

Durham E-Theses

New Techniques for Coherence Imaging Fusion Plasmas

JOSEPH STEVEN