

Absolute distance interferometry capable of long-term high frequency measurements of fast targets

Andrew James Lancaster

Balliol College, Oxford



Thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2015

Absolute distance interferometry capable of long-term high frequency measurements of fast targets

Andrew James Lancaster
Balliol College, Oxford

Thesis submitted in partial fulfilment of the requirements for the
degree of Doctor of Philosophy at the University of Oxford
Trinity Term, 2015

Abstract

This thesis presents two novel absolute distance measurement techniques which expand the recently developed dynamic frequency scanning interferometry method [1][2][3]. Preliminary work to improve the frequency referencing of the techniques will also be reported.

The first method was developed to eliminate gaps in the distance measurement which were present in previous approaches and is referred to as continuous frequency scanning interferometry (continuous FSI). This technique has been successfully demonstrated on stationary targets at a range of 2.5–5 m and moving targets at speeds around 2 mm/s and 8 mm/s.

The second method is referred to as enhanced continuous FSI (enhanced CFSI). It addresses one of the main limitations of the continuous FSI technique by greatly increasing the allowable target motion speed. Proof of principle measurements of enhanced continuous FSI have been performed at speeds of up to 16 mm/s and are presented here.

The improved frequency referencing technique uses an optical nanofibre to perform saturated absorption spectroscopy of acetylene-13. The equipment, experimental procedure and preliminary results are presented.

Acknowledgements

I would like to acknowledge those who have helped me throughout the course of my DPhil. Starting with the research group I would like to thank John Dale for teaching me a huge amount of programming and practical know how, Patrick Brockill for his help with CUDA, Paul Coe for his help with doppler-free spectroscopy and of course Armin Reichold for his supervision and encouragement.

From NPL I would like to thank Ben Hughes, Andrew Lewis, Mike Campbell and Matt Warden. Gilberto Brambilla, from the ORC at Southampton, generously allowed the use of his nanofibre drawing rig, and his students Imran Mustafa and Rand Ismaeel who patiently taught me how to make the fibres. I would also like to thank the physics workshop and support staff for manufacturing the various custom pieces of equipment. Specific thanks go to James Lynn for his nanofibre cell designs and Russell Allcock for his patience dealing with an incredibly unreliable server.

Heinrich von Jagwitz-Biegnitz, Muhammad Kasim and Jimmy Holloway have put up with my constant fidgeting while sharing the office, and also provided much needed breaks. Ian Coulter, Nathan Hinton, Rahul Nath, Martha Imprialou, Duncan Money, Mikey Davis, Seb Pike, Stephen Gaw, Phoebe Downing, Georgina Lang and Gabi de Wit have kept me sane when things started to break and made Oxford a fun place to be, while Phil Rhind and Adam Baggett provided much needed distractions while I was away.

Finally I would like to thank my family, who have been incredibly supportive and put up with a disproportionate amount of grumpiness.

Thank you all.

The author gratefully acknowledges financial support for his work from STFC and the National Physical Laboratory under the auspices of studentship ST/I000526/1.

This thesis was written using the $\text{\LaTeX} 2_{\epsilon}$ typesetting package. Graphs and diagrams were produced using Matlab, Scavis and Inkscape. This thesis and the research it describes are original work carried out solely by the named author.

Copyright, Andrew Lancaster, 2015.

All rights reserved, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the express permission of the author.

Published at the University of Oxford, United Kingdom.

Contents

1	Introduction	1
1.1	Optical absolute distance measurement	1
1.1.1	Frequency Scanning Interferometry	2
1.1.2	Multi-wavelength Interferometry	4
1.1.3	Time of flight measurements	5
1.1.4	Modulation rangefinding	5
1.1.5	Frequency combs	6
1.2	Applications in Particle Physics	6
1.2.1	Monitoring of the ATLAS semiconductor tracker	7
1.2.2	Surveying the reference network for the ILC	8
1.2.3	Monitoring the final focus quadrupole magnets at the ILC.	10
1.3	Recent FSI developments at Oxford	11
2	Continuous FSI - Theory	13
2.1	The limitations of dynamic FSI	13
2.2	Continuous FSI	16
2.2.1	Schematic layout of a continuous FSI system	16
2.2.2	Derivation of the algorithm for length calculation	18
2.2.3	Laser frequency calculation using a gas cell	20
2.2.4	Performing a hand-over and moving to continuous FSI	22
2.2.5	Using continuous FSI for fast feedback	23

3	Continuous FSI - Implementation	24
3.1	Experimental layout	24
3.1.1	Frequency scanning lasers	25
3.1.2	Optical switches	25
3.1.3	Master splitter box	27
3.1.4	Front-end splitter tree	28
3.1.5	Interferometers	29
3.1.6	Gas cell	31
3.1.7	Photodiodes and data acquisition	32
3.1.8	Computing	34
3.2	Implementation of the analysis	36
3.2.1	Overview	37
3.2.2	Phase extraction	39
3.2.3	Gas cell normalisation and filtering	48
3.2.4	Absorption peak finding	51
3.2.5	Length calculation	55
3.3	Parameter selection	57
3.3.1	Separability constraint	57
3.3.2	Timing constraint	59
4	Continuous FSI - Verification	61
4.1	Absolute calibration	62
4.1.1	The reference system	62
4.1.2	Experimental layout	63
4.1.3	Data integrity	66
4.1.4	Calculation of the calibration constants	67
4.1.5	Experimental uncertainties	70
4.2	ADC sampling frequency synchronisation	78

4.2.1	Experimental layout	78
4.2.2	Time axis and sampling frequency synchronisation	80
4.3	Differential length calibration	88
4.3.1	Experimental layout	88
4.3.2	Data alignment in time	89
4.3.3	Comparison of the measurements	90
5	Enhanced Continuous FSI	97
5.1	Speed limitations	97
5.1.1	Fringe frequency	98
5.1.2	Motion tolerance without round trips in the interferometer	99
5.1.3	Round trips in the interferometer	100
5.1.4	Motion tolerance with broadened frequency spectra	104
5.2	Incorporating a fixed frequency laser	107
5.2.1	Implementation of enhanced CFSI - equipment	110
5.2.2	Implementation of enhanced CFSI - analysis	113
5.2.3	Enhanced CFSI experimental demonstration	115
6	Improved frequency referencing	126
6.1	Overview of doppler-free saturated absorption spectroscopy	126
6.2	Choice of target gas	128
6.3	Optical nanofibres	130
6.3.1	Hollow core fibres	132
6.4	Equipment for saturated absorption	133
6.4.1	Acetylene storage cylinder	133
6.4.2	Nanofibre gas cell	136
6.4.3	Nanofibre fabrication and transfer	139
6.5	Experimental results	143
6.5.1	Doppler broadened acetylene absorption features	143

6.5.2	Doppler-free saturated acetylene absorption features	145
6.5.3	Nanofibre fragility	148
7	Summary and future work	152
7.1	Interferometric techniques	152
7.1.1	Experimental developments	153
7.1.2	Analysis developments	155
7.2	Saturated absorption spectroscopy	156
A	Length difference plots	158

List of Figures

1.1	A schematic of a Michelson interferometer.	3
1.2	The layout of the International Linear Collider.	9
1.3	The LiCAS RTRS system.	10
1.4	The MONALISA measurement interferometer launch head.	11
2.1	A typical laser scanning pattern for dynamic FSI.	14
2.2	A typical laser scanning pattern for continuous FSI.	15
2.3	Schematic layout of a continuous FSI measurement system.	16
2.4	Schematic layout of a gas cell based frequency measurement system.	17
3.1	Schematic of the full experimental setup for continuous FSI.	26
3.2	The two TLM-8700 lasers.	27
3.3	The MEMS switch box.	28
3.4	Schematic of the front end splitter tree.	29
3.5	A schematic of a fibre-fed Fizeau interferometer.	30
3.6	The gas cell as mounted for continuous FSI.	31
3.7	The DAQ crate configured for continuous FSI.	33
3.8	Gas cell data taken using the LiCAS photodiodes and fibre circulators.	35
3.9	Gas cell data taken using the Thorlabs photodiodes and fibre circulators.	35
3.10	Gas cell data taken using the Thorlabs photodiodes and PLC splitters.	36
3.11	Schematic of a continuous FSI analysis iteration.	38
3.12	A subsection of an interference buffer.	40

3.13	Magnitude of the FFT of the interference data buffer.	42
3.14	Threshold values and peak limits as found by the auto-filter algorithm.	44
3.15	Combined filter and Hilbert windows.	44
3.16	Wrapped phase calculated using the Hilbert transform.	45
3.17	Unwrapped phase.	46
3.18	Spurious oscillations at the start and end of the extracted phase.	47
3.19	Chaining the current phase results to those in the previous iteration.	47
3.20	Difference in phase between the current and previous iterations.	48
3.21	Raw gas cell and power data, along with the predicted gas cell data.	51
3.22	Normalised and filtered gas cell data over a pair of laser scans.	52
3.23	Normalised and filtered gas cell data over a single gas cell feature.	52
3.24	Example fit region for the gaussian gas cell fits, highlighted in red.	54
4.1	Experimental layout for the absolute calibration experiments.	63
4.2	Average fringe frequency spectra for a series of interferometers.	65
4.3	Example fringe frequency spectrum for position 7, calibration 1.	68
4.4	Standard deviations of the continuous FSI distance measurements.	70
4.5	Standard deviations of the LaserTRACER distance measurements.	71
4.6	Measured environmental parameters during the calibration experiments.	73
4.7	Changes in the refractive index ratio.	74
4.8	Residuals from the straight line fits for all calibration sets.	75
4.9	Residuals from the straight line fits, ignoring the end of of set 5.	75
4.10	Calibration values using the available data as described in table 4.3.	76
4.11	Calibration values calculated using data which ignored position 7 values.	77
4.12	Experimental layout for the time base calibration experiment.	79
4.13	Scan of central time and stretch parameter for calibration 1.	83
4.14	Scan of central time and stretch parameter for calibration 2.	83
4.15	Scan of central time and stretch parameter for calibration 3.	84

4.16	Scan of central time and stretch parameter for calibration 4.	84
4.17	Oscillation measurements aligned using the fit parameters.	86
4.18	Alignment of the oscillation data at activation.	87
4.19	Alignment of the oscillation data at the end of a set.	87
4.20	Combined result for the stretch parameter.	88
4.21	Subset of the region used to align the FSI and LaserTRACER data. . . .	90
4.22	T_{centre} values for the forward experiments.	91
4.23	T_{centre} values for the backward experiments.	91
4.24	Alignment of the FSI and LaserTRACER data in a forward experiment. . . .	92
4.25	Alignment of the FSI and LaserTRACER data in a backward experiment. . . .	92
4.26	Difference between FSI and LaserTRACER measurements (forward).	94
4.27	Difference between FSI and LaserTRACER measurements (backward).	95
4.28	Drift in forward experiment 4 as the transfer length is adjusted.	96
5.1	Idealised spectra from a 2 m interferometer without round trips.	101
5.2	Idealised spectra from a 2 m interferometer with two round trips.	103
5.3	Simulated spectra from a 2 m interferometer with two round trips.	106
5.4	Simulated spectra including a 1540 nm fixed-frequency laser.	108
5.5	Schematic of the full experimental setup for enhanced continuous FSI.	111
5.6	Clarification of the beat signal.	114
5.7	Examples of the beat signals observed after filtering.	115
5.8	Laser frequency results calculated during an enhanced CFSI run.	116
5.9	A histogram of the laser frequency results during an enhanced CFSI run.	116
5.10	Laser frequencies calculated during the FFI demonstration experiment.	118
5.11	Behaviour of the CFSI, FFI and LaserTRACER measurements.	118
5.12	Difference between the FSI / FFI systems and the LaserTRACER.	119
5.13	Gradients from fits while monitoring a moving target.	121
5.14	Demonstration of switching between continuous FSI and FFI systems.	123

5.15	Demonstration of scanning laser deactivation.	125
6.1	Saturated absorption features observed in the ^{85}Rb	128
6.2	The ν_1 and ν_3 vibration modes of acetylene.	129
6.3	Unnormalised intensity of light guided by an optical nanofibre.	131
6.4	Schematic of an optical nanofibre.	132
6.5	The acetylene storage cylinder and associated vacuum system.	134
6.6	Schematic of the acetylene storage system and associated vacuum layout.	134
6.7	External view of the nanofibre gas cell.	136
6.8	Internal view of the nanofibre gas cell.	137
6.9	Overview drawing of the nanofibre gas cell	138
6.10	Clamps used to hold the optical nanofibre in place.	139
6.11	The ORC nanofibre drawing rig with the Oxford clamps in place.	140
6.12	An overview of the nanofibre drawing process.	142
6.13	The transfer bar used to move the nanofibre to the gas cell.	143
6.14	Schematic of the experiment to observe doppler-broadened features.	144
6.15	The P16 acetylene-13 absorption feature, observed with a nanofibre.	146
6.16	The P16 acetylene-13 absorption feature at range of pressures.	146
6.17	Schematic of the experiment for doppler-free saturated absorption.	147
6.18	Scans across a potential saturated absorption feature.	149
6.19	A broken nanofibre illuminated using a red laser	150
A.1	Difference between FSI and LaserTRACER measurements (forward).	159
A.2	Difference between FSI and LaserTRACER measurements (backward).	160
A.3	Difference between the FSI / FFI systems and the LaserTRACER.	161

List of Tables

3.1	Laser tuning parameters chosen for continuous FSI.	58
4.1	Atmospheric sensor parameters.	66
4.2	Error propagation into the refractive index ratio.	72
4.3	Number of measurement points in the absolute calibration data sets. . .	76
4.4	Number of measurement points in the final absolute calibration data sets.	77
4.5	Final absolute calibration constant.	78
4.6	Piezo actuator parameters during the ADC synchronisation experiments.	80
4.7	Stretch parameters and associated uncertainties.	82
4.8	Central times and associated uncertainties.	83
4.9	Final stretch parameter and associated uncertainties.	85
5.1	Results from the differential gradient calculations.	120

Chapter 1

Introduction

This chapter will provide a brief overview of optical distance measurement techniques. It will focus on Frequency Scanning Interferometry (FSI), the technique which is utilised in the remainder of this thesis, but will also cover competing techniques and describe the applications of optical absolute distance measurement. The applications of FSI which led to interest in the method in Oxford will also be described. Recent advances will be described in order to show where the work presented in this thesis will widen the uses of FSI.

1.1 Optical absolute distance measurement

Optical systems capable of measuring absolute distance, rather than displacement of a target, have been studied since at least 1898 when Benoit [4] used the method of excess fractions to calculate the optical path difference of an interferometer illuminated by light from a series of spectral lines. The accuracy of the technique was a fraction of the wavelength of the light over an optical path difference of up to 10 cm [5]. More details of the technique are given in section 1.1.2. Since that time many other techniques have been developed including, but not limited to, Frequency Scanning Interferometry [6] (see section 1.1.1), time of flight measurements [7] (see section 1.1.3) and modulation

rangefinding [8] (see section 1.1.4). Other techniques are detailed in [7][9][10].

Absolute distance measurements are of interest because the target does not have to be moved from a known position to its destination, as is the case with displacement measurements, making them suitable for measuring the distance between objects. The measurements can recover quickly when interrupted, whereas a differential system would require the target to be returned to the known position before restarting the measurements. Absolute measurements also allow a target to be monitored over long periods without continuous measurement, because they do not need to record changes from the initial position.

Within an industrial environment absolute distance measurements can play an important role. One common use is to re-establish the absolute distance after a beam break for differential systems such as laser trackers [10]. Other applications include long term monitoring of large structures and scientific applications (see section 1.2).

1.1.1 Frequency Scanning Interferometry

The absolute distance measurement technique used throughout this thesis is Frequency Scanning Interferometry (FSI), which determines the optical path difference L of an interferometer such as the Michelson interferometer [11] shown in figure 1.1.

The light in each arm of the interferometer is coherent with the light in the other but experiences a different optical path. When the light recombines at the beam splitter it will interfere and so the detector will observe an irradiance I of the form given in equation 1.1 [12].

$$I = I_1 + I_2 + \sqrt{I_1 I_2} \cos\left(\frac{2\pi\nu L}{c}\right) \quad (1.1)$$

Here I_1 and I_2 are the irradiances in each arm of the interferometer, ν is the frequency of the light and c is the speed of light. The term $2\pi\nu L/c$ is known as the phase of the interferometer and will be denoted by ϕ .

The principle of FSI is that by scanning the frequency of the light by an amount $\Delta\nu$

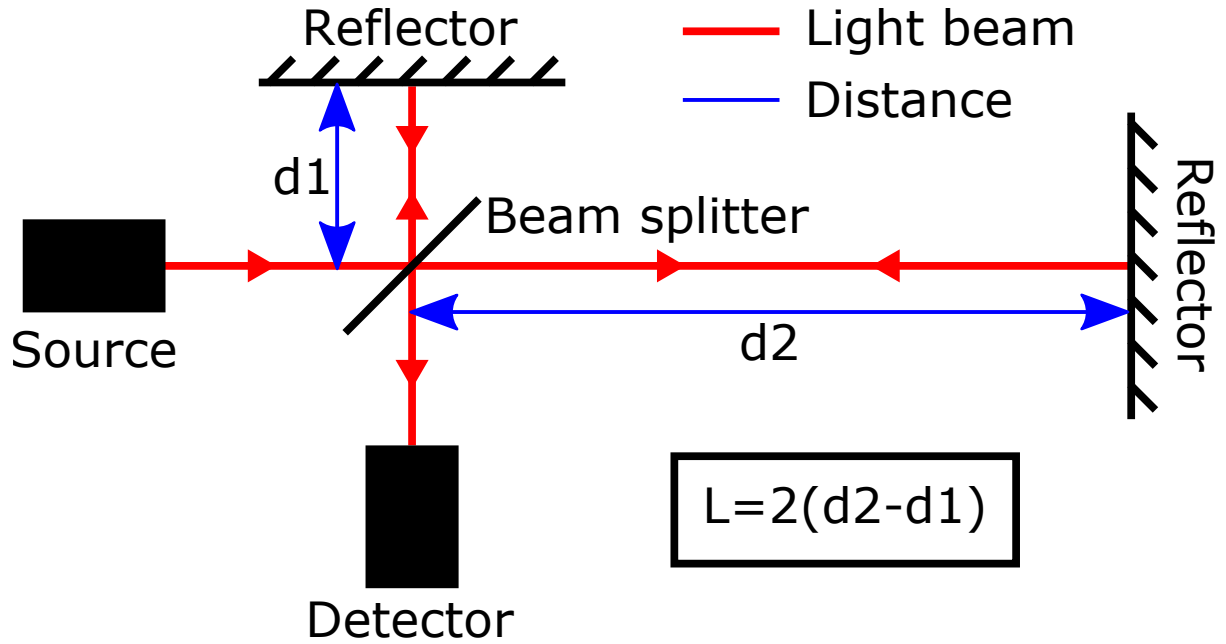


Figure 1.1: A schematic of a Michelson interferometer.

we can induce a phase change $\Delta\phi$ in the interferometer and so calculate the optical path difference from equation 1.2.

$$\Delta\phi = \frac{2\pi}{c} L \Delta\nu \quad (1.2)$$

There are two main approaches to calculate the optical path difference. The first is to measure the phase and frequency changes and calculate the OPD directly. This was the approach taken by Barwood et. al., who measured the optical path length of a stabilised 1 m Fabry-Pérot etalon with a relative accuracy of approximately 3 parts in 10^{10} [6]. The measurement was performed by locking a laser to one of the etalon peaks and measuring the frequency of its beat with a second laser locked to a saturated rubidium D_2 line. The lasers were then scanned over a range of approximately 7 THz, the first locked to another etalon peak and its frequency measured again with respect to the second laser, now locked to the rubidium D_1 line. These measurements provide the differences in phase and laser frequency required. When FSI is performed in this way the length scale is derived from the measurement of the laser frequency.

The second, more common [13][14][15], method of calculating the unknown optical

path difference is to simultaneously measure two interferometers, one of which (known as the reference interferometer) has a pre-determined OPD. By dividing the phase difference equations for the two interferometers we can cancel the frequency change term and so retrieve equation 1.3.

$$\frac{\Delta\phi_M}{\Delta\phi_R} = \frac{L_M}{L_R} \quad (1.3)$$

Here subscript M refers to the (unknown) measurement interferometer and subscript R refers to the reference interferometer. L_M is the only unknown in equation 1.3, and so can be directly extracted. In this implementation the length scale of the system comes from the length of the reference interferometer, and so is dependant on the technique used to determine its length.

Frequency scanning interferometry as described above is vulnerable to drifts of the interferometer length over the measurement period. Two methods of correcting for this drift are the inclusion of a second, wavelength stabilised, laser to act as a counting interferometer which can correct the FSI phase change [16], or to include a second laser scanning its frequency in the opposite direction to the first [17][18]. The two scanning lasers experience opposite phase drifts from interferometer length changes and so can be averaged to provide a drift-free measurement.

1.1.2 Multi-wavelength Interferometry

Multi-wavelength interferometry is the technique which was proposed by Benoit [4]. It works by illuminating the measurement interferometer with a series of wavelengths λ_i and measuring the fractional excess ϵ_i , as defined by equation 1.4 [19] where a_i are integers, for each.

$$L = a_i \left(\frac{\lambda_i}{2} \right) + \epsilon_i \quad (1.4)$$

With a series of N wavelengths we have N+1 unknowns (L and a_1 - a_N), in a system of equations which are solvable because all a_i are integers.

The light for multi-wavelength systems can be generated in a variety of ways. Benoit

used a series of spectral lines, but modern systems such as [20] often use frequency stabilised lasers. The wavelengths of light used in the measurement provide the scale for the measurement.

1.1.3 Time of flight measurements

Time of flight measurements are performed by observing the time taken for a laser pulse to travel from a transmitter to the target and then return to a detector. The distance travelled by the pulse is then calculated using the speed of light. It is a technique which was initially developed for surveying and military uses [7].

While the time of flight measurement technique is conceptually simple it can be technically challenging to achieve high accuracies. This is because the speed of light demands high-accuracy timing systems. Amann et. al. describe several causes of error in the timing system including noise-based jitter and variations in the shape and amplitude of the laser pulse when it returns from the target. The timing system provides the scale for the measurement.

1.1.4 Modulation rangefinding

In modulation rangefinding a laser beam is subjected to an external modulation, for example of its intensity [8]. The modulated beam propagates to the target and returns, with the modulation acquiring a phase shift relative to the outgoing beam of $(2d/\lambda_{mod}) \times 2\pi$, where d is the distance to the target and λ_{mod} is the wavelength of the applied modulation. By measuring this phase shift it is possible to accurately determine the distance to the target reflector modulo an integer number of half-wavelengths, because the phase will wrap every 2π radians. As an example, the 1.5 GHz modulation frequency used in [8] causes the phase to repeat for every 10 cm to the target. This ambiguity could be resolved by a change of modulation frequency. The length scale of the measurement comes from the timing in the phase measurement system.

1.1.5 Frequency combs

The use of frequency combs for absolute distance measurement is an area of active research, trying to take advantage of the large number of accurately determined optical frequencies they provide. Frequency combs are being used both to monitor the frequency of lasers used in frequency modulated continuous wave LADAR [21] (a modulation ranging technique in which the optical frequency is varied, rather than the intensity as described in section 1.1.4) and as a source for multi-wavelength interferometry [22].

Many of the multi-wavelength techniques take advantage of other features of the frequency comb to enhance their measurements. Examples include using time of flight measurements of the comb pulses to break the half-wavelength ambiguity [23], using the phase measurements for a large number of frequencies to break the ambiguity by calculating the distance using spectral interferometry [24] and using widely spaced frequency modes to correct for the refractive index of the air [25]. It has been predicted that frequency comb systems could provide measurement accuracies of 2 parts in 10^{13} at long (≈ 30 km) ranges [23].

1.2 Applications in Particle Physics

Frequency Scanning Interferometry in a variety of forms has been used in Oxford for a series of high-energy and accelerator physics applications. These are monitoring the geometry of the semiconductor tracker in the ATLAS detector [26], a robotic surveying device for the tunnel reference network in the International Linear Collider (ILC) [27] and a system for monitoring the position and orientation of the final focus quadrupole magnets, again for the ILC [28]. The following sections will briefly describe those applications and explain the evolution of the FSI technique.

1.2.1 Monitoring of the ATLAS semiconductor tracker

The ATLAS detector is one of the four main particle physics experiments at the Large Hadron Collider [29]. The semiconductor tracker is one of the sub-detectors within the experiment, using 4,088 independent silicon detector modules to measure the trajectories of ionizing particles [26]. More information on the semiconductor tracker can be found in [30].

To maintain the performance of the tracker changes in the relative positions of the silicon modules must be recorded. A FSI system capable of measuring individual distances between points on the trackers support structure to $1\ \mu\text{m}$ precision and thus providing 3-dimensional tracker shape corrections for the detection elements with $10\ \mu\text{m}$ precision was implemented [26]. The system provides simultaneous measurements of the 842 interferometers arranged as a set of interlinked geodetic grids within the high-radiation detector environment. The measurement interferometers consist of two fibres with a beam splitter mounted on the front, and are pointed at a retroreflector. One fibre feeds the interferometer and the second returns the pW-level interference signal to the multiplexed read-out system.

The length measurements were performed using two lasers which scan in opposite directions to correct for length drifts, one of which can tune over 818–851 nm and the second can fine-tune near 836 nm, providing a single measurement per interferometer every 8 minutes. Optical choppers, phase locked at 50Hz, are used to ensure that only one laser illuminates the interferometers during any given data acquisition period (lasting for 10 ms). The frequency scanning range was extended by linking a series of smaller frequency scans together with larger and faster coarse tuning scans.

The reference scale for the measurements is provided by a pair of evacuated reference interferometers with a Super-Invar rod frame because of its low coefficient of thermal expansion [31]. The reference interferometers have been designed to have stable lengths over long time scales to enable high precision measurements, although high absolute ac-

curacy in the length is not a requirement because the objective is to measure changes in the detector shape. The reference interferometers have a four-fibre readout system to perform phase measurements. The length of the measurement interferometers is computed by fitting their interference signals against the reference phase.

The ATLAS FSI system began development in 1995 [32], with first in-situ data and commissioning in 2008. The system is still in use and being upgraded, for example with *vibrato* FSI which allows more rapid changes in the detector to be monitored [33] and is sensitive to sub-50 nm changes in interferometer length.

1.2.2 Surveying the reference network for the ILC

The International Linear Collider is a proposed particle accelerator which will collide electrons and positrons at a centre-of-mass energies between 200–500 GeV to perform a precision study of the Higgs boson and its couplings to other particles, along with studies of the top quark and searches for new physics such as supersymmetry. The machine is expected to have a total footprint of approximately 31 km [34], with the components required to be aligned within tight tolerances (ranging from 0.1 mm RMS over 150 m in the electron and positron sources to 0.02 mm over 200 m in the beam delivery system). A diagram of the ILC is given in figure 1.2. To achieve this the ILC will be aligned absolutely against a network of reference markers in the accelerator tunnel and relatively against itself, and later to the particle beams. The tunnel markers would necessarily take a significant effort to manually survey [35]. The Linear Collider Alignment and Survey (LiCAS) Rapid Tunnel Reference Surveyor (RTRS) was proposed as a method to automate the survey process, and is described in detail in [27][35][36][37].

The LiCAS RTRS consists of a series of linked measurement cars which move along a rail on the tunnel wall. These cars measure their positions relative to each other and also measure the position of the reference markers (one marker per car). When the measurements at one RTRS position have been completed the system moves on by one reference




Figure 1.2: The layout of the International Linear Collider. Image reproduced with permission from the ILC Technical Design Report [34].

marker and measures again, which provides a series of overlapping measurements. A diagram of the LiCAS RTRS is shown in figure 1.3

Two FSI systems are used in the prototype RTRS, the *internal* and *external* systems. The external system measures the position of the reference markers with respect to the measurement car using uncollimated fibre-fed Fizeau interferometers. The interferometers are uncollimated to allow the target retroreflector to be significantly offset from the pointing axis of the fibre. The 3-dimensional position of the reference markers are found using multilateration. The internal FSI system measures the distance, and two rotations around the transverse axis, between measurement cars. The internal system uses collimated Fizeau interferometers to perform the measurements. Note that in both internal and external FSI systems the interferometers use only a single fibre, rather than the two fibre system used by ATLAS.

There are many other differences between the LiCAS and ATLAS FSI systems. The LiCAS system uses only a single laser, scanning from 1535 nm to 1565 nm at 40 nm/s. The frequency scan was contiguous, and did not require the linking of smaller scans




Figure 1.3: The LiCAS RTRS system. Image courtesy of Dr. A Reichold. LSM stands for laser straightness monitor.

as was the case in the ATLAS system. All measurement and reference interferometers, again constructed from Invar, were sampled simultaneously at 2.77 MHz. The reference interferometers use only one fibre rather than four for readout, and the measurement length is calculated using a Lomb periodogram of filtered measurement interferometer data. Overall the LiCAS FSI system is significantly simpler than the ATLAS system. The external FSI precision was below $2.43 \mu\text{m}$, and the internal FSI precision below $0.521 \mu\text{m}$.

1.2.3 Monitoring the final focus quadrupole magnets at the ILC.

The Monitoring Alignment and Stabilisation with high Accuracy (MONALISA) project developed a system to monitor the final focus quadrupole magnets for the International Linear Collider [28]. These magnets require nanometre-level relative vertical stability in order to focus and collide the nominally 5.9 nm high particle bunches [34]. When particle beams are present beam-based position corrections of the beam are anticipated to correct for misalignment of the final focus quadrupoles, but the alignment could change when the particle beams are not present. The MONALISA system aims to provide a way to

monitor displacements of the magnets which is independent of the particle beams.

The MONALISA system planned to incorporate both a FSI system, which measures the absolute distance between components, and a fixed frequency interferometry (FFI) system, to provide displacement measurements. In an FFI system displacements are measured by monitoring the interferometric phase from a fixed frequency laser, which can be converted into changes in distance through the definition of phase given in section 1.1.1.

To enable the FFI measurements the measurement interferometer was designed to have a multi-fibre readout system, shown in figure 1.4. Using such a system the FFI measurements achieved a 0.2 nm RMS over 1 second, rising to 1 nm over 1 minute [38].




Figure 1.4: The MONALISA measurement interferometer launch head.

The FSI system, which used a similar analysis to that used in the LiCAS experiments [39] and was capable of distance measurements with a 70 nm standard deviation (again of a 400 mm interferometer) [28], was incorporated to provide absolute length measurements. The combination of the two measurement systems provides the powerful combination of continuous, high frequency displacement measurements with the ability to redetermine the absolute target position in the case of measurement interruption.

1.3 Recent FSI developments at Oxford

Towards the end of the MONALISA project Warden developed dynamic FSI [1][2][3]. This is a technique which uses a pair of frequency scanning lasers, tuning in opposite

directions, to calculate the length of the measurement interferometer at every point in the data acquisition period. The technique allows the measurement interferometer to change its length throughout the acquisition period and provides the user with a much greater density of information about the interferometer length compared to the methods described previously, which return only a single length measurement per laser scan. Dynamic FSI will be described in detail in chapter 2 as it forms the basis of the work presented in this thesis.

The dynamic FSI technique was used extensively throughout the AMULET project, which sought to apply the FSI expertise developed in Oxford to industrial applications. The project developed a compact fibre reference interferometer system, incorporated Dynamic FSI into a laser tracer and developed a commercial multi-line FSI system in conjunction with Etalon A.G. [40].

During the AMULET project it became apparent that dynamic FSI could be applied to more applications, such as real-time feedback for the control of coordinate measuring machines, if its measurement period could be extended and target motion tolerance increased. This work aims to address these problems. Chapters 2 to 4 describe the theory and testing of the continuous FSI technique which potentially enables measurements over extended time periods. Chapter 5 describes the enhanced continuous FSI technique which increases the motion tolerance of the measurements through the inclusion of a fixed frequency laser without requiring a change in interferometer layout as was necessary in the MONALISA project. Chapter 6 details work to improve the frequency referencing of the measurements through the use of a saturated acetylene absorption system and chapter 7 summarises the progress which has been made and suggests areas for future work.

Chapter 2

Continuous FSI - Theory

This chapter will give the motivation and theoretical overview of continuous frequency scanning interferometry (continuous FSI). The first section will detail the limitations of dynamic frequency scanning interferometry (dynamic FSI) as described by Warden [1][2] and Dale [3], and explain how continuous FSI addresses these problems. The second section will describe an idealised continuous FSI system and the theory needed to calculate the measured length.

2.1 The limitations of dynamic FSI

As discussed in chapter 1, dynamic FSI is a technique which uses a pair of frequency scanning lasers to measure the unknown optical path difference (OPD) of an interferometer. The two lasers scan their frequencies in opposite directions simultaneously, and the technique can determine the OPD of interest at every data sample, with typically millions of samples per laser scan. Dynamic FSI has been demonstrated to reach relative uncertainties of 0.41×10^{-6} at the two-sigma ($k=2$) level, with a distance resolution of 40 nm within each scan [3].

Dynamic FSI is clearly a powerful technique, but it has some disadvantages. The first of these is that at the end of each laser scan the measurement of the OPD stops,

giving the end user a series of high-resolution measurements with gaps in the information while the lasers reset. This is shown in figure 2.1. Throughout this thesis I will refer to each period of contiguous length measurements as a *run*, so figure 2.1 shows three runs each lasting for the duration of a single pair of laser scans. The lasers scan at different speeds, and throughout this thesis the slow laser will be referred to as laser 1 and the fast laser as laser 2. In each run the length is independently recalculated, meaning that the measured lengths vary between runs with a spread determined by the absolute accuracy of the system, which is less precise than the resolution.




Figure 2.1: A typical laser scanning pattern for dynamic FSI. The thick blue and thin green lines show the wavelengths of laser 1 and laser 2 respectively as a function of time. The red regions highlight times at which the measurement is interrupted.

Continuous FSI, as proposed by Dale [41], addresses this disadvantage by changing the timing of the laser scans, as shown in figure 2.2. This change means that at no time are both lasers resetting and so, by taking advantage of a subtlety in the dynamic FSI equations, there do not need to be any gaps in the measurement and a single run can

last for multiple laser scans. The exact analysis technique to achieve this is described in section 2.2.




Figure 2.2: A typical laser scanning pattern for continuous FSI, with the thick blue and thin green lines showing the wavelengths of laser 1 and laser 2 respectively. The striped grey region (S) is the set up period, which is effectively a dynamic FSI run. The white regions (b1 etc.) have only one scanning laser and the length is extracted using the technique described in section 2.2. In orange regions (h1 etc.) both lasers are scanning again and it is possible to swap between them in what is called a hand-over, again described in section 2.2. This enables measurement to continue into the next white region and so the length can be calculated at all times.

The second disadvantage is that the dynamic FSI measurement requires information from the entirety of each run to calculate the length during the run, meaning that analysis must be performed after all of the data from the full run has been recorded. This prevents dynamic FSI from being used in fast feedback applications due to the necessary time between data acquisition and calculation of the length. With continuous FSI, once the initial set up period (the grey striped region labelled S in figure 2.2) has elapsed, it is possible to calculate the length on a point by point basis. If the raw data can be retrieved

as it is taken during the run then this can be used to calculate the length at that point using only data that has been taken in the past, dramatically reducing the time between data acquisition and the availability of the calculated length. If the analysis program could calculate the length quickly enough this would enable real-time applications such as low-latency feedback systems, needed in applications such as computer numerical control (CNC) machine control.

