability density function. One very common sampling algorithm is the Metropolis-Hastings algorithm but there are also others, for example the Gibbs sampling.

To determine the posterior probability distribution for the parameters $\vec{\theta}$ given the measured data \vec{y} , we make use of Bayes theorem [Kar14]

$$p(\vec{\theta}|\vec{y}) = \frac{p(\vec{y}|\vec{\theta})p_0(\vec{\theta})}{p(\vec{y}|\mathcal{M})}. \quad (\text{A.4})$$

Here, $p(\vec{\theta}|\vec{y})$ is the posterior probability distribution of the parameters $\vec{\theta}$ given the data \vec{y} which we would like to sample with the MCMC. $p(\vec{y}|\vec{\theta})$ is the probability density for the data given the parameters. It is determined by the previously determined likelihood function [HBL10]. $p_0(\vec{\theta})$ designates the prior probability distribution of the parameters. In the case of the actuator parameter estimation, we obtain a guess for the shape of this prior probability distribution from the estimated covariance of the results of the Nelder-Mead simplex method. In addition, it is affected in case a possible parameter range is specified at the beginning of the sampling. The denominator $p(\vec{y}|\mathcal{M})$ is also called marginal likelihood and is usually neglected in the case of parameter estimation where it is only a normalisation constant [Kar14]. As the MCMC is insensitive to this parameter [HFM18], it will not be discussed in more detail here.

The Bayes theorem is the central part of the Bayesian approach to data analysis, in contrast to what is often called the ‘frequentist approach’ [Kar14]. A key difference of these approaches is the concept of probabilities. Whereas in the ‘frequentist approach’, the probability is taken to be the ‘relative frequency of occurrence of a given event in sequential repetitions of the same experiment’ the Bayesian approach treats probability as ‘the assignment of a degree of belief for the given event, based on the evidence at hand’ [Kar14]. A more detailed discussion is beyond the scope of this thesis. However, we would like to emphasise the advantage of the Bayesian approach for the case of the actuator parameter estimation as it allows us to estimate confidence intervals for a relatively small data set [Kar14] as the 6-8 laser frequency modulation per loop characterisation experiment by making use of our prior knowledge.

A.2.2 Definition of the term MCMC

As the name indicates, the Markov Chain Monte-Carlo (MCMC) methods are a combination of Markov Chain and Monte Carlo methods. In heuristic terms, a Markov Chain designates a random walker whose next step only depends on the previous one, not on the chain history [Kar14]. If the walker can freely move in the parameter space, the kernel can be constructed in such a way that a desired stationary probability distribution is represented if the walker’s chains are also ergodic [Kar14]. The Monte Carlo term makes clear that it is one of the many Monte Carlo methods which, roughly speaking, replace an integral, which may be hard to compute numerically, with adequate random sampling. This can be seen if we look at the original Metropolis algorithm. As summarised in [HFM18], this algorithm draws a proposal from the proposal probability distribution function. In the next step, a random number between 0 and 1 is generated. Then, the ratio of the function to be sampled, in cases like ours the log-likelihood function, at the proposed new sample to the current sample is computed. If this ratio is larger than the random number, this new sample is accepted and else it is rejected. The ratio of accepted samples per proposed samples is the acceptance ratio. As further explained below, this is one indicator for a good sampling of the posterior probability density function.

A.2.3 Important technical aspects

In this paragraph, we want to summarise two technical aspects that should be considered when using the MCMC algorithm. One aspect is that the first few samples depend on the starting point and that samples closely together, especially at the beginning of the sampling process, may be correlated [Kar14]. These two points result in the requirement that the MCMC needs to run long enough. However, it is difficult to quantify ‘long enough’ in general. But there are a number of criteria that can help to answer this question for the case at hand. [HFM18] recommend as heuristic criteria to ensure that the random walker has crossed the areas of high probability in the parameter space many times or equivalently, that subsets of the chain give comparable results for the posterior probability density function. This is similar to the potential scale reduction factor explained in [Kar14]. Alternatively, they claim a MCMC has run long enough when the result does not change with different starting values and different random number seeds used for the computation. This list of criteria is by no means exhaustive, already [Kar14] also proposes further possible checks.

A second technical aspect to be considered is that the MCMC user should ensure that a good fraction of the probability distribution is sampled. One common way to simplify this process for the case of complicated distributions is the so-called ‘heating profile’ [Kar14]. This approach is smoothing the posterior distribution such that for example valleys can be crossed more easily and is one example for a MCMC annealing procedure [Kar14]. Naturally, this occurs at the beginning of the sampling and these samples will be discarded later on for the estimation of the final posterior probability density distribution [Kar14]. This beginning phase is called the burn-in and these samples can also be discarded if the user is afraid of having started the sampling with non-typical parameters [HFM18].

Another important diagnostic is the acceptance ratio of the MCMC [HFM18]. If nearly all samples are accepted, this can often indicate that the step sizes of the chains, determined by the proposal distribution, are small. Hence, this can be an indicator for the chain being confined to a small region of the parameter space for a given amount of run time. Contrarily, if we accept too few samples, this could be an indicator of the chain being in a low-probability region of the posterior probability distribution such that it may take a long time to reach the high probability regions from there. The acceptance ratio can be influenced via the proposal probability distribution function. One option to do so is the adaptive proposal scheme, where the proposed probability distribution function is updated based on the history of the chain, compare for example [Wis17]. Again, these parts have to be discarded for the calculation of the posterior probability distribution function such that the Markov character of the chains remains [HFM18].

A.3 MCMC uncertainty estimation results and residuals for the actuator parameters

In this section, we present the actuator fit results in more detail. For each of the five experiments, we show the MCMC parameter chains to increase the confidence in the MCMC estimate. We also show the resulting histograms and estimated covariances. The chains and the histograms are a heuristic indicator that the number of MCMC samples is sufficient for our purpose as explained in the previous section. The MCMC was run for $15 \cdot 10^3$ samples. We show the results after the burn-in phase which lasted $2 \cdot 10^3$ samples and where the adaptive proposal scheme has been used. As a part of this scheme, the covariance matrix for the proposal distribution from which

A.3. MCMC UNCERTAINTY ESTIMATION RESULTS AND RESIDUALS FOR THE ACTUATOR PARAMETERS

the samples are drawn, was updated after every 200 samples. Here, we show plots of the relative error in each experiment as

$$r = \left| 100 \cdot \frac{\vec{y} - \mathcal{M}(\vec{\theta}_{\text{fit}})}{\vec{y}} \right| \quad (\text{A.5})$$

with r being the relative error in % which is calculated from the difference between the data \vec{y} and the model \mathcal{M} evaluated using the fit results $\vec{\theta}_{\text{fit}}$ at the modulation frequencies of the experiment. Let us mention that r measures fit quality which is a different quantity than the parameter uncertainties reported in Table 3.3.