2.2 Continuous FSI

2.2.1 Schematic layout of a continuous FSI system

Continuous FSI uses a hardware configuration that closely resembles that of dynamic FSI. The difference is the addition of a pair of optical switches between the lasers and the rest of the optical system. The switches are a purely practical consideration and are used to stop light emitted from the lasers during their reset periods disrupting the length measurements. The generic layout described in figure 2.3 could be used without alteration as a dynamic FSI system.

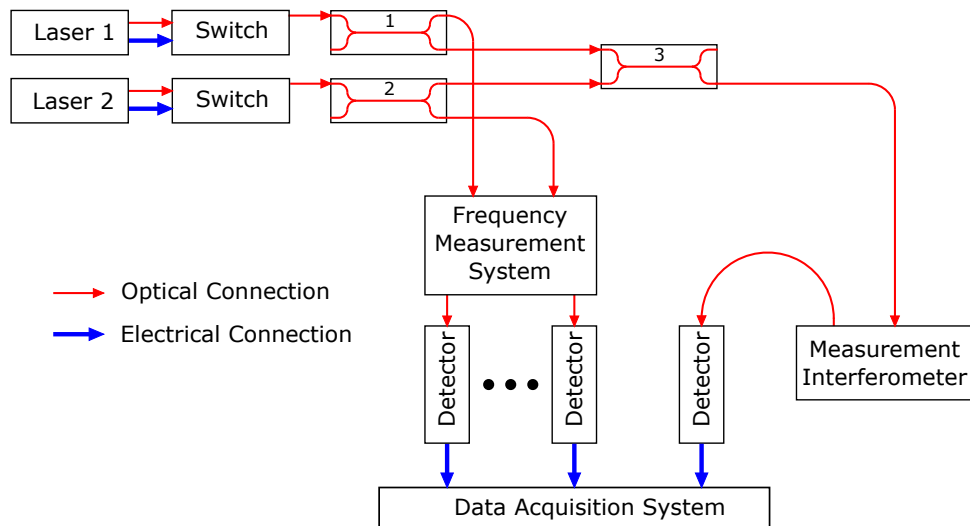


Figure 2.3: Schematic layout of a continuous FSI measurement system.

A continuous FSI system is built around a pair of frequency-scanning lasers, referred

to as laser 1 and laser 2. These are lasers that can scan their frequency between two limits without experiencing a mode-hop. A small fraction of the light from each laser is split off and fed in to a frequency measurement system. The frequency measurement system and its associated data acquisition channels measure the laser frequency every time the measurement interferometer intensity pattern is sampled. An example of such a system is a gas absorption cell, such as [42][43], used in conjunction with a stable reference interferometer. This type of frequency measurement system was developed in [3], and is used throughout this thesis. The gas cell system is shown schematically in figure 2.4, with an explanation of how it works given in section 2.2.3.

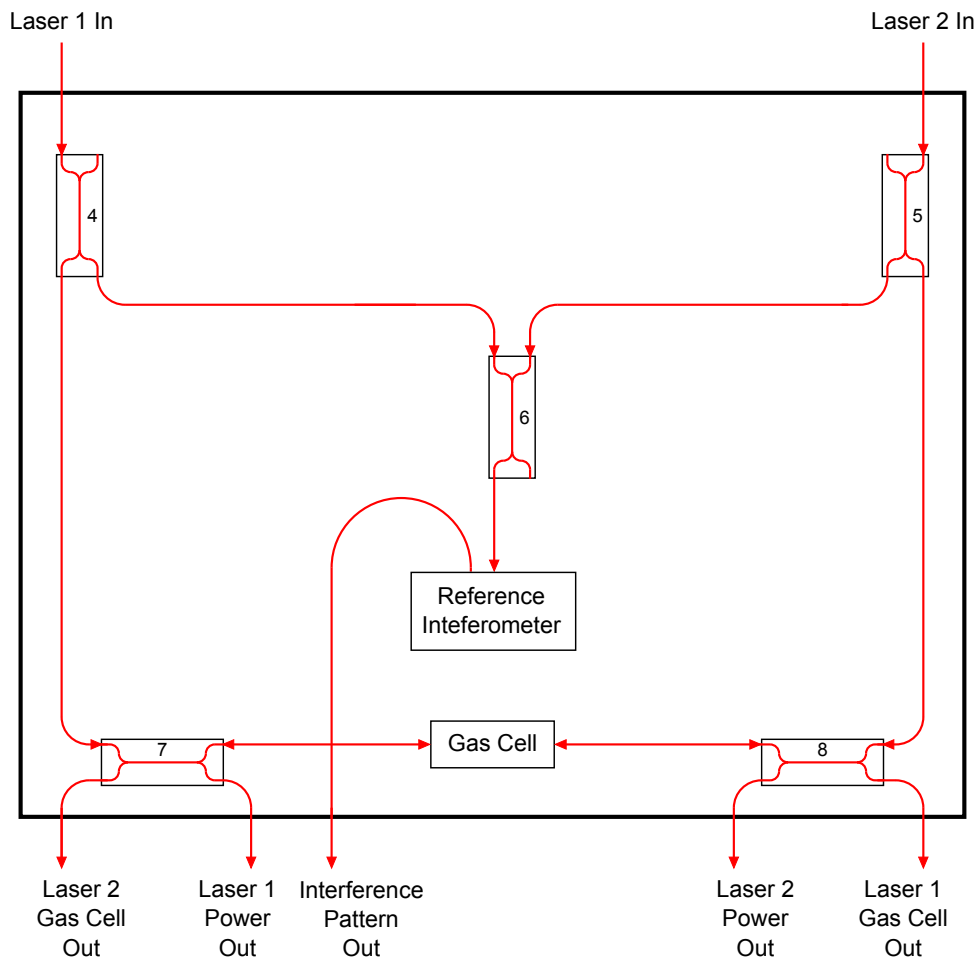


Figure 2.4: Schematic layout of a gas cell based frequency measurement system.

The remaining light from both lasers is combined in a fibre splitter (labelled 3 in figure 2.3) and used to illuminate the measurement interferometer. Continuous FSI will

calculate time-resolved measurements of the OPD of this interferometer, and the OPD is allowed to vary over the run.

The intensities from the measurement interferometer and the information from the frequency measurement system are recorded simultaneously by a high-speed data acquisition (DAQ) system at discrete times, which will be referred to by index.

2.2.2 Derivation of the algorithm for length calculation

The length calculation in continuous FSI is initially identical to that of dynamic FSI, with the difference being the handovers which are described in section 2.2.4. Continuous FSI uses the phase information extracted from the recorded interference pattern of the measurement interferometer to calculate the length. At each time index we can extract two distinct pieces of phase information from the measurement interferometer, one from each laser. There are several methods to separate the two laser signals and to extract the phases from the interference pattern, which can depend quite heavily on the exact layout of the interferometer. The technique to extract phase from an intensity measurement that has been used throughout this thesis is described in section 3.2.2, but for an idealised system the assumption is simply that the phases can be extracted without going into the exact details. The dependence of the two measurement interferometer phases on the interferometer length and laser frequencies at an arbitrary time index i are given in equations 2.1 and 2.2.

$$\phi_{i,1,M} = \frac{2\pi}{c}\nu_{i,1}L_{i,M} \quad (2.1)$$

$$\phi_{i,2,M} = \frac{2\pi}{c}\nu_{i,2}L_{i,M} \quad (2.2)$$

Here ϕ is the phase, ν is the laser frequency, L is the optical path difference of the interferometer and c is the speed of light. The subscript 1 or 2 is the laser number and M indicates that the value refers to the measurement interferometer. Of these parameters the phases can be extracted from the intensity of the interference pattern, modulo an

arbitrary offset, and the laser frequencies are measured by the frequency measurement system. This leaves the length of the interferometer at each point as an unknown to be calculated. It should be emphasized that the interferometer length used in these calculations is the optical path difference, and so must be corrected using the absolute refractive index to provide a measurement of the physical distance.

An arbitrary time index t , within the grey setup period S in figure 2.2, is chosen to be the *transfer index* at which the first length is calculated. Differences in phase between index t and any other index i are then calculated, as shown in equations 2.3 and 2.4.

$$\phi_{i,1,M} - \phi_{t,1,M} = \Delta\phi_{i,1} = \frac{2\pi}{c} (\nu_{i,1}L_{i,M} - \nu_{t,1}L_{t,M}) \quad (2.3)$$

$$\phi_{i,2,M} - \phi_{t,2,M} = \Delta\phi_{i,2} = \frac{2\pi}{c} (\nu_{i,2}L_{i,M} - \nu_{t,2}L_{t,M}) \quad (2.4)$$

The calculation of this difference is important because the extracted phases are typically calculated modulo an arbitrary offset, which is the same for all indices in a laser scan and so cancels out. These equations can be rearranged to give the measurement length at index i , if we can calculate the measurement length at index t . This is shown in equations 2.5 and 2.6. However, the length of the measurement interferometer at the index t , known as the *transfer length* $L_{t,M}$, is still unknown.

$$L_{i,M} = \frac{1}{\nu_{i,1}} \left(\frac{c}{2\pi} \Delta\phi_{i,1} + \nu_{t,1}L_{t,M} \right) \quad (2.5)$$

$$L_{i,M} = \frac{1}{\nu_{i,2}} \left(\frac{c}{2\pi} \Delta\phi_{i,2} + \nu_{t,2}L_{t,M} \right) \quad (2.6)$$

To calculate the transfer length we subtract equations 2.5 and 2.6, which removes $L_{i,M}$, and rearrange to give $L_{t,M}$ as shown in equation 2.7.

$$L_{t,M} = \frac{\frac{\nu_{i,1}}{\nu_{i,2}} \Delta\phi_{i,2} - \Delta\phi_{i,1}}{\frac{2\pi}{c} \left(\nu_{t,1} - \frac{\nu_{i,1}}{\nu_{i,2}} \nu_{t,2} \right)} \quad (2.7)$$

Equation 2.7 means that the transfer length can be calculated using any index except

index t , which would require division by zero. The method used throughout this thesis is to calculate the transfer length using a series of indices away from the transfer index and then to take an average. With the knowledge of the transfer length, the length at every other point can be calculated using equation 2.5 or 2.6. The transfer index and the indices used to calculate the transfer length are all taken from the set up period (S) of figure 2.2.

2.2.3 Laser frequency calculation using a gas cell

In the previous section the laser frequencies were assumed to be known. It is important to know how the laser frequency is calculated before the process of a hand-over in the orange sections of figure 2.2 can be described. The system shown in figure 2.4 provides five measurements. These are the interference pattern from the reference interferometer, the power of each laser and the transmitted power of each laser through the gas cell. The laser power is simply used to normalise the gas cell transmission data to give the transmission function and so is not referred to here. The detailed algorithm for normalisation is described in section 3.2.3.

The core of the frequency measurement system shown in 2.4 is the gas cell. The gas cell provides a series of points of known frequency for each laser by virtue of the calibrated gas absorption features. Here ‘point’ rather than ‘index’ has been used deliberately because the minimum of an absorption feature, which is the point of known frequency, is unlikely to coincide with a time index.

To extrapolate the laser frequency at each time index from these absorption features we need to use the phases from the reference interferometer, which are strictly proportional to the laser frequency providing the optical length of the interferometer is stable with time over the period in which the laser frequency needs to be extrapolated, and is independent of laser frequency. As in equations 2.1 and 2.2, we can extract one phase for each laser from the recorded reference interferometer intensity pattern. These phases

are given in equations 2.8 and 2.9.

$$\phi_{i,1,R} = \frac{2\pi}{c}\nu_{i,1}L_R \quad (2.8)$$

$$\phi_{i,2,R} = \frac{2\pi}{c}\nu_{i,2}L_R \quad (2.9)$$

As with the measurement interferometer equations ϕ is the phase, ν is the laser frequency, L is the optical path difference of the interferometer and c is the speed of light. The subscript i refers to the time index, subscripts 1 or 2 indicate the laser number and subscript R identifies the reference interferometer.

The first step to calculate the frequency at each time index is to measure the reference phase at all of the absorption features that have been observed in the “set up period”. This is achieved by fitting the transmission function of each gas absorption feature against phase from the reference interferometer. These fits provide a series of points at which we know both interferometer phase and laser frequency. In the second step these points are used to calculate the reference length L_R by performing a linear fit of phase against frequency for each laser, which will give a gradient of $\frac{2\pi}{c}L_R$. The average of the calculated length values from the fits of both lasers is taken to be the reference length.

The final step is to combine the knowledge of the reference length with the knowledge of the phase and frequency of a gas cell feature to calculate the laser frequencies at each time index. This is achieved by taking the difference between the phases at the index of interest (i) and the phase at one of the gas cell features, denoted as p for laser 1 and q for laser 2. Using equations 2.8 and 2.9, this gives equations 2.10 and 2.11.

$$\phi_{i,1,R} - \phi_{p,1,R} = \frac{2\pi}{c}L_R(\nu_{i,1} - \nu_{p,1}) \quad (2.10)$$

$$\phi_{i,2,R} - \phi_{q,2,R} = \frac{2\pi}{c}L_R(\nu_{i,2} - \nu_{q,2}) \quad (2.11)$$

Re-arranging for $\nu_{i,1}$ and $\nu_{i,2}$ gives equations 2.12 and 2.13, allowing calculation of the frequency at each time index for each laser. This knowledge of the frequency can be used

to complete the determination of the measurement interferometer lengths as described in section 2.2.2.

$$\nu_{i,1} = \left(\left(\frac{c}{2\pi L_R} \right) (\phi_{i,1,R} - \phi_{p,1,R}) \right) + \nu_{p,1} \quad (2.12)$$

$$\nu_{i,2} = \left(\left(\frac{c}{2\pi L_R} \right) (\phi_{i,2,R} - \phi_{q,2,R}) \right) + \nu_{q,2} \quad (2.13)$$

We note here that the optimal choice of absorption feature from which to extrapolate the laser frequency remains open and is explained in section 3.2.5.2.

2.2.4 Performing a hand-over and moving to continuous FSI

So far the analysis described here is identical to that used in dynamic FSI as described by Dale [3]. We now note that in both equations 2.5 and 2.6 only the phase information from one of the two lasers is required to calculate the length at any point, excluding the transfer length. The transfer length only needs to be calculated using equation 2.7 once, in the grey set up period in figure 2.2, so after it has been calculated there is no longer a need for both lasers to be active simultaneously at all times. Importantly the equations that are used to calculate the laser frequencies, 2.12 and 2.13, also require information from only one of the two lasers.

The consequence of this laser-independence is that after the transfer length is calculated one of the lasers can stop scanning and reset without disrupting the measurement of the length (period b1 in figure 2.2). In figure 2.2 the first laser to stop scanning is laser 1. If laser 1 can restart and have its frequency determined, for example by crossing one of the gas cell features, before laser 2 stops scanning then there is a period during which we know the length of the measurement interferometer from laser 2 and the reference interferometer phase and frequency of laser 1. This occurs in section h1 in figure 2.2, which is referred to as a hand-over period.

Once we are in a hand-over period the only parameter needed to begin calculating

the length of the measurement interferometer using the new laser is knowledge of a new transfer length. This is acquired by choosing an arbitrary time index in the past and using the length as calculated by the old laser at that index as the transfer length for the new laser. Using this length in conjunction with the phase and frequency of the new laser at that same index, the new laser can be used to calculate the length of the measurement interferometer from then on. The above procedure is called a *hand-over*. It should be noted that by simply taking the length as calculated by laser 2 at the new transfer index, rather than performing an average as was done for the first transfer length, we may become susceptible to random walk over the run because uncertainties in the phase extraction are not suppressed by the averaging of the first calculation, and so could build up over successive hand-overs.

Once the hand-over is completed laser 2 is no longer needed for length measurement and can reset (period b2). When laser 2 starts to scan again we enter into a second hand-over period (h2) and the hand-over procedure repeats, allowing laser 1 to reset. This can happen indefinitely, as shown in figure 2.2, allowing contiguous length measurement over an arbitrary time period. Hand-overs are the core principle behind continuous FSI.

2.2.5 Using continuous FSI for fast feedback

The ability of continuous FSI to measure the length on a point by point basis after the initial set up period makes it potentially useful for fast-feedback applications. Once the set up period, which would roughly correspond to a full dynamic FSI run, is passed the length and frequency calculations do not need any further information beyond the new phase values from the reference and measurement interferometers. If these phases can be extracted with low latency then the length can also be calculated with low latency, opening up applications such as position feedback on CNC machines.

Chapter 3

Continuous FSI - Implementation

This chapter will describe the hardware implementation and software algorithms of a continuous FSI system. The first section will detail the equipment used to implement the system. The second section describes the software used to calculate the interferometer length. The third section describes the reasoning behind the choice of the various system parameters, such as laser scanning speeds and wavelength ranges.

3.1 Experimental layout

The experimental layout of the continuous FSI system differs superficially from the layout of the schematic diagrams given in figures 2.3 and 2.4. This difference arises from the fact that the functions of both splitters 1 to 3 in figure 2.3, and splitters 4 to 6 from figure 2.4, are performed by just three splitters in the experimental system. These three splitters are grouped together in the master splitter box, described below in section 3.1.3. A final difference is that the interferometers are not illuminated straight from splitter 3 (figure 2.3, for the measurement interferometer) or 6 (figure 2.4, for the reference interferometer). The light is instead fed into a splitter tree which can illuminate up to 24 interferometers at any one time, although we only use two during continuous FSI experiments. This splitter tree is known as the front-end splitter tree, and is described in

section 3.1.4. The splitter tree arrangement is used to enable compatibility with the multi-interferometer measurement system as described by Dale [3], for which it was designed. A final difference is the technology used in the splitters which feed the gas cell. These are not fused-taper splitters but back-to-back planar lightwave circuit (PLC) splitters, which are described in detail in section 3.1.6.

Figure 3.1 shows the full schematic layout of the continuous FSI system as implemented for this work, with differences from the theoretical schematics highlighted by shading the relevant components in green.

3.1.1 Frequency scanning lasers

The two lasers used in the experiment were both New Focus TLM-8700-H-CL lasers [44]. These are fibre-coupled lasers with an output power of up to 8 mW, a mode hop free wavelength range of 1520 nm to 1630 nm at scanning speeds between 2 nm/s and 2000 nm/s. The coherence length of the lasers is not given in the manufacturer specifications, but is known to enable interferometry over an OPD of at least 40 m, as described by Dale [3]. The lasers provide a series of voltage outputs which provide information about their current state. One output, called ‘Sync out’, provides a ‘high’ voltage level for the entire duration of the laser scan and swaps to a ‘low’ voltage level when the laser is resetting. The ‘Sync out’ output is used to control the optical switches, and the first rising edge from laser two is used as the trigger for the data acquisition system. The lasers are controlled using RS-232 communication, and are shown in figure 3.2.

3.1.2 Optical switches

The optical switches are non-latching Seralco SXNT-2x2-9N MEMS switches [45]. These are fibre-coupled switches and have been mounted on a circuit board by the Central Electronics Group at Oxford. The switches are set up to allow light to pass when the input from the laser sync out is high, which indicates that the lasers are scanning, and

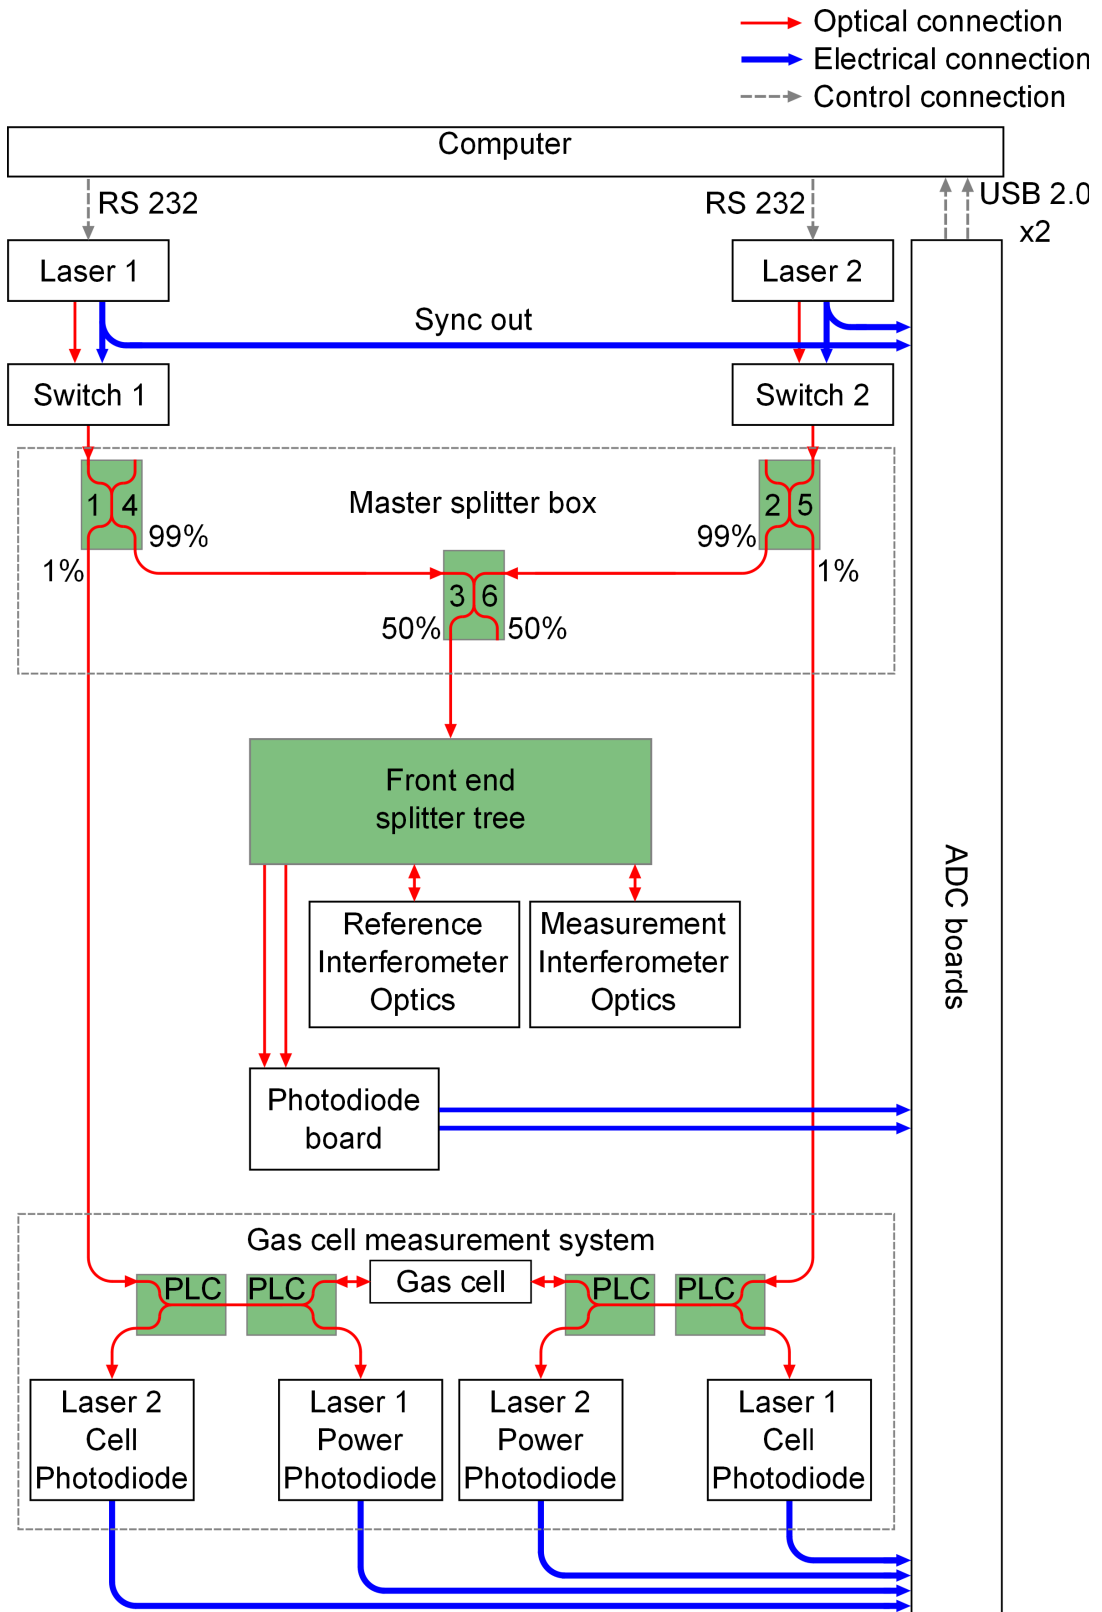


Figure 3.1: Schematic of the full experimental setup for continuous FSI. Boxes shaded green highlight superficial differences from the theoretical schematics (see figures 2.3 and 2.4). Merged splitters are labelled with both numbers.




Figure 3.2: The two TLM-8700 lasers.

to block the light when the input is low, when the lasers are resetting or not scanning. Figure 3.3 shows the inside of the fibre switch box.

3.1.3 Master splitter box

The master splitter box is the point at which the experimental system shown in figure 3.1 superficially starts to deviate from the generic schematic given in figures 2.3 and 2.4. The master splitter box contains a pair of 99:1 fused-taper splitters. The 1% outputs are used to feed the gas cell system described in section 3.1.6. The 99% outputs are combined in a 50:50 fibre-taper splitter, one of the outputs of which is used to feed the front-end splitter tree. The three splitters are securely attached to a sturdy metal mesh, which is in turn mounted in a metal casing for added stability and protection from accidental damage.

For the sake of completeness, the unused output from the 50:50 fibre splitter is passed through a 99:1 splitter before it leaves the master splitter box. The 1% output was used for previous experiments [3] but is not required for continuous FSI. The extra splitter




Figure 3.3: The MEMS switch box.

would have no effect on the behaviour of the continuous FSI system, and is left in place to allow the previous experiments to be repeated with minimal effort.

The master splitter box is highlighted in figure 3.1.

3.1.4 Front-end splitter tree

The measurement and reference interferometers are illuminated by light from the front-end splitter tree, which is a series of fused-taper splitters spliced together. The splitter tree used in this case has two inputs and 24 outputs. As shown in figure 3.4, each of these 24 outputs has an extra 50:50 splitter, with the second output terminated. This is required to allow light that is returning from the interferometer to be tapped off and sent into the DAQ system. The extra 50:50 splitter means that the outputs each emit $1/48^{th}$ of the power from each of the inputs, assuming a loss free system.

One of the two available inputs to the splitter tree is used to inject the combined

infra-red laser light from the master splitter box. The second input is used to insert red (635 nm) light from an Intelite FRM635-10S-APC 13.8 mW laser [46] which has been locked inside the splitter tree. The red light is included as an alignment aid and is switched off when continuous FSI measurements are being performed. The front end splitter tree is shown schematically in figure 3.4.

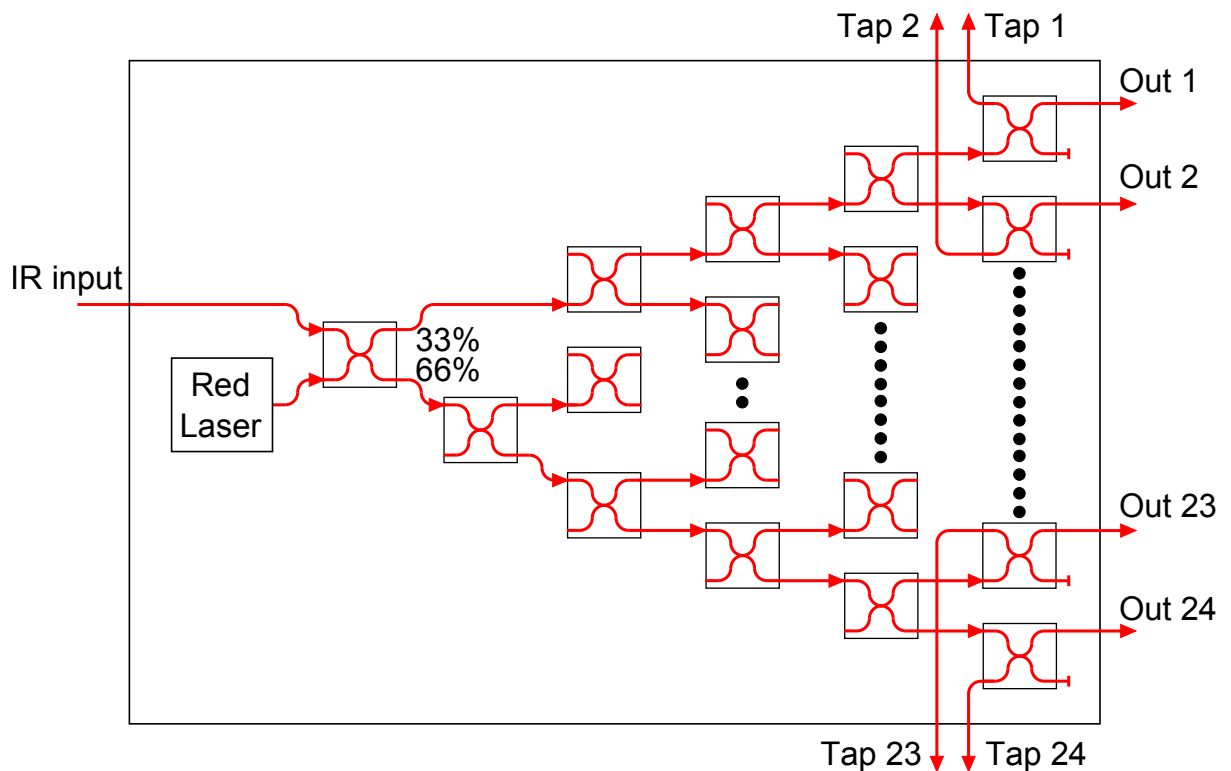


Figure 3.4: Schematic of the front end splitter tree. Arrows show completed fibre connections and terminated lines show terminated connections. The fibre splitters are 50:50 unless otherwise labelled.

3.1.5 Interferometers

The interferometers used in this continuous FSI implementation are all Fizeau interferometers [47], shown in figure 3.5. The reference surface is the end of the optical fibre that is used to supply the laser light. The fibre has been re-polished to provide good reflection back down the fibre while mitigating the effects of reflections from the back of the collimator lens. The second reflective surface is typically a corner cube retro-reflector. The

interferometers have an OPD of twice the optical distance between the fibre tip and the apex of the retro-reflector. The measurement interferometers used in this thesis are typ-




Figure 3.5: A schematic of a fibre-fed Fizeau interferometer, with an example mounted on an optical table.

ically mounted on an optical table using standard optical mounts, as shown in figure 3.5. The collimators used are Thorlabs F280FC-1550 [48], and the retroreflectors are PLX 0.5" Stainless Steel Ball Mounted Hollow Retroreflectors [49]. However, the reference interferometers are custom built Invar interferometers with an OPD of approximately 6.5 m. The thermal expansion of the interferometers is further compensated by a counter expansion mechanism, reducing the relative thermal expansion to $5.25 \times 10^{-7} \text{ mK}^{-1}$ [35]. The interferometers are sealed in an insulated metal tube and are expected to be much more stable due to the low thermal expansion co-efficient and the reduction in macroscopic air flow provided by the container. These interferometers will be referred to as *LiCAS* interferometers, after the project that they were constructed for. Details of the construction of these interferometers are available in [3][35]. The reference interferometers were not evacuated during the continuous FSI experiments, which means that their length varies with optical frequency due to changes in the refractive index of the air.

This was not taken into account in section 2.2.3 because the maximum effect on the measured laser frequencies is expected to be negligible (approximately 40 kHz compared with optical frequencies of order 194 THz).

3.1.6 Gas cell

The gas cell used is a hydrogen cyanide gas cell from Wavelength References [50], which is based on NIST SRM 2519a [42][43]. The gas cell is 165 mm long and filled to a pressure of approximately 25 Torr (3.3 kPa). The gas cell provides 54 absorption features with vacuum wavelengths between 1527.6 nm and 1564.4 nm, which covers most of the available scanning range of the TLM 8700 lasers. To simplify the fitting of the features only the central and most pronounced 44 peaks are used, with vacuum wavelengths in the range 1529.8 nm to 1559.8 nm.

The gas cell is mounted in a similar way to the splitters in the master splitter box, on a sturdy mesh in a metal casing. The temperature of the cell is monitored using a National Instruments NI USB-TC01 with J-Type grounded probe thermocouple [51]. The metal casing is isolated from vibration using thick foam and an optical breadboard, as shown in figure 3.6.




Figure 3.6: The gas cell as mounted for continuous FSI.

The splitters that are used in the gas cell system are back-to-back custom built 50:50

1×2 PLC splitters from OpticReach Ltd. [52]. This is an important feature as it results in a more consistent division of optical power over the wavelength range of the lasers when compared with fused-taper splitters or even 2×2 PLC splitters. The disadvantage of the 1×2 PLC layout is the dramatic reduction in power, to approximately 6.25% of the input power on the gas cell photodiode. For comparison a system based on fused-taper or 2×2 PLC splitters would provide 25% of the input power to the gas cell photodiode, which is a factor of four higher. However, the continuous FSI system is not currently power limited, and so can take advantage of the 1×2 PLC splitter layout.

The previous iteration of the gas cell system used fibre circulators [3]. The circulators were found to have some small internal reflections that superimpose a slow interference pattern onto the gas cell readouts, reducing the overall precision of the reference length measurements, and so were swapped for the PLC system, which suffers from this problem to a significantly lesser degree and is also cheaper. Further improvements were also made by changing the photodiodes, and so full details of all of the improvements are given in section 3.1.7.

3.1.7 Photodiodes and data acquisition

The data acquisition (DAQ) system was custom built in Oxford over the course of the LiCAS and MONALISA projects [1][35][36][37]. The DAQ system used in the continuous FSI experiments consists of a pair of analogue-to-digital conversion (ADC) boards, each of which samples up to 16 channels simultaneously at a sampling speed of 50/18 MHz (≈ 2.77 MHz). These are mounted in a DAQ crate, which distributes triggers and power to the cards. Analogue voltage signals are sent to the board from either a photodiode board or BNC input board placed adjacent to the ADC board in the crate, and data is transferred to a computer via a USB 2.0 connection. The DAQ system is shown in figure 3.7.

The ADC boards have two modes of operation, burst mode and continuous mode.

In burst mode the board will store the data in on-board dynamic random-access memory (DRAM) and transfer the entire data set to the host computer at the end of the acquisition period. This limits the recording time. In continuous mode, the ADC boards transfer the data to the PC in small packets as it is taken. However, the board cannot access the DRAM in this mode, using a large FIFO instead. The maximum number of channels which can be continuously read out from an ADC is limited to 9 channels per board by one of the on-board FPGAs.

The continuous FSI system uses eight channels, four on each ADC board. Splitting the channels over two boards allows the data to be transferred to the computer at double the rate of a single board. On the left ADC board in figure 3.7 the signals from the laser power and gas cell transmission photodiodes are recorded using a BNC input board, positioned to the right of that ADC board. The ADC board on the right of figure 3.7 records the interferometer signals using a custom photodiode board on its left as well as the sync output from each laser using a BNC input board on its right.




Figure 3.7: The DAQ crate configured for continuous FSI.

The custom photodiode boards have eight fibre-coupled InGaAs photodiodes. The boards provide a range of remotely controlled gains for each photodiode readout channel.

The electronics has been optimised for fast FSI-style signals. The BNC input boards simply route their input signals to the backplane of the DAQ crate, and so to the ADC cards. The positioning of the various input boards is important because of the way the remotely controlled gains for the photodiodes are set. The settings are passed via the back plane of the DAQ crate and are lost if they pass through the BNC input board before reaching the photodiode boards. The BNC board must be positioned on the right of the ADC board to ensure the settings are not lost.