An example of how the MCMC chains should not look are the plots of experiment 2. This is because, as can be seen in Figure A.7, the gains and delays are not only drifting at the beginning of the sampling but all along the sampling process. This is most likely because there is only a single data point but 4 parameters to determine. Perhaps the result could have been improved with adapting the prior probability density function more strongly using the results from experiment 1. However, to ensure this measurement can be compared to the others, it was decided to employ the same prior probability density function. The chains are fluctuating around a fixed value as expected, compare Section A.2.3, in the other experiments. Of course, it is still possible that only a valley of the posterior probability density function is sampled but together with a reasonable acceptance ratio found, these two heuristic criteria were considered sufficient given that we have found no evidence for a complicated multimodal distribution. On the contrary, the histograms have a Gaussian-like shape for all experiments with the exception of experiment two. This increases the confidence in the reported simple parameter uncertainties. For each experiment, we also show a covariance matrix. With the exception of experiment 2, due to the ambiguity explained above, the covariances are similar. The fast actuator gain estimate is correlated to both the fast and the slow actuator delay. This can for example be seen in the bottom left panel of Figure A.5. The fast actuator delay is correlated to the fast actuator gain as well as the slow actuator delay. The slow actuator gain is the only parameter which appears significantly less correlated than the other parameters. The cause of this observation could not be identified yet. The slow actuator delay estimate is affected by the fast actuator gain and delay. The relative error is mostly around 1% and sometimes smaller but never above 10%.

A.3.1 FM test campaign experiment




Figure A.1: The MCMC chains using the data from the FM test campaign.




Figure A.2: The covariances and histograms of the parameters for the ground test results.

A.3. MCMC UNCERTAINTY ESTIMATION RESULTS AND RESIDUALS FOR THE ACTUATOR PARAMETERS




Figure A.3: The relative error of the fit to transfer function measurement $T_{ff_{out} \rightarrow \Psi_F}$ for ground test.

A.3.2 Loop characterisation experiment 1




Figure A.4: The MCMC chains of experiment 1.




Figure A.5: The covariances and histograms of the parameters for experiment 1.




Figure A.6: The relative error of the fit to transfer function measurement $T_{ff_{out} \rightarrow \psi_F}$ for experiment 1.

A.3.3 Loop characterisation experiment 2




Figure A.7: The MCMC chains of experiment 2.




Figure A.8: The covariances and histograms of the parameters for experiment 2.




Figure A.9: The relative error of the fit to transfer function measurement $T_{ff_{out} \rightarrow \psi_F}$ for experiment 2.

A.3.4 Loop characterisation experiment 3




Figure A.10: The MCMC chains of experiment 3.




Figure A.11: The covariances and histograms of the parameters for experiment 3.




Figure A.12: The relative error of the fit to transfer function measurement $T_{\text{ff}_{\text{out}} \rightarrow \Psi_F}$ for experiment 3.

A.3.5 Loop characterisation experiment 4




Figure A.13: The MCMC chains of experiment 4.




Figure A.14: The covariances and histograms of the parameters for experiment 4.




Figure A.15: The relative error of the fit to transfer function measurement $T_{\text{ff}_{\text{out}} \rightarrow \Psi_F}$ for experiment 4.

A.3.6 Loop characterisation experiment 5




Figure A.16: The MCMC chains of experiment 5.




Figure A.17: The covariances and histograms of the parameters for experiment 5.




Figure A.18: The relative error of the fit to transfer function measurement $T_{\text{ff}_{\text{out}} \rightarrow \Psi_F}$ for experiment 5.

Appendix B

List of OMS telemetry channels

This list provides technical information on the OMS measurements on LPF used in the analysis of this thesis. The LTPDA telemetry names of the measurement channels are given along with a list of the parameter aliases under which the objects can be found in databases. A short description as it is implemented in LTPDA is also provided along with the terms under which this measurement can be found in the other chapters. The parameters of Tables B.21 to B.28 are part of the laser amplitude control loop, see [Aud14].

Parameter	Alias
DMU_OMS_o12	LAT00005
	GST00103
	GST10103
	GST20103
	GST30103
	GST50103
	LAT10005

Table B.1: Table of parameter alias for DMU_OMS_o12: The x position of TM2 w.r.t. TM1 as measured by the X12 interferometer. This measurement is denoted as o12 in the thesis.

Parameter	Alias
DMU_OMS_o1	LAT00003
	GST00101
	GST10101
	GST20101
	GST30101
	GST50101
	LAT10003

Table B.2: Table of parameter alias for DMU_OMS_o1: The x position of the SC w.r.t. TM1 as measured by the X1 interferometer at 10Hz. This measurement is denoted as o1 in the thesis.

Parameter	Alias
DMU_OMS_PSLF	LST12406
	LAT00070
	GST00112
	GST10112
	GST20112
	GST30112
	GST50112
	LAT10070

Table B.3: Table of parameter alias for DMU_OMS_PSLF: The output of the frequency interferometer as given by the DMU. This measurement is denoted as Ψ_F in the thesis.

Parameter	Alias
DMU_OMS_PSLR	LST12407
	LAT00071
	GST00113
	GST10113
	GST20113
	GST30113
	GST50113
	LAT10071

Table B.4: Table of parameter alias for DMU_OMS_PSLR: The output of the reference interferometer as given by the DMU. This measurement is denoted as Ψ_R in the thesis.

Parameter	Alias
DMU_OMS_FAST_FREQ_CTRL_ERROR	LST17359

Table B.5: Table of parameter alias for DMU_OMS_FAST_FREQ_CTRL_ERROR: The error signal of the frequency interferometer fast loop as given by the DMU. This measurement is denoted as ff_{err} in the thesis.

Parameter	Alias
DMU_OMS_FAST_FREQ_CTRL_OUT	LST17360

Table B.6: Table of parameter alias for DMU_OMS_FAST_FREQ_CTRL_OUT: The output signal of the frequency interferometer fast loop as given by the DMU. This measurement is denoted as ff_{out} in the thesis.

Parameter	Alias
DMU_OMS_SLOW_FREQ_CTRL_ERROR	LST17362

Table B.7: Table of parameter alias for DMU_OMS_SLOW_FREQ_CTRL_ERROR: The error signal of the frequency interferometer slow loop as given by the DMU. This measurement is denoted as sf_{err} in the thesis.

Parameter	Alias
DMU_OMS_SLOW_FREQ_CTRL_OUT	LST17363

Table B.8: Table of parameter alias for DMU_OMS_SLOW_FREQ_CTRL_OUT: The output signal of the frequency interferometer slow loop as given by the DMU. This measurement is denoted as sf_{out} in the thesis.

Parameter	Alias
DMU_OMS_DWS_eta1	LAT00011
	GST00109
	GST10109
	GST20109
	GST30109
	GST50109
	LAT10011

Table B.9: Table of parameter alias for DMU_OMS_DWS_eta1: The eta attitude of TM1 w.r.t. the SC as measured by the X1 interferometer and calibrated in the DMU (sampled at 10Hz). This measurement is denoted as η_1 in the thesis.

Parameter	Alias
DMU_OMS_DWS_eta2	LAT00013
	GST00111
	GST10111
	GST20111
	GST30111
	GST50111
	LAT10013

Table B.10: Table of parameter alias for DMU_OMS_DWS_eta2: The eta attitude of TM2 w.r.t. the SC estimated by combining the measurements of the X12 and X1 interferometers and calibrated in the DMU (sampled at 10Hz). This measurement is denoted as η_2 in the thesis.

Parameter	Alias
DMU_OMS_DWS_phi1	LAT00010
	GST00108
	GST10108
	GST20108
	GST30108
	GST50108
	LAT10010

Table B.11: Table of parameter alias for DMU_OMS_DWS_phi1: The phi attitude of TM1 w.r.t. the SC as measured by the X1 interferometer and calibrated in the DMU (sampled at 10Hz). This measurement is denoted as ϕ_1 in the thesis.

Parameter	Alias
DMU_OMS_DWS_phi2	LAT00012
	GST00110
	GST10110
	GST20110
	GST30110
	GST50110
	LAT10012

Table B.12: Table of parameter alias for DMU_OMS_DWS_phi2: The phi attitude of TM2 w.r.t. the SC estimated by combining the measurements of the X12 and X1 interferometers and calibrated in the DMU (sampled at 10Hz). This measurement is denoted as ϕ_2 in the thesis.

Parameter	Alias
DMU_OMS_DC_eta1	LAT00007
	GST00105
	GST10105
	GST20105
	GST30105
	GST50105
	LAT10007

Table B.13: Table of parameter alias for DMU_OMS_DC_eta1: The eta attitude of TM1 w.r.t. the SC as given by the DC quadrant powers of the X1 interferometer (sampled at 10Hz). This measurement is denoted as η_1^{DC} measurement in the thesis.

Parameter	Alias
DMU_OMS_DC_eta2	LAT00009
	GST00107
	GST10107
	GST20107
	GST30107
	GST50107
	LAT10009

Table B.14: Table of parameter alias for DMU_OMS_DC_eta2: The eta attitude of TM2 w.r.t. the SC as given by the DC quadrant powers of the X1 and X12 interferometers. This measurement is denoted as η_2^{DC} measurement in the thesis.

Parameter	Alias
DMU_OMS_DC_phi1	LAT00006
	GST00104
	GST10104
	GST20104
	GST30104
	GST50104
	LAT10006

Table B.15: Table of parameter alias for DMU_OMS_DC_phi1: The phi attitude of TM1 w.r.t. the SC as given by the DC quadrant powers of the X1 interferometer. This measurement is denoted as ϕ_1^{DC} measurement in the thesis.

Parameter	Alias
DMU_OMS_DC_phi2	LAT00008
	GST00106
	GST10106
	GST20106
	GST30106
	GST50106
	LAT10008

Table B.16: Table of parameter alias for DMU_OMS_DC_phi2: The phi attitude of TM2 w.r.t. the SC as given by the DC quadrant powers of the X1 and X12 interferometers. This measurement is denoted as ϕ_2^{DC} measurement in the thesis.

Parameter	Alias
LA_RLU_TEMP	LLT10003

Table B.17: Table of parameter alias for LA_RLU_TEMP: Internal operating temperature of the RLU. It is a 16-bit word which covers the range -40degC to 129degC.

Parameter	Alias
LA_RLU_Pump_Current	LLT10001

Table B.18: Table of parameter alias for LA_RLU_Pump_Current: RLU monitor of the pump diode current. It is a 16-bit word which covers the range 0A to 5.3A.

Parameter	Alias
LA_RLU_PWR_OUT	LLT10002

Table B.19: Table of parameter alias for LA_RLU_PWR_OUT: RLU monitor of the laser power from the internal photodiode. It is a 16-bit word which covers the range 0mW to 65mW.

Parameter	Alias
LA_LMU_AOM_TEMP	LLT10008

Table B.20: Table of parameter alias for LA_LMU_AOM_TEMP: Monitor of the temperature of the Laser modulator unit between the AOMs. It is a 16-bit word which covers the range -40degC to +169degC.

Parameter	Alias
LA_LM_RF_AMPLITUDE_1	LLT10050

Table B.21: Table of parameter alias for LA_LM_RF_AMPLITUDE_1: Monitor of the power amplifier 1. It is a 16-bit word which covers the range -10V to 10V.

Parameter	Alias
LA_LM_RF_AMPLITUDE_2	LLT10051

Table B.22: Table of parameter alias for LA_LM_RF_AMPLITUDE_2: Monitor of the power amplifier 2. It is a 16-bit word which covers the range -10V to 10V.

Parameter	Alias
LA_LM_PWR1	LLT10052

Table B.23: Table of parameter alias for LA_LM_PWR1: Monitor of the power signal of beam 1 after TIA. It is a 16-bit word which covers the range 1000uA to -1000uA.

Parameter	Alias
LA_LM_PWR2	LLT10053

Table B.24: Table of parameter alias for LA_LM_PWR2: Monitor of the power signal of beam 2 after TIA. It is a 16-bit word which covers the range 1000uA to -1000uA.

Parameter	Alias
LA_LM_PWR_STAB_ERR1	LLT10046

Table B.25: Table of parameter alias for LA_LM_PWR_STAB_ERR1: Monitor of the power stabilization signal after the Error amplifier 1. It is a 16-bit word which covers the range -10V to 10V.

Parameter	Alias
LA_LM_PWR_STAB_ERR2	LLT10047

Table B.26: Table of parameter alias for LA_LM_PWR_STAB_ERR2: Monitor of the power stabilization signal after the Error amplifier 2. It is a 16-bit word which covers the range -10V to 10V.

Parameter	Alias
LA_LM_PWR_STAB1_Feedback	LLT10048

Table B.27: Table of parameter alias for LA_LM_PWR_STAB1_Feedback: Monitor of the power stabilization signal after the Servo amplifier 1. It is a 16-bit word which covers the range -10V to 10V.

Parameter	Alias
LA_LM_PWR_STAB2_Feedback	LLT10049

Table B.28: Table of parameter alias for LA_LM_PWR_STAB2_Feedback: Monitor of the power stabilization signal after the Servo amplifier 2. It is a 16-bit word which covers the range -10V to 10V.

Appendix C

Free-Fall Experiment

In this chapter, we briefly summarise the free-fall experiment on LPF. The focus is on parameter estimation work, which has been undertaken during the mission preparation phase. The applied procedures follow Appendix A.

C.1 Concept of the experiment

In the course of the LPF mission development and testing, the expected performance could be improved below the originally required residual acceleration level because of four main reasons [A+12a]. First, a precise estimate of the expected gravitational imbalances allowed to derive the compensation forces which the GRS would have to apply. In combination with laboratory measurements of electrostatic actuation noise, a smaller acceleration noise contribution from the actuation system was expected. The second reason summarised in [A+12a] was that with a proposed redesign of attitude control loops, the coupling of star tracker noise into the science measurement could be minimised. In addition, the laboratory measurements of the OMS sensing noise were below the required level. The fourth contribution to a lower noise level was expected to result from dedicated experiments to measure a different coupling of the two TMs to satellite motion and subsequent system adjustments. These factors led to an estimated acceleration noise budget as shown in Figure C.1. From this noise budget, it is apparent that the residual acceleration below 3 mHz was expected to be dominated by electrostatic actuation noise. This, however, was undesired because, with such a limitation, the technology demonstrator mission LPF would be less representative of the LISA mission where there is no actuation along the measurement axis. However, the electrostatic actuation on TM2 in the science mode (see Section 1.2.1) cannot be switched off easily because the expected remaining gravitational imbalance on the satellite, which leads to a static acceleration along the measurement axis of 0.65 nm s^{-2} (see for example [A+12a]), would move the TM away from the desired operating point. Hence, a compensation force is mandatory.