The photodiodes used in the gas cell system are external to the DAQ crate. In the current implementation Thorlabs PDA10CS-EC photodiodes [53] have been chosen. The photodiode boards are not used because the photodiode packages mounted on them contain small ball lenses, which have internal reflections leading to a slowly varying interference pattern superimposed on each output. This slow signal is filtered away during analysis in the interferometer lines, but this is not possible on the gas cell system which observes signals which also vary slowly with time, such as laser intensity. The improvements made by changing to a PLC splitter layout and the photodiodes are shown in figures 3.8, 3.9 and 3.10. It can clearly be seen that figures 3.9 and 3.10 have less overall variation than figure 3.8. It can also be seen that the power data has less rapid variation in the PLC system than in the circulator system, with the local range dropping from approximately 140 ADC counts to approximately 60 ADC counts. More information on the DAQ system is given in [54][55].

3.1.8 Computing

The DAQ and analysis of the continuous FSI experiments are performed using a server with 2 CPUs, each providing 16 cores clocked at 2 GHz, and 64 GB RAM. The server has two separate USB controllers, which is necessary to enable maximum data transfer speed from the ADC boards. The server also contains a series of NVIDIA graphics processing units (GPUs) to enable parallel programming using CUDA [56], a GPU programming

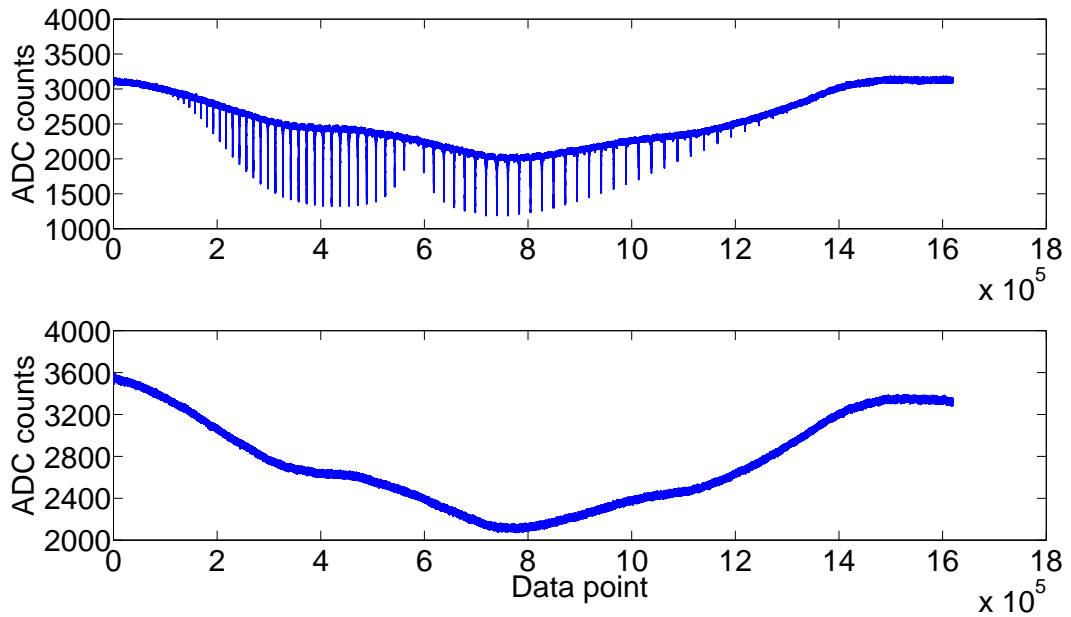


Figure 3.8: Example raw gas cell (upper plot) and power data (lower plot) using the LiCAS photodiode boards and fibre circulators.




Figure 3.9: Example raw gas cell (upper plot) and power data (lower plot) using the Thorlabs PDA10CS-EC photodiodes and fibre circulators.

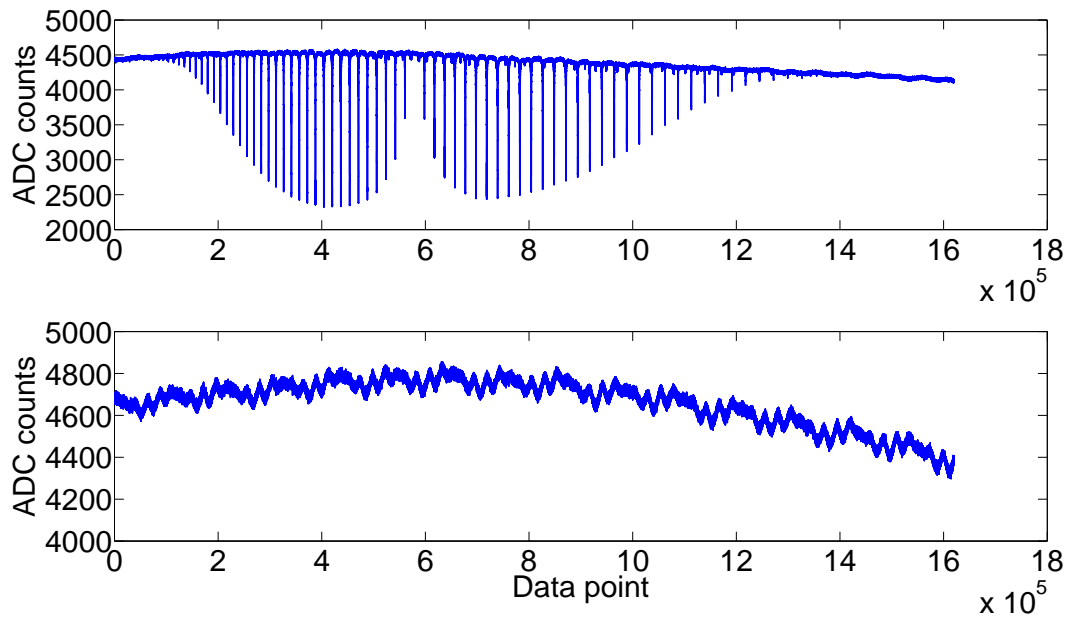


Figure 3.10: Example raw gas cell (upper plot) and power data (lower plot) using the Thorlabs PDA10CS-EC photodiodes and PLC splitters.

language from NVIDIA. The GPU used is a Tesla K20c card [57], which has 4800 MBytes of onboard RAM, 2496 CUDA cores and are clocked at 706 MHz. The Tesla card is able to perform host-device memory copies concurrently with kernel execution (processing). By taking advantage of the GPU it is possible to perform the analysis much quicker than with only CPUs. The current CPU code takes approximately 45 minutes to analyse 3 seconds of continuous FSI data, whereas the GPU code analyses 1 minutes worth of continuous FSI data in approximately 2 1/2 minutes, a speed increase of a factor of 360. The speed increase made offline analysis of long data sets feasible, and with further development the code may reach sufficient speed for real-time analysis.

3.2 Implementation of the analysis

This section will describe how the results from a continuous FSI run are analysed, starting with an overview of the analysis structure before going into more detail on each individual process.

3.2.1 Overview

The analysis of the continuous FSI experiments was designed from the outset to run in real time, as described in section 2.2.5, although the current implementation has not yet reached this benchmark. The data is supplied to the analysis in *packets*. Each packet contains a set of 16,384 (2^{14}) data points for each of the 8 signals (two interferometer signals, two gas cell transmission signals, two laser power signals and two laser sync output signals), giving 131,072 data points per packet. The analysis is forced to use only information from the current or previous packets, even when performing the analysis offline which could provide more information.

The 8 signals are retrieved from the data packet, with each added to a separate buffer holding 65,536 (2^{16}) data points (4 packets worth of data). This is the amount of data that is required to perform the Fourier transforms used in the phase extraction and gas cell normalisation processes of the analysis with a satisfactory frequency resolution. A power of 2 was chosen for the data length to increase computation speed of our Fast Fourier Transforms (FFTs). Other processes that are performed are gas cell peak finding and, once sufficient data has been collected, reference and measurement length calculation. Each process provides information about the third packet of data in the buffer. The data from the first and fourth packets is cut away after filtering to remove spurious oscillations which are an artefact of the hard windows used in the Fourier transform process (see figure 3.18 for details). The information from the second buffer is required to link the phase extracted in the current iteration with the phase extracted in the previous iteration. Once a set of buffers has been analysed the first packet is removed and the remaining packets are shifted forwards. The next data packet is then added and the analysis run again on the new buffers. A schematic of the analysis process is shown in figure 3.11.

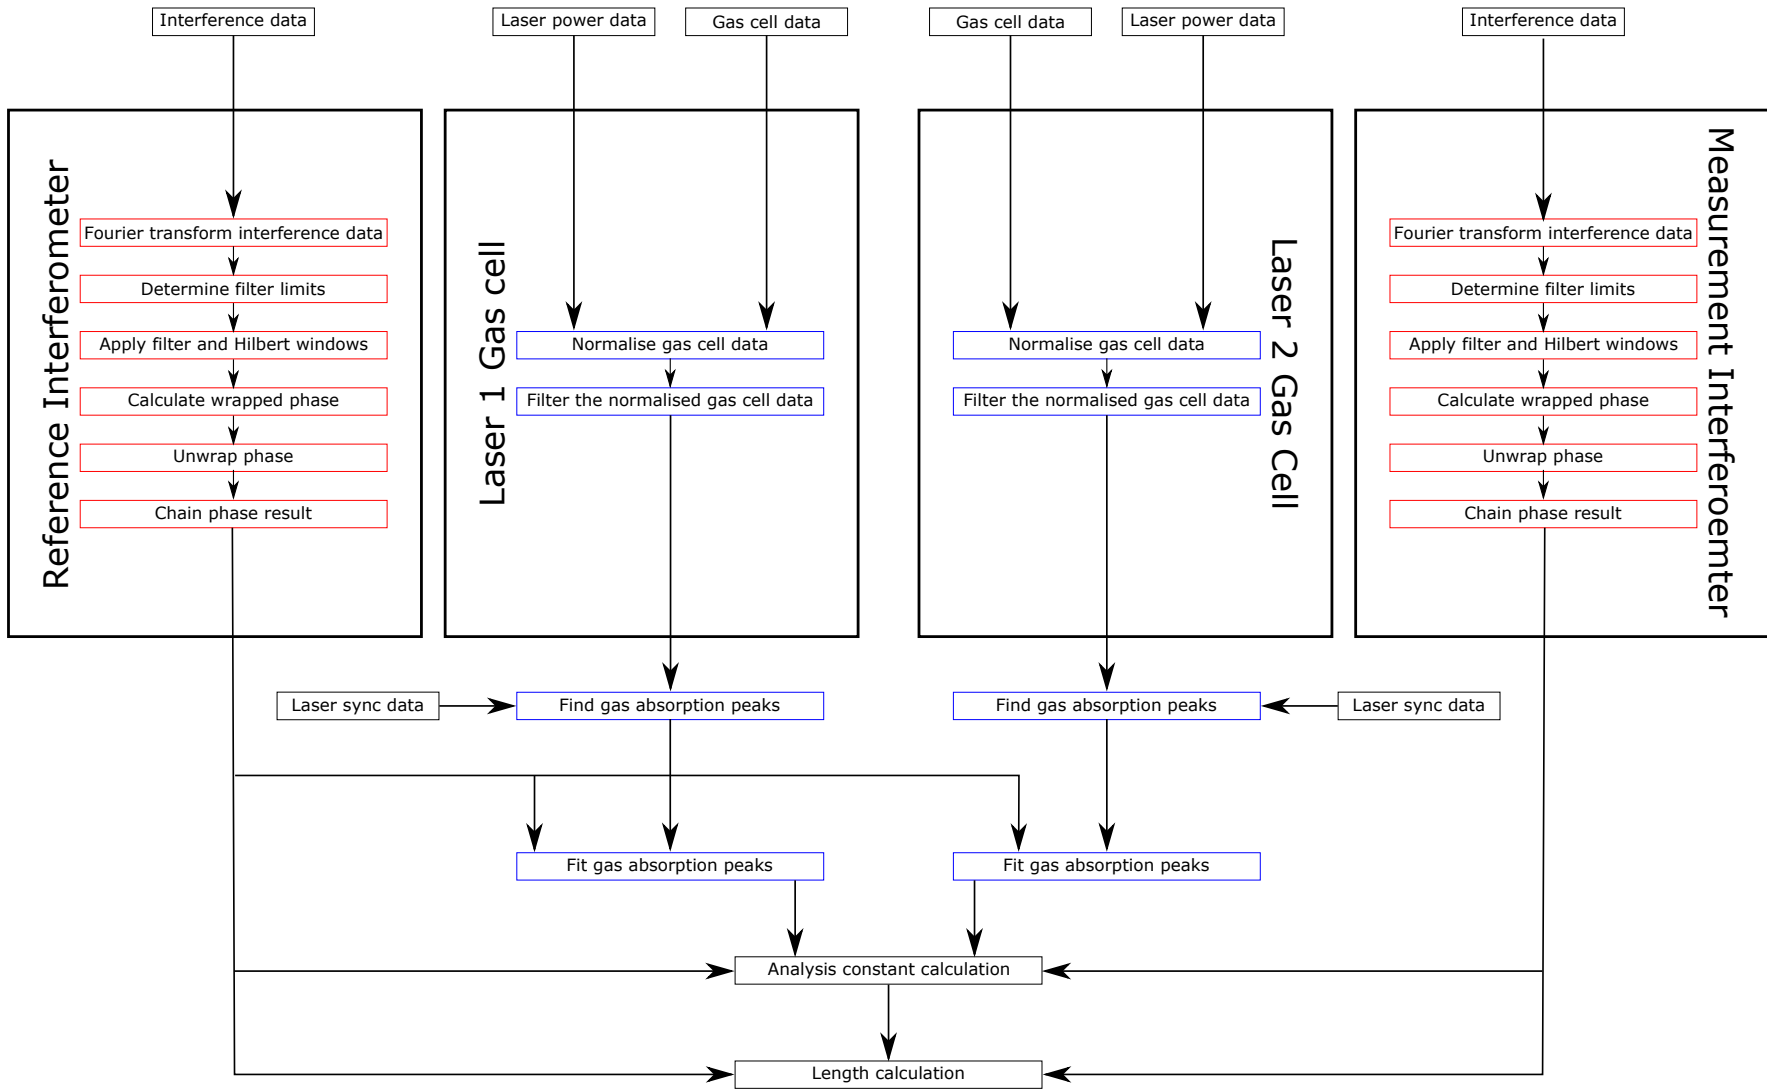


Figure 3.11: Schematic of a continuous FSI analysis iteration.

3.2.2 Phase extraction

A Hilbert transform is used for phase extraction, as described by Warden [1] and Dale [3] for earlier FSI studies. The Hilbert transform works on real signals with a rate of change of phase (phase velocity) that does not change sign. This has implications for laser tuning speeds which are discussed in section 3.3. The Hilbert transform applies a window in frequency space which when used on a co-sinusoidal signal, such as an interference pattern, results in a complex return value with the original cosinusoidal signal as the real component but with the corresponding sinusoidal signal as the imaginary component. The phase of the signal can then be calculated using the arctangent of the two components. A detailed explanation of the technique is available in [58][59], although Warden [1] provides a clear derivation in chapter 2. The Hilbert window ($H(\omega)$) applied to the data is defined by equation 3.1 where $sgn(\omega)$ is -1 for negative ω , zero for $\omega = 0$ and $+1$ for positive ω , as given in equation 2.2.7 of Warden [1].

$$H(\omega) = 1 + sgn(\omega) \quad (3.1)$$

The Hilbert transform technique will not work with data that contains two interference signals, such as the continuous FSI interferometer data, because it is not cosinusoidal. This problem is overcome by separating the two signals by fringe frequency using software level Butterworth band pass filters [60]. The filtered data is then cosinusoidal and the Hilbert transform will work as intended. The positioning of these filters is determined by an auto-filter algorithm. The steps required to perform the phase extraction using a Hilbert transform are:

1. Fourier transform the raw interference data.
2. Determine the filter limits to separate the two laser signals using the auto-filter algorithm.
3. Apply Hilbert and filter windows to the separated data sets.

4. Calculate the wrapped phase for each laser.
5. Unwrap the phase for each laser.
6. Chain the phases to previous results.

The steps are described in detail in sections 3.2.2.1 to 3.2.2.6 below.

3.2.2.1 Fourier transforming the raw interference data

The raw interference data buffer, an example of which is shown in figure 3.12, is transferred onto the Tesla K20c GPU card. Once transferred to the GPU the data is converted




Figure 3.12: A subsection of an interference buffer showing the form of the interference fringes.

from short format (16 bit) into complex double format (two 64-bit numbers per value) with zero imaginary components. This change maintains numerical accuracy throughout the analysis. The Fourier transform is performed using the cuFFT library [61], which is included in the standard CUDA package, producing a complex double results (65,536 pairs of doubles).

3.2.2.2 Determining the filter limits using the auto-filter algorithm

The auto-filter algorithm separates the signals from the two lasers using the difference in their interferometric fringe frequencies. The algorithm follows the following steps:

1. Calculate the magnitude of the Fourier transform.
2. Average the magnitudes over several data points.
3. Find potential spectral peaks.
4. Assign the spectral peaks to the correct lasers.
5. Find the filter limits for both lasers.

Calculating the magnitude of the Fourier transform The magnitude of the first half of the complex data, corresponding to the positive frequency data, is calculated on the GPU. The negative frequency data can be ignored because it will be zeroed by the Hilbert window at a later stage. At this point the magnitude data consists of 32,768 doubles.

Average the magnitudes The magnitude data is averaged using the GPU, typically over bins of 32 data points. This gives 1,024 bins to analyse. The algorithm is not well optimised, simply using one CUDA-thread per bin and calculating the average of the corresponding 32 points. More complex algorithms were not deemed necessary because of the small number of calculations required, although this could be worth revisiting in future work. The magnitude of the Fourier transform for the raw data is shown in figure 3.13, along with its averaged form.

Finding potential spectral peaks The auto-filter peak finding algorithm is performed on the CPU, so the 1,024 averaged values are transferred back to standard memory. The nature of the spectral peak finding algorithm involves many comparisons between data points which, coupled with the small size of the data set, led to the decision to avoid the added complexity of GPU programming.

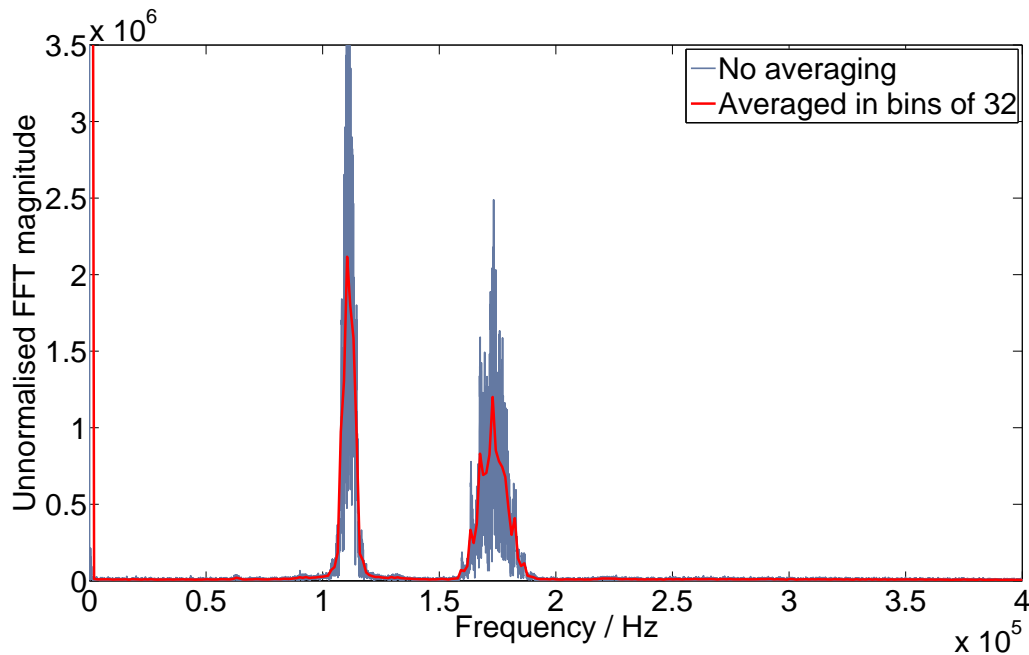


Figure 3.13: Magnitude of the FFT of the interference data buffer.

A spectral peak in the spectrum is defined as two or more bins (containing averages of 32 points) above a threshold value. The threshold value is calculated as the mean of all bins plus a multiple of their standard deviation. The multiple is known as the *cut-off* value, and is typically 0.2.

The peak finding algorithm starts at the 21st bin. The first 20 bins are ignored in order to avoid the DC peak, which is often larger than the interference peaks. This prevents peaks below 27 kHz from being detected. The algorithm advances through the data marking any set of two or more bins above threshold. Once it has found a peak the algorithm requires two bins to fall below threshold before ending the peak, and starting to search for the next candidate. The two bin requirement is designed to prevent spurious noise from being misidentified as a peak, or from ending a peak prematurely. The algorithm will therefore return peaks which are at least 2 bins wide, and separated by at least 2 bins.

Once all of the peaks have been found the two which have the highest mean values are taken to be the interference peaks.

Assigning the spectral peaks to the correct lasers The algorithm then compares the current peaks with those found in the previous buffer. This is necessary because when one of the lasers is not active the round trips (or other noise features) from the active laser could be interpreted as peak from the missing laser. Comparing with the previous results allows this problem to be avoided. If a feature is found to be missing the previous filter values are used as a place-holder.

Finding the filter limits for both lasers. If the start and end values of the peaks were used as the filter window edges there would be a risk that some data could be filtered away incorrectly as the peaks often have wide shoulders. To avoid this the windows are expanded symmetrically around their central values by a factor known as the *proportional width*, typically 6. If this expansion causes the filter windows to overlap the proportional width is reduced (in steps of 1) until they have been separated, or the proportional width reaches 2. If there is still a clash with a proportional width of 2 then the clashing limits are set to be the value equidistant from the initial peak edges $\pm 10\%$ of the size of the gap between them. Figure 3.14 shows an example of the peak finding algorithm process, with the threshold level, original peak boundaries and expanded peak boundaries shown.

3.2.2.3 Application of the Hilbert and filter windows

The two filter windows are now applied to the data on the GPU, producing two filtered spectra. The filters are Butterworth band-pass filters [60] with order 50 at the high pass edge and order 100 at the low pass edge. The Hilbert window is applied simultaneously by setting values corresponding to negative frequencies to zero, and doubling those corresponding to positive frequencies. The zero-frequency point has always been filtered away and so can be ignored when applying the Hilbert window. The windows that are applied to the data are shown in figure 3.15.

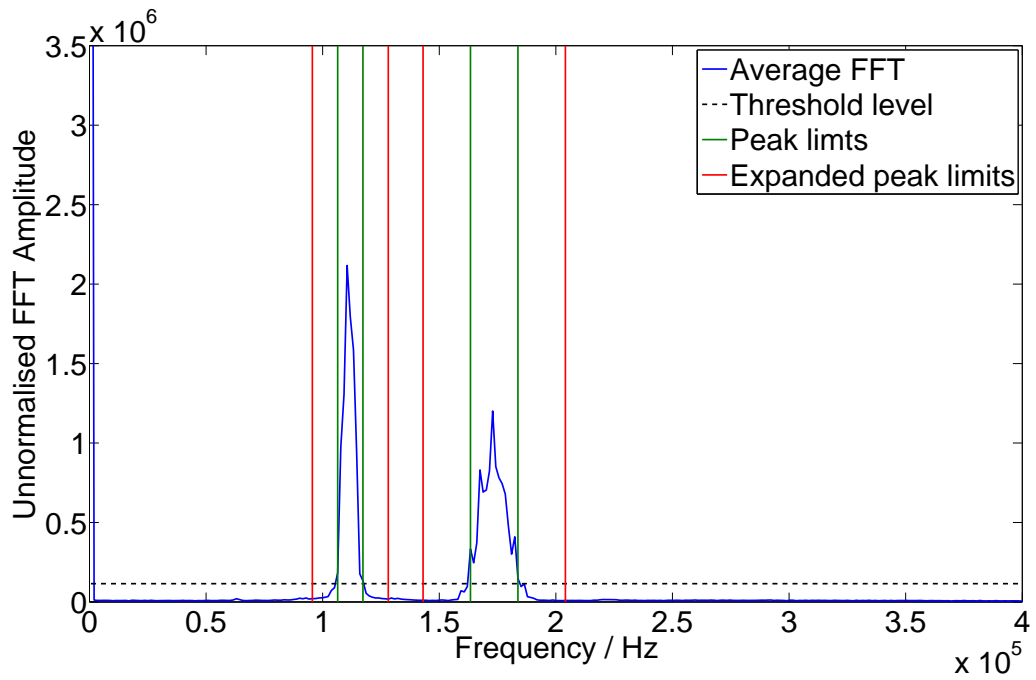


Figure 3.14: Threshold values and peak limits as found by the auto-filter algorithm.

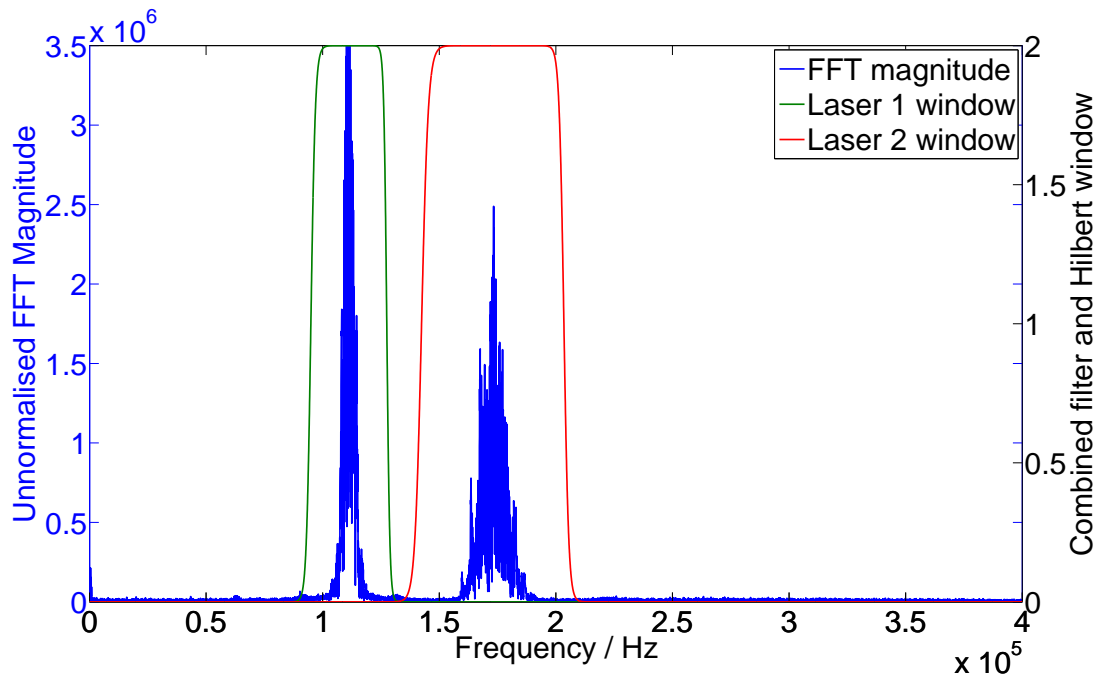


Figure 3.15: Combined filter and Hilbert windows.

3.2.2.4 Calculation of the wrapped phase

The two filtered spectra are transformed back into the time domain, again using cuFFT. The phase at each point is calculated using the arctangent of the imaginary component divided by the real component for that point, as described above. This produces phase in the range $\pm\pi$, and is known as the wrapped phase. This is shown in figure 3.16.




Figure 3.16: Wrapped phase calculated using the Hilbert transform.

3.2.2.5 Unwrap the phase

To generate contiguous phase data the jumps in the wrapped phase must be removed in a process called *unwrapping*. The jumps evident in the wrapped phase are removed by adding integer multiples of 2π . Unwrapping is performed on the GPU. The phase jumps are identified by calculating point-by-point differences in the unwrapped phase. If the difference is below -1.5π then all subsequent values are increased by 2π . If the difference is above 1.5π then all subsequent values are reduced by 2π . This produces the unwrapped phase as shown in figure 3.17. This is the point at which the knowledge of the direction




Figure 3.17: Unwrapped phase.

of the lasers tuning is applied to the data. If the lasers optical frequency is decreasing the unwrapped phase is subtracted from 0 so that the phase changes have the correct sign.

This is also the most suitable point to show the spurious oscillations described in section 3.2.1. By calculating the point-by-point phase differences, as shown in figure 3.18, large spikes and oscillations at the start and end of the data regularly become apparent. These must be cut before the phase can be used reliably.

3.2.2.6 Chain the phase to previous results

The unwrapped phase must now be aligned with phases that have been extracted in previous iterations of the analysis. This is possible because the phase in the second packet of the buffer should be equivalent to the phase in the third packet of the previous buffer. One point in the region of overlap is chosen and the difference between the new and old data added on to the new data. This aligns the data and gives the final extracted phase values, as shown in figure 3.19.

This method of alignment is vulnerable to noise on the phase extraction, as can be




Figure 3.18: Spurious oscillations at the start and end of the extracted phase shown by looking at the phase differences.

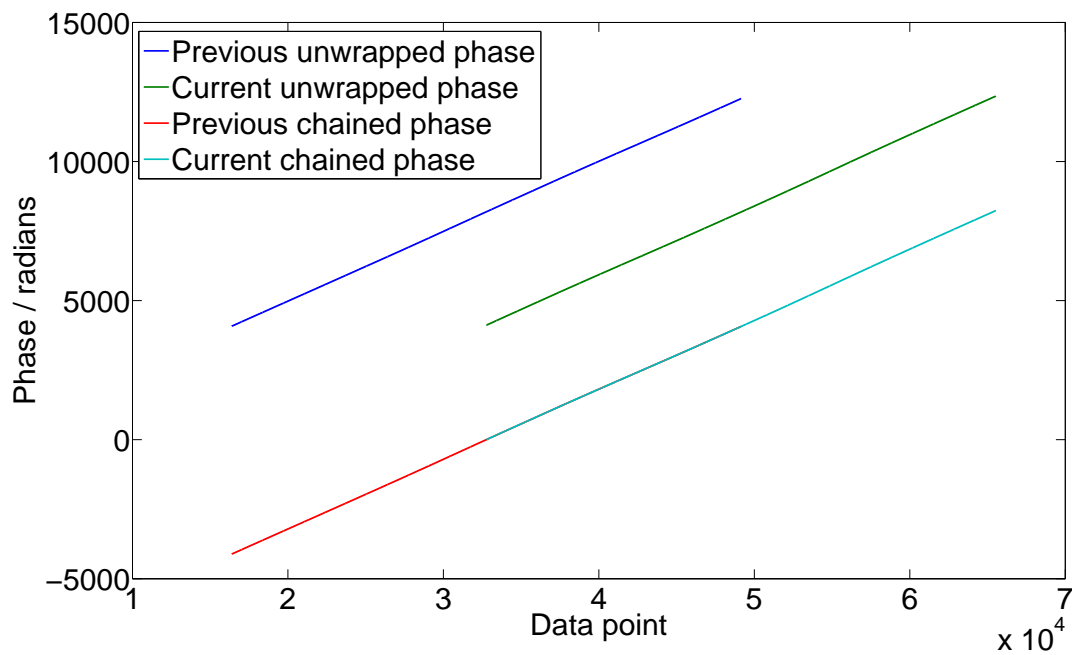


Figure 3.19: Chaining the current phase results to those in the previous iteration. The phase from the previous iteration has been centred on zero to aid visibility. Only the phases from the central two packets have been shown.

seen by the 6 mrad mean deviation in figure 3.20, which is a plot of the differences between the current and previous analysis iterations over the entire overlapping phase region. These offsets are likely to lead to random walk over time. Further study is required to determine the source of this discrepancy, and the implementation of a more complex chaining algorithm has been delayed until this has been performed.




Figure 3.20: Difference in phase between the current and previous iterations for the entire overlapping packet. Note that the final point is exactly zero.

3.2.3 Gas cell normalisation and filtering

The gas cell normalisation and filtering process removes noise and laser-induced power variations from gas cell data, which could cause phase errors during the fitting process described in section 3.2.4 It is performed on the Tesla GPU, and requires information from both the gas cell and laser power photodiodes. The normalisation process can be broken down into two steps:

1. Correction for changes in laser power.
2. Filtering the normalised data.

3.2.3.1 Laser power corrections

The first step in the analysis removes variations in laser power from the gas cell data and so calculates the transmission through the gas cell. Assuming that the detection chains of photodiodes, amplifiers and ADCs provide linear signals proportional to the detected optical power we record signals of the forms given in equations 3.2 and 3.3, where $S_{i,c}$ and $S_{i,p}$ are the signals recorded by the ADC at time index i for the gas cell (subscript c) and laser power (subscript p), I_i is the true laser intensity at time index i , g_c and g_p are the overall gains (whether from the optics, photodiodes or electronics) for each signal, o_c and o_p are the two offsets in the detection chain and finally T_i is the fractional transmission through the gas cell. T_i ranges from 0, which is total absorption, to 1, which is no absorption. The gains and offsets are assumed to be constant in all circumstances.

$$S_{i,c} = I_i g_c T_i + o_c \quad (3.2)$$

$$S_{i,p} = I_i g_p + o_p \quad (3.3)$$

Equation 3.3 is then re-arranged to give the laser intensity, as shown in equation 3.4.

$$I_i = \frac{S_{i,p} - o_p}{g_p} \quad (3.4)$$

This result can then be substituted into equation 3.2, giving equation 3.5.

$$S_{i,c} = \frac{S_{i,p} - o_p}{g_p} g_c T_i + o_c \quad (3.5)$$

By choosing a position in the data where there is no absorption in the gas cell, so $T_i = 1$, equation 3.5 can be rewritten as equation 3.6 giving a linear transformation between the

recorded power signal and the recorded gas cell signal.

$$S_{i,c} = \left(\frac{g_c}{g_p} \right) S_{i,p} + \left(o_c - o_p \frac{g_c}{g_p} \right) \quad (3.6)$$

Equation 3.6 can then be used in sections of the data where T_i is not known to calculate what the gas cell signal would be if $T_i = 1$. This prediction is referred to as $S'_{i,c}$ and is summarised in equations 3.7 to 3.8.

$$S'_{i,c} = \left(\frac{g_c}{g_p} \right) S_{i,p} + \left(o_c - o_p \frac{g_c}{g_p} \right) \quad (3.7)$$

$$S'_{i,c} \equiv I_i g_c + o_c \quad (3.8)$$

The final step is to subtract o_c from equations 3.2 and 3.8 and then divide them, giving equation 3.9 which simplifies to give T_i as shown in equation 3.10.

$$\frac{S_{i,c} - o_c}{S'_{i,c} - o_c} = \frac{I_i g_c T_i}{I_i g_c} \quad (3.9)$$

$$\frac{S_{i,c} - o_c}{S'_{i,c} - o_c} = T_i \quad (3.10)$$

To perform this calculation only 3 parameters (formed from the offsets and gradients of the two channels) are required to characterise and normalise the readout channels: o_c , (g_c/g_p) and $(o_c - o_p g_c/g_p)$. The parameters are calculated before starting to analyse the continuous FSI data, typically using a different data set. The offset of the gas cell channel, o_c , is calculated by taking an average of 1,000 data points from gas cell data taken using the same optical and electronic systems, while the laser has been blocked by the optical switches. The second two parameters are the gradient and offset of the straight line described in equation 3.6, and are calculated by fitting a straight line to 1,000 points of gas cell and power data taken during a change of switch state in a region where $T = 1$, providing a set of data running from the minimum to maximum possible values. An example of raw gas cell, power and corrected power data is given in figure 3.21.