The subsequent need for control in the absence of electrostatic actuation is fulfilled by the drift mode or free-flight experiment. Here, the continuous control of the science mode is replaced by an intermittent control scheme. This means that periods of control impulses are interleaved with uncontrolled periods, the so-called free-flights. Accordingly, the force applied during the impulses is significantly larger than during the same amount of time in the standard science mode. Here,




Figure C.1: The LPF acceleration noise budget current best estimate before launch. Electrostatic actuation noise was expected to dominate below 3 mHz. Reprint of [A⁺12a].

a little more detail on the electrostatic actuation is necessary. On LPF, the actuation electronics have two operational modes: the wide-range and high resolution mode. As the names indicate, the first is used if a comparatively high actuation force of up to $0.94 \mu\text{N}$ is required but is relatively noisy [Bra09]. If a maximum actuation force of 2 nN is sufficient, the high-resolution mode can be set [Bra09]. During LPF operations, sensing noise levels of 1.2 and $2.4 \text{ nm}/\sqrt{\text{Hz}}$ in displacement have been measured in this mode [A⁺17b]. In the nominal design of the free-fall experiment, the force impulse was applied with the actuation electronics in the wide range mode and for the free-flights, the actuation electronics switched to the high resolution mode. This implies that the switches between these two modes have to be fast and reliable over longer periods of time.

The intermittent control scheme holds true for the control along the sensitive axis of LPF while all other degrees of freedom are controlled as in science mode. Other intermittent control modes, where for example also the ϕ degree of freedom of TM2 is subject to intermittent control, have also been designed [Sch14][Giu17]. The underlying rationale is that the same electrodes are used for the control along x and ϕ such that a constant actuation along ϕ implies a certain level electrostatic actuation noise along x and vice versa. However, as these control modes were not applied in flight, they will not be further discussed in this work.

Let us move now from the description of the experiment to the related data analysis. This intermittent control scheme results in two features which are not present in the estimation of the residual acceleration noise spectrum during a nominal measurement in science mode.

First, the force impulses result in a quasi-parabolic differential TM motion, as shown for a simulated example in Figure C.2. The force impulse of typically 1 s is calculated such that TM2 returns to its starting position within a given time, typically approximately 350 s . Thus, the amplitude of the quasi-parabolas depends on the residual static acceleration along the measurement axis. Additionally, the motion of each of the TMs is slightly affected by the satellite environment




Figure C.2: An example of the simulated force impulses and free-flights where the typical quasi parabolic motion is discernible.

via a force gradient due to gravitational, electrostatic and other forces. These force gradients are often called ‘stiffness’ and are measured in s^{-2} [A⁺18c]. The relative motion of TM2 during the free-fall periods has to be modelled and subtracted from the data to yield a measurement of residual acceleration which can be compared to the measurements obtained in the standard science mode.

The second difference to the nominal science mode is that the differential acceleration data taken during the force impulses exhibits a significantly increased noise level due to the use of the wide-range mode. Consequently, it is too noisy to be used for residual acceleration measurements on the levels of the science mode measurements or below and therefore has to be discarded. This leads to the challenge of estimating a PSD below 1 mHz from data segments of several hundreds of seconds. In other words, the task at hand can be described as estimating the PSD of data with periodic gaps.

C.2 Preparations for operations

The free-flight experiment had been carefully prepared prior to LPF operations, including tests on the torsion balance facility at University of Trento. It was verified experimentally that the intermittent control scheme allowed for a precise torque noise estimation in the presence of a large DC torque. The reduction of actuation noise during the actuation-free periods, however, could not be shown within the measurement accuracy which was limited by aliasing [R⁺18].

In addition, several approaches to the spectral estimation in the presence of periodic gaps had been tested, as summarised in [A⁺15].

One of these was the Constrained-Gaussian Gap Patching which aims at filling the gaps with noise which has a proper correlation to the noise used to fill the other gaps and to the measurement data during the free-fall periods. Therefore, a model of the underlying noise spectrum is needed.

The second approach, summarised in [A⁺15], relies on mitigating the effect of the force impulses by multiplying the data with a window function which smoothly goes to 0 during the force impulses. For more details, we refer for example to [Gry09].

In addition, another approach where the data is low-pass filtered using a Blackman-Harris window, then downsampled and set to 0 during the force impulses has also been developed, see for example [Giu17]. Work at the AEI has been done to estimate the parameters of the quasi-parabolas of the free-flights to subtract the large deterministic motion. This work is described in the following paragraph.

C.2.1 The local simulation environment

The local simulation environment is based on the modelling of LPF as an assembly of linear and time-invariant state space models for each of the subsystems and noise models implemented in LTPDA. For more details about the state space models, we refer to [Pac14]. This simulation has been developed by A. Grynagier, see for example [GFV10].

Clearly, the free-flight experiment is non-linear due to the switch between force impulses and the free-flight periods. This behaviour is approximated by constructing a linear model both for the kick periods and the free-flight periods. The system state at the end of the kick period is then used as the initial state for the free-flight simulation and the system state at the end of this period is used as the initial state of the next kick period, as illustrated in Figure C.3. In this




Figure C.3: The concept of the state-space model free-flight simulation. The system state at the end of either the kick or the free-flight period is used as the initial state for the simulation of the next period.

simulation, many system parameters can be set, such as the stiffnesses and the static acceleration level. In addition, different noise models are available and the dimension of the simulation can be chosen.

We will now show an example simulation and the corresponding parameter estimation results. In the example one-dimensional simulation under study, the static differential acceleration $\Delta DCacc$ is set to 1.04 nm s^{-2} and the stiffness of TM2 per unit mass ω is set to $2 \mu\text{s}^{-2}$. With the default settings for the stiffness of TM1, this results in a difference in stiffness between the two TMs of $\approx 64.9 \text{ ns}^{-2}$. This is significantly smaller than the stiffness of TM2 and thus has not been included in the model used in this example. In this simulation, coloured noise is used. It results from the dedicated LPF state space models which have been used in the version ‘Best Case June 2011’ if available and else in the ‘DFACS ICD v1.4’ version. The force noise acting on the TMs

arising from outside the satellite and the force noise acting between the satellite and the TMs has been reduced by 50%. The capacitive actuation noise is increased for the force impulse periods. After a short initialisation phase, the o12 and the force impulses time series data looks like the example data stretch that has been shown in Figure C.2. In total, 349 free-flight quasi-parabolas have been simulated and used for parameter estimation.