Figure 3.21: Raw gas cell and power data, along with the predicted ($T=1$) gas cell data S' . The data are negative because the BNC input boards invert input signals, meaning that higher input voltages lead to lower recorded values. This does not affect the normalisation routine.

3.2.3.2 Filtering the normalised data

A Butterworth low-pass filter is applied to the data [60], to remove unwanted high frequency noise. The data is transformed between time and frequency domains with real-to-complex / complex-to-real FFTs using `cuFFT`, and the Butterworth filter has order 20 and is set at 11 kHz. Figure 3.22 shows the unfiltered and filtered normalised data, with a single feature highlighted in figure 3.23.

3.2.4 Absorption peak finding

The absorption peak finding algorithm searches for and performs fits of the gas absorption features. The algorithm uses the transmission function (from the normalisation process described in section 3.2.3), the laser sync out data and the reference interferometer phase (extracted as described in section 3.2.2), and has 4 steps:

1. Check laser light was present in the system.




Figure 3.22: Normalised and filtered gas cell data over a pair of laser scans. The large blue regions occur because the laser is blocked by the MEMS switches, providing low powers and so the division in equation 3.10 can produce largely varying values of the transmission with small noise-based variations in the raw gas cell and laser power data.




Figure 3.23: Normalised and filtered gas cell data over a single gas cell feature.

2. Search for an absorption feature.
3. Perform a fit of the gas cell feature.
4. Identify the gas cell feature.

The algorithm is performed on the CPU.

3.2.4.1 Checking for laser light

The peak finding algorithm begins by checking that the laser light was being transmitted by the MEMS switch, which is necessary to prevent a zero signal region being misidentified as a gas cell feature. This is done by looking at the laser sync out signal, which is also used to change the state of the switch. If the laser sync out signal has indicated that the laser had been blocked at any time in the last 24,576 data points then the transmission function is set to 2. The effect of this check was shown in figure 3.22. The effect of setting the transmission function to a high value is that the peak finding algorithm will ignore the data. We wait for 24,576 data points to allow the MEMS switch to change state, which takes longer than it takes for the laser sync out signal to change.

3.2.4.2 Search for an absorption feature

The peak finding algorithm breaks the normalised transmission data into small sections, typically of 4,096 data points (4 sections per packet) and finds the minimum value in each section. If this value falls below a certain cut-off level (typically 0.9) then the algorithm decides that it has found a gas cell feature. The algorithm will continue to search for lower values in subsequent 4,096 data point sections, until the minimum value rises. The time index of the minimum value is recorded and passed to the peak fitting part of the algorithm, which uses it to select the reference interferometer phase and normalised gas cell transmission data to fit as described below.

3.2.4.3 Fitting the gas cell feature

To use the peak in the length calculation we need to know the reference phase at its centre, as described in section 2.2.3. This is found by fitting a Gaussian function to a plot of the transmission function versus reference phase. The fit is performed using MnMigrad from Minuit2 [62] to minimise the sum of the square of the vertical residuals in a region around the minimum value from the search algorithm. The fit region consists of 1,201 points for laser 1 and 851 points for laser 2, which uses less points because of the difference in tuning speeds. The fit region is limited to improve the reliability of the fitting procedure, which is trying to fit a Gaussian function to a Voigt signal, which is not used because previous work has found it difficult to fit. An example fit region is shown in figure 3.24.




Figure 3.24: Example fit region for the gaussian gas cell fits, highlighted in red.

3.2.4.4 Identifying the gas cell feature

Individual gas cell features are difficult to distinguish from each other. Identification of each feature therefore requires external information. The technique used in the current

implementation of continuous FSI is to carefully select the start and end wavelengths of each laser scan, and so be able to predict the first peak that will be encountered in the scan. Using the current start wavelength values (described in section 3.3) the peaks that are first encountered are P1 or P22 for laser 1, with laser 2 encountering either R21 or P22 first. Once the first peak has been identified the algorithm simply counts the number of peaks and labels new peaks accordingly.

This method of peak identification does not work for laser 1 during the initial set-up period. This is because variations in the laser scan start times mean that one gas peak for laser 1 is not always recorded. In this period we need to count backwards from the last peak found. This is described in section 3.2.5.1.

3.2.5 Length calculation

The complexity in length calculation comes from calculating the various constants, such as the reference and transfer lengths, and performing handovers. This section works through the constants and describes how they are calculated and evolve over the course of a continuous FSI run. I will refer to the laser currently being used to calculate the length as the primary laser. The length calculation is performed on the CPU. The constants which are described are as follows:

1. Reference interferometer length.
2. Choice of gas cell absorption feature.
3. Initial transfer length.

3.2.5.1 Reference length calculation

The first constant to be calculated is the reference length. The analysis waits until laser 1 turns off, marking the end of the set-up period in figure 2.2, and then identifies the gas cell features that have been found for laser 1 up to that point by counting backwards

from the last feature (P1). This is necessary because the start time of laser 1 can vary with respect to the start time of laser 2 (which is used as the DAQ trigger) and so the first feature cannot be reliably predicted. Once the laser 1 gas cell features are assigned their correct frequencies a linear fit of phase versus frequency is performed for each laser and the reference length is calculated from the gradient as described in section 2.2.3. The laser 1 fit uses 10 gas features (P1 – P10) and laser 2 uses 22 gas features (R21 – R0). The fit is again performed in Minuit2 using MnMigrad to minimise the sum of the squared vertical residuals. The average of the reference lengths from the two lasers is taken as the true value.

3.2.5.2 Choice of gas cell absorption feature

The gas cell frequencies and phases are calculated in the gas absorption feature fitting process described in section 3.2.4. In principle any of the identified absorption frequencies could be used to calculate the laser frequencies. The length calculation routine always uses the gas cell feature which was encountered most recently in the calculations. The current gas cell feature is used to calculate all laser frequencies in the calculation, even the frequency at the transfer length which will have to be recalculated when a new feature is found.

3.2.5.3 Initial transfer length calculation

Finally the initial transfer length is calculated. The first 524,288 extracted phases in the run (2^{19} points, 32 packets) are recorded for both lasers in both the reference and measurement interferometers. This is the only time that phases are stored after the data used to generate them has passed from the analysis buffers. The stored phase data is used to calculate the transfer length at the first point in the third packet of the current buffer as per equation 2.7.

3.2.5.4 Length calculation and hand-overs

Once the transfer length is calculated and the phase and frequency of laser 2 at the same index (the first point in the third packet of the current buffer) are stored it is possible to begin length calculation as per equation 2.6. The transfer length needs to be re-evaluated each time a hand-over occurs over the course of the run. The handover occurs in the same packet in which the final primary laser gas cell feature is found. As soon as the primary laser encounters the final feature, one index in the **previous** packet is chosen to act as the new transfer index. All four phases (both lasers in both interferometers) are known at that index, as is the length of the measurement interferometer as calculated in the previous analysis iteration. This is all of the information needed to perform a hand-over, and so the analysis can be swap between the two lasers as described in section 2.2.4. The analysis is now able to continue until the next hand-over occurs.

3.3 Parameter selection

Continuous FSI as implemented in this chapter sets two constraints on the available laser parameters. These are that the interference signals from the two lasers must be separable, and that the time taken for each laser to perform a pair of forward and backward scans must be the same, with laser down-time small enough that handovers can be performed. These constraints are known as the separability constraint and the timing constraint respectively. For the preliminary measurements the data was only taken for up to a minute, relaxing the timing constraint slightly. The laser parameters chosen to satisfy the constraints are given in table 3.1.

3.3.1 Separability constraint

The separability constraint can be satisfied by setting the laser scanning speed ratio, which determines how well separated the interference signals are in fringe-frequency space.

Parameter	Laser 1 (slow)	Laser 2 (fast)
Scanning speed / nm/s	32	48
Low wavelength / nm	1542.15	1529.20
High wavelength / nm	1560.40	1560.26
Scan start time / s	-1.05	0
Dwell time / ms	10	0
Deceleration coefficient	3	3
Acceleration coefficient	3	3

Table 3.1: Laser tuning parameters chosen for continuous FSI.

The fringe frequency of the lasers is directly proportional both to the tuning speed and to the OPD of the interferometer. Initially the aim is to maximise the separation of the interference signals when the interferometer OPD is constant. Limits on moving interferometers are discussed in chapter 5.

The complicating factor is that with a very well aligned interferometer it is possible that a second round trip will be observed. A second round trip is caused by some of the light in the open part of the interferometer reflecting back from the surface of the optical fibre and so travelling through the interferometer for a second time. This generates a new interference signal at double the original fringe frequency, although at a much reduced intensity.

The presence of round trips means that the faster laser should not have a speed which is near to an integer multiple of the speed of the slow laser, especially not near to the doubled fringe-frequency. A suitable compromise has been found to have the fast laser scan at 1.5 times the speed of the slow laser. This positions the interference pattern of the second laser half-way between the first and second round trips of the slow laser, giving maximal separation of the interference signals. The spaces between higher order round trips were not used because the round trip interference signals are wider in fringe-frequency space, reducing the separability of the signals. By selecting the space between the first and second round trips we also maximise the motion tolerance of the measurement, which is described in detail in chapter 5.

This is a suitable point to explain the choice of laser tuning speeds. The fast laser

was the determining factor, and its speed was chosen to be 48 nm/s so that the optical switch can change state while the laser wavelength is between any pair of gas absorption features and then have enough data recorded at full laser power to give reliable phase information at the first gas absorption feature encountered. This is not a true constraint, but simplifies the absorption feature identification algorithm. The tuning ratio then sets the slow laser tuning speed to 32 nm/s.

3.3.2 Timing constraint

There are several parameters which affect the time taken for a laser to complete a pair of scans, other than the laser scanning speeds. These are the laser scanning range, the acceleration and deceleration coefficients, and the dwell time. The timing constraint is not sufficient to uniquely determine all of these parameters, allowing a certain amount of flexibility when setting them. The process used to determine them for the current implementation is described below.

3.3.2.1 Wavelength scanning ranges

The first scanning range to be set is the fast laser scanning range. This was set so that only the gas absorption features used in the analysis (R21 to P22) are covered. This choice simplifies identification of the features, although other choices could have been made.

The slow laser was set to cover only the P-branch of the Hydrogen Cyanide absorption features. The P-branch is almost two-thirds of the R21 to P22 wavelength range and so provided a suitable starting point, given the laser scanning speed ratio. The exact limits were fine tuned to satisfy the timing constraints once the other parameters had been set.

3.3.2.2 Acceleration and deceleration coefficients

The acceleration and deceleration coefficients determine how quickly the lasers slow down and restart the scanning process. The longer the time allowed to do this the more uniform the laser scanning curve. The limit is that the lasers must be able to restart within the time that the other laser is scanning, in order to allow handovers to be performed. The parameters for the fast laser were set to 3, which was the largest parameter that made its turnaround quickly enough that the handover procedure can occur. The slow laser parameters were then set to the same values so that its behaviour would be comparable with the fast laser.

3.3.2.3 Laser dwell

The dwell parameter introduces a wait at both the upper and lower wavelengths of the scan. A dwell of 10 ms for the slow laser was found to correct most of the timing discrepancy, with fine tuning of the wavelength scanning ranges finalising the parameters which meet the timing constraint.

Chapter 4

Continuous FSI - Verification

This chapter will describe the experiments performed to evaluate the performance of the continuous FSI system. The experiments were performed in three stages:

1. Absolute calibration, designed to check that the system calculates the correct length during the set up period.
2. Time synchronisation studies of the continuous FSI data acquisition system and reference system's ADC.
3. Differential calibration, to check that the length measurements remain accurate after the initial set up period.

We are unaware of any previous work attempting to perform a differential calibration of an interferometer in dynamic mode in the way described in stages 2 and 3. As such the methodologies developed in those sections (4.2 and 4.3) will be of use for the calibration of other systems capable of performing such measurements, rather than being specific to the continuous FSI system.

4.1 Absolute calibration

The absolute calibration provides the overall length scale of the technique. This calibration removes systematic effects from gas cell pressure uncertainty or fit bias caused by background effects, and ensures that the continuous FSI distance measurements are accurate. The scale is established by comparing the distance measurements from the continuous FSI system with those from a reference system, which provides measurements which are traceable to the SI metre.

4.1.1 The reference system

An Etalon A.G LaserTRACER [63] has been used as the reference system throughout this thesis. The LaserTRACER provides calibrated and traceable differential distance measurements using a stabilised He-Ne laser in a Michelson-style interferometer [11]. The He-Ne laser frequency has been calibrated against an Iodine stabilised laser which is traceable to frequency standards at the Physikalisch-Technische Bundesanstalt (PTB)¹. The LaserTRACER target tracking motors were disabled for the entire experimental period, except during the initial alignment procedure, in order to reduce vibration during the distance measurements. The disadvantage of this choice is that the interferometers automatic gain control will not work unless observing a moving target, so was manually deactivated when the target was stationary. The LaserTRACER environmental measurement station was also disabled, because it was reporting inaccurate pressures. External measurements of the atmospheric temperature, pressure and relative humidity were made using a custom system provided by Dr. Andrew Lewis from the National Physical Laboratory (NPL)², described in section 4.1.2.

¹Physikalisch-Technische Bundesanstalt, Bundesallee 100, 38116 Braunschweig.

²National Physical Laboratory, Hampton Road, Teddington, Middlesex, TW11 0LW.

4.1.2 Experimental layout

The absolute calibration experiment was performed using a motion stage with a range of travel of approximately 2 m. The experimental layout is shown in figure 4.1.




Figure 4.1: Experimental layout for the absolute calibration experiments.

The LaserTRACER is mounted at one end of an optical table and adjusted so that the measurement beam is horizontal, at which point the motors were deactivated using the emergency stop. The LaserTRACER beam provides the optical axis for the experiment. A dichroic mirror, with coatings that transmit the LaserTRACERs light (with wavelengths around 633 nm) but reflect light from the FSI system (with wavelengths around 1550 nm), is placed into the beam. This mirror is used to combine the LaserTRACER and continuous FSI beams.

The motion stage, with a 0.5" retroreflector mounted on the carriage via a magnetic mount which was glued into place, is then aligned to the LaserTRACER beam over the entire range of travel. The continuous FSI beam is then added into the experiment,

launching from a standard fibre-collimator system as described in section 3.1.5. The infra-red light is made co-linear with the LaserTRACER beam using a gold coated mirror to position the IR beam at the same point through which the LaserTRACER beam exits the dichroic mirror, and then using the dichroic mirror to adjust the pointing of the infra-red beam.

The continuous FSI beam is launched from a position close to the far end of the motion stage. This was done because at slow laser tuning speeds, such as those chosen in section 3.3, smaller interferometers generate interference fringes relatively slowly. This can lead to difficulty separating the signals from the two lasers because of the frequency resolution of the Fourier transformations. This is shown in figure 4.2, which shows fringe frequency spectra for a series of interferometers with lengths between 1 m and 4 m (2 m to 8 m OPD). The 1 m interferometer almost experiences a peak clash and would be difficult for the automatic filtering to separate.

The usual remedy to this problem would be to increase the laser tuning speeds, but this was not attempted with the current continuous FSI system, in order to maintain agreement with the tuning parameter limits described in section 3.3. The approach was instead to avoid short measurement distances, which also increases the fringe frequency, leading to the layout described above. Future work should aim to address this problem (see chapter 7). In previous experiments it was found to be difficult to perform LaserTRACER measurements at the longer distances required for continuous FSI while using a dichroic mirror, which led to the decision to allow the two systems to see different optical paths. The difference in path between the two systems will lead to variations between the two measured distances, as the properties of the air that the measurements are performed in will be slightly different. If possible future experiments should avoid this atypical layout.

Seven calibration experiments were performed. In each experiment distance measurements were taken at up to seven positions on the motion stage, in the order labelled in figure 4.1. The order of the measurements was intended to reduce any effects from the




Figure 4.2: Average fringe frequency spectra for a series of interferometers. The peaks are labelled with the interferometer length (1–4) and S for the slow laser or F for the fast laser. The 1 m interferometer (2 m OPD) is close to a peak clash.

direction of travel of the motion stage. The motion stage was controlled via RS-232, with the stage position adjusted manually. This means that the motion stage was not stopped in exactly the same position in each calibration experiment, except at positions 1 and 4 which are defined by the ends of the motion stage.

Thirty distance measurements were taken at each position. The continuous FSI system recorded 0.319 seconds of data (corresponding to 54 packets) in each measurement, which is enough to cover the set up period. The transfer length as defined in equation 2.7 was used as the measured length for the purposes of calibration. The LaserTRACER was triggered using the laser 2 sync output, and recorded a single measurement per continuous FSI run.

The environment monitor, developed by Dr. Andrew Lewis at the National Physical Laboratory [64], was based on a Raspberry Pi computer. It had two temperature sensors, one pressure sensor and a humidity sensor. The measured values are used to correct for the refractive index of air. The specifications of the sensors are given in table 4.1. One atmospheric measurement was recorded per position in each calibration, with the sensors positioned centrally near to the FSI-only beam path as shown in figure 4.1.

Sensor	Maximum error
Temperature	0.14 C
Pressure	2.1 hPa
Relative humidity	1.4 %

Table 4.1: Atmospheric sensor parameters [64]

4.1.3 Data integrity

Before the analysis was performed some of the measurements had to be removed. In one case this was done during the measurement process, due to a hardware issue, but otherwise measurements were cut during the analysis process, which revealed more subtle effects. These effects are described in sections 4.1.5.1 and 4.1.5.2.

4.1.3.1 Hardware failure

Data was not recorded for position 7 in calibration 7. This was because an electrical noise spike, caused by deactivating the motion stage power supply, interrupted the LaserTRACER USB communications. This disrupted the differential distance measurements that the LaserTRACER provides, and so any measurements would have an arbitrary offset from the data recorded at positions 1–6. Any such offset would prevent useful analysis, so measurements at position 7 for calibration 7 were abandoned.

4.1.3.2 Analysis-level data cuts

The data removed at a software level was position 7 in calibration 1, and one point in the position 7 data in calibration 2. The effect which required these cuts was similar in both cases, and came from the automatic filtering in the phase extraction of the measurement interferometer. In the case of position 7 in calibration 1, laser 2 was providing reduced power. This means that the laser 2 primary peak was smaller than the round trip of laser 1, causing it to be misidentified. It would technically be possible to adjust the filtering algorithm to avoid this problem, but this has not been done because failures of this nature highlight problems with laser tuning in an obvious way. If the filtering had been corrected these tuning problems could make the distance measurements unreliable without being obviously inaccurate, causing further problems in later analysis. An example of the problem is shown in figure 4.3.

The same problem recurred in calibration 2 at one point in position 7. Only one of the automatic filter attempts in the run failed in this way, but it was deemed safer to simply remove the run entirely.

4.1.4 Calculation of the calibration constants

As observed during dynamic FSI measurements performed by Dale [3], a linear relation between the continuous FSI measurements and the LaserTRACER measurements




Figure 4.3: Example measurement interferometer fringe frequency spectrum for position 7 in calibration 1. The third spectral feature (1'), which is a double round trip of the first spectral feature (1), both from the slow laser, is misidentified as the second spectral feature (2), which is the primary feature of the fast laser.

was expected. The form of the relation is given in equation 4.1, where L_{Tracer} is the LaserTRACER distance measurement, L_{CFEI} is the continuous FSI distance measurement, M_{cal} is the calibration constant and C_{cal} is the offset. The offset has no physical significance in the current experiment, because the LaserTRACER provides differential measurements and the offset depends on the position which was taken to have zero displacement.

$$L_{Tracer} = M_{cal}L_{CFEI} + C_{cal} \quad (4.1)$$

The calibration constant was calculated in two stages. The first stage is to perform a linear fit to the length measured without correction for refractive indices (raw length). The fit was performed in Matlab [65] using a non-linear least squares algorithm [66]. As shown in equations 4.2 to 4.4, the gradient of this fit (m_{raw}) is related to the calibration constant by the ratio of the refractive indices at the two wavelengths. Lower case letters indicate raw length values, with n_{633} and n_{1545} the refractive indices at 633 nm and 1545 nm respectively.

$$\frac{l_{Tracer}}{n_{633}} = M_{cal} \frac{l_{CFEI}}{n_{1545}} + C_{cal} \quad (4.2)$$

$$l_{Tracer} = \frac{M_{cal}n_{633}}{n_{1545}}l_{CFEI} + C_{cal}n_{633} \quad (4.3)$$

$$m_{raw} = \frac{M_{cal}n_{633}}{n_{1545}} \quad (4.4)$$

The second stage of the analysis applies the effect of the refractive index to the measurements. The Ciddor equation [67] was used to calculate the refractive indices, which were applied as a ratio. This approach was chosen in place of applying the correction to the raw data because the two refractive indices will have many correlated errors as they are based on the same atmospheric measurements.

4.1.5 Experimental uncertainties

There are several uncertainties which need to be taken in to account in the calculation of the calibration constants. The first pair of these are the uncertainties in the measured length from the continuous FSI system and from the LaserTRACER. These were calculated at each position by taking the standard deviation of the 30 measured values. These uncertainties are used as inputs to the non-linear least squares fit so that it produces an accurate error estimate on the gradient. The uncertainties calculated for the continuous FSI and LaserTRACER measurements are shown in figures 4.4 and 4.5 below. One interesting feature of the continuous FSI uncertainties is that the standard deviations increase at shortest distances. This could be indirect evidence for the overlapping spectral features problem described in section 4.1.2 and shown in figure 4.2.




Figure 4.4: Standard deviations of the continuous FSI distance measurements during the calibration experiments.

The remaining sources of uncertainty influence the refractive index ratio, and include the atmospheric parameters (temperature, pressure, humidity and the concentration of CO_2) along with variations caused by the changing wavelength of the continuous FSI




Figure 4.5: Standard deviations of the LaserTRACER distance measurements during the calibration experiments.

lasers. No attempt was made to analytically propagate the errors through the ratio of Ciddor refractive index equations, so the effect of uncertainties were estimated numerically by varying the parameter of interest while keeping the other parameters constant. A linear fit of the parameter against the refractive index ratio then gives the sensitivity which can be used in standard error propagation calculations.

The values of the atmospheric measurements, averaged over all positions in a given calibration experiment, are shown in figure 4.6. The refractive index ratio changes caused by changes in the various parameters are shown in figure 4.7, with gradients and error propagation described in table 4.2.

The concentration of CO_2 was not measured and is assumed to be 450 parts per million (ppm), which is the nominal value used by Ciddor [67]. An uncertainty range of ± 150 ppm for the concentration of CO_2 was assumed, as Ciddor also uses 600 ppm as a maximum value for comparisons between alternative refractive index calculations for precise length interferometry and geodetic surveying. These values seemed reasonable as upper and lower bounds for the concentration values. The upper and lower bounds for

the wavelength values were set to be the start and end wavelengths for the laser 2 scans.

Parameter	Uncertainty	Sensitivity coefficient	Resultant uncertainty in $\frac{n_{1545}}{n_{633}}$
Temperature	± 0.1 °C	1.016×10^{-8} / °C	1.006×10^{-9}
Pressure	± 210 Pa	-3.137×10^{-11} / Pa	6.589×10^{-9}
Relative humidity	± 1.4 %	-2.607×10^{-10} / %	3.650×10^{-10}
Wavelength	± 16 nm	-0.824 / m	1.319×10^{-8}
Concentration of CO ₂	± 150 ppm	-1.687×10^{-12} / ppm	2.531×10^{-10}
n_{IR}/n_{633}	-	-	1.478×10^{-8}

Table 4.2: Error propagation into the refractive index ratio. The temperature was an average of two sensors, so is reduced compared with table 4.1. The wavelength and CO₂ concentration uncertainties were taken as half of the range they are expected to cover.

4.1.5.1 Discrepancies in calibration 5

The measurement length data from calibration 5 was not self consistent, which became apparent when residuals from the straight line fit, before refractive index correction, were plotted (see figure 4.8). The residuals for the other calibrations do not exceed $16.6 \mu\text{m}$, whereas the calibration 5 data shows excursions of the order 1.93 mm .

The cause of this discrepancy is a jump in the length data between position 5 and position 6. By using only data from the first 5 positions in the calibration, the residuals become comparable with the other calibrations. This improvement is shown in figure 4.9. It is not possible to conclude which of the two distance measurement systems records the distance jump because the calibration experiments were performed at the National Physical Laboratory but the analysis was performed in Oxford after the equipment had been shipped back. This prevented further experiments to determine the cause of the problem from being performed.

4.1.5.2 Calibration constant results

Once the erroneous data had been removed the calibration constants were calculated. A summary of the measurements used in the calculation are described in table 4.3. The

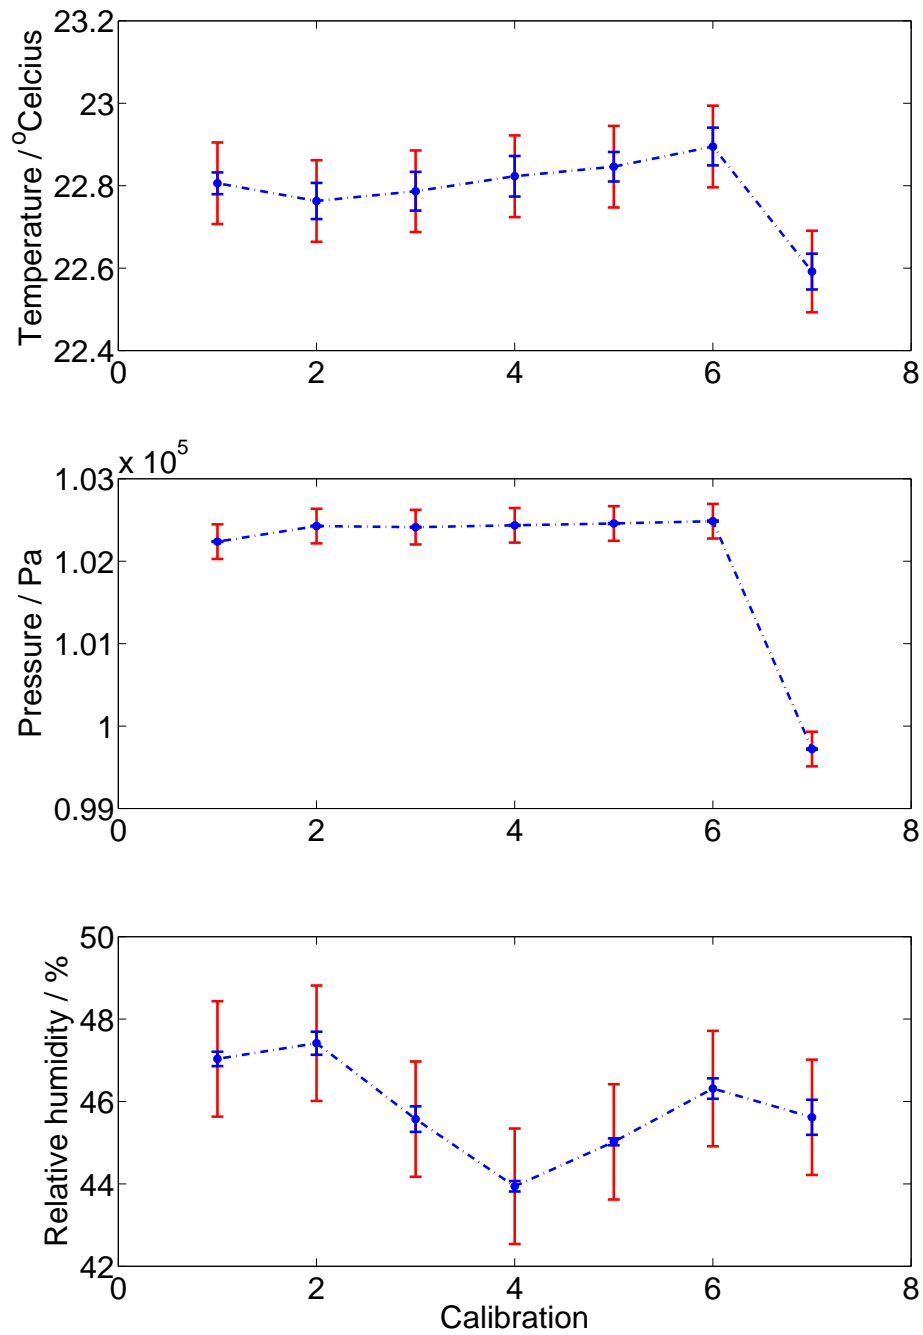


Figure 4.6: Measured environmental parameters during the calibration experiments. Blue error bars show the standard deviation over the 7 positions, with the red bars showing the errors described in table 4.1, which are much larger.




Figure 4.7: Changes in the refractive index ratio with environmental parameters and IR laser wavelength.




Figure 4.8: Residuals from the straight line fits for all calibration sets, before refractive index corrections were made.




Figure 4.9: Residuals from the straight line fits between L_{Tracer} and L_{CFSI} for all calibration sets, before refractive index corrections were made. The final two positions of calibration 5 have been ignored, leaving only the data from positions 1–5.

Calibration	Positions	Total measurement points
1	1-6	180
2	1-7	209
3	1-7	210
4	1-7	210
5	1-5	150
6	1-7	210
7	1-6	180

Table 4.3: Number of measurement points in the data sets used to perform the absolute calibration of the continuous FSI system.

seven calibration constants derived from the reliable data and their weighted mean are shown in figure 4.10. The error bars show the one-sigma errors.




Figure 4.10: Calibration values using the available data as described in table 4.3.

It is apparent that the calibration constants are neither consistent with each other nor with the weighted mean value. The values broadly fall into two sets, with calibrations 1, 5 and 7 forming one consistent set and calibrations 2, 3, 4 and 6 forming another. The only obvious difference between these two sets is that the first data set does not include any data from position 7, whereas the second does. If the position 7 data is removed from all sets the calibration constants become consistent. The data used in this case is

detailed in table 4.4 and the calibration constant results shown in figure 4.11.

Calibration	Positions	Total measurement points
1	1-6	180
2	1-6	180
3	1-6	180
4	1-6	180
5	1-5	150
6	1-6	180
7	1-6	180

Table 4.4: Number of measurement points in the data sets used to perform the final absolute calibration of the continuous FSI system.




Figure 4.11: Calibration values calculated using data which ignored position 7 values. The relevant data is described in table 4.4.

This is an entirely unsatisfactory approach but, as described in section 4.1.5.1 above, it was not possible to perform further experiments to determine the cause of the discrepancy. Those measurements which could be performed, such as checking the straightness of the motion stage, have not provided a conclusive explanation for the discrepancy. The weighted mean of the calibration constants calculated without the position 7 data, along with the one-sigma uncertainty, is given in table 4.5. This value was taken to be the

absolute calibration constant for the continuous FSI system. The range of the seven calibration constant measurements is also given to provide a more conservative estimate of the uncertainty.

Calibration constant	0.999998496
Uncertainty	8.919×10^{-8}
Range of measurements	8.668×10^{-7}

Table 4.5: Final absolute calibration constant, calculated without the position 7 data as described in table 4.4 and figure 4.11.

4.2 ADC sampling frequency synchronisation

The ADC synchronisation is performed to ensure that the continuous FSI and Laser-TRACER ADC sampling speeds are consistent with each other. This is important for the differential calibration which will be described in section 4.3, because differences in timing would be interpreted as differences in the distance travelled by the target retroreflector, and so lead to an incorrect determination of the differential calibration constant.

4.2.1 Experimental layout

To perform ADC synchronisation between the FSI system and the reference system they must simultaneously observe a target with a time-varying position. The sampling frequency of one measurement (in this case the FSI system) and the time offset between the two measurements can then be adjusted until they provide a consistent observation of the moving target. Of those two parameters only the sampling frequency adjustment is applicable to future measurements as it should depend on the two sets of ADC hardware and so remain constant over all runs. The time offset will vary on a run by run basis, even when using triggers, because of the variations in CFSI laser timing but is a necessary parameter in any successful fit. One further parameter, the distance offset, is also determined to facilitate visual checking of the data alignment.

The ADC synchronisation experiments were performed by observing periodic motion of the measurement interferometer. The two measurement beams are combined in the same manner as in the absolute calibration experiments, but in this case the target retroreflector was mounted on a Thorlabs PAS001 piezoelectric actuator [68] rather than a motion stage. The actuator was driven by a TTI TG550 function generator [69] amplified by a power supply custom built by the central electronics group in Oxford. The system layout is shown in figure 4.12.




Figure 4.12: Experimental layout for the time base calibration experiment.

Four experiments were performed, each recording continuous FSI data for 59.8 seconds (10,145 packets). The LaserTRACER recorded 32,768 measurements at a sampling frequency of 781.250 Hz, giving a data acquisition period of 41.9 seconds. Both systems were triggered by the sync output from laser 2. The actuator was driven sinusoidally by a voltage varying between 10 V and 30 V. During each experiment the actuator was either activated or deactivated within the acquisition period, giving oscillatory motion at either the end or the beginning of the run respectively. The change of state was implemented in order to provide an unambiguous time at which to align the two measurements, which would be difficult to do with a purely sinusoidal signal. See table 4.6 for details.

ADC synchronisation experiment	Drive frequency / Hz	State change
1	70.94	Deactivation
2	70.97	Activation
3	30.40	Deactivation
4	30.40	Activation

Table 4.6: Piezo actuator parameters during the ADC synchronisation experiments.

4.2.2 Time axis and sampling frequency synchronisation

The time axis of the measurements was determined by the LaserTRACER. For a given LaserTracer data point j the time T_j is given by equation 4.5, where f_{tracer} is the nominal LaserTRACER sampling frequency.

$$T_j = \frac{j}{f_{tracer}} \quad (4.5)$$

The time T_i of a CFSI data point i is given by equation 4.6 where f_{CFSI} is the nominal CFSI sampling frequency, s is the stretch parameter used to adjust the sampling frequency, and i_{centre} is the index at which the offset time T_{centre} is determined. In the work described in this chapter i_{centre} has been chosen to be the central data point of the continuous FSI data set, but it could in principle be any data point.

$$T_i = \frac{(i - i_{centre})s}{f_{CFSI}} + T_{centre} \quad (4.6)$$

Of the parameters in equation 4.6 s and T_{centre} need to be determined by comparing the two length measurements. To perform the comparison we need to determine which CFSI point corresponds to each LaserTRACER point, which is achieved by equating T_i and T_j and rearranging for i resulting in equation 4.7.

$$i = \frac{(j/f_{tracer}) - T_{centre}}{s/f_{CFSI}} + i_{centre} \quad (4.7)$$

It should be noted that when calculated using equation 4.7 index i is often non-integer. When this was the case the continuous FSI length was interpolated between the mea-

surements at the adjacent indices. With the lengths from both systems determined at each LaserTRACER index the difference between the measurement systems was calculated according to equation 4.8, where D_o is the distance offset between the measurement systems.

$$L_{tracer}(j) = L_{CFSI}(i(j)) - D_o \quad (4.8)$$

The difference between the two lengths was minimised by varying s , T_{centre} and D_o . The procedure used was as follows:

1. Fit the continuous FSI data to the LaserTRACER data using Minuit.
2. Perform a parameter scan to verify that the best minimum has been reached, and to provide uncertainty estimates.
3. Combine the measurements of the stretch parameter to give the final calibration constant.

4.2.2.1 Fit the oscillation data using Minuit

The Minuit fitter minimises the difference between $L_{tracer,j}$ and $L_{CFSI,i}$ using the simplex routines. During a fit the CFSI time axis is changing, and so any interpolations between continuous FSI data points are recalculated in each iteration. The fit is performed several times, with the three parameters fixed and released in turn as described below. All of the LaserTRACER data is used in the fit.

1. Fit T_{centre} . D_o and s are fixed.
2. Fit s . D_o and T_{centre} are fixed.
3. Fit D_o . T_{centre} and s are fixed.
4. Fit all parameters.
5. Fit T_{centre} . D_o and s are fixed.

6. Fit s . D_o and T_{centre} are fixed.
7. Fit all parameters.

4.2.2.2 Perform a parameter scan and extract a new minimum with error bounds

The fitting procedure would often struggle to reach the best minimum in the data. To improve the results a two-dimensional parameter scan, centred on the fit results, is performed by varying T_{centre} and s while holding D_o constant. The χ^2 per degree of freedom was calculated and the minimum value taken as the optimal result. The 1σ uncertainty was calculated by looking for the contour where the χ^2 per degree of freedom increased by one from this value. The central time is varied by ± 2 ms and the stretch is varied by $\pm 5 \times 10^{-5}$ in calibrations 1, 2 and 4. Calibration 3 required a much larger region of ± 6 ms around the central time and $\pm 15 \times 10^{-5}$ in stretch to provide a useful result.

The χ^2 per degree of freedom calculated in these parameter scans are shown for each calibration in figures 4.13 to 4.16. The best values for the stretch are given, along with uncertainties, in table 4.7. The corresponding values and uncertainties for the central time parameter are given in table 4.8, which shows that the uncertainties are typically below the LaserTRACER time increment (1.28 ms).

Synchronisation experiment	Stretch	σ^-	σ^+
1	0.99996875	1.85×10^{-5}	4.68×10^{-5}
2	0.99997981	4.70×10^{-6}	6.51×10^{-6}
3	0.99997345	4.65×10^{-5}	6.19×10^{-5}
4	0.99998008	1.19×10^{-5}	1.28×10^{-5}

Table 4.7: Stretch parameters and associated uncertainties calculated from the parameter scans.

The length measurements from both systems, with the time axis and offset of the continuous FSI measurements corrected, are shown in figure 4.17. The alignment of the calibration 4 data at the time that the actuator was activated is shown in figure 4.18. The continued consistency of the continuous FSI and LaserTRACER timing up until the




Figure 4.13: Scan of central time and stretch parameter for calibration 1. χ^2 per degree of freedom is denoted by colour, and capped at the minimum value plus 1 to visualise the 1σ uncertainty region.




Figure 4.14: Scan of central time and stretch parameter for calibration 2. χ^2 per degree of freedom is denoted by colour, and capped at the minimum value plus 1 to visualise the 1σ uncertainty region.

Synchronisation experiment	Central time / s	σ^- / s	σ^+ / s
1	29.3999	0.46×10^{-3}	1.15×10^{-3}
2	29.4002	0.05×10^{-3}	0.06×10^{-3}
3	29.3971	1.15×10^{-3}	1.54×10^{-3}
4	29.4002	0.13×10^{-3}	0.12×10^{-3}

Table 4.8: Central times and associated uncertainties calculated from the parameter scans.




Figure 4.15: Scan of central time and stretch parameter for calibration 3. χ^2 per degree of freedom is denoted by colour, and capped at the minimum value plus 1 to visualise the 1σ uncertainty region.




Figure 4.16: Scan of central time and stretch parameter for calibration 4. χ^2 per degree of freedom is denoted by colour, and capped at the minimum value plus 1 to visualise the 1σ uncertainty region.

end of the LaserTRACER data is shown in figure 4.19. These figures confirm that the fit results provide a realistic description of the time axis.

Several points should be noted about the data. Firstly, the two calibrations with the largest uncertainty in the stretch parameter, calibrations 1 and 3, have the shortest period of time with an oscillating target. Future experiments should aim to record data with oscillations present for as long as possible, as in calibrations 2 and 4. Secondly, the distance offset does not appear to be accurate in all of the calibrations. This is not a concern, because the offset is specific to each individual experiment and has no physical significance beyond it. The offset discrepancy is also likely to have been influenced by the final observation, which is that in most cases the continuous FSI data drifts with respect to the LaserTRACER over the course of the run. This effect will be discussed in section 4.3 and chapter 5.

4.2.2.3 Combination of the stretch measurements

The asymmetric uncertainties on the stretch parameter, as given in table 4.7, make the calculation of a weighted mean and error propagation more involved than would be the case with symmetric uncertainties. The procedure for combination of results described by Barlow [70] has been used, using a variable-width Gaussian parametrisation of the likelihood function, with σ varying linearly with stretch. The Barlow approach requires the solution to be found using an iterative approach, so again a Minuit2 fitter was used. The combined result and uncertainties are given in table 4.9, and a comparison with the individual measurements shown in figure 4.20.

Stretch	σ^-	σ^+
0.99997957	4.32×10^{-6}	5.5×10^{-6}

Table 4.9: Final stretch parameter and associated uncertainties.

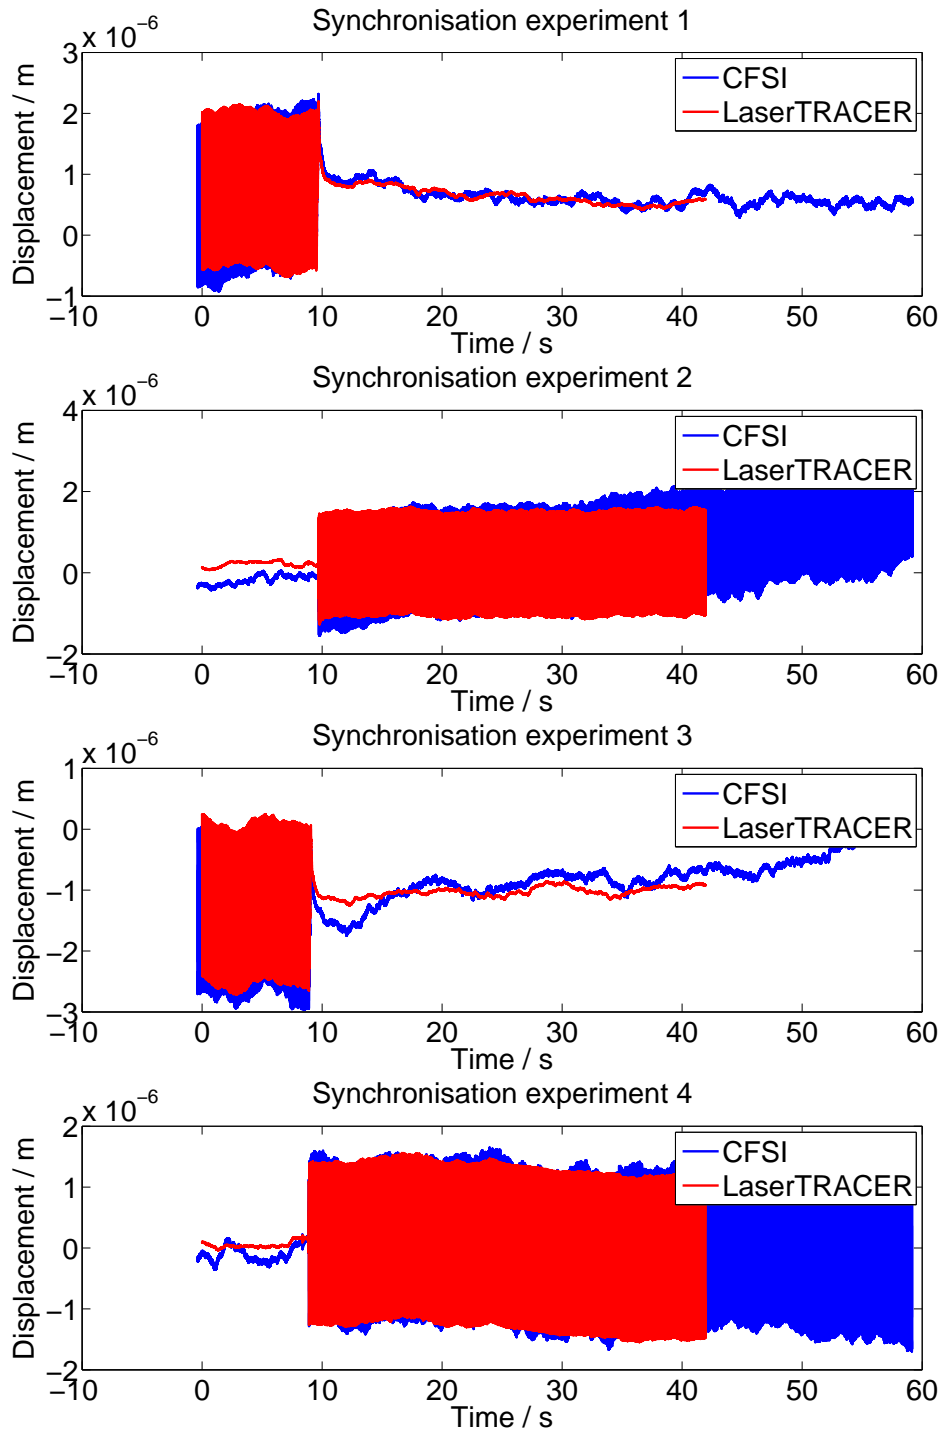


Figure 4.17: Oscillation measurements aligned using the fit parameters. All of the Laser-TRACER data is used in the fit.




Figure 4.18: Alignment of the oscillation data for calibration 4 at the time that the oscillations were activated.




Figure 4.19: Alignment of the oscillation data for calibration 4 at the end of the Laser-TRACER data.




Figure 4.20: Combined result for the stretch parameter, with the individual measurements included to show the consistency between the results.

4.3 Differential length calibration

The differential length calibration procedure is designed to ensure that the continuous FSI distance measurement remains accurate after the initial transfer length calculation. The procedure looks for deviations from the reference system measurements. Many effects such as random walk from the phase extraction or successive handovers should contribute to this.

4.3.1 Experimental layout

The experimental layout used for the differential length calibration experiments was identical to that of the absolute length calibration experiments (figure 4.1). The equipment was not altered between the two calibrations. The differential calibration was performed by observing the target retroreflector with both continuous FSI and the LaserTRACER as it was moved along the motion stage. As was the case in the time synchronisation experiments, the continuous FSI system recorded 10,145 packets of data (59.8 seconds)

and the LaserTRACER recorded 32,768 measurements at 781.250 Hz (41.9 seconds) per run, with both systems triggered by the laser 2 sync output. As described in section 4.2.1, this triggering does not remove the need to align the two measurements in time because of variations in CFSI laser timing.

Two styles of experiment were performed. In the first, known as “forward” experiments, the target was initially stationary and was then moved in the direction of increasing OPD. The forward experiments typically started near stage position 2 and moved towards position 6. Measurements in which the OPD was reduced are known as “backward” experiments. The backward experiments started near position 2 and moved towards position 5. In all cases the stage moved at approximately 2 mm/s, and the motion was started manually in order to overcome communication problems with its controller which leads to variation in the start time of the stage movement relative to the trigger time.

4.3.2 Data alignment in time

To observe deviations between the continuous FSI and reference system measurements they must first be aligned in time. This is achieved using a similar routine to that described in section 4.2.2. The key differences are that the stretch parameter is fixed at the value calculated by the time synchronisation experiments, leading to a simplified fitting procedure, and that only a fraction of the data before the start of the motion is used for the fits. The new fit procedure is described below.

1. Fit D_o . T_{centre} is fixed.
2. Scan T_{centre} and set limits at the minimum ± 5 ms. D_o is fixed.
3. Fit T_{centre} . D_o is fixed.
4. Fit D_o . T_{centre} is fixed.
5. Fit T_{centre} . D_o is fixed.

The fit region was determined by looking for the point at which the LaserTRACER data first increases or decreases by $1 \mu m$, and fitting from the start of the data up to 300 points before that point. Fitting this region was possible because of background vibrations in the measurement interferometer, which provide sufficient variation to extract the required parameters. A subsection of this fit region (after alignment) is shown in figure 4.21, with the calculated T_{centre} values and errors shown in figures 4.22 and 4.23. Examples of the overall alignment of forward and backward data are shown in figures 4.24 and 4.25 respectively.




Figure 4.21: A subset of the fit region used to align the continuous FSI and LaserTRACER data, showing the small oscillation as measured by both systems.

4.3.3 Comparison of the measurements

Once the two data sets were aligned in time it was possible to make direct comparisons between the continuous FSI and LaserTRACER reference measurements. The distance measured by the continuous FSI system at each LaserTRACER measurement index was interpolated as described in the fitting procedure, and the difference between the two




Figure 4.22: T_{centre} values for the forward experiments.




Figure 4.23: T_{centre} values for the backward experiments.




Figure 4.24: An example of the alignment of the continuous FSI and LaserTRACER data in a forward experiment.




Figure 4.25: An example of the alignment of the continuous FSI and LaserTRACER data in a backward experiment.

measurements calculated. A series of examples of distance differences from the forward and backward experiments are shown in figures 4.26 and 4.27 respectively.

The difference plots show two main effects: drift and fluctuations. The fluctuations, of which experiment 15 in figure 4.26 is a clear example, are caused by physical length changes in the interferometer. This is best demonstrated by the results in chapter 5, so will be discussed there.

The drift between the two measurements varies on a run by run basis. Some runs, such as experiment 25 in figure 4.26 and experiment 20 in figure 4.27, experience only minor drifts and are dominated by the fluctuations. Other experiments observe drifts of almost $2 \mu\text{m}$, with the CFSI measurements observing both excess length (as in experiment 26 in figure 4.26) and length deficits (as in experiment 4 in figure 4.26 and experiment 13 in figure 4.27).

The two likely causes of such drifts within the continuous FSI technique are errors in the calculation of the transfer length (through its interaction with laser frequency in equations 2.5 and 2.6) or the reference length (which determines frequency changes), both of which have been investigated. Experiment 4 in figure 4.26 was chosen for the investigation, in which the transfer length or reference length was adjusted to eliminate the observed drift.

The transfer length was found to require a change of $20\text{--}22 \mu\text{m}$ ($40\text{--}44 \mu\text{m}$ OPD) to remove the overall drift, as shown in figure 4.28. This change corresponds to a change of 10–11 standard deviations from the calculated value. While the drift has been successfully removed a new effect becomes apparent, as the length measurement acquires a significant zig-zag structure. This structure is caused by the changing laser frequency scanning direction throughout the run and indicates that the transfer length is incorrect.

A similar calculation for the reference length found that it was necessary for an adjustment of order millimetres to be made to remove the gradient, which is thousands of standard deviations away from the original value, therefore neither the transfer length nor the reference length provide a suitable explanation for the variations in drift observed.




Figure 4.26: A series of examples of the difference between the continuous FSI and LaserTRACER measurements for the forward experiments.




Figure 4.27: A series of examples of the difference between the continuous FSI and LaserTRACER measurements for the backward experiments.




Figure 4.28: Drift in forward experiment 4 as the transfer length is adjusted.

The drift therefore remains unexplained, but physical drifts between the optical paths of the two measurement systems have not been excluded. It was not possible to observe the effect of any scale error because of the influence of the drift on the measurements.

Chapter 5

Enhanced Continuous FSI

The previous chapters in this thesis have described and demonstrated a method of chaining dynamic FSI scans together to give contiguous distance measurements over a theoretically indefinite period. However, the maximum motion speed of the target is limited by the implementation given in chapter 3. The mechanism which leads to the speed limitation is described in the following chapter, and a method to overcome the problem is proposed and demonstrated.

5.1 Speed limitations

The interferometer design described in chapter 3 results in the interference signals for both scanning lasers being observed by the same photodiode. This places a series of constraints on the form of the recorded signal, and so on the maximum motion tolerance in the measurement interferometer. Firstly, the interference signals from each laser must be separable in fringe frequency in order for the Hilbert transform to extract the correct interferometric phases. The Hilbert transform also demands that the phase velocity does not change sign, as described in section 3.2.2, so the fringe rate must remain above zero. Finally, a minimum number of points per interference fringe are required in order for phase unwrapping to work reliably, which provides a maximum fringe frequency which

should not be exceeded. This work uses a conservative figure of 8 points per fringe, leading to a maximum fringe frequency of $50/(18 \times 8)$ MHz = 347 kHz with the current DAQ system. The maximum frequency allowed by the points per fringe limit will be denoted by f_{ppf} .

5.1.1 Fringe frequency

The formula for the instantaneous fringe rate, referred to as fringe frequency, can be determined by differentiating a phase equation (for example equation 2.1) with respect to time, giving the rate of change of phase $\dot{\phi}$ as shown in equation 5.1.

$$\dot{\phi} = \frac{2\pi}{c} (L\dot{\nu} + \nu\dot{L}) \quad (5.1)$$

One fringe is given by a change of phase of 2π radians, and so the fringe frequency F can be found by dividing equation 5.1 by 2π and taking the absolute value. The absolute value must be taken because the fringe rate is identical in the case of both increasing and decreasing phase. This gives equation 5.2.

$$F = \left| \frac{L\dot{\nu}}{c} + \frac{\nu\dot{L}}{c} \right| \quad (5.2)$$

The first term, containing the contribution to the fringe rate from the changing laser frequency, and will be referred to as the FSI term and denoted by I_1 or I_2 , where the subscript number is the laser number. As has been the convention in previous chapters the slower laser will be referred to as laser 1 and the faster laser as laser 2. When the interferometer is stationary, as is universally the case in the reference interferometer, only the FSI term contributes to the fringe frequency. This demands that the two laser scanning speeds be different to prevent them producing signals with equal fringe frequency. This limitation also means that only laser 1 can impinge upon the zero fringe rate limit, and only laser 2 can impinge upon the points per fringe limit.

The second term of equation 5.2 adds in the effect of changing interferometer length, and will be referred to as the motion term. For the calculation of the motion term we will assume that both lasers are on average emitting light at the same optical frequency, which simplifies the problem significantly by making the motion term equal for both lasers. This is true to a good approximation, as the optical frequency of each laser only varies by approximately 2% and both lasers cover a similar range. The motion term will be denoted by M .

5.1.2 Motion tolerance without round trips in the interferometer

The first scenario which will be examined is that of a system which does not experience a second round trip in its interferometers. When using the assumption of equal optical frequencies the three constraints on the fringe frequencies in such a system can be described using equations 5.3 to 5.5.

$$|I_1| - |M| > 0 \quad (5.3)$$

$$|I_2| + |M| < f_{ppf} \quad (5.4)$$

$$|I_1| + |M| < |I_2| - |M| \quad (5.5)$$

Equations 5.3 and 5.4 show that the optimal FSI term (that which maximises the allowable size of the motion term) for laser 1 is the same amount above zero as the FSI term for laser 2 is below f_{ppf} .

At this point only the separability condition remains to be satisfied. When both laser 1 and laser 2 are scanning in the same direction interferometer motion affects both fringe frequencies in the same way, so that if the fringe frequencies are separable at the start then they will remain separable as the target moves. When the lasers are scanning in opposite directions the effect of taking the absolute value of the phase velocity becomes

relevant because target motion causes equal but opposite changes in the observed fringe frequencies. In the case of increasing interferometer length the laser which is increasing its optical frequency sees an increased fringe rate but the other laser will see a decrease of the same magnitude, and vice versa for shrinking interferometer length. The result of this is that the separability limit will be approached twice as fast as the zero or points per fringe limits, as described by equation 5.5, and so the optimal position of the fringe frequencies in a stationary interferometer are $|I_1| = f_{ppf}/4$ and $|I_2| = 3f_{ppf}/4 = 3|I_1|$. In this case the separability limit is broken at equal but opposite target speed to the speed which breaks the zero and points per fringe limits, giving maximum motion tolerance.

The maximum motion tolerance achievable under this set of assumptions is found by equating $|M|$ and $|I_1|$. This is done in equation 5.6, where the optimum value for $|I_1|$ has been used, and which can be trivially rearranged to give equation 5.7.

$$\frac{\nu |\dot{L}|}{c} = \frac{f_{ppf}}{4} \quad (5.6)$$

$$|\dot{L}| = \frac{\lambda f_{ppf}}{4} \quad (5.7)$$

Using $f_{ppf} = 347$ kHz, as calculated above, and taking the average λ to be 1540 nm the maximum value of $|\dot{L}|$ is calculated to be 134 mm/s. A series of idealised fringe frequency spectra for a 2 m interferometer with optimised laser tuning speeds and a range of target motion speeds is given in figure 5.1 to demonstrate this result. The simulation was performed with perfectly monochromatic lasers, and as such a real system would have wider features because of uneven frequency tuning curves and non-uniform length changes in the interferometer. These effects will be discussed in section 5.1.4.

5.1.3 Round trips in the interferometer

We will now take into account the effect of multiple round trips in the interferometer on the motion tolerance, which is a more realistic model of our observed data. As described in

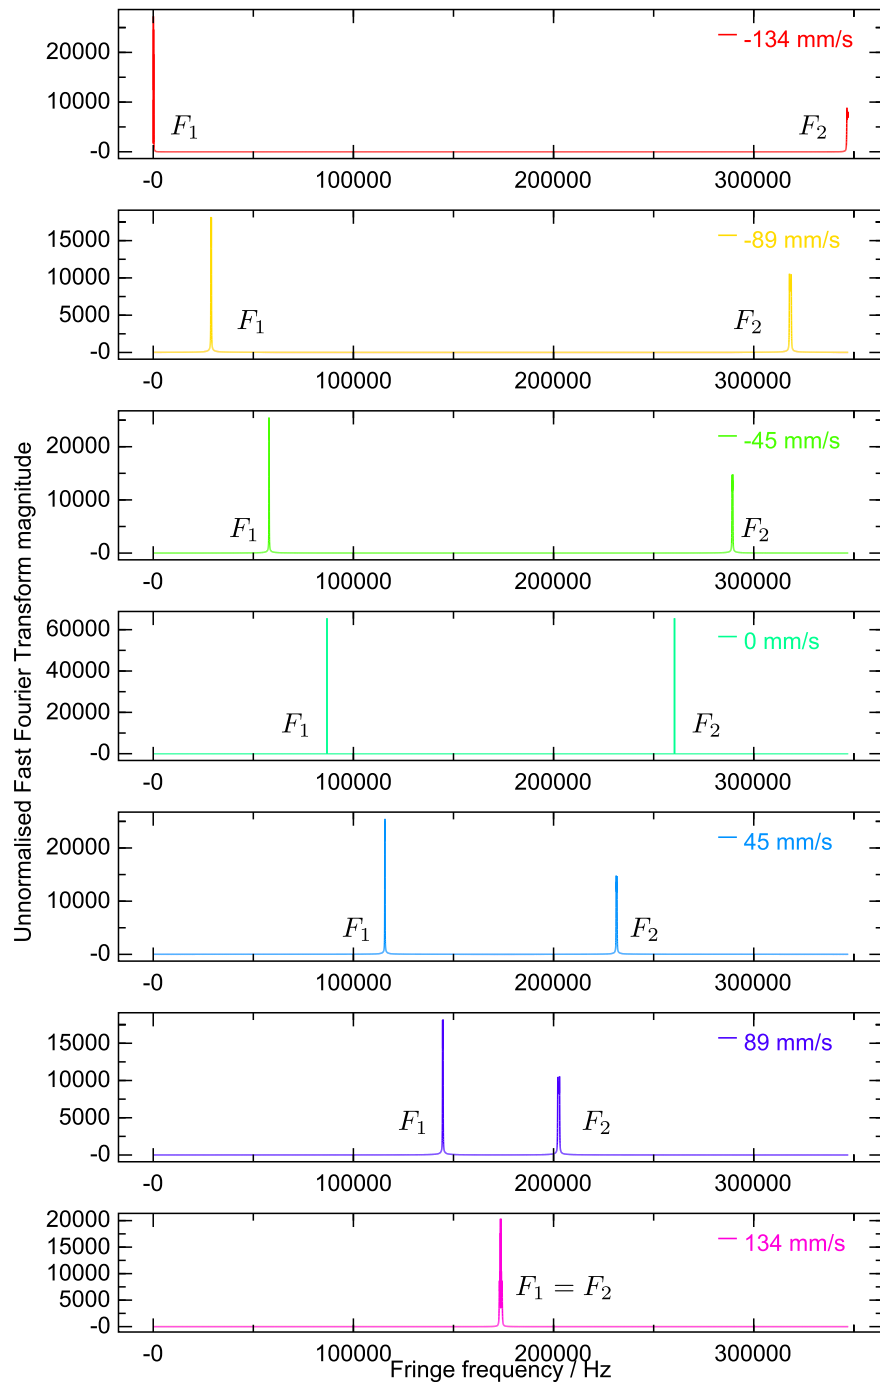


Figure 5.1: Idealised spectra from a 2 m interferometer without round trips. The maximum of the frequency axis is the points-per-fringe limit. The slow laser is scanning at 13 THz/s, with the fast laser scanning at -39 THz/s. The scan range of both lasers is centred at 194.67 THz (1540 nm). The peaks from the two lasers clash at 134 mm/s (bottom plot), and hit the extreme limits at -134 mm/s (top plot).

section 3.3.1 these round trips result in extra interference patterns with fringe frequencies at integer multiples of the single trip fringe frequency, when the interferometer is not changing its OPD. When the OPD is changing the round trips will therefore experience length changes which are integer multiples of the actual OPD changes, and the width of the fringe frequency spectrum also scales up.

The fringe rate for laser 2 (F_2) should lie between the fringe rate for laser 1 (F_1) and its second round trip, which reformulates the constraint equations as 5.8 to 5.11.

$$|I_1| - |M| > 0 \quad (5.8)$$

$$|I_2| + |M| < 2 * (|I_1| - |M|) \quad (5.9)$$

$$|I_1| + |M| < |I_2| - |M| \quad (5.10)$$

$$|I_2| + |M| < f_{ppf} \quad (5.11)$$

The combination of equations 5.9 and 5.10 determine that the optimal value of $|I_2|$ to be $2|M|$ above $|I_1|$ and $3|M|$ below $2|I_1|$, and so $|I_2| = \frac{7}{5}|I_1|$. At this position $|M|$ can reach $0.2|I_1|$ before the laser 2 feature clashes with either the laser 1 feature or its second round trip. As such we must maximise $|I_1|$ in order to maximise the motion tolerance.

The maximum fringe frequency which can be achieved by laser 2 before it merges with the second round trip of laser 1 is $1.6|I_1|$. As such, any space in the fringe frequency spectrum above this frequency is wasted. It therefore makes sense to equate this frequency with the points per fringe limit, giving $|I_1| = f_{ppf}/1.6$ and so $|I_2| = 1.4f_{ppf}/1.6$. The maximum tolerable size of the FFI term ($0.2|I_1|$) is then $f_{ppf}/8$, which is half that achieved in the no round trip case and gives a maximum value of $\left|\dot{L}\right|$ of 67 mm/s (assuming a wavelength of 1540 nm). Idealised fringe frequency spectra including round trips are given in figure 5.2. Note that laser 1 never approaches the zero-fringe frequency limit, with a minimum value at $0.8|I_1|$, because of the tight limits imposed by the laser 1 round trip.

Even this reduced figure will overestimate the achievable motion tolerance because it

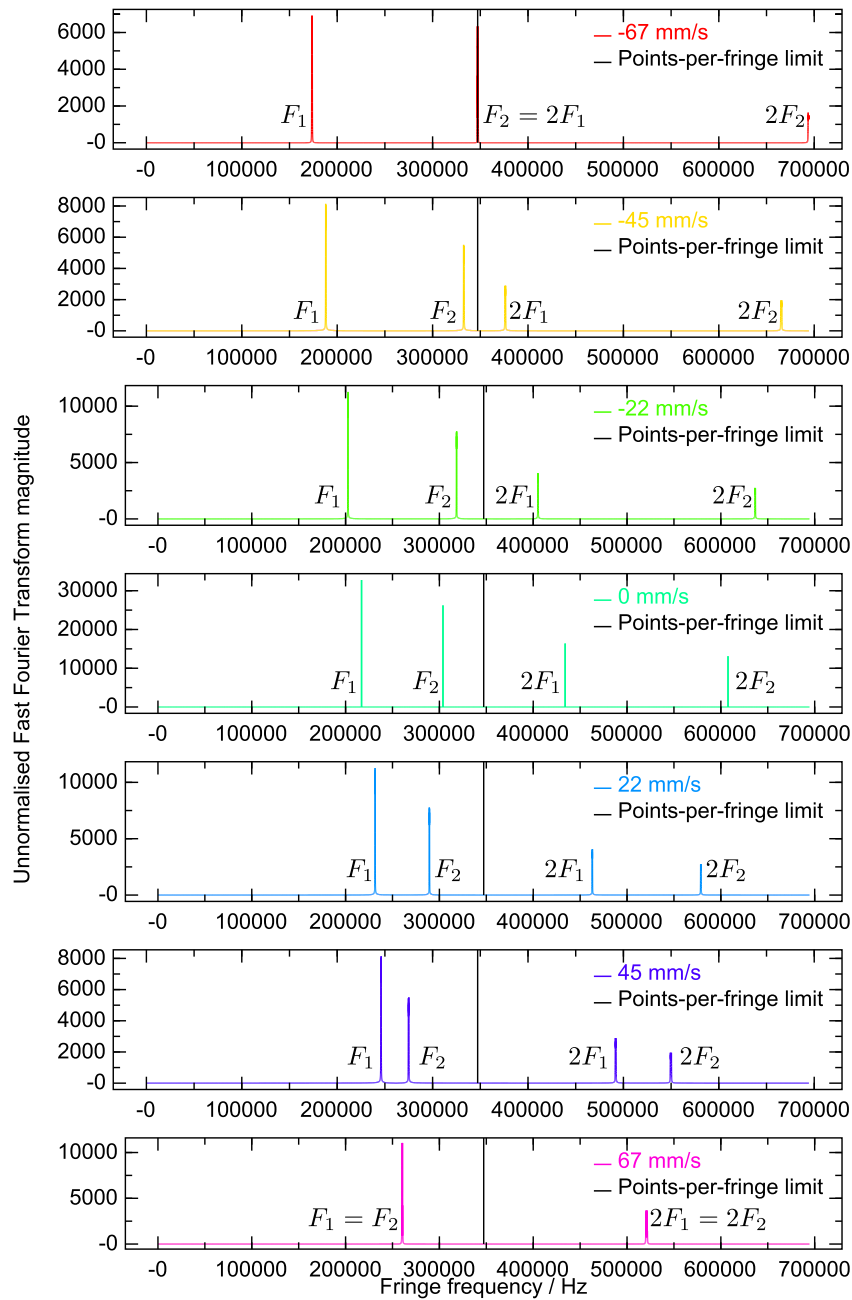


Figure 5.2: Idealised spectra from a 2 m interferometer with two round trips. The lasers have been given different intensities to help distinguish between them. The the points-per-fringe limit is shown in black. The slow laser is scanning at 32.5 THz/s, with the fast laser scanning at -45.5 THz/s. The scan range of both lasers is centred at 194.67 THz (1540 nm). The peaks from both second round trips are all above the points-per-fringe limit. The peaks from the two lasers clash at 67 mm/s (bottom plot), and the peak from laser 2 clashes with the second round trip from laser 1 at -67 mm/s (top plot), which results in a fringe rate equal to the points-per-fringe limit.

does not take laser tuning non-linearities in to account. The next section will provide a basic treatment to estimate the reduced limits.

5.1.4 Motion tolerance with broadened frequency spectra

To estimate the changes to the motion tolerance of the system due to the width of the fringe frequency spectra we will introduce a new parameter, α , to describe the width of the spectral peak due to the time variation of the tuning rate. If the peak is centred around fringe frequency F , as per equation 5.2, then the peak limits will be defined by equations 5.12 and 5.13.

$$F_{upper} = F(1 + \alpha) \quad (5.12)$$

$$F_{lower} = F(1 - \alpha) \quad (5.13)$$

The parameter α can obviously vary between lasers, and so the parameters will be identified through a subscript number. The addition of the peak width term modifies the constraint equations (5.8 to 5.11), leading to equations 5.14 to 5.17 below.

$$(|I_1| - |M|)(1 - \alpha_1) > 0 \quad (5.14)$$

$$(|I_2| + |M|)(1 + \alpha_2) < 2 * (|I_1| - |M|) (1 - \alpha_1) \quad (5.15)$$

$$(|I_1| + |M|)(1 + \alpha_1) < (|I_2| - |M|)(1 - \alpha_2) \quad (5.16)$$

$$(|I_2| + |M|)(1 + \alpha_2) < f_{ppf} \quad (5.17)$$

To determine the maximum motion tolerance of the system we equate the two sides of equation 5.16 and rearrange, giving equation 5.18 below.

$$|M| = \frac{|I_2|(1 - \alpha_2) - |I_1|(1 + \alpha_1)}{2 - \alpha_2 + \alpha_1} \quad (5.18)$$

By substituting the result for $|M|$ into equation 5.15, we can rearrange and find the optimal ratio of the FSI terms, as shown in equation 5.19.

$$|I_2| = |I_1| \left(\frac{7 - \alpha_1 - \alpha_2 + 3\alpha_1\alpha_2 - 4\alpha_1^2}{5 - \alpha_1 - \alpha_2 + 3\alpha_1\alpha_2 - 2\alpha_2^2} \right) \quad (5.19)$$

Note that equations 5.18 and 5.19 are both consistent with our idealised calculations, as setting $\alpha_1 = \alpha_2 = 0$ yields $|I_2| = 1.4 |I_1|$ and $|M| = 0.2 |I_1|$ as before. To find the maximum value of $|I_1|$ and so maximise the expected motion tolerance of the system we substitute equations 5.18 and 5.19 into equation 5.17, giving equation 5.20.

$$|I_1| = \frac{f_{ppf}(2 + \alpha_1 - \alpha_2)(5 - \alpha_1 - \alpha_2 + 3\alpha_1\alpha_2 - 2\alpha_2^2)}{16 - 4(3\alpha_1^2 - 2\alpha_1\alpha_2 + 3\alpha_2^2) + 4(\alpha_1\alpha_2^2 - \alpha_1^2\alpha_2 - \alpha_1^3 + \alpha_2^3) - 4(\alpha_1^3\alpha_2 - 2\alpha_1^2\alpha_2^2 + \alpha_1\alpha_2^3)} \quad (5.20)$$

With this final equation we can determine the maximum motion tolerance that we expect to be achievable. Using the green peak limits (determined by where the peak falls below threshold as described in section 3.2.2.2) from figure 3.14 to calculate the α values yields $\alpha_1 = 0.0485$ and $\alpha_2 = 0.0586$. This makes the optimal value of $|I_1| = 0.61f_{ppf}$, slightly reduced from $0.625f_{ppf}$ from section 5.1.3. $|I_2|$ is found to be $1.41 |I_1| = 0.86f_{ppf}$, again reduced from $0.875f_{ppf}$ in the previous section. The maximum value of $|\dot{L}|$ is 45 mm/s in this case. Simulated fringe frequency spectra including round trips are given in figure 5.3. Note that laser 1 again never approaches the zero-fringe frequency limit because of the tight limits imposed by the round trip.