For each of these, the stiffness per unit mass ω and the static differential acceleration $\Delta DCacc$ previously set has been estimated. For clarity, we will replace o12 in the following equations with x_{12} . That is we use a simple model

$$\ddot{x}_{12} = \omega x_{12} + \Delta DCacc, \quad (C.1)$$

with x_{12} as the relative displacement. This model is used in two different ways.

The first option tested is to fit this model to the direct displacement simulation data, o12. Therefore, the ordinary differential equation C.1 is solved for x_{12} and the result is compared to the simulated o12 measurement. The resulting data-model expression is then used to find the parameter values for ω and $\Delta DCacc$. For the minimisation, we build a log-likelihood function and apply the Nelder-Mead simplex method as described in Appendix A. This allows for significantly more complex and non-linear models. This first option is called ‘displacement fit’ in Table C.1. However, this option directly implies that the ordinary differential equation has to be solved many times during the minimisation which is far from ideal in terms of runtime.

Therefore, the second approach to the parameter estimation for the simulated data is to differentiate the simulated results for o12 to then minimise

$$\ddot{x}_{12} - (\omega x_{12} + \Delta DCacc) . \quad (C.2)$$

This second option is called ‘acceleration fit’ in Table C.1 and the minimisation technique for parameter estimation is the same as for the ‘displacement fit’.

Each of these two approaches is applied to each of the simulated 349 free-flight quasi-parabolas that have been simulated. The free-flight segments have been treated individually, that is without taking into account the result of the fit to the previous free-flight. Accordingly, the same initial guess is used for each of these. An uncertainty is estimated for the results of each quasi-parabola using the inverse of the derivative of the cost function with respect to its parameters at the simplex results. The thus estimated uncertainties are denoted as Fisher Matrix (FM) uncertainty in the following. This approach to uncertainty estimation is also used to obtain a starting covariance matrix for MCMC sampling, as explained in Appendix A.

The parameter estimation results and the uncertainty estimates are shown in Table C.1. This table exhibits a good agreement between the fit results and the parameters chosen for the simulation as well as the expected agreement of the two different expression that have been used for minimisation and uncertainty estimation.

Figure C.4 shows the time series of the residuals from the displacement fit for each free-flight period. We find that already with this simple model, we can subtract the quasi-parabolic motion and approach the noise level throughout. Together with Table C.1, this confirms the two approaches to parameter estimation of simulated free-fall data.

parameter	fitting method	mean value	standard deviation	rms of FM uncertainty
Δ DC acc	acceleration fit	$1.03999 \cdot 10^{-9}$	$1.2 \cdot 10^{-14}$	$1.3 \cdot 10^{-14}$
Δ DC acc	displacement fit	$1.03999 \cdot 10^{-9}$	$2 \cdot 10^{-14}$	$2.5 \cdot 10^{-14}$
ω	acceleration fit	$1.9999 \cdot 10^{-6}$	$8.1 \cdot 10^{-10}$	$8.3 \cdot 10^{-10}$
ω	displacement fit	$1.9999 \cdot 10^{-6}$	$1.4 \cdot 10^{-9}$	$1.7 \cdot 10^{-9}$

Table C.1: Parameter estimation results for simulated free-flight experiment data. We find that the parameters set in the one-dimensional simulation are recovered by the parameter estimation and the uncertainty estimates from different approaches agree.




Figure C.4: The time series data of the fit residuals of the displacement fit for the simulated free-fall data. We find that already with this simple model, we can subtract the quasi-parabolic motion and approach the noise level.

C.3 The experiment during operations

As stated in Section C.1, the static differential acceleration along the measurement axis was estimated to be 0.65 nm s^{-2} before launch. In flight, this value was found to be reduced to 20 pm s^{-2} and below due to the gravitational balancing of the satellite components achieved [A+16c]. Along the measurement axis, the smaller static differential acceleration implies that the suspension control loop (see Section 1.2.1) needs to apply less force during a standard science measurement. This was an important contribution to the small residual acceleration noise levels measured on LPF because the force and torque noise resulting from the electrostatic actuation depends not only on the actual applied force or torque but also on the currently chosen maximum force/torque limit, the so-called ‘actuation authority’ [A+16c]. After some initial measurements

taken with the nominal actuation authorities, the actuation authorities were lowered during the subsequent nominal science mode measurements. As a consequence, the residual acceleration below 1 mHz was not limited by electrostatic actuation noise as anticipated.

However, the free-fall experiment could be performed several times during the LPF mission with varying settings. During the experiment’s analysis, it became apparent that putting the gaps to zero as it was done by two approaches to the free-flight data analysis (see Section C.2) introduces a significant bias in the spectrum. This can easily be seen by applying the free-fall experiment data analysis procedure to a standard noise run and comparing the results. This bias is mitigated in several steps which we summarise here from [Giu17]. At first a parametrised spectral model of the true underlying noise is assumed. Then, it is convolved with a model of the window applied, the downsampling and the gaps introduced. Next, the parameters of the spectral model are fit to match the recorded data. Then, the original spectrum and the processed spectrum are compared to find the corresponding scaling factor. In the last step, this factor is applied to the experimental data. This technique has been applied for the results summarised below.

In the first run, the experiment was performed using the nominal actuation authorities. Compared to a standard science mode measurement with these actuation authorities, the expected reduction of roughly a factor 2 between 0.1 and 1 mHz could be achieved [A+19e]. This first run thus showed that the intermittent control mode is working in orbit and reduces the low frequency noise as was aimed for.

In another run of the experiment, the actuation authority during the drift phases had been reduced as in the standard science measurements to 50 pN along x_2 and 1 pN m along ϕ_2 , respectively. This configuration was called ‘URLA’. Moreover, in this run, the application of a constant out-of-loop force of 11.2 pN was applied to further reduce the static differential acceleration and thus the amplitude of the quasi parabolas. This avoided difficulties in the data recording. A third difference in this version of the experiment compared to other runs is that the high resolution actuation mode could also be used during the force impulse periods by extending them to 5 s duration.

In the URLA actuation authority configuration, the noise levels during the free-fall periods are no longer expected to be below those of a nominal science measurement in the same configuration. The reason is that in the URLA configuration, the actuation noise is dominated by the contribution from ϕ actuation which remains the same in the applied version of the free-fall mode. This expectation was confirmed by the measurements, see [A+19e]. Even though this version of the free-fall experiment did not reduce the residual acceleration noise level, it provides a highly valuable confirmation of the previously measured LPF performance in a different control mode.

Acknowledgements

Many people have supported me in preparing this thesis in various ways and I sincerely hope that with this attempt I will do everyone justice.

First, I would like to thank Prof. Dr. Danzmann for accepting me as a PhD student at the AEI and for continuously spreading enthusiasm. Thank you as well for standing up for your students and their concerns. I am no less grateful to Prof. Dr. Gerhard Heinzel for sharing his immense knowledge via an open office door. It is a pleasure to work in the space interferometry group because of the friendly and open-minded atmosphere you create and combine with different international collaborations.

I would like to offer my thanks and appreciation to Martin Hewitson for a constant guidance throughout my PhD work on top of an everlasting immense pile of work. Thank you especially for your interest in all questions which arose along the years, helpful discussions, endless cross-checking of slides and for reading this thesis!

A huge thank you to the whole space interferometry group, and especially to everyone who read a part or more of this work. I also enjoyed working with our former colleagues Markus Otto and Heather Audley. Furthermore, I would also like to explicitly thank Vitus Händchen for reading the thesis in parallel to a new job. Excellent LTPDA support, endless patience when necessary and a large amount of fresh fruit has been contributed by our former office mate Ingo Diepholz. Nikos Karnesis, thank you for many interesting discussions on parameter estimation and LPF questions during your time at the AEI. For a sketch, the ComponentLibrary (<http://www.gwoptics.org/ComponentLibrary/>) by A. Franzen was used.

I would also like to thank Michael Born for sharing the timing correction code used especially to obtain the results of Chapter 3. Thank you also for many useful discussions. Thanks Gudrun Wanner for lots of advice and helpful discussions much beyond what can be expected from a busy colleague next door.

I am very much obliged to my fellow AEI companions along the LPF PhD track: Brigitte Kaune, Andreas Wittchen and Lennart Wissel. It is a pleasure to share the work with friends. Thank you for an infinite amount of helpful discussions, great breaks and support in the everyday business. Thereby, special thanks go to Andreas Wittchen for performing and discussing the AEI laboratory measurements shown in Chapter 6.

Being part of the LISA Pathfinder collaboration during the mission operations phase was a unique experience. Thank you all for including also the PhD students into this team and many helpful meetings, comments and discussions. I would like to specifically thank Roberta Giusteri, Giuliana Russano, Oliver Jennrich, James Ira Thorpe and William Joseph Weber for very interesting

ACKNOWLEDGEMENTS

discussions about the free-flight experiment as well as David Robertson and Paul McNamara for answering a number of OMS-related questions.

Graduating would not have been possible without a committee who have readily agreed to spend their valuable time on this thesis and the defense. Thank you.

Finally, I would like to use this opportunity to humbly thank my family, my friends and especially my better half for your immense support throughout the years. Everyone of you knows best for what and even an incomplete list is certainly beyond the scope of this thesis!

List of publications

- [P1] G. Anderson, J. Anderson, M. Anderson, G. Aveni, D. Bame, P. Barela, K. Blackman, A. Carmain, L. Chen, M. Cherng, S. Clark, M. Connally, W. Connolly, D. Conroy, M. Cooper, C. Cutler, J. D’Agostino, N. Demmons, E. Dorantes, C. Dunn, M. Duran, E. Ehrbar, J. Evans, J. Fernandez, G. Franklin, M. Girard, J. Gorelik, V. Hruby, O. Hsu, D. Jackson, S. Javidnia, D. Kern, M. Knopp, R. Kolasinski, C. Kuo, T. Le, I Li, O. Liepack, A. Littlefield, P. Maghami, S. Malik, L. Markley, R. Martin, C. Marrese-Reading, J. Mehta, J. Mennela, D. Miller, D. Nguyen, J. O’Donnell, R. Parikh, G. Plett, T. Ramsey, T. Randolph, S. Rhodes, A. Romero-Wolf, T. Roy, A. Ruiz, H. Shaw, J. Slutsky, D. Spence, J. Stocky, J. Tallon, J. I. Thorpe, W. Tolman, H. Umfress, R. Valencia, C. Valerio, W. Warner, J. Wellman, P. Willis, J. Ziemer, J. Zwahlen, M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesai, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jenntich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I Mateos, P. W. McNatnara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. Robertson, I, F. Rivas, G. Russano, C. F. Sopena, T. Sumner, D. Texier, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, P. Zweifel, ST7 Team, and LISA Pathfinder Collaboration. Experimental results from the ST7 mission on LISA Pathfinder. *PHYSICAL REVIEW D*, 98(10), NOV 14 2018.
- [P2] M. Armano, H. Audley, G. Auger, J. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, R. De Rosa, L. Di Fiore, I. Diepholz, G. Dixon, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, R. Flatscher, M. Freschi, C. Garcia Marirrodriaga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, A. Grado, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, S. Madden, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, N. Meshksar, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, D. I. Robertson, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, R. Stanga, T. J. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, M. Troebs, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass,

- D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanoni, T. Ziegler, P. Zweifel, LISA Pathfinder Collaboration, and LISA Pathfinder Collaboration. Capacitive sensing of test mass motion with nanometer precision over millimeter-wide sensing gaps for space-borne gravitational reference sensors. *PHYSICAL REVIEW D*, 96(6), SEP 26 2017.
- [P3] M. Armano, H. Audley, G. Auger, J. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, M. de Deus Silva, R. De Rosa, L. Di Fiore, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, R. Flatscher, M. Freschi, C. Garca Marrirrodrga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, A. Grado, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killo, A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, S. Madden, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, D. I. Robertson, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, D. Shaul, J. Slutsky, C. F. Sopena, R. Stanga, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, M. Troebs, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanoni, T. Ziegler, P. Zweifel, and LISA Pathfinder Collaboration. LISA Pathfinder: First steps to observing gravitational waves from space. In *11TH INTERNATIONAL LISA SYMPOSIUM*, volume 840 of *Journal of Physics Conference Series*. LISA, 2017.
- [P4] M. Armano, H. Audley, G. Auger, J. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, M. de Deus Silva, R. De Rosa, L. Di Fiore, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, R. Flatscher, M. Freschi, C. Garcia Marrirrodrga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, A. Grado, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killo, A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, S. Madden, I. Mateos, P. W. McNamara, J. Mendes, L. Mendel, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, D. I. Robertson, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, D. Shaul, J. Slutsky, C. F. Sopena, R. Stanga, T. Sumner, D. Texier, J. I. Thorpe, C. Trenke, M. Troebs, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanoni, T. Ziegler, P. Zweifel, and LISA Pathfinder Collaboration. LISA Pathfinder closed-loop analysis: a model breakdown of the in-loop observables. In *11TH INTERNATIONAL LISA SYMPOSIUM*, volume 840 of *Journal of Physics Conference Series*. LISA, 2017.
- [P5] M. Armano, H. Audley, G. Auger, J. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, M. de Deus Silva, R. De Rosa, L. Di Fiore, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, R. Flatscher, M. Freschi, C. Garca Marrirrodrga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, A. Grado, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé,

- O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, S. Madden, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, D. I. Robertson, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, D. Shaul, J. Slutsky, C. F. Sopena, R. Stanga, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, M. Troebs, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanoni, T. Ziegler, P. Zweifel, and LISA Pathfinder Collaboration. Laser Frequency Noise Stabilisation and Interferometer Path Length Differences on LISA Pathfinder. In *11TH INTERNATIONAL LISA SYMPOSIUM*, volume 840 of *Journal of Physics Conference Series*. LISA, 2017.
- [P6] M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, C. Cutler, K. Danzmann, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, J. Gallegos, C. Garcia Marirrodriga, R. Gerndt, L. I. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, P. Maghami, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. Free-flight experiments in LISA Pathfinder. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P7] M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, M. de Deus Silva, D. Desiderio, E. Piersanti, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, R. Flatscher, M. Freschi, J. Gallegos, C. Garcia Marirrodriga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, R. Tomlinson, C. Trenkel, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, C. Warren, P. J. Wass, D. Wealthy, W. J. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. Constraints on LISA Pathfinder's self-gravity: design requirements, estimates and testing procedures. *CLASSICAL AND QUANTUM GRAVITY*, 33(23), DEC 8 2016.
- [P8] M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, I. Diepholz,