The discussion up to this point has assumed that I_1 and I_2 can be set to any position in fringe frequency space by the user adjusting the tuning speeds of the two lasers. While this is true in dynamic FSI it is not the case in the current continuous FSI implementation, which has its laser tuning speeds limited by timing constraints, as described in section 3.3. The current settings give $|I_2| = 1.5 |I_1|$, which is above the optimal value calculated above and so makes equation 5.15 the relevant limit. Substituting the actual ratio of 1.5 into equation 5.15 provides the relation between $|I_1|$ and $|M|$, given in equation 5.21, and so

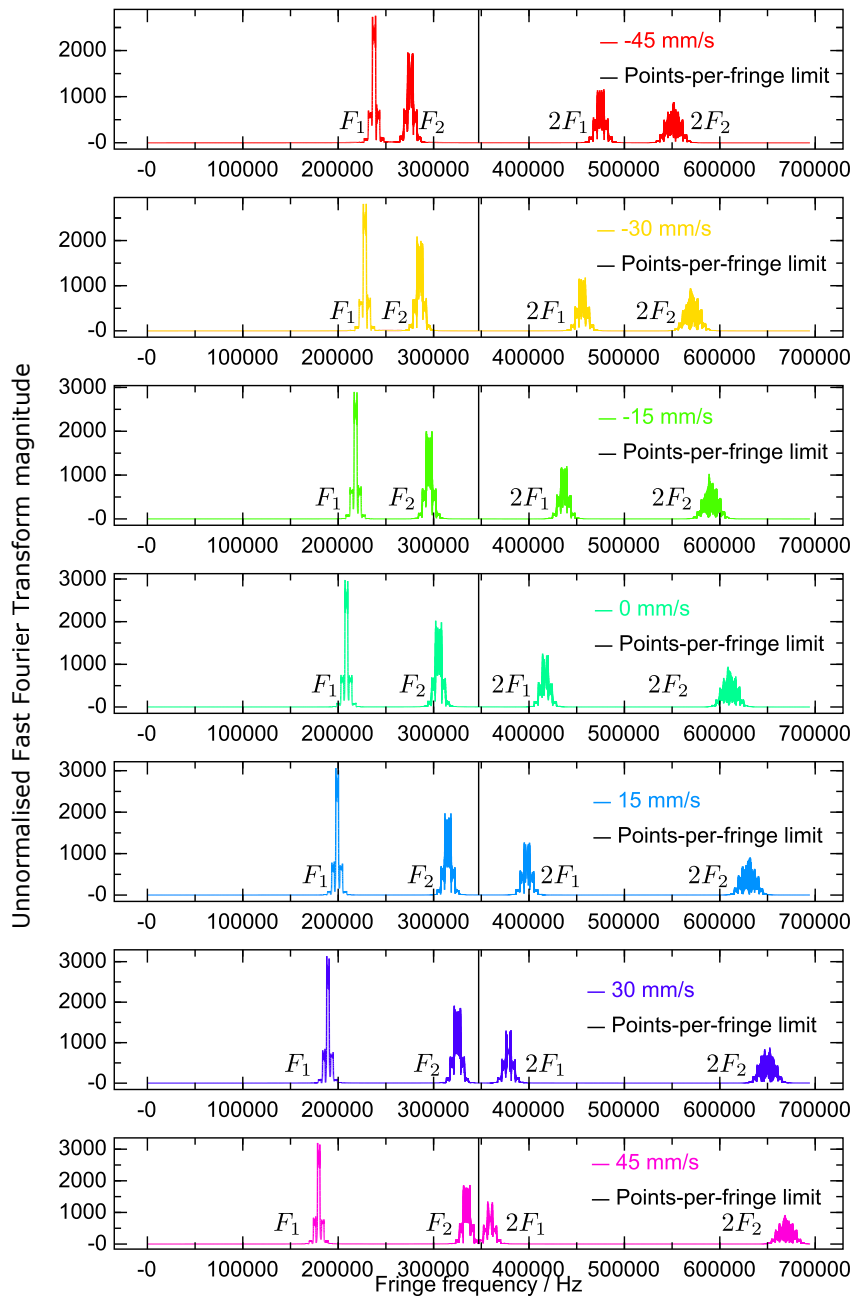


Figure 5.3: Simulated spectra from a 2 m interferometer with two round trips. The lasers have been given different intensities to help distinguish between them. The the points-per-fringe limit is shown in black. The slow laser is scanning at -31.7 THz/s, with the fast laser scanning at 44.7 THz/s. The scan range of both lasers is centred at 194.67 THz (1540 nm). The peaks from the second round trips are all above the points-per-fringe limit. The peaks from the two lasers clash at 45 mm/s (red), and the peak from laser 2 clashes with the second round trip from laser 1 at 45 mm/s (pink), which results in a fringe rate equal to the points-per-fringe limit.

leads to the length-dependent value given in equation 5.22.

$$|M|_{max} = |I_1| \left(\frac{1 - 4\alpha_1 - 3\alpha_2}{6 + 2\alpha_2 - 4\alpha_1} \right) \quad (5.21)$$

$$\left| \dot{L} \right|_{max} = \frac{L \left| \dot{\lambda} \right|}{\lambda} \left(\frac{1 - 4\alpha_1 - 3\alpha_2}{6 + 2\alpha_2 - 4\alpha_1} \right) \quad (5.22)$$

Using $\dot{\lambda} = 32 \text{ nm/s}$, $\lambda = 1540 \text{ nm}$ and the alpha values estimated above we can simplify equation 5.22 to $\left| \dot{L} \right| = 0.0022 L s^{-1}$ for the current continuous FSI system.

5.2 Incorporating a fixed frequency laser

In the round trip case described in section 5.1.3 the lowest fringe frequency which can be reached by laser 1 before laser 2 clashes with its first round trip is $0.8 |I_1|$ ($f_{ppf}/2$ in the ideal case). This effectively leaves half of the available fringe frequency space, below $0.8 |I_1|$, unused. Adding a third laser to the system, with a fixed frequency rather than scanning frequency, it is possible to take advantage of this unused fringe frequency region.

A fixed frequency laser does not have the FSI term in the fringe rate equation (equation 5.2), and so generates fringes at a rate which is proportional to the motion speed of the target. This means that the Hilbert transform cannot be used to extract the phase of the fixed-frequency lasers interference pattern when the target motion is slow because the laser linewidth and any frequency jitter will cause the fringe spectrum to impinge upon the zero fringe rate limit. Once the target is moving quickly enough for the fringe frequency to always be non-zero the Hilbert transform could be applied as before. This emergence of the fixed frequency laser peak is simulated in figure 5.4.

Once the phase can be extracted a differential length measurement can be performed. A start index is chosen to be the base position, which will be labelled with a subscript s . The phase at this index is given in equation 5.23, where $\phi_{s,ff}$ is the fixed frequency laser phase at index s , L_s is the length calculated at that index by the CFSI system and ν_{ff}

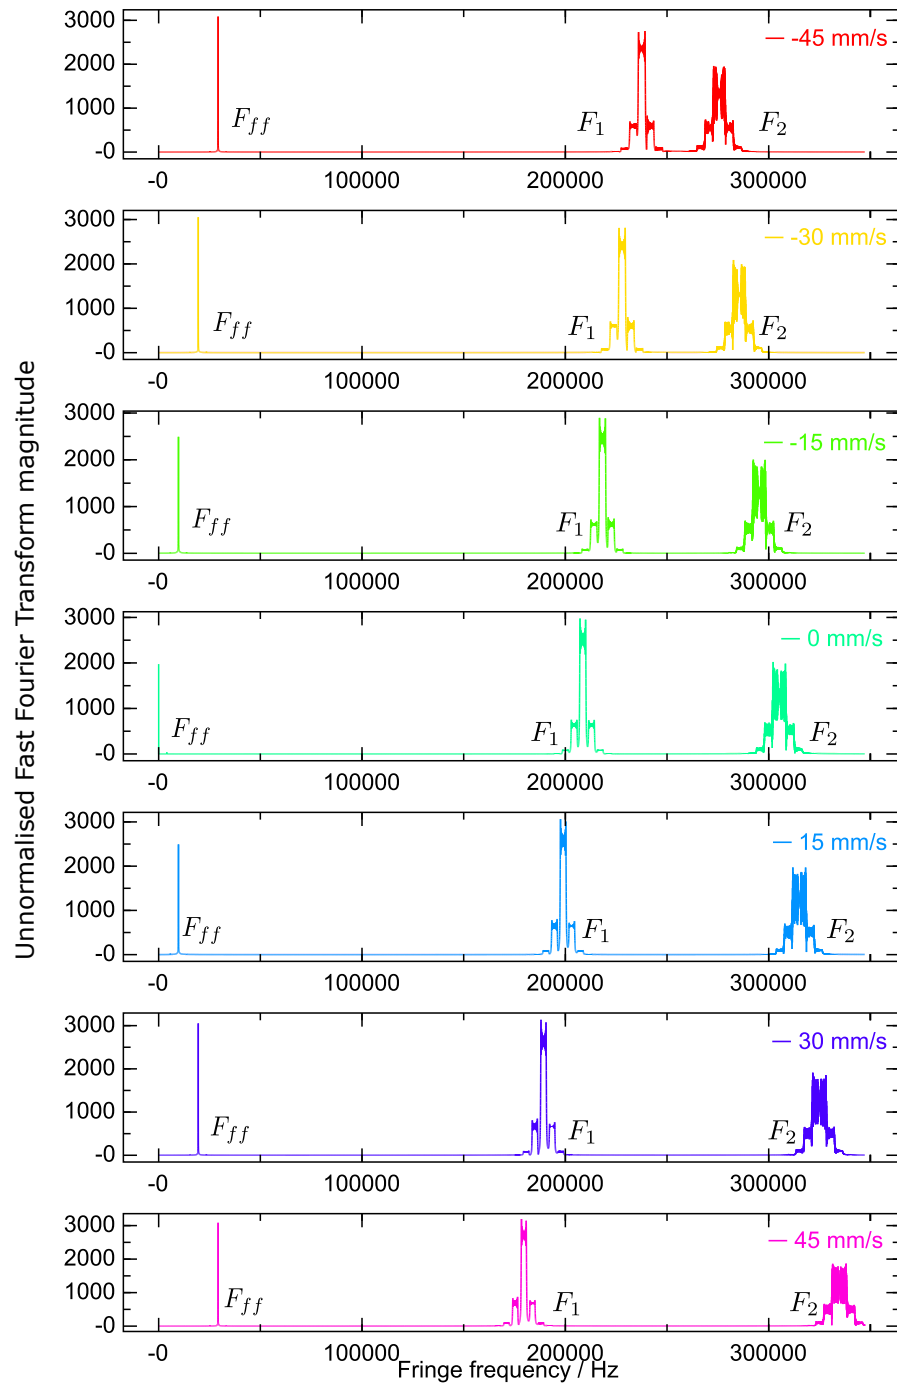


Figure 5.4: Simulated spectra from a 2 m interferometer including a 1540 nm fixed-frequency laser. Round trips have been removed for clarity. The FFI laser peak can be seen at lower frequencies. The maximum of the frequency axis is the points-per-fringe limit.

is the optical frequency of the fixed frequency laser.

$$\phi_s = \frac{2\pi}{c} L_s \nu_{ff} \quad (5.23)$$

At a later index i the interferometer length can be calculated using phase differences as shown in equations 5.24 and 5.25. The ambiguity in the direction of phase changes inherent in the Hilbert transform can be broken because we have knowledge of the motion direction from the CFSI signal.

$$\phi_i - \phi_s = \frac{2\pi\nu_{ff}}{c} (L_i - L_s) \quad (5.24)$$

$$L_i = \frac{c}{2\pi\nu_{ff}} (\phi_i - \phi_s) + L_s \quad (5.25)$$

This shows that the length of the interferometer can be tracked using only data from the fixed frequency laser, and so the scanning lasers can be ignored in periods in which the motion direction is known and does not change sign. If they are shut off from the interferometer entirely, for example using a MEMS switch, then the fringe frequency of the fixed frequency laser (F_{ff}) can rise all the way up to f_{ppf} which gives the maximum possible motion tolerance (535 mm/s assuming a wavelength of 1540 nm and using the current equipment).

An important point to note is that the scanning lasers will never clash in the reference interferometer because this does not change its length appreciably. This means that the reference interferometer can always be used to track the optical frequencies of the scanning lasers. When $\left| \dot{L} \right|$ drops the fixed frequency laser will eventually be unable to track the OPD, and so the system must swap back to continuous FSI. As the scanning laser frequencies are continuously monitored the length of the interferometer from the FFI laser can be used as the transfer length to perform a handover as in ordinary continuous FSI. The continuous FSI system would then monitor the length until $\left| \dot{L} \right|$ increases again. This technique is known as enhanced continuous FSI (enhanced CFSI).

The method proposed above provides absolute length measurements with the maximum possible motion tolerance. However, it adds a subtler limit on the acceleration of the measurement interferometer OPD. The system needs time to swap between FSI and FFI lasers systems, and so if the target acceleration is too great at this point the analysis could fail before the swap is completed. This limit has not yet been studied in detail and is likely to be system dependent so will be left as future work (see chapter 7). As an estimate of the limit if 7 packets are allowed before the motion stops (2 packets for detection, 1 for switching and 4 for the CFSI phase extraction) the current system needs approximately 41 ms to respond. The maximum motion tolerance is $|\dot{L}| = 0.0022Ls^{-1}$ as described in section 5.1.4, which gives a maximum allowable acceleration of $0.054Ls^{-2}$ when switching back to the CFSI system.

5.2.1 Implementation of enhanced CFSI - equipment

An enhanced continuous FSI system requires a different system layout to continuous FSI, as shown in figure 5.5 below. This section will detail the new equipment and reasoning for the changes in experimental layout. Those sections which remain unchanged, such as the master splitter box and gas cell measurement systems, will not be reviewed.

5.2.1.1 Fixed frequency laser

The fixed frequency laser used in the enhanced CFSI experiments was a New Focus TLB6728 Velocity laser [71]. This laser has a tunable wavelength range of 1520 nm to 1570 nm, a linewidth of 200 kHz over a 50 ms integration time and an output power of 30 mW. The power output of the laser was one of the main reasons for the experimental layout changes, as it could easily damage the photodiodes in the system. To prevent this from happening it was decided to reduced the power going into the enhanced CFSI system from the laser by sending it through the front end splitter tree first. This reduced the optical power to ≈ 0.625 mW throughout the rest of the system, which keeps the

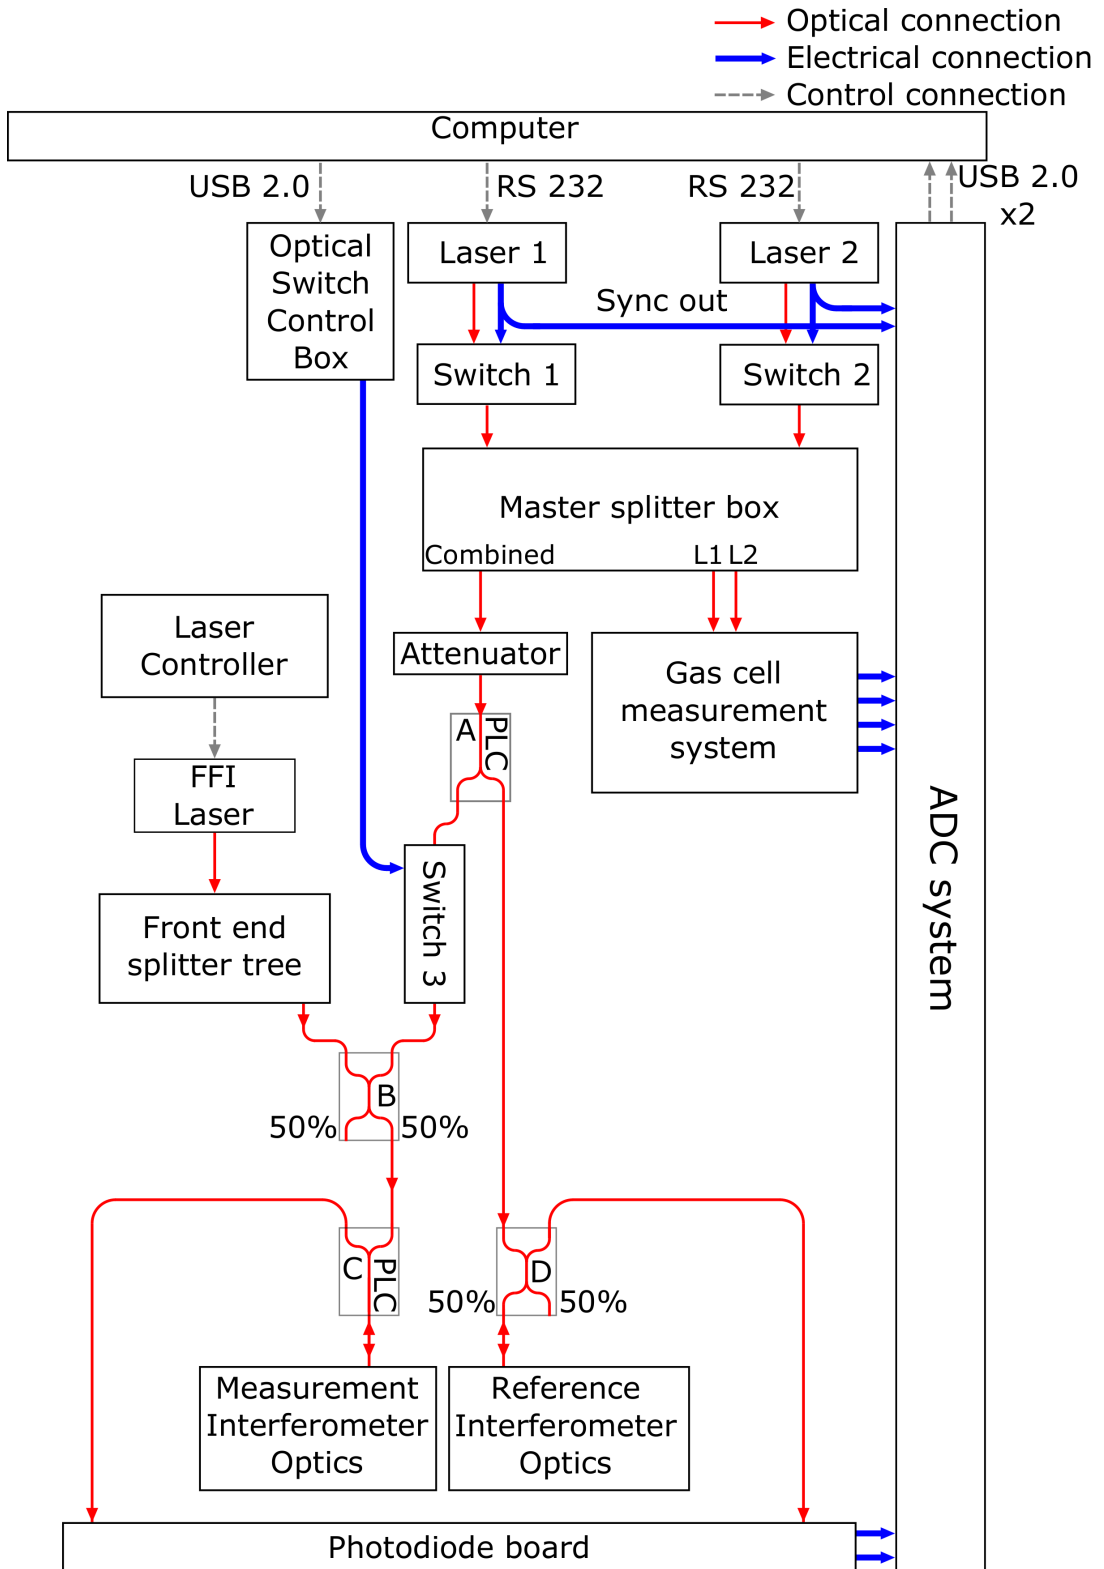


Figure 5.5: Schematic of the full experimental setup for enhanced continuous FSI. The master splitter box, front-end splitter tree and gas cell measurement system are described in detail in chapter 3. Splitters C and D only use different technologies because of the limited availability of equipment.

collimated output below the 10 mW Class 1 limit [72]. The front end splitter tree was used rather than a variable attenuator to prevent the accessible laser power from being accidentally increased during use.

5.2.1.2 Fibre splitter layout

Using the front end splitter tree to attenuate the Velocity laser meant that the system to feed the interferometers had to be updated. The first change is the introduction of a variable attenuator straight after the master splitter box combined output. The attenuator, which was a Thorlabs VOA50-APC single-mode variable attenuator [73], was needed to prevent the scanning lasers from overloading the photodiodes and was tuned to give the best possible signal in the reference interferometer. The attenuator is followed by a 1×2 PLC splitter, labelled as A, which divides the light between the measurement and reference interferometers.

The measurement interferometer section is then passed through an extra optical switch (switch 3). The switch is used to remove the scanning lasers from the measurement interferometer as required, and was controlled via a control box as described in section 5.2.1.3. The light was then combined with the light from the Velocity laser using a 50:50 fibre taper splitter (labelled B) before being passed into a 1×2 PLC splitter labelled C. Splitter C feeds the measurement interferometer and provides a tap to send the returning light to the photodiode board.

The reference interferometer section is simpler than the measurement interferometer section, simply feeding a 50:50 fibre taper splitter (labelled D) which serves an analogous function as splitter C. This is because the fixed frequency laser was not used in the reference interferometer.

5.2.1.3 Optical switch control box

The extra optical switch in the enhanced CFSI system was controlled using a National Instruments NI USB-6501 Digital I/O Device [74]. This device was used to send either a

high or low signal to the switch and so change its state.

5.2.2 Implementation of enhanced CFSI - analysis

The analysis of the fixed frequency interference signal is significantly simpler than that for the overall continuous FSI system. It requires two processes to happen: laser frequency determination and phase extraction. Phase extraction is performed using a Hilbert transform in the same manner as for continuous FSI, although the filter limits are set manually rather than using the autofilter routine, due to the low FFI signal relative to the optimised CFSI signals.

Determination of the FFI laser frequency was necessary because the laser did not report it with sufficient precision. In an ideal system the laser frequency would be known *a priori*, for example by locking it to a gas absorption feature, but this was not the case with the current system.

The laser frequency was calculated by looking for the optical beat signal between the fixed and scanning lasers. The beat signal is observable when the difference between the two optical frequencies is less than the bandwidth of the DAQ system. The frequency of the scanning laser at the point at which the beat is detected is then assigned to the fixed-frequency laser. This calculation is currently performed offline. The fixed frequency laser was placed between the R2 and R1 HC^{13}N absorption peaks, at a wavelength of approximately 1541 nm. This wavelength was chosen because it is only covered by laser 2 in the current system, removing the need to determine the source of the beat signal as further beat signals are generated when laser 1 and laser 2 cross.

A dedicated beat detection photodiode was not used in the enhanced CFSI system. The beat signal was instead extracted from the measurement interferometer, and was clarified by removing the interference signals from the recorded data with a spectral filter. Removal of the interference patterns leaves only background noise and the beat signal, as shown in figure 5.6. Figure 5.7 shows a series of examples of the observable

detail of the beat signals.




Figure 5.6: Clarification of the beat signal between laser 2 and the fixed frequency laser.

With the interference signals removed the beat point is found by looking for the largest magnitude change in signal amplitude between two time indices. The frequency at the second index, as calculated by laser 2 in the usual way, is taken to be the frequency of the fixed frequency laser. The differences between points, rather than the absolute magnitude, was used to find the beat signal because background variations that were larger than the beat were observed in some cases. The laser frequency used in the enhanced CFSI analysis is the average of the laser frequencies found in the run.

There are some difficulties with finding the laser frequency in this way. The most severe of these is that the beat signal is not always present in the data, due to its rapid nature. The algorithm will then decide that a background fluctuation is in fact the beat signal and provide an incorrect result. These results are cut before the averaging procedure is performed. The use of index by index differences for beat finding introduces a series of difficulties as it depends heavily upon the form of the beat signal, which varies with each measurement, and also upon noise fluctuations at the base of the signal. These




Figure 5.7: Examples of the beat signals observed after filtering.

effects were expected to shift the beat location by a few indices but, as can be seen in the plot of typical laser frequency results over a run (after cutting) shown in figure 5.8 and the histogram of the results shown in figure 5.9, the calculated laser frequency was found to vary by less than 20 MHz over a run. The 20 MHz variation corresponds to approximately 1×10^{-7} of the mean frequency (194.546171 THz), which was deemed to be an acceptable level of variation for the experiments.

Once the laser frequency has been calculated we have all of the information required to analyse the fixed frequency portion of an enhanced CFSI system.

5.2.3 Enhanced CFSI experimental demonstration

The demonstration experiments for enhanced CFSI are divided into three sections. In the first section the experiments to confirm that the fixed frequency portion of the system is working correctly are described. This section will also demonstrate the problems caused by the difference in optical path between the LaserTRACER reference system and the continuous FSI systems, as described in section 4.3.3. The second experimental section




Figure 5.8: Laser frequency results calculated during an enhanced CFSI run.




Figure 5.9: A histogram of the laser frequency results calculated during an enhanced CFSI run.

will demonstrate the enhanced CFSI system switching between continuous FSI and the fixed frequency interferometry system, and also the change back to continuous FSI. The final section will detail experiments which demonstrate that enhanced continuous FSI can return uninterrupted measurements even as the two scanning lasers are isolated from the measurement interferometer, and then continue as they are reintroduced.

5.2.3.1 Fixed frequency interferometry demonstration

The fixed frequency interferometry demonstration experiments were forward-style experiments as described in section 4.3. The fixed-frequency system was used to determine the length changes once the continuous FSI measurements confirmed that the OPD was changing sufficiently rapidly. The data was aligned in time against the LaserTRACER in the same manner as before, with the addition of a further fit of T_{centre} including an extra 600 data points to cover the first period of the stage motion. This additional fit was found to improve the data alignment in the enhanced CFSI experiments. The calculated laser frequency throughout these experiments is shown in figure 5.10, which covers a period of nearly 3 hours.

An example of the length measured by all three distance measurement systems is shown in figure 5.11, showing macroscopic consistency between all three systems. Figure 5.12 shows a series of examples of the differences of both continuous FSI and FFI systems from the LaserTRACER results. The FFI measurement is a differential measurement, and was equated to the continuous FSI measurement at one index to allow direct comparison. The index chosen was the LaserTRACER index which was 300 points after the LaserTRACER exceeded a $1 \mu\text{m}$ point-by-point change.

The first point of note is that in all cases the same rapidly changing deviations of the continuous FSI result from the LaserTRACER, as observed in figures 4.26 and 4.27, are still present in the data. These deviations are also present of the FFI differences. This is an important point because the FFI measurement is largely independent from the continuous FSI measurement, as the filtering procedures in each routine entirely remove




Figure 5.10: Laser frequency results calculated during the FFI demonstration experiment. The first value (194.546171 THz) has been subtracted to aid visibility. The error bars show the standard deviations of the calculated laser frequencies during the run.




Figure 5.11: Macroscopic behaviour of the CFSI, FFI and LaserTRACER measurements during a FFI demonstration run. The target was moved at approximately 8 mm/s.




Figure 5.12: A series of examples of the difference between the continuous FSI and FFI systems from the LaserTRACER measurements for the enhanced CFSI demonstration experiments. Further examples are provided in appendix A.

the range of signal frequencies used in the other, although the use of the reference length and gas cell measurements from the continuous FSI analysis could lead to correlations in the overall scale factor. The common factor between the two measurements is that they have an identical optical path, suggesting the conclusion that the rapidly varying elements of the difference between the continuous FSI or FFI measurements and the LaserTRACER reference measurements are actual changes in the OPD of the system that are not common to the LaserTRACER and FSI / FFI paths. Possible contributory factors include changes in refractive index and changes due to the physical mounting of the two measurement systems. In either case it is not possible to calibrate the continuous FSI or FFI systems against the LaserTRACER using these measurements.

The difficulties with the differential calibration experiments is illustrated by figure 5.13, which shows the gradient of a straight line fit of continuous FSI and FFI distance measurements versus LaserTRACER distance measurements for each of the motion experiments that have been performed (forward, backward, enhanced CFSI (CFSI) and enhanced CFSI (FFI)). The mean gradient values and associated standard deviations are given in table 5.1.

Experiment style	Mean gradient - 1	Standard deviation
Forward	-5.16×10^{-6}	9.46×10^{-6}
Backward	7.19×10^{-6}	5.73×10^{-6}
Enhanced CFSI (CFSI)	-9.55×10^{-6}	1.62×10^{-6}
Enhanced CFSI (FFI)	-2.76×10^{-6}	1.60×10^{-6}

Table 5.1: Results from the differential gradient calculations.

The first point to note is that the gradients from the forward experiments are consistently below 1 for the early experiments but then tend to increase after experiment 14. This effect can be partially explained by fluctuations within the measurement, as is the case in experiment 15 (shown in figure 4.26) which experiences a rapid drift that results in the high gradient, although other experiments generate the high gradient through consistent drift. As discussed previously this inconsistency casts doubt upon any conclusions drawn from the gradient experiments.

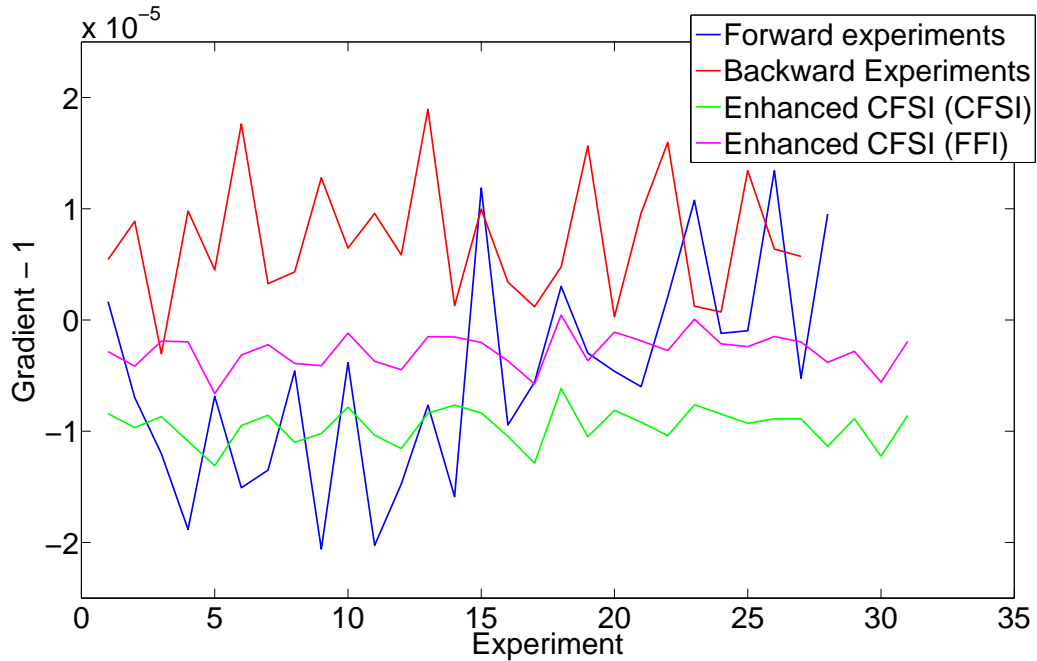


Figure 5.13: Gradients from fits of CFSI/FFI versus LaserTRACER monitoring a moving target.

The second point is that the backward experiments, which provide gradients that are consistently above 1, and the enhanced CFSI (CFSI) experiments, which have gradients consistently below 1 (and are also consistent with the early forward measurements), were performed with the target retro-reflector travelling in opposite directions. The cause of this discrepancy is that the continuous FSI system gradually measures a shorter distance than the LaserTRACER over the course of each run, independent of target motion. This can be seen in figures 4.27 and 5.12, with further examples provided in appendix A. It is unclear whether the length deficit is atmospheric in nature or comes from the continuous FSI technique, and this would require further experiments to determine.

Without explaining the continuous FSI length deficit it would not be appropriate to attempt to calibrate the FFI length measurements. If the effect comes from the atmosphere then the calibration will be inaccurate, and if the effect comes from the continuous FSI algorithm then it could also affect the FFI measurement, which uses the same phase extraction algorithm and has a scale dependence on the reference interferometer length

and gas cell peak fitting results, which are used to calculate the laser frequency. FFI measurements were only performed in one direction, meaning it is not possible to check for a FFI length deficit when the stage is moving backwards.

While calibrating the FFI measurements would not be appropriate a comparison with the continuous FSI measurement provides a measure of the effect of atmospheric variations on the repeatability of the gradient calculation. This is because the two systems operate in the same wavelength range, through exactly the same air and are perfectly co-linear, removing any misalignment concerns. Fits of FFI displacement versus continuous FSI displacement provide a mean value of $1 + 6.79 \times 10^{-6}$ with a standard deviation of 6.66×10^{-7} , showing a significant improvement in consistency over the measurements against the LaserTRACER. Future experiments should aim to minimise the difference in optical path between the CFSI and reference systems.

With the experimental layout preventing reliable calibration of the enhanced CFSI system, the remainder of this chapter will be dedicated to proof-of-principle measurements designed to demonstrate successful changes between CFSI and FFI measurements.

5.2.3.2 Enhanced CFSI switching demonstration

Switching between the two measurement systems was initially performed using the same data that was used by the FFI demonstration described above. Switching to the FFI system was achieved by choosing a data point and equating the length measured by the FFI system at that point to the equivalent CFSI length measurement. Future length measurements were calculated using the FFI system. To change back to the continuous FSI system a point was chosen to act as the transfer point, and the length as calculated by the FFI system at that point used as the transfer length. The complete measurement, along with the initial and intermediate stages are shown in figure 5.14 as differences from the LaserTRACER.

As in previous figures, the difference between the basic continuous FSI measurement and the LaserTRACER is shown in blue, and difference between the FFI measurement




Figure 5.14: A measurement to demonstrate switching between continuous FSI and FFI systems. The *CFSI* and *FFI* lines are the original measurements. The *To FFI* measurement has switched from CFSI to FFI. The *Combination* measurement has switched from CFSI to FFI, and then back again at a later time.

and the LaserTRACER in red. These traces show our best current measurements of the OPD. The green plot, labelled “To FFI”, shows a measurement where we have switched from continuous FSI to FFI at approximately 10 seconds. The black plot, labelled “Combination”, shows the full measurement, switching to FFI at the same point and then back to continuous FSI at approximately 33.6 seconds. Note that the black plot covers the blue until the switch to FFI, and obscures the green until the switch back to continuous FSI. Overall the switched measurements do not deviate from the LaserTRACER measurements to a greater extent than the basic continuous FSI system, so the switching technique has been successfully demonstrated.

5.2.3.3 CFSI removal demonstration

The final demonstration measurement required is one in which the scanning lasers are deactivated for part of the measurement. The experiments performed were a variation

on forward-style experiments with a speed of approximately 8 mm/s. After the motion was started the scanning lasers were removed using the extra optical switch. The motion stage was speeded up to approximately 16 mm/s for a few seconds, and then slowed to its original speed. The scanning lasers were reintroduced, and the motion stopped. The results of an example experiment are shown in figure 5.15.

The lower plot shows the measurement interferometer interference pattern recorded by the ADC system and it can be seen that the period of low intensity where the two scanning lasers were isolated, running from 9.8–21.7 seconds, encompasses the region of fast motion visible in the upper plot. The difference in measured displacement between the enhanced CFSI data and LaserTRACER data is similar to that which had been observed in earlier experiments. This was the final demonstration of enhanced CFSI, and has successfully shown that it is a technique which can measure fast moving targets by isolating the scanning lasers from a measurement interferometer and using a fixed frequency laser, although the system still needs to be calibrated. Further development and work required for enhanced CFSI to be a deployable measurement system is described in chapter 7.