- R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, J. Gallegos, C. Garcia Marirrodriga, R. Gerndt, L. I. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. A noise simulator for eLISA: Migrating LISA Pathfinder knowledge to the eLISA mission. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P9] M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, J. Gallegos, C. Garcia Marirrodriga, R. Gerndt, L. I. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. Disentangling the magnetic force noise contribution in LISA Pathfinder. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P10] M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, J. Gallegos, C. Garcia Marirrodriga, R. Gerndt, L. I. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. The LISA Pathfinder Mission. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P11] M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, J. Gallegos,

- C. Garcia Marirrodriga, R. Gerndt, L. I. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreux-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. A Strategy to Characterize the LISA-Pathfinder Cold Gas Thruster System. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P12] M. Armano, H. Audley, G. Auger, J. T. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, L. Carbone, A. Cavalleri, A. Cesarini, G. Ciani, G. Congedo, A. M. Cruise, K. Danzmann, M. de Deus Silva, R. De Rosa, M. Diaz-Aguilo, L. Di Fiore, I. Diepholz, G. Dixon, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, W. Fichter, E. D. Fitzsimons, R. Flatscher, M. Freschi, A. F. Garcia Marin, C. Garcia Marirrodriga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, F. Guzman, A. Grado, C. Grimani, A. Grynagier, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, U. Johann, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, S. Madden, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, A. Monsky, D. Nicolodi, M. Nofrarias, S. Paczkowski, M. Perreux-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, B. Rais, J. Ramos-Castro, J. Reiche, D. I. Robertson, H. Rozemeijer, F. Rivas, G. Russano, J. Sanjuan, P. Sarra, A. Schleicher, D. Shaul, J. Slutsky, C. F. Sopena, R. Stanga, F. Steier, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, M. Troebs, H. B. Tu, D. Vetrugno, S. Vitale, V. Wand, G. Wanner, H. Ward, C. Warren, P. J. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanoni, T. Ziegler, and P. Zweifel. Sub-Femto-g Free Fall for Space-Based Gravitational Wave Observatories: LISA Pathfinder Results. *PHYSICAL REVIEW LETTERS*, 116(23), JUN 7 2016.
- [P13] M. Armano, H. Audley, G. Auger, J. T. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, R. Flatscher, M. Freschi, J. Gallegos, C. Garcia Marirrodriga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, C. J. Killow, N. Korsakova, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreux-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. I. Robertson, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. Sopena, T. J. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, D. Wealthy, W. J. Weber, A. Wittchen, C. Zanoni,

- T. Ziegler, and P. Zweifel. Charge-Induced Force Noise on Free-Falling Test Masses: Results from LISA Pathfinder. *PHYSICAL REVIEW LETTERS*, 118(17), APR 26 2017.
- [P14] M. Armano, H. Audley, G. Auger, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, C. Garcia Marirrodriga, R. Gerndt, L. I. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzel, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, E. Mitchell, A. Moroni, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanon, T. Ziegler, and P. Zweifel. Bayesian statistics for the calibration of the LISA Pathfinder experiment. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P15] M. Armano, H. Audley, J. Baird, M. Bassan, S. Benella, P. Binetruy, M. Born, D. Bortoluzzi, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, M. Fabi, L. Ferraioli, V. Ferroni, N. Finetti, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzel, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, M. Laurenza, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, F. Sabbatini, J. Slutsky, C. F. Sopena, T. Sumner, D. Telloni, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanon, and P. Zweifel. Characteristics and Energy Dependence of Recurrent Galactic Cosmic-Ray Flux Depressions and of a Forbush Decrease with LISA Pathfinder. *ASTROPHYSICAL JOURNAL*, 854(2), FEB 20 2018.
- [P16] M. Armano, H. Audley, J. Baird, S. Benella, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, M. Fabi, L. Ferraioli, V. Ferroni, N. Finetti, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzel, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, K. Kudela, M. Laurenza, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. Robertson, I. F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Telloni, D. Texier, J. Thorpe, I. D. Vetrugno, M. Villani, S. Vitale, G. Wanner, H. Ward, P. Wass, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Forbush Decreases and < 2 Day GCR Flux Non-recurrent Variations Studied with LISA Pathfinder. *ASTROPHYSICAL JOURNAL*, 874(2), APR 1 2019.

- [P17] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, N. Finetti, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshkar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Measuring the Galactic Cosmic Ray flux with the LISA Pathfinder radiation monitor. *ASTROPARTICLE PHYSICS*, 98:28–37, MAR 2018.
- [P18] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, M-S Hartig, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. Robertson, I. F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, P. Zweifel, and Lisa Pathfinder Collaboration. Novel methods to measure the gravitational constant in space. *PHYSICAL REVIEW D*, 100(6), SEP 20 2019.
- [P19] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, C. Mansanet, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, N. Meshksar, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, J. Sanjuan, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Temperature stability in the sub-milliHertz band with LISA Pathfinder. *MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY*, 486(3):3368–3379, JUL 2019.
- [P20] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson,

- D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, P. Zweifel, and LISA Pathfinder Collaboration. LISA Pathfinder micronewton cold gas thrusters: In-flight characterization. *PHYSICAL REVIEW D*, 99(12), JUN 28 2019.
- [P21] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, P. Zweifel, and LISA Pathfinder Collaboration. LISA Pathfinder platform stability and drag-free performance. *PHYSICAL REVIEW D*, 99(8), APR 16 2019.
- [P22] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, D. Giardini, F. Gibert, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, L. Liu, I. Lloro, J. A. Lobo, J. P. López-Zaragoza, R. Maarschalkerweerd, F. Mailland, D. Mance, V. Martín, L. Martin-Polo, F. Martin-Porqueras, J. Martino, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, N. Meshksar, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, M. Pfeil, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. Robertson, I. F. Rivas, G. Russano, G. Santoruvo, P. Sarra, D. Shaul, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, P. Zweifel, and LISA Pathfinder Collaboration. Precision charge control for isolated free-falling test masses: LISA pathfinder results. *PHYSICAL REVIEW D*, 98(6), SEP 11 2018.
- [P23] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias,