Figure 5.15: A measurement to demonstrate switching between continuous FSI and FFI systems including removal of the scanning lasers from the measurement interferometer. The top plot shows the displacement measurement. The central plot shows the difference from the laser tracer. The bottom plot shows the measurement interferometer interference pattern, with the period without the scanning lasers clearly visible.

Chapter 6

Improved frequency referencing

Continuous FSI as implemented in this thesis relies heavily on a gas absorption cell, which provides the frequency scale for the measurements. Given this reliance on the gas cell significant effort has been put in to improving the gas cell system with the aim of moving towards using doppler-free saturated absorption features. The frequency stability and repeatability of the fixed frequency laser used in chapter 5 could also be improved by locking it to a gas absorption feature, which could also remove the need for an additional frequency determination routine.

This chapter will give an overview of doppler-free saturated absorption spectroscopy, motivate the choice of target gas cell species, describes the experimental approach taken and present the progress made towards implementing the new system.

6.1 Overview of doppler-free saturated absorption spectroscopy

Throughout the experiments described in chapters 3, 4 and 5 all of the gas absorption features have been doppler-broadened features. Doppler broadening is a well understood mechanism [75] and is caused by the thermal motion of molecules in the gas. Those molecules which have a component of their velocity along the direction of the laser path

see light with either a lower frequency, if travelling in the same direction as the laser beam, or higher frequency, if travelling in the opposite direction. This links the shape of the absorption features to the velocity profile of the gas, with the central transition frequency of a moving molecule shifting according to equation 6.1 where ν' is the shifted frequency, ν is the stationary absorption frequency and v is the speed of the molecule parallel to the light (direction determines the sign used from the \pm) and c is the speed of light. We do not consider the effect of second-order Doppler broadening caused by the transverse motion of the molecules, as it is expected to be several orders of magnitude smaller [76].

$$\nu' = \nu \left(1 \pm \frac{v}{c} \right) \quad (6.1)$$

Doppler-free saturated absorption spectroscopy is a common technique used to remove the effect of molecular motion [77]. The technique works by counter-propagating a pair of laser beams derived from the same laser source, and so containing light of identical frequency. As the beams are counter-propagating any given molecule of the gas will see the frequencies with equal and opposite doppler shifts. This difference in doppler shift means that the light in each beam interacts with molecules in different regions of the velocity profile of the gas.

Molecules with only transverse velocity, which will be referred to here as stationary molecules, are a special case because they experience no doppler shifts. They will interact with both beams when the laser frequency is the same as the absorption frequency of the transition. If the intensity of one of the beams can be increased to a level where most of the stationary molecules are kept in their excited state then the molecules are said to be saturated. When the molecules are saturated the second beam will see an increase in transmission as there are fewer molecules available to absorb it. This increase in transmission provides a measurement of the underlying gas cell transition which is free from the effects of doppler broadening. An example of such a feature, from a doppler broadened Rubidium peak made out of three transitions and observed using pump-probe

spectroscopy with crossed polarisations, is shown in figure 6.1.




Figure 6.1: Saturated absorption features observed in the ^{85}Rb D2 F=2 atomic transition using the fibre laser frequency stabilisation system (FLFS) [78]. Tuning voltage is proportional to wavelength.

The reduced spectral width of the gas cell feature could lead to improved laser frequency measurement for a FSI system, as has already been demonstrated with iodine [79], and also provide a suitable target for a laser frequency locking system for the enhanced CFSI fixed frequency laser.

6.2 Choice of target gas

Throughout the previous chapters a hydrogen cyanide gas cell has been used. For the saturated absorption experiments acetylene-13 ($^{13}\text{C}_2\text{H}_2$) was chosen instead. The primary reason for this change is that the P16 rotational transition accompanying the $\nu_1 + \nu_3$ vibrational transition in acetylene-13 is the International Committee for Weights and Measures recommended radiation for a stabilised laser to realise the definition of the

metre in the wavelength range of our lasers [80]. A fixed frequency laser locked to this saturated transition would be an incredibly useful tool.

The two modes which make up the $\nu_1 + \nu_3$ vibration are shown in figure 6.2, with the other modes described in figure 6.20 in [81]. Each transition between the vibrational levels also results in a change in the rotation of the molecule, with the selection rule $\Delta J = \pm 1$. A change of $\Delta J = +1$ leads to a transition in the R branch, with transitions in the




Figure 6.2: The ν_1 and ν_3 vibration modes of acetylene [81].

P branch corresponding to $\Delta J = -1$, again described in chapter 6 of [81]. Acetylene-13 has 61 features in the $\nu_1 + \nu_3$ band with vacuum wavelengths between 1520.11 nm and 1552.26 nm [82], allowing reference interferometer length measurement as previously performed with the hydrogen cyanide. A final advantage of acetylene-13 is that it was available at 99% atomic purity from Sigma Aldrich [83], whereas hydrogen cyanide could not be sourced commercially.

Working with acetylene raises several challenges for the implementation of a saturated spectroscopy system. Acetylene is highly reactive, which limits the choice of materials used in the construction of its containers, and explosive at higher pressures. These concerns will be described in more detail and addressed in section 6.4 which describes the equipment used. A more fundamental challenge is the small lifetime of the acetylene excitations, which makes them difficult to saturate. The parameter used to describe the level of saturation is the saturation intensity. A numerical formula to estimate this for

acetylene is given in [84], and repeated in equation 6.2 where I_s has units of W/m^2 and p has units of Torr. C_0 is zero in a purely collisional regime.

$$I_s = 22 \times 10^6 (C_0 + 10p)^2 \quad (6.2)$$

Some example calculations to illustrate the difficulties saturating acetylene transitions, using equation 6.2 and a 1 mm diameter beam, give a saturation power of 1.1 MW at 3.33 kPa (25 Torr, the same pressure as the HCN cell) and 875 mW for a pressure of 3 Pa (22.5 mTorr) as specified in [80]. To reach these high powers a Fabry-Perot cavity is often used to increase the power transmitted through the gas cell [82][84][85], although it is possible to achieve saturation without such a cavity as demonstrated in [86], from which the 1 mm beam diameter value used above was taken. Even without the cavity all of these approaches require components to be mounted on a stable optical table. A solution where this is not required would be preferable, as portability is a requirement when looking towards wide applicability and commercial applications where such a table may not be available.

6.3 Optical nanofibres

An alternative approach to increase the intensity of the light is to confine it into a smaller area. This is the approach taken by Takiguchi [87], who achieved saturation of acetylene using laser power of order 10 mW coupled into an optical nanofibre. An optical nanofibre is an optical fibre with a diameter which is smaller than the wavelength of the light. The result of this small diameter is that the light is guided mainly as an evanescent field, rather than being confined within the optical fibre. The evanescent field extends for only a short distance from the optical fibre, which leads to high intensities even with low laser powers.

Solutions of the Maxwell equations for a guided wave travelling down a nanofibre are

available in [88], and an overview of the fabrication process and applications of nanofibres is given in [89], with the fabrication process used in this work described in section 6.4.3. An example intensity distribution, calculated using the Optical Fibre Toolbox written using Matlab by Karapetyan et. al. [90], is shown in figure 6.3 and a schematic of the diameter of an optical nanofibre along its length is given in figure 6.4.




Figure 6.3: Unnormalised intensity of light guided by an optical nanofibre. The nanofibre, shown by the mesh, is 410 nm diameter and the wavelength of the light was 1550 nm.

The system built to implement saturated absorption spectroscopy was based on the nanofibre approach because of the possibility to have a self-contained gas cell fed by connectorised optical fibres. The surrounding optics could also be fibre based and so the entire experiment would be both portable and easily reconfigurable.

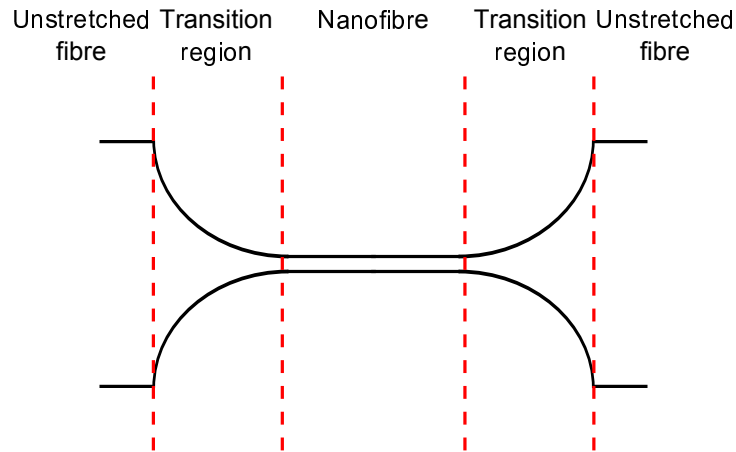


Figure 6.4: Schematic of an optical nanofibre, based on [91]. The nanofibre manufacturing process is described in section 6.4.3.

6.3.1 Hollow core fibres

An alternative fibre-based approach to saturated absorption spectroscopy of acetylene is to use hollow core fibres. These are fibres which contain a series of holes running along their length which can be filled with acetylene gas. Saturated absorption features have been observed with such fibres [92][93].

This approach was not taken in this work because of the lack of flexibility of such systems. If the hollow core fibres are spliced onto ordinary optical fibres the acetylene gas is sealed within and so its pressure cannot be changed without breaking the splices and risking damage to the fibre. Leaving the hollow core fibres unspliced means that a stable, high precision alignment system is required to couple light into the fibre, eliminating one of the primary benefits of a fibre-based system.

A system based on an optical nanofibre avoids these problems as the evanescent field is located outside of the fibre, meaning that the acetylene pressure can be varied without breaking any splices which have been made.

6.4 Equipment for saturated absorption

This section details the equipment used to store the acetylene gas and the nanofibre gas cell. The mechanical drawings for the custom pieces were created by James Lynn in the physics department design office. Acetylene is flammable, with a lower explosion limit of 2.5 % and upper explosion limit of 82 %. It is incompatible with a wide range of materials, including brass and silver. Acetylene will also form explosive compounds with copper [94], requiring copper to be entirely avoided during the construction of the equipment. Vacuum grease was also carefully checked for acetylene compatibility. Internally sharp edges were avoided wherever possible, as they can catalyse spontaneous dissociation of acetylene, and the gas transfer processes, both into the storage vessel from the source flask and from the storage vessel into the gas cell, were designed to minimise rapid expansion of the gas, as shock waves can also lead to dissociation of the gas.

6.4.1 Acetylene storage cylinder

The first component in the acetylene spectroscopy experiment is the storage cylinder for the acetylene gas. In most commercial storage systems, such as those for welding, acetylene is stored at high pressure in an acetone solution but this is not suitable for the frequency referencing experiment because of the risk of contamination of the gas and low pressures required. The pipes and valves needed to transfer gas into and out of the cylinder are fixed in place on the same base to minimise the number of vacuum seals that are broken and remade when the system is moved. A picture of the acetylene storage system is given in figure 6.5, with a schematic given in figure 6.6.

The acetylene cylinder is a 3 L stainless steel vessel with a CF-16 fitting to connect it to the vacuum system. The volume of the acetylene breakseal flask in which the acetylene is delivered is 500 mL, giving an approximate total volume of 3.5 L (ignoring the extra volume from the vacuum pipes). The breakseal flask contains acetylene at a pressure of one atmosphere. The relative volumes mean that the pressure of the acetylene




Figure 6.5: Picture showing the acetylene storage cylinder and associated vacuum system.




Figure 6.6: Schematic of the acetylene storage system and associated vacuum layout.

drops from atmospheric pressure in the breakseal flask down to below $1/7^{\text{th}}$ atmospheric pressure. Approximately $6/7^{\text{th}}$ of the acetylene is contained in the storage cylinder. The storage cylinder gas pressure is monitored using a MKS Instruments Type 626B Baratron Absolute Pressure transducer pressure gauge with a maximum readout pressure of 133 kPa (1,000 Torr) [95], which is P_1 in figure 6.6.

Once the gas pressure has fallen below atmospheric pressure the risk of explosion is low [96], but the initial filling was performed in a shielded fume cupboard. The procedure begins by connecting the break-seal flask to the system using a rubber pipe. All of the vacuum valves are opened and the system is evacuated. The vacuum pump is then isolated by closing valve A. Valves B and C are also closed. The glass seal on the break-seal flask is then broken which fills the rubber piping with acetylene gas. Valve B is then slowly opened to fill the pipe network with the acetylene gas while avoiding a rapid expansion. Valve C is then opened in the same way, gradually filling the storage cylinder. The pressure in the cylinder is monitored and allowed to stabilise. Once this has happened valves C and D are closed, the vacuum disconnected and valve A opened to vent the pipe network. The break-seal flask is removed and the pipes re-evacuated before all of the valves are closed.

The reverse of this procedure is followed when filling the gas cell, which is evacuated at the same time as the pipe network. Valve D is opened in place of breaking the glass seal. The gas cell pressure is not allowed to rise above 133 Pa (1 Torr), which is the maximum pressure that can be measured using the gas cell pressure gauge as described in section 6.4.2.

The two components in figure 6.6 not yet described are the second pressure gauge (P_2) and gate valve F. The second pressure gauge was used to control the vacuum pump. Valve F was used to isolate the gauge from the system when acetylene was present, because of concerns that it would interact.

The cylinder itself was finite-element modelled to ensure that it could contain any explosion of the acetylene, and the clamp around the edge of the base plate, which is

used to seal a plastic sheet over the cylinder when it is not in use. The volume contained by the sheet is filled with a pure nitrogen atmosphere so that any leakage into the cell contains no oxygen, again reducing the risk of explosion.

6.4.2 Nanofibre gas cell

The nanofibre gas cell was designed to keep the nanofibre intact for as long as possible, with the minimum possible volume to limit the amount of acetylene that has to be used when it is filled. A series of gas cells were manufactured as the design evolved, and were made of either aluminium or stainless steel. An outside view of the final cell, including the pressure gauge and vacuum valves, is shown in figure 6.7, with internal details shown in figure 6.8. An overview drawing, produced by James Lynn, is given in figure 6.9 to add clarity.




Figure 6.7: External view of the nanofibre gas cell. The cell is 700 mm long.

The main feature of the gas cell which is visible on the external view is the position




Figure 6.8: Internal view of the nanofibre gas cell.

of the vacuum “tower” used for pumping and filling. This is positioned at one end of the cell so as to minimise any rapid flow of gas across the core of the nanofibre, as would be the case if the tower was mounted centrally. The tower consists of a MKS Instruments Type 626B Baratron Absolute Pressure transducer pressure gauge, with a maximum readout pressure of 133 Pa (1 Torr) and good acetylene compatibility [95], a gate valve to seal the gas cell and a needle valve to slow any gas flow into or out of the system. The fibre feedthroughs at the end of the gas cell are custom made O-ring sealed flanges with an 8 fibre ribbon cable sealed into it with Stycast [97]. Ribbon cables were chosen to allow multiple splice attempts when connecting the nanofibre to the gas cell.

The internal view shows the features which contain the nanofibre. At each end of the cell are spools which hold down excess fibre, with movable and lockable arms to aid with the coiling process. Running down the centre of the cell are a pair of guide rails which are used to align and mount the clamps which hold the nanofibre in place. The alignment of the clamps is important so as to minimise stress on the nanofibre. The clamps could




Figure 6.9: Overview drawing of the nanofibre gas cell, with extra details of the clamps (top left) and spooling mechanisms (bottom right). Drawing created by James Lynn.

be mounted anywhere upon the 460 mm long guide rails to enable a wide range of fibre lengths to be supported, and were used both to secure the fibre in the cell and during the drawing process described in section 6.4.3. An example clamp is shown in figure 6.10. The lid contains a hollow which was added to stop the nanofibre from touching it, which would strip off the light from the evanescent field.




Figure 6.10: Clamps used to hold the optical nanofibre in place. The clamp holds the untapered part of the fibre.

It was found to be incredibly difficult to make the gas cell vacuum tight. It was necessary to simplify the shape of the main O-ring groove over the various design iterations, and to try to minimise trapped volumes. A further problem was that it was not possible to bake the gas cell to improve the quality of the vacuum. This is because the gas cell needed to be open when the nanofibre was mounted, and baking with the nanofibre in-situ was deemed too risky.

6.4.3 Nanofibre fabrication and transfer

The optical nanofibres were manufactured at the Optoelectronics Research Centre (ORC) at the University of Southampton, using the equipment and software of Gilberto Brambilla's group.

The nanofibres were created by stretching standard Corning SMF28 optical fibre. The region which will be stretched has its cladding removed using a mechanical fibre stripper, and is then cleaned using lens tissue and isopropanol. The fibre was then mounted in a drawing rig, shown in figure 6.11, with mechanical guide clamps to ensure the fibre alignment is correct and magnetic clamps to secure the fibre in place. The clamps are secured on unstripped sections of the optical fibre.




Figure 6.11: The ORC nanofibre drawing rig with the Oxford clamps in place.

At this point the Oxford clamps were secured and the guide clamps released. The Oxford clamps were themselves mounted on a supporting bar fastened to the ORC drawing rig.

To stretch the fibre it must be heated to a temperature which allows it to flow slowly. The heating and transfer process is illustrated in figure 6.12. The heater used in the ORC drawing rig was a ceramic microheater approximately the size of a 2 cm cube. A narrow groove allows access to the central bulb, which provides the region of uniform

temperature. The heater is mounted on a tower with a series of micrometer controlled mounts to allow some adjustment of the angle and height of the heater. The heater tower can be moved over the optical fibre along a rail mounted on the optical table. All tower positioning and movement is performed manually.

Once the heater is positioned over the fibre the motion stages are started (diagram a in figure 6.12). These are controlled using custom software to draw the fibre to the specified shape, following the model described in [91]. This is achieved by simultaneously moving the fibre side to side through the heater, to provide a growing region of heated fibre, and stretching the optical fibre, which reduces the diameter through conservation of mass. The software provides an estimated fibre diameter throughout the drawing process.

While the fibre is being stretched it tends to arch upwards, requiring the operator to adjust the height of the heater throughout the process. If the fibre touches the heater at any point it will be destroyed. The fibre will also arch when the stages stop, so the ceramic heater must be withdrawn as the fibre is being stretched (and so is in motion). The stages are then manually stopped before the fibre cools and snaps. Withdrawing the heater is the most challenging aspect of the drawing procedure due to the difficulty of viewing a sub-500 nm optical fibre.

With the heater removed the Oxford claps are secured to a transfer bar (diagram b in figure 6.12). The transfer bar is shown in figure 6.13. This jig will keep the clamps in their current relative orientation as the fibre is moved into the gas cell. The magnetic clamps are only removed after the transfer bar is secured, and the move to the gas cell rails is performed as quickly and safely possible. The clamps are secured onto slot stones between the rails, and the transfer bar removed along with the L-pieces on the clamps (diagram c in figure 6.12). A clear shield is positioned over the nanofibre to protect it from falling dust.

The tails of unstretched optical fibre are spliced on to the feed-through fibres in the gas cell, the shield is removed and the gas cell sealed. The gas cell is then evacuated through a needle valve to avoid rapid air movements and prevent degradation of the nanofibre.




Figure 6.12: An overview of the nanofibre drawing process.




Figure 6.13: The transfer bar used to move the nanofibre from the ORC drawing rig to the gas cell.

Once the cell was evacuated it was sealed using the gate valve, and transported back to Oxford for the experiments.

6.5 Experimental results

Several preliminary experiments were performed using the optical nanofibres before work was stopped. These experiments were designed firstly to see a doppler broadened acetylene absorption feature (see section 6.5.1), and then to observe a doppler-free feature in saturated gas (see section 6.5.2). Further experiments, such as locking a laser to the saturated feature, were not performed because of difficulties handling the nanofibres, as described in section 6.5.3.

6.5.1 Doppler broadened acetylene absorption features

The primary motivation for observations of the doppler-broadened gas absorption features was to check that the evanescent field of the nanofibre was able to interact with the acetylene gas. The experimental layout is shown in figure 6.14.

The New Focus TLB6728 Velocity laser [71], as used in the Enhanced CFSI experiments described in chapter 5, was used for all of the nanofibre experiments. The laser is split using a 1×2 PLC splitter, with half of the light sent through the nanofibre gas

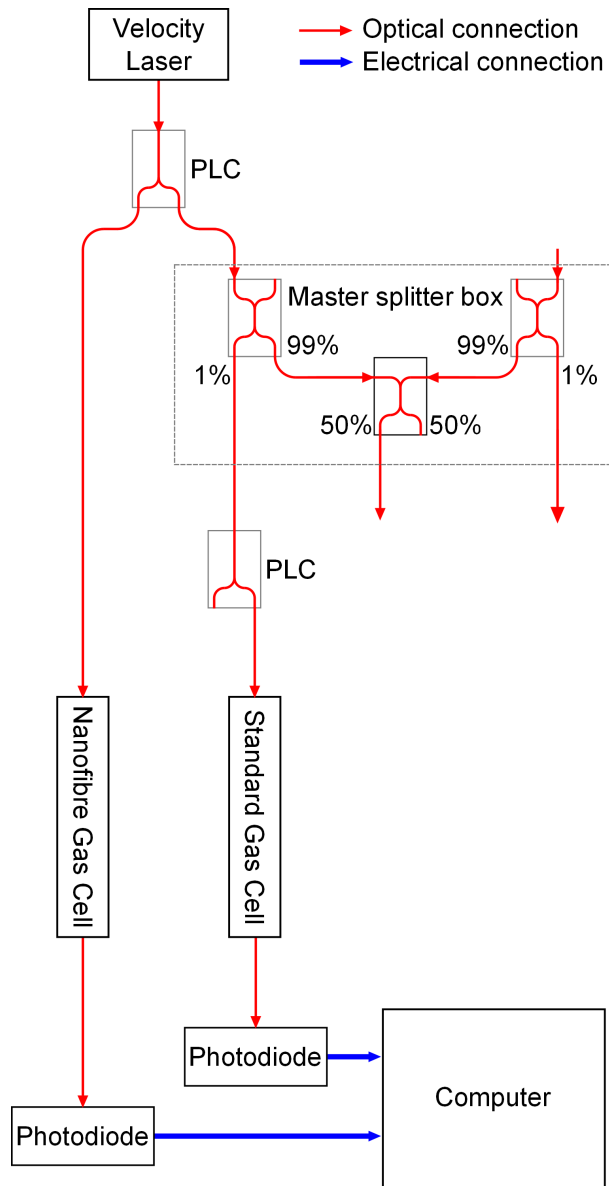


Figure 6.14: Schematic of the experiment to observe doppler-broadened acetylene-13 absorption features using a nanofibre. Both the nanofibre and standard gas cells contain acetylene-13.

cell and half of the light sent into the master splitter box described in section 3.1.3. The master splitter box allows us to feed the front-end splitter tree and a reference interferometer, with the intention of using it to generate a frequency axis, although the limited extent of the experiments meant this was not done. The output of the master splitter box is split again using another 1×2 PLC splitter, feeding both the hydrogen cyanide gas cell used previously (see section 3.1.6) and an extra acetylene-13 gas cell. The acetylene-13 gas cell, which is a Wavelength References product [98], is a fibre-coupled, 165 mm long and filled to a pressure of 6.67 kPa (50 Torr). The hydrogen cyanide cell was included to help with frequency axis generation, but was not used. The light was detected using Thorlabs PDA10CS-EC photodiodes [53] and recorded using a BNC input board and the DAQ system described in section 3.1.7.

The experiments were performed with the laser scanning its wavelength between 1541.5 nm and 1543 nm. Scans of both increasing and decreasing wavelength were performed, with a speed of 1 nm/s. The laser current was set at 200 mA. Doppler broadened absorption features were observed at a range of acetylene pressures, with an example at 130 Pa shown in figure 6.15.

The visibility of the P16 absorption feature increased with gas pressure. This can be seen by comparing the depth of the features in figure 6.16.

6.5.2 Doppler-free saturated acetylene absorption features

After the successful observation of doppler-broadened absorption features the next stage was to attempt to saturate the acetylene gas. This required a significant change in the experimental layout, shown in figure 6.17.

The Velocity laser is split using a 90:10 fibre taper splitter. The 10% output is used to feed the normal gas cell system as described in section 6.5.1. The 90% output is used to feed a Keopsys 23 dBm Erbium-doped fibre amplifier (EDFA) which is now obsolete and replaced by [99]. The amplified light is sent into a 2×1 PLC splitter, which feeds the




Figure 6.15: The P16 acetylene-13 absorption feature, observed with a nanofibre gas cell pressure of 130 Pa. The standard gas cell feature is also shown to confirm the positioning of the feature. Data point is proportional to wavelength, but an accurate wavelength axis was not generated. The filtered data has a 1 kHz low-pass filter applied.




Figure 6.16: The P16 acetylene-13 absorption feature, observed at range of nanofibre gas cell pressures. Data point is proportional to wavelength, but an accurate wavelength axis was not generated. Filtered with a 1 kHz low-pass filter.

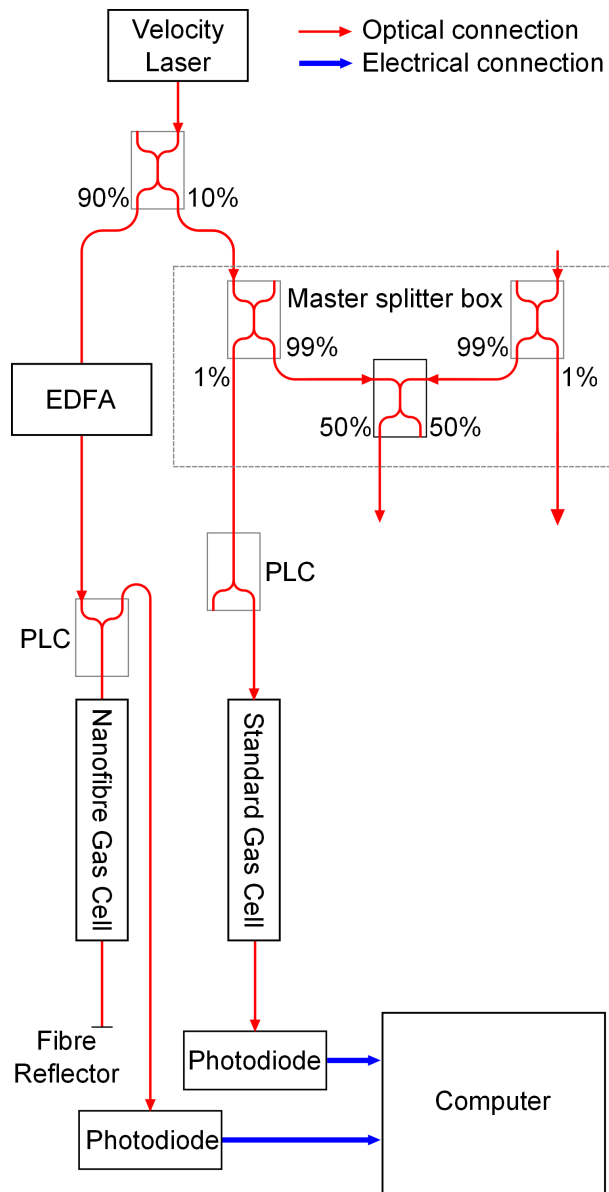


Figure 6.17: Schematic of the experiment to observe doppler-free saturated acetylene-13 absorption features using a nanofibre. Both the nanofibre and standard gas cells contain acetylene-13.

nanofibre gas cell. The input power to the nanofibre was ≈ 100 mW, which is significantly higher than the 16.2 mW used to achieve saturation with a nanofibre of similar diameter using acetylene at a higher pressure in [87], so saturated absorption features were expected to be visible. At the far side of the gas cell the light is reflected using a Thorlabs fibre retro-reflector [100], with the returning light tapped off by the PLC and observed using a PDA10CS-EC photodiode [53]. The data was recorded using a National Instruments PCI-6115 multifunction DAQ board [101], which has reduced resolution compared with the previous DAQ system (12 bit resolution compared with 14 bit resolution) but this is more than offset by the reduction in ADC noise.

The Velocity laser was not set up to scan its wavelength in this experiment. Instead, it was set to a “fixed” wavelength near to the acetylene-13 P16 absorption feature and set to wavelength tracking mode. The wavelength tracking mode of the laser oscillates the wavelength around the fixed point to enable the laser to stabilise the central wavelength. The wavelength excursion is sufficiently large to encompass the absorption feature. The benefit of this change is that the gas cell absorption feature is observed more frequently, but this comes at the cost of losing information about the direction of the laser tuning as the feature is observed.

It is unclear whether doppler-free saturated absorption features have successfully been observed. The potential saturated features, shown in figure 6.18, are too small (on the order of a few ADC digitisation levels) to be confidently identified. Other possible causes are wavelength dependent changes in laser power or an interference pattern superimposed on top of the absorption feature. Further experiments to clarify the data were curtailed due to the destruction of the nanofibre, as described in section 6.5.3.

6.5.3 Nanofibre fragility

The nanofibre gas cell experiments were stopped because the nanofibre broke. This was not a one-off occurrence, with seven nanofibres breaking in Oxford over an experimental

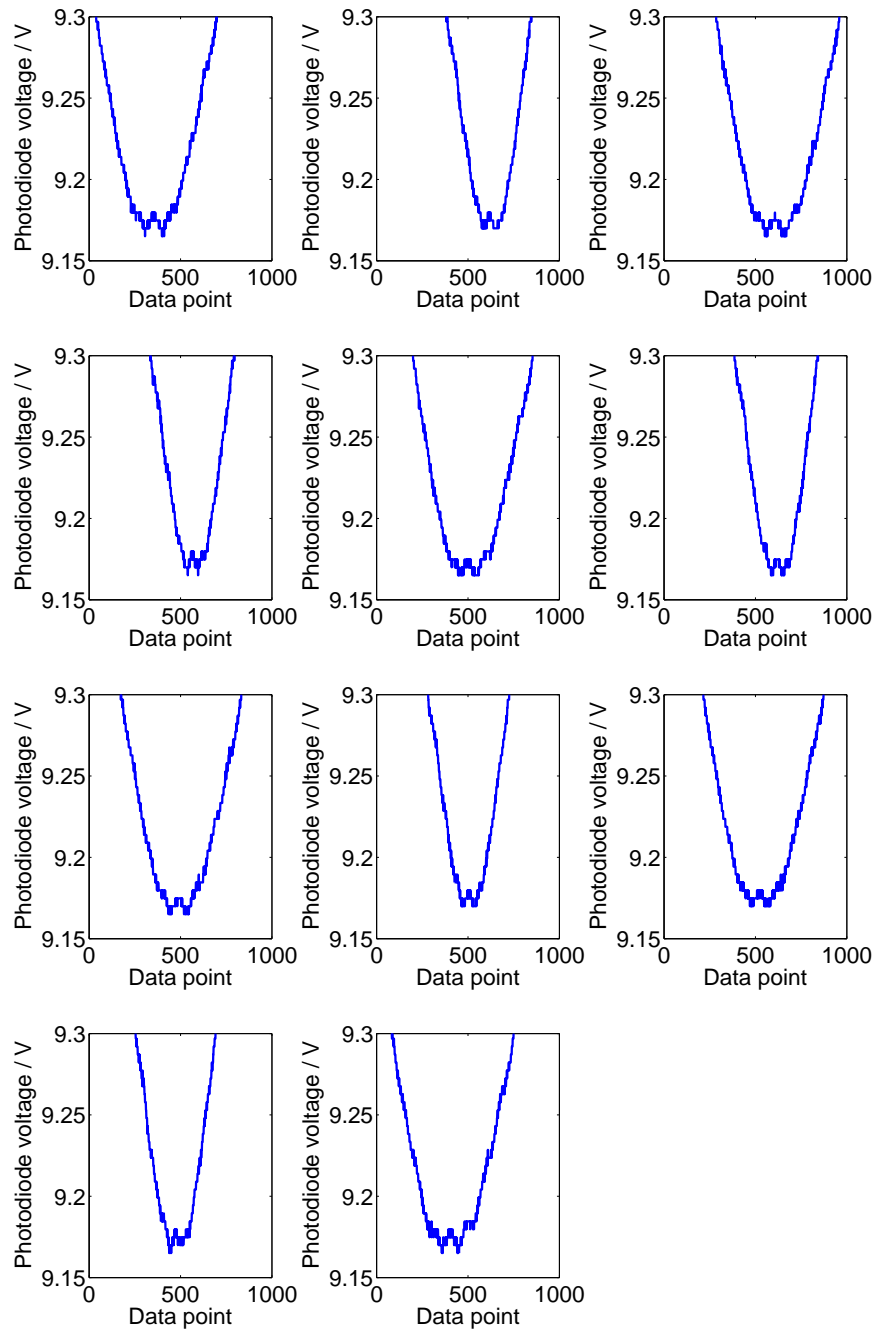


Figure 6.18: 11 consecutive scans across a potential saturated absorption feature as observed using the nanofibre gas cell. Data point indicates increasing or decreasing wavelength, which is ambiguous.

period lasting nine months. The fibres were made during a series of visits (approximately 30 days in total) to the ORC, with several attempts at fabrication during each visit. The fibre breaks were confirmed by injecting red light from the front-end splitter tree (see section 3.1.4), which escapes from the nanofibre and makes it visible. An example of a broken fibre is shown in figure 6.19.




Figure 6.19: A broken nanofibre illuminated using a red laser

Several factors combined to bring an end to the nanofibre experiments. The first of these was the difficulty in manufacturing new nanofibres. When a nanofibre broke it was necessary to organise further visits to the ORC, and the difficulty of the manufacturing process meant that it often required multiple visits to produce a single fibre. The longest period of time that a nanofibre survived in Oxford was around three weeks, during which all of the measurements presented were performed, but most nanofibres broke after a couple of days. The short lifetime of the fibres meant that it would typically take significantly longer to manufacture the fibre than it would survive in operation.

The second factor is the ambiguous cause of the nanofibre failures. In some cases the cause of the breakage was obvious, such as an inrush of gas causing instantaneous

destruction. In other cases the problem was not obvious. The nanofibres would often experience a gradual drop in transmission over a period of days, despite being kept under vacuum, until no transmission was observed. The time between nanofibre failures, caused by the manufacturing delays, and the individual nature of the failures meant that improving the equipment to prevent failures was almost impossible.

The third factor was the time taken to implement changes to the equipment when they could be made. The gas cell described in this chapter was the third iteration of the gas cell, with each iteration taking a significant amount of time to redesign and manufacture. The fibre transfer equipment also required several iterations, but did seem to be reasonably reliable by the end of the experimental period.

The experiments were beginning to show promise. If they were to be continued it would be a great improvement to build a specialised fibre tapering rig. A new rig could be tailored to produce fibres for the specific gas cell experiments, for example by enabling easier fibre transfer and to specialise in the manufacture of the specific fibres needed, rather than a general purpose manufacturing rig such as the set up at ORC. A new rig could be placed in a clean room, reducing the exposure of the fibre and vacuum system to dust. Other changes which may be worth consideration are drawing the fibre vertically, rather than horizontally, as this could reduce transverse forces, for example from air currents and sag, on the nanofibre. If a microscopy system could be implemented it would aid viewing of the nanofibre during the drawing process, and integration of the final fibre mounts into the drawing rig could reduce the number of large movements that the nanofibre would have to undergo. A simpler and more accessible heater geometry, possibly even a split heater which could be separated, would also help with fibre removal but temperature uniformity requirements could prevent this from being a suitable change. However, the largest benefit of a tapering rig in Oxford would be a significant reduction in the time between equipment iterations. The gas cell could also be improved, with a major aim being to reduce the vacuum sealing surface area, for example by loading the nanofibre from one end rather than from above which requires a large lid.

Chapter 7

Summary and future work

This chapter will summarise the work in this thesis, highlight areas where results require clarification and suggest suitable approaches for further investigation. The techniques developed fall into two categories which will be summarised separately: the interferometric techniques (continuous FSI and enhanced continuous FSI) in section 7.1 and the acetylene saturated absorption spectroscopy work in section 7.2.

7.1 Interferometric techniques

Both continuous FSI and enhanced continuous FSI have been shown to be promising techniques in a series of demonstration experiments performed with modified dynamic FSI hardware based around a hydrogen cyanide gas cell and using single-channel, non-polarization maintaining optics. Continuous FSI measurements have been performed in a range of 2.8–4.9 m with standard deviations below $3 \mu\text{m}$, and compared against a reference system giving an absolute calibration constant (scale factor) of $0.999998496 \pm 8.919 \times 10^{-8}$, calculated from seven measurements with a range of 8.668×10^{-7} . Measurements have been performed for up to 1 minute and compared against the reference system for 40 seconds, remaining consistent to within a few microns over the course of each measurement. The measurements have successfully monitored a variety of target motions at a range of

over 4 m (8 m OPD), covering micron-scale oscillations at up to 70 Hz, distance increases at approximately 2 mm/s and 8 mm/s and distance decreases at 2 mm/s.

It has been demonstrated that it is possible to swap between continuous FSI and differential fixed frequency interferometry measurements to increase the motion tolerance of the measurement system, from an estimated 0.0022L/s to an estimated 268 mm/s (535 mm/s OPD change), and to isolate the scanning lasers from the measurement interferometer. Target motions of between 8 mm/s and 16 mm/s have been observed using the FFI system, with the direction ambiguity eliminated using the continuous FSI distance measurements. Graphics processing units have been used in the analysis of the data and demonstrated dramatically reduced processing times when compared with previous implementations, approaching real-time speeds in the case of basic continuous FSI (which have reached 2 1/2 minutes for analysis of a 1 minute measurement).

Although both FSI techniques qualitatively appear to work significant questions remain unanswered. Answering these will require both changes to the experimental layout of the measurement system and calibration experiments, along with further study of the analysis process and its sources of drift over time. Some of the suggestions made here are already being implemented.

7.1.1 Experimental developments

The most important change required in the experimental procedure is to move to a calibration experiment layout with a significantly reduced difference in optical path between the FSI and reference systems. This would remove one of the suspected sources of the measurement drifts, common to both FSI and FFI systems. While this adjustment sounds trivial it could be difficult to implement and require a change in reference system and a larger optical bench to enable the long-range measurements currently required for continuous FSI. As an extra observation the FFI system should be included in all future calibration experiments as it has proven to be invaluable for highlighting external con-

tributions to differences between the continuous FSI and reference systems. The mature analysis code would now permit results to be generated while the experiments are being performed and so allow unexpected effects to be investigated in-situ rather than being detected months later.

The implementation of the enhanced continuous FSI system should be adjusted to equalise the fringe visibility of the FSI and FFI signals. This would enable the FFI component to be detected by the automatic filtering algorithm, removing the need for manual intervention in the analysis. It may also be beneficial to include a dedicated beat detection photodiode, which would provide an improved signal to noise ratio for the measurement and remove the need for heavy filtering of the measurement interferometer signal to extract the beat signal. If a beat detection photodiode is not implemented then the fixed frequency laser light should be introduced into the reference interferometer. This should not interfere with its measurement but would provide a beat signal even when the scanning lasers are removed from the measurement interferometer, which is not currently the case. If the fixed frequency laser were locked to a saturated absorption feature its frequency would be accurately known and determining its frequency via the FSI system would be unnecessary. Using its beat signal as an extra point of known frequency (effectively an extra gas absorption feature) may also prove beneficial.

The gas cell referencing system could be improved by including more gas cells containing a variety of different gases. These are commercially available, and would provide additional independent measurements of the reference interferometer length and scanning laser frequencies.

The synchronisation of the two scanning lasers should also be improved. In the current implementation they are set to perform a certain number of scanning cycles and started with a fixed time delay. The level of synchronisation provided by this system was enough for the minute-long experiments in this thesis, but degrades over time. A system which restarts the lasers at regular intervals or lasers accepting external triggers could help to alleviate the problem. It may be possible to improve the level of synchronisation provided

by the current system by further optimisation of the scanning ranges of the two lasers.

An important improvement of the continuous FSI system would be to reduce the minimum measurable length, currently limited by signal separation requirements as described in section 4.1.2. The obvious solution is to increase the laser tuning speeds at small distances, but this would raise a series of issues. Firstly it may be necessary to increase the sampling speed of the data acquisition system in order to allow the analysis to cope with the increased speeds while meeting the requirements described in section 3.3. Other potential difficulties arise from the data requirements for the peak finding algorithms and Fourier transformations used throughout the analysis. Faster optical switches may be required to isolate the lasers quickly enough at the end of each scan. The analysis will also have to be performed at a quicker rate.

The absolute calibration constant, as described in section 4.1, is also an area of concern as it represents something unknown about the FSI system. Further effort should be made to understand its source, which may in turn change the point in the analysis at which it is applied or even eliminate it entirely. One possible source of the calibration constant is the gas cell pressure uncertainty, which will also affect the laser frequency calculation, and so effort should be made to determine the pressure more accurately to separate pressure effects from the global absolute calibration constant. It may be possible to determine the gas cell pressure by performing a series of absolute calibration experiments with the gas cell held at a range of temperatures, which will in turn modify the gas cell pressure.

7.1.2 Analysis developments

The most obvious improvement needed by the analysis is that it needs to be faster. Moving to a real-real time system greatly increases the utility of the techniques. This is especially true of enhanced continuous FSI, which would need to decide when to change to the fixed frequency system and isolate the scanning lasers, processes which are currently controlled manually. While the GPU approach could feasibly attain real-time analysis

of a single measurement channel it may be necessary to introduce new approaches, such as the use of field programmable gate arrays for some signal processing or changing the interferometer design to simplify phase extraction (for example by separating the two scanning lasers by polarization and using multi-fibre readout for phase determination), in a multi-channel system.

The analysis algorithm may itself require significant revision. There are many potential sources for error which have not yet been studied such as random walk in the phase extraction procedure (which could be introduced through the phase chaining procedure described in section 3.2.2.6), the optimal choice of gas absorption features to use for laser frequency determination and the ability to update the length of the reference interferometer as a run progresses. Recalculation of the transfer length using the averaging method used for the initial calculation could also be of benefit.

7.2 Saturated absorption spectroscopy

The acetylene doppler-free saturated absorption experiments are difficult to draw conclusions from because of their premature end. Significant progress has been made in terms of equipment design and manufacture. An acetylene storage cylinder with an associated vacuum system which has been successfully used transfer acetylene both from the break-seal flask in which it was supplied into the cylinder and also into the nanofibre gas cell for use in experiments. Several iterations of the nanofibre gas cell have been manufactured, along with mechanisms to improve the transfer of nanofibres from the drawing rig used to manufacture them into the gas cell. Seven nanofibres were successfully manufactured and transported to Oxford for experiments.

Doppler broadened acetylene absorption features have been successfully observed using the nanofibre gas cell at pressures below 130 Pa. It may also be the case that doppler-free saturated absorption features have been observed, but the current measurements are ambiguous and would benefit from clarification measurements that were not

possible due to the destruction of the optical nanofibres as described in section 6.5.3.

The experiments were beginning to show promise but will require significant redevelopment, with recommendations to improve the experiments already described in section 6.5.3. The nanofibre-based system implemented here was proposed for its flexibility and mobility, but with more experience of the fragility of nanofibres it may prove difficult to make a portable system and other approaches, such as use of a Fabry-Perot cavity, may need to be considered.

Appendix A

Length difference plots

This appendix provides further examples of the difference between continuous FSI or FFI distance measurements from the LaserTRACER reference system. Figures A.1 and A.2 provide additional examples to complement figures 4.26 and 4.27, with the experiments described in section 4.3. Figure A.3 provides additional examples to complement figure 5.12, with the experiments described in section 5.2.

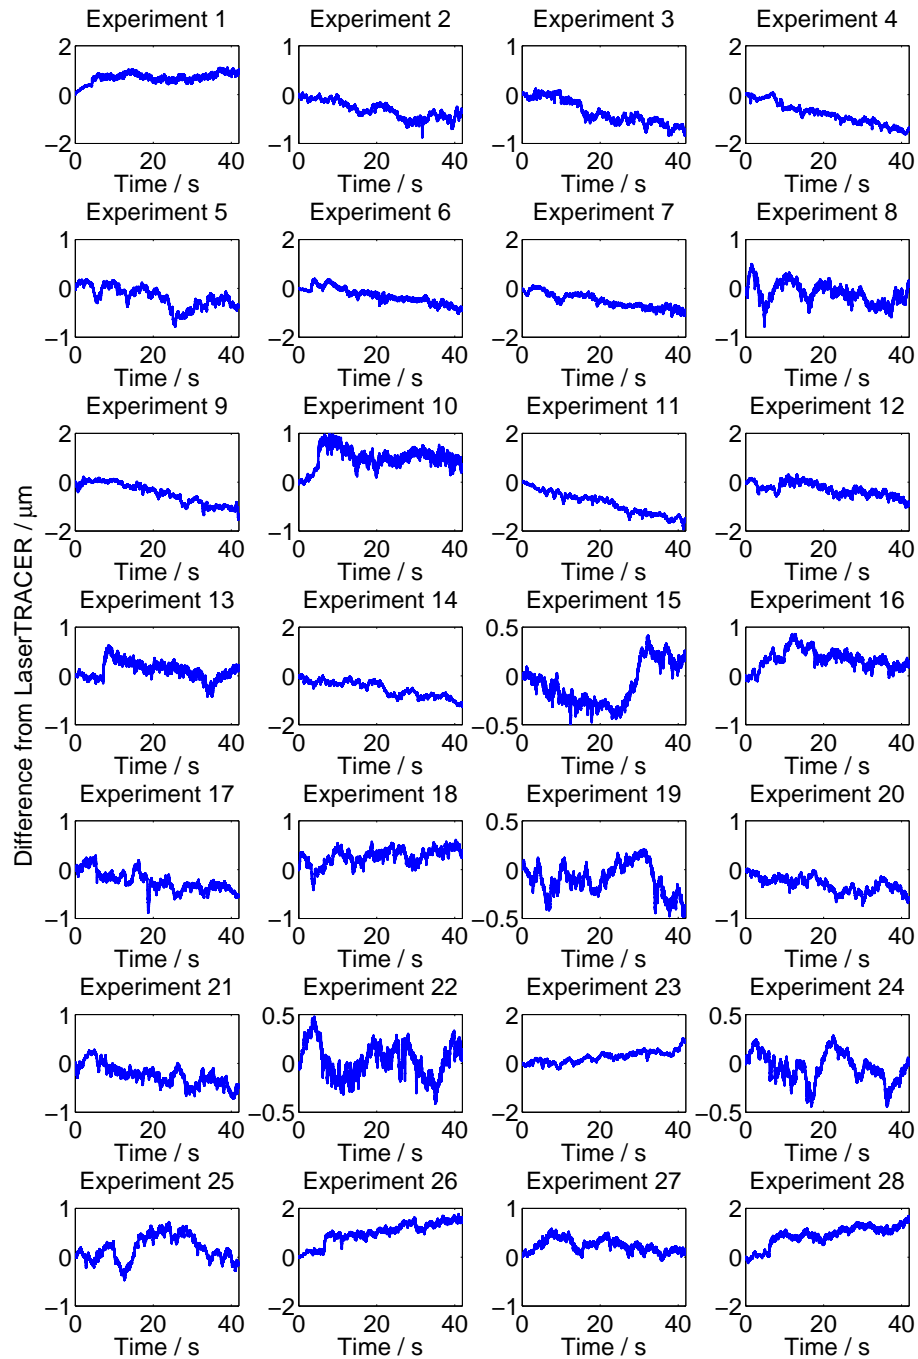


Figure A.1: Examples of the difference between the continuous FSI and LaserTRACER measurements for the forward experiments.

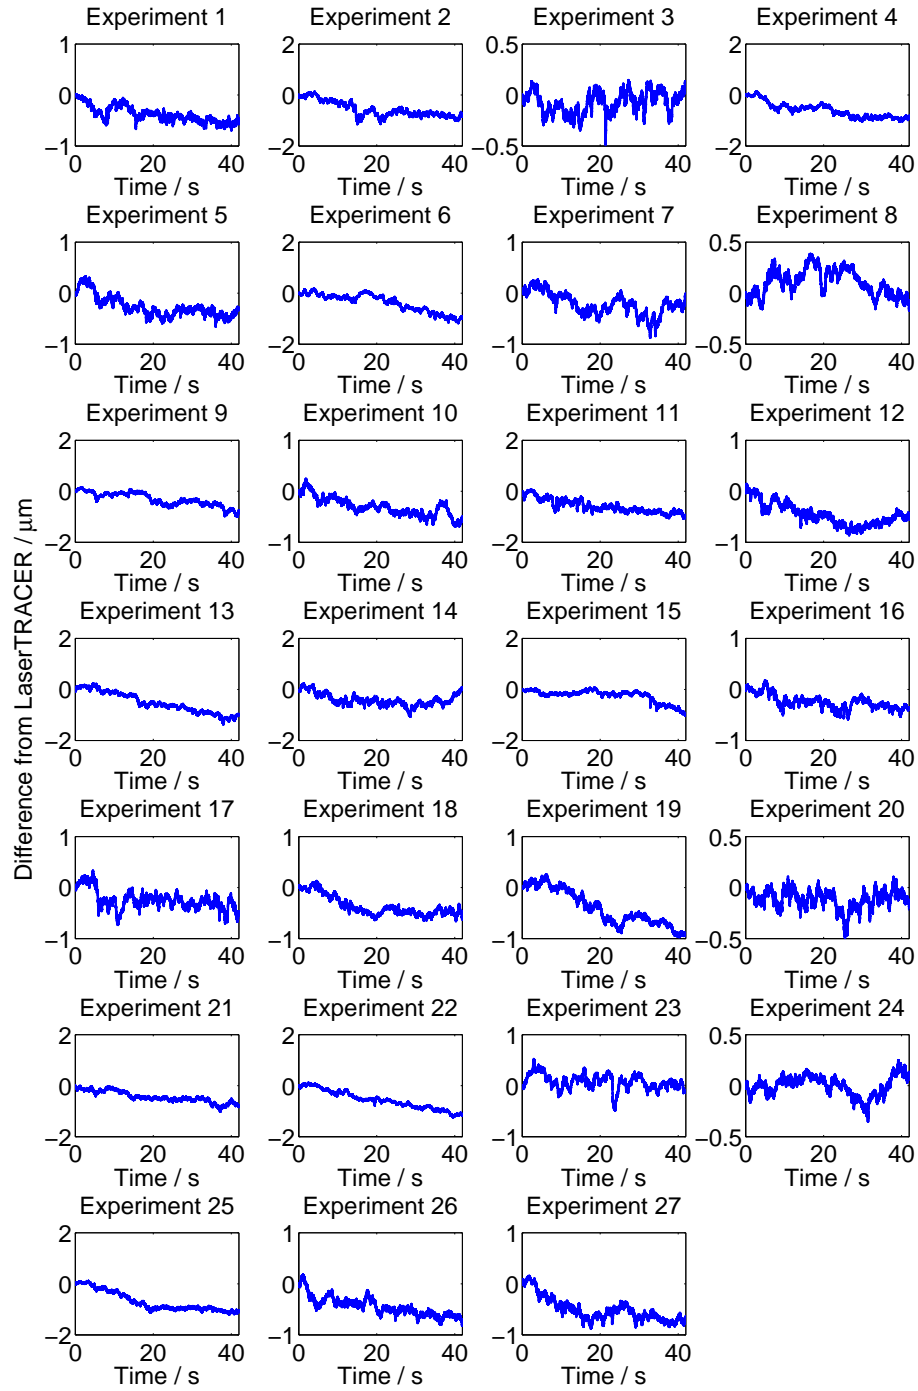


Figure A.2: Examples of the difference between the continuous FSI and LaserTRACER measurements for the backward experiments.




Figure A.3: Examples of the difference between the continuous FSI and FFI systems from the LaserTRACER measurements for the enhanced CFSI demonstration experiments.

Bibliography

- [1] M S Warden. *Absolute distance metrology using frequency swept lasers*. D.Phil. thesis, University of Oxford, 2011.
- [2] D Urner. M Warden. Apparatus and method for measuring distance. International Patent Number WO2012022956 (A1).
- [3] John Dale, Ben Hughes, Andrew J. Lancaster, Andrew J. Lewis, Armin J. H. Reichold, and Matthew S. Warden. Multi-channel absolute distance measurement system with sub ppm-accuracy and 20 m range using frequency scanning interferometry and gas absorption cells. *Opt. Express*, 22(20):24869–24893, Oct 2014.
- [4] René Benoît. Application des phénomènes d’interférence à des déterminations métrologiques. *J. Phys. Theor. Appl.*, 7(1):57–68, 1898.
- [5] F. Bien, M. Camac, H. J. Caulfield, and S. Ezekiel. Absolute distance measurements by variable wavelength interferometry. *Appl. Opt.*, 20(3):400–403, Feb 1981.
- [6] G P Barwood, P Gill, and W R C Rowley. High-accuracy length metrology using multiple-stage swept-frequency interferometry with laser diodes. *Measurement Science and Technology*, 9(7):1036, 1998.

- [7] Markus-Christian Amann, Thierry Bosch, Marc Lescure, Risto Myllyla, and Marc Rioux. Laser ranging: a critical review of usual techniques for distance measurement. *Optical Engineering*, 40(1):10–19, 2001.
- [8] J. M. Payne, D. Parker, and R. Bradley. Rangefinder with fast multiple range capability. *Review of Scientific Instruments*, 63(6):3311–3316, 1992.
- [9] W.T. Estler, K.L. Edmundson, G.N. Peggs, and D.H. Parker. Large-scale metrology an update. *CIRP Annals - Manufacturing Technology*, 51(2):587 – 609, 2002.
- [10] G N Peggs, P G Maropoulos, E B Hughes, A B Forbes, S Robson, M Ziebart, and B Muralikrishnan. Recent developments in large-scale dimensional metrology. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, 223(6):571–595, 2009.
- [11] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, Series 3 Vol. 34(203):333–345, 1887.
- [12] Eugene Hecht. *Optics*. Addison Wesley, 4th edition, 2002.
- [13] Karl-Heinz Bechstein and Werner Fuchs. Absolute interferometric distance measurements applying a variable synthetic wavelength. *Journal of Optics*, 29(3):179, 1998.
- [14] Jack A. Stone, Alois Stejskal, and Lowell Howard. Absolute interferometry with a 670-nm external cavity diode laser. *Appl. Opt.*, 38(28):5981–5994, Oct 1999.
- [15] Th Kinder and K-D Salewski. Absolute distance interferometer with grating-stabilized tunable diode laser at 633 nm. *Journal of Optics A: Pure and Applied Optics*, 4(6):S364, 2002.

- [16] Florian Pollinger, Karl Meiners-Hagen, Martin Wedde, and Ahmed Abou-Zeid. Diode-laser-based high-precision absolute distance interferometer of 20 m range. *Appl. Opt.*, 48(32):6188–6194, Nov 2009.
- [17] Richard Schneider, Peter Thurmel, and Michael Stockmann. Distance measurement of moving objects by frequency modulated laser radar. *Optical Engineering*, 40(1):33–37, 2001.
- [18] J.J. Martinez, M.A. Campbell, M.S. Warden, E.B. Hughes, N.J. Copner, and A.J. Lewis. Dual-sweep frequency scanning interferometry using four wave mixing. *Photonics Technology Letters, IEEE*, 27(7):733–736, April 2015.
- [19] G. L. Bourdet and A. G. Orszag. Absolute distance measurements by CO2 laser multiwavelength interferometry. *Appl. Opt.*, 18(2):225–227, Jan 1979.
- [20] A Lewis. Measurement of length, surface form and thermal expansion coefficient of length bars up to 1.5 m using multiple-wavelength phase-stepping interferometry. *Measurement Science and Technology*, 5(6):694, 1994.
- [21] Esther Baumann, Fabrizio R. Giorgetta, Ian Coddington, Laura C. Sinclair, Kevin Knabe, William C. Swann, and Nathan R. Newbury. Comb-calibrated frequency-modulated continuous-wave ladar for absolute distance measurements. *Opt. Lett.*, 38(12):2026–2028, Jun 2013.
- [22] Ruitao Yang, Florian Pollinger, Karl Meiners-Hagen, Jiubin Tan, and Harald Bosse. Heterodyne multi-wavelength absolute interferometry based on a cavity-enhanced electro-optic frequency comb pair. *Opt. Lett.*, 39(20):5834–5837, Oct 2014.
- [23] I. Coddington, W. C. Swann, L. Nenadovic, and N. R. Newbury. Rapid and precise absolute distance measurements at long range. *Nature Photonics*, 3(6):351–356, May 2009.

- [24] S. A. van den Berg, S. T. Persijn, G. J. P. Kok, M. G. Zeitouny, and N. Bhattacharya. Many-wavelength interferometry with thousands of lasers for absolute distance measurement. *Phys. Rev. Lett.*, 108:183901, May 2012.
- [25] Kaoru Minoshima and Hirokazu Matsumoto. High-accuracy measurement of 240-m distance in an optical tunnel by use of a compact femtosecond laser. *Appl. Opt.*, 39(30):5512–5517, Oct 2000.
- [26] P A Coe, D F Howell, and R B Nickerson. Frequency scanning interferometry in ATLAS: remote, multiple, simultaneous and precise distance measurements in a hostile environment. *Measurement Science and Technology*, 15(11):2175, 2004.
- [27] A. Reichold et al. The LiCAS Rapid Tunnel Reference Surveyor—the status after commissioning. *10th IWAA, KEK, Tsukuba, Japan.*, Feb 2008.
- [28] M. Warden, P.A. Coe, D. Urner, and A. Reichold. MONALISA: A precise system for accelerator component position monitoring. *10th IWAA, KEK, Tsukuba, Japan.*, Feb 2008.
- [29] G. Aad et al. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.
- [30] M Turala. The ATLAS semiconductor tracker. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 466(2):243 – 254, 2001. 4th Int. Symp. on Development and Application of Semiconductor Tracking Detectors.
- [31] S. M. Gibson, P. A. Coe, M. Dehchar, J. Fopma, D. F. Howell, R. B. Nickerson, and G. Viehhauser. First Data from the ATLAS Inner Detector FSI Alignment System. *10th IWAA, KEK, Tsukuba, Japan.*, Feb 2008.
- [32] A.F. Fox-Murphy, D.F. Howell, R.B. Nickerson, and A.R. Weidberg. Frequency scanned interferometry (FSI): the basis of a survey system for ATLAS using fast

- automated remote interferometry. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 383(1):229 – 237, 1996. Development and Application of Semiconductor Tracking Detectors.
- [33] S.M. Gibson, M. Dehchar, K. Horton, A. Lewis, Z. Liang, et al. A novel method for ATLAS FSI alignment based on rapid, direct phase monitoring, 2010.
- [34] C Adolphsen et al. The International Linear Collider Technical Design Report, 2013.
- [35] J Dale. *A Study of Interferometric Distance Measurement Systems on a Prototype Rapid Tunnel Reference Surveyor and the Effects of Reference Network Errors at the International Linear Collider*. D.Phil. thesis, University of Oxford, 2009.
- [36] J R Green. *Development of a Prototype Frequency Scanning Absolute Distance Measurement System for the Survey & Alignment of the International Linear Collider*. D.Phil. thesis, University of Oxford, 2007.
- [37] G R Moss. *A Laser Based Straightness Monitor for a Prototype Automated Linear Collider Tunnel Surveying System*. D.Phil. thesis, University of Oxford, 2013.
- [38] D. Urner, M.S. Warden, P.A. Coe, and A. Reichold. The Compact Distance Meter Interferometer Prototype. http://www.eurotev.org/reports_presentations/eurotev_reports/2008/e37072/EUROTeV-Report-2008-097.pdf, Feb 2009.
- [39] D. Urner, P.A. Coe, and A. Reichold. Stabilization of the ILC final focus using interferometers. <http://accelconf.web.cern.ch/AccelConf/e06/PAPERS/THPCH090.PDF>.
- [40] Etalon A.G. Absolute Multiline Tecnology. http://www.etalon-ag.com/wp-content/uploads/pdf/PR_Absolute-Multiline-Technologie-Etalon-2014-02.pdf.

- [41] John Dale. Private communication.
- [42] William C. Swann and Sarah L. Gilbert. Line centers, pressure shift, and pressure broadening of 1530-1560 nm hydrogen cyanide wavelength calibration lines. *J. Opt. Soc. Am. B*, 22(8):1749–1756, Aug 2005.
- [43] Sarah L. Gilbert William C. Swann and Chih-Ming Wang. Hydrogen Cyanide $H^{13}C^{14}N$ Absorption Reference for 1530 nm to 1565 nm Wavelength Calibration SRM 2519a. *NIST Special Publication 260-137*, 2005 Edition.
- [44] Newport Spectra-Physics. New Focus TLM-8700 OEM Tunable Laser Modules data sheet. http://www.newport.com/TLM-8700-OEM-Tunable-Laser-Modules/917478/1033/info.aspx#tab_Specifications.
- [45] Sercalo microtechnology. Miniature fiber optic MEMS switch datasheet.
- [46] Intelite Inc. FRM635–10S-APC fiber pigtailed laser module specifications.
- [47] D.R. Williams M. Bass, E.W. van Stryland and W.L. Wolfe, editors. *Handbook of Optics*, volume 2, chapter 21 - Interferometers. McGraw-Hill, 2nd edition, 1995.
- [48] Thorlabs. FC/PC fixed aspheric lens fiber collimation packages. <http://www.thorlabs.de/catalogpages/V21/1098.PDF>.
- [49] PLX inc. Plx stainless steel ball mounted hollow retroreflectors product page. <http://www.plxinc.com/products/stainless-steel-ball-mounted-hollow-retroreflectors-bmrs>.
- [50] Wavelength References. Hydrogen cyanide fiber coupled gas cell for remote optical gas sensor and calibration c-band. <http://www.wavelengthreferences.com/pdf/HCN%20fibercoupled%20-%202014.pdf>.
- [51] National Instruments corporation. NI USB-TC01 - data sheet. <http://sine.ni.com/ds/app/doc/p/id/ds-215/lang/en>.

- [52] OpticReach Ltd. PLC splitter modules data sheet. http://www.opticreach.com/images/stories/Opticreach_Datasheets/DATASHEET%20-%20%20PLC%20Splitter.pdf.
- [53] Thorlabs. PDA10CS InGaAs Switchable Gain Detector User Guide. <http://www.thorlabs.de/thorcat/13000/PDA10CS-EC-Manual.pdf>.
- [54] Roy Wastie. ADC Board VHDL Firmware development for Mona Lisa. <http://www-pnp.physics.ox.ac.uk/~monalisa/open/pages/projects/SummerProjects2008/FPGA/ADCBoardVHDLFirmware.ppt>.
- [55] Jack Hickish. MONALISA ADC/FPGA Documentation. <http://www-pnp.physics.ox.ac.uk/~monalisa/open/pages/projects/SummerProjects2008/FPGA/Documentation.pdf>.
- [56] NVIDIA Corporation. CUDA programming language. http://www.nvidia.com/object/cuda_home_new.html.
- [57] NVIDIA. Nvidia Tesla GPU accelerators. <http://www.nvidia.co.uk/content/tesla/pdf/NVIDIA-Tesla-Kepler-Family-Datasheet.pdf>.
- [58] B. Boashash. Estimating and interpreting the instantaneous frequency of a signal. i. fundamentals. *Proceedings of the IEEE*, 80(4):520–538, Apr 1992.
- [59] B. Boashash. Estimating and interpreting the instantaneous frequency of a signal. ii. algorithms and applications. *Proceedings of the IEEE*, 80(4):540–568, Apr 1992.
- [60] S. Butterworth. On the Theory of Filter Amplifiers. *Experimental Wireless and the Wireless Engineer*, 7:536–541, Oct 1930.
- [61] NVIDIA Corporation. cuFFT library. <https://developer.nvidia.com/cuFFT>.

- [62] F. James and M. Roos. Minuit - a system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343 – 367, 1975.
- [63] Etalon AG. LaserTRACER. <http://www.etalon-ag.com/?lang=en>.
- [64] Andrew Lewis. Private communication.
- [65] The MathWorks Inc. MATLAB. <http://uk.mathworks.com/products/matlab/>.
- [66] D.E. Wells and E.J. Krakiwsky. *The Method of Least Squares*. Geometry and Geomatics Engineering University of New Brunswick, 1971.
- [67] Philip E. Ciddor. Refractive index of air: new equations for the visible and near infrared. *Appl. Opt.*, 35(9):1566–1573, Mar 1996.
- [68] Thorlabs. Thorlabs piezoelectric actuator pas001, discontinued. Thorlabs catalogue V18.
- [69] Thurlby Thandar Instruments Ltd. Aim & Thurlby Thandar Instruments TG550 data sheet. <http://www.ttid.co.uk/products-tti/pdf-brochure/gen-tg550-4p.pdf>.
- [70] Roger Barlow. Asymmetric Statistical Errors <http://arxiv.org/abs/physics/0406120>. *arXiv:physics/0406120*, June 2004.
- [71] Newport Spectra-Physics. New Focus Velocity 6700 Widely Tunable Lasers data sheet. http://assets.newport.com/webDocuments-EN/images/DS_041104_Velocity_Datasheet.pdf.
- [72] British Standards Institution. BS EN 60825-1:2007 Safety of laser products - Part 1: Equipment classification and requirements, 2007.
- [73] Thorlabs. VOA50-APC - SM Variable Attenuator, 1310/1550 nm, 50 dB, in-line, FC/APC connectors. <http://www.thorlabs.de/thorproduct.cfm?partnumber=VOA50-APC>.

- [74] National Instruments corporation. NI USB-6501 data sheet. <http://www.ni.com/datasheet/pdf/en/ds-135>.
- [75] S. J. Blundell and K. M. Blundell. *Concepts in Thermal Physics*. Oxford University Press, 2006.
- [76] Guang S. He. *Nonlinear Optics and Photonics*. Oxford University Press, 2015.
- [77] Christopher J. Foot. *Atomic Physics*. Oxford master series in Atomic, Optical and Laser Physics. Oxford University Press, 2005.
- [78] P.A. Coe, A.J. Lancaster, G. Hammet, R. Wastie, H. Dabiri Khah, J. Fopma, P. Nandanwar, M.S. Warden, and A.J.H. Reichold. Fibre Laser Frequency Stabilisation system. Unpublished.
- [79] Günther Prellinger, Karl Meiners-Hagen, and Florian Pollinger. Spectroscopically in situ traceable heterodyne frequency-scanning interferometry for distances up to 50m. *Measurement Science and Technology*, 26(8):084003, 2015.
- [80] T J Quinn. Practical realization of the definition of the metre, including recommended radiations of other optical frequency standards (2001). *Metrologia*, 40(2):103, 2003.
- [81] J. Michael Hollas. *Modern Spectroscopy*. Wiley, 4th edition, 2004.
- [82] C.S. Edwards, H.S. Margolis, G.P. Barwood, S.N. Lea, P. Gill, and W.R.C. Rowley. High-accuracy frequency atlas of $^{13}\text{C}_2\text{H}_2$ in the $1.5\ \mu\text{m}$ region. *Applied Physics B*, 80(8):977–983, 2005.
- [83] Sigma Aldrich. Product 617520-500ML-EU. Discontinued.
- [84] M. de Labachellerie, K. Nakagawa, and M. Ohtsu. Ultranarrow $^{13}\text{C}_2\text{H}_2$ saturated-absorption lines at $1.5\ \mu\text{m}$. *Opt. Lett.*, 19(11):840–842, Jun 1994.

- [85] A. Czajkowski, J.E. Bernard, A.A. Madej, and R.S. Windeler. Absolute frequency measurement of acetylene transitions in the region of 1540 nm. *Applied Physics B*, 79(1):45–50, 2004.
- [86] Petr Balling, Marc Fischer, Philipp Kubina, and Ronald Holzwarth. Absolute frequency measurement of wavelength standard at 1542nm: acetylene stabilized DFB laser. *Opt. Express*, 13(23):9196–9201, Nov 2005.
- [87] Masato Takiguchi, Yutaka Yoshikawa, Takayuki Yamamoto, Kazuyuki Nakayama, and Takahiro Kuga. Saturated absorption spectroscopy of acetylene molecules with an optical nanofiber. *Opt. Lett.*, 36(7):1254–1256, Apr 2011.
- [88] Limin Tong, Jingyi Lou, and Eric Mazur. Single-mode guiding properties of subwavelength-diameter silica and silicon wire waveguides. *Opt. Express*, 12(6):1025–1035, Mar 2004.
- [89] Gilberto Brambilla, Fei Xu, Peter Horak, Yongmin Jung, Fumihito Koizumi, Neil P. Sessions, Elena Koukharenko, Xian Feng, Ganapathy S. Murugan, James S. Wilkinson, and David J. Richardson. Optical fiber nanowires and microwires: fabrication and applications. *Adv. Opt. Photon.*, 1(1):107–161, Jan 2009.
- [90] Konstantin Karapetyan, Fabian Bruse, Cristian Dan, Dimitri Pritzkau, Ulrich Wiedemann, Wolfgang Alt, and Dieter Meschede. MATLAB Optical Fibre Toolbox. <http://uk.mathworks.com/matlabcentral/fileexchange/27819-optical-fibre-toolbox>.
- [91] T.A. Birks and Y.W. Li. The shape of fiber tapers. *Lightwave Technology, Journal of*, 10(4):432–438, Apr 1992.

- [92] Jes Henningsen, Jan Hald, and Jan C. Peterson. Saturated absorption in acetylene and hydrogen cyanide in hollow-core photonic bandgap fibers. *Opt. Express*, 13(26):10475–10482, Dec 2005.
- [93] R. Thapa, K. Knabe, K. L. Corwin, and B. R. Washburn. Arc fusion splicing of hollow-core photonic bandgap fibers for gas-filled fiber cells. *Opt. Express*, 14(21):9576–9583, Oct 2006.
- [94] R. J. Lewis and N.I. Sax. *Sax's dangerous properties of industrial materials*, volume 2. Wiley-Interscience, 11th edition, 2004.
- [95] MKS Instruments. Type 626b absolute baratron ambient-temperature capacitance manometer data sheet. <http://www.mksinst.com/docs/UR/626B-profibus-DS.pdf>.
- [96] James Lynn. Private communication.
- [97] Henkel AG & Co. KGaA. Loctite stycast 2850ft. <http://www.henkel-adhesives.com/product-search-1554.htm?nodeid=8802585018369>.
- [98] Wavelength References. C-band wavelength calibrator $^{13}\text{C}_2\text{H}_2$ acetylene gas cell.
- [99] Keopsys SA. CEFA-C-BO-HP series data sheet. <http://www.keopsys.com/uploads/images/PDF/datasheet/CEFA-C-BO-HP.pdf>.
- [100] Thorlabs. P5-SMF28ER-P01-1 - single mode retroreflector, 1260 - 1620 nm, FC/APC connector. <http://www.thorlabs.de/thorproduct.cfm?partnumber=P5-SMF28ER-P01-1>.
- [101] National Instruments corporation. S series multifunction DAQ data sheet. http://www.ni.com/pdf/products/us/4daqsc197-198_ETC_212-213.pdf.