- S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. Robertson, I. F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Calibrating the system dynamics of LISA Pathfinder. *PHYSICAL REVIEW D*, 97(12), JUN 12 2018.
- [P24] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzysimisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I. Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20μ Hz. *PHYSICAL REVIEW LETTERS*, 120(6), FEB 5 2018.
- [P25] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzysimisch, I. Harrison, M-S Hartig, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, N. Meshksar, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, F. Rivas, D. Robertson, I. D. Roma-Dollase, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Telloni, D. Texier, J. I. Thorpe, C. Trenkel, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Spacecraft and interplanetary contributions to the magnetic environment on-board LISA Pathfinder. *MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY*, 494(2):3014–3027, MAY 2020.
- [P26] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzysimisch, I. Harrison, M-S Hartig, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. I. Robertson, F. Rivas, G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. LISA Pathfinder Perfor-

- mance Confirmed in an Open-Loop Configuration: Results from the Free-Fall Actuation Mode. *PHYSICAL REVIEW LETTERS*, 123(11), SEP 11 2019.
- [P27] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, M. de Deus Silva, I. Diepholz, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, A. Lobo, I. Lloro, L. Liu, J.-P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, N. Meshksar, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, D. Robertson, F. Rivas, G. Russano, J. Slutsky, C. Sopena, T. Sumner, D. Texier, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, W. Weber, L. Wissel, A. Wittchen, and P. Zweifel. LISA Pathfinder. *arXiv e-prints*, page arXiv:1903.08924, 2019.
- [P28] M. Armano, H. Audley, J. Baird, M. Born, D. Bortoluzzi, N. Cardines, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, G. Dixon, R. Dolesi, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, M. Freschi, L. Gesa, D. Giardini, F. Gibert, R. Giusteri, C. Grimani, J. Grzymisch, I. Harrison, M-S Hartig, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, N. Karnesis, B. Kaune, C. J. Killow, N. Korsakova, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, I. Mateos, P. W. McNamara, J. Mendes, L. Mendes, N. Meshksar, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, J. Ramos-Castro, J. Reiche, F. Rivas, D. Robertson, I. G. Russano, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. ten Pierick, J. I. Thorpe, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. J. Wass, W. J. Weber, L. Wissel, A. Wittchen, and P. Zweifel. Analysis of the accuracy of actuation electronics in the laser interferometer space antenna pathfinder. *REVIEW OF SCIENTIFIC INSTRUMENTS*, 91(4), APR 1 2020.
- [P29] F. Gibert, M. Nofrarias, M. Armano, H. Audley, G. Auger, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, A. Bursi, M. Caleno, A. Cavalleri, A. Cesarini, M. Cruise, K. Danzmann, I. Diepholz, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. Fitzsimons, M. Freschi, J. Gallegos, C. Garcia Marirrodrga, R. Gerndt, L. I. Gesa, D. Giardini, R. Giusteri, C. Grimani, I. Harrison, G. Heinzl, M. Hewitson, D. Hollington, M. Hueller, J. Huesler, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. Killow, I. Lloro, R. Maarschalkerweerd, S. Madden, P. Maghami, D. Mance, V. Martín, F. Martin-Porqueras, I. Mateos, P. McNamara, J. Mendes, L. Mendes, A. Moroni, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, J. A. Romera Perez, D. Robertson, H. Rozemeijer, G. Russano, P. Sarra, A. Schleicher, J. Slutsky, C. F. Sopena, T. Sumner, D. Texier, J. I. Thorpe, C. Trenkel, H. B. Tu, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, S. Waschke, P. Wass, D. Wealthy, S. Wen, W. Weber, A. Wittchen, C. Zanoni, T. Ziegler, and P. Zweifel. In-flight thermal experiments for LISA Pathfinder: Simulating temperature noise at the Inertial Sensors. In *10TH INTERNATIONAL LISA SYMPOSIUM*, volume 610 of *Journal of Physics Conference Series*, 2015.
- [P30] Gerhard Heinzl, Martin Hewitson, Michael Born, Lennart Wissel, Brigitte Kaune, Gudrun Wanner, Karsten Danzmann, Sarah Paczkowski, Michael Tröbs, Andreas Wittchen, Nikos

Karnesis, Heather Audley, and Jens Reiche. LPF final report for the German contribution to the nominal mission. Technical report, Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Hannover, 2018.

- [P31] J. I. Thorpe, J. Slutsky, John G. Baker, Tyson B. Littenberg, Sophie Hourihane, Nicole Pagane, Petr Pokorny, Diego Janches, M. Armano, H. Audley, G. Auger, J. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, A. Cavalleri, A. Cesarini, A. M. Cruise, K. Danzmann, M. de Deus Silva, R. De Rosa, L. Di Fiore, I. Diepholz, G. Dixon, R. Dolesi, N. Dunbar, L. Ferraioli, V. Ferroni, E. D. Fitzsimons, R. Flatscher, M. Freschi, C. Garcia Marirrodriga, R. Gerndt, L. Gesa, F. Gibert, D. Giardini, R. Giusteri, A. Grado, C. Grimani, J. Grzymisch, I Harrison, G. Heinzl, M. Hewitson, D. Hollington, D. Hoyland, M. Hueller, H. Inchauspé, O. Jennrich, P. Jetzer, B. Johlander, N. Karnesis, B. Kaune, N. Korsakova, C. J. Killow, J. A. Lobo, I Lloro, L. Liu, J. P. López-Zaragoza, R. Maarschalkerweerd, D. Mance, V. Martín, L. Martin-Polo, J. Martino, F. Martin-Porqueras, S. Madden, I Mateos, P. W. McNamara, J. Mendes, L. Mendes, M. Nofrarias, S. Paczkowski, M. Perreur-Lloyd, A. Petiteau, P. Pivato, E. Plagnol, P. Prat, U. Ragnit, J. Ramos-Castro, J. Reiche, D. Robertson, I, H. Rozemeijer, F. Rivas, G. Russano, P. Sarra, A. Schleicher, D. Shaul, C. F. Sopena, R. Stanga, T. Sumner, D. Texier, C. Trenkel, M. Troebs, D. Vetrugno, S. Vitale, G. Wanner, H. Ward, P. Wass, D. Wealthy, W. J. Weber, L. Wissel, A. Wittchen, A. Zambotti, C. Zanoni, T. Ziegler, P. Zweifel, P. Barela, C. Cutler, N. Demmons, C. Dunn, M. Girard, O. Hsu, S. Javidnia, I Li, P. Maghami, C. Marrese-Reading, J. Mehta, J. O'Donnell, A. Romero-Wolf, J. Ziemer, LISA Pathfinder Collaboration, and ST7-DRS Operations Team. Micrometeoroid Events in LISA Pathfinder. *ASTROPHYSICAL JOURNAL*, 883(1), SEP 20 2019.

Curriculum Vitae

Personal Details

Name Sarah Katharina Paczkowski
Nationality German
Date of birth 30.03.1988
Place of birth Hannover

Education and Qualification

2014 – 2020 **PhD student**, *AEI & Leibniz Universität Hannover*.
2011 – 2014 **MSc Physics**, *Leibniz Universität Hannover*, Thesis: Modelling the S/C and TM Motion for LISA.
2007 – 2011 **BSc Physics**, *Leibniz Universität Hannover*, Thesis: Schwarze Löcher in Galaxien.
June 2007 **Abitur (German general qualification for university entrance)**, *Gymnasium Neustadt am Rübenberge*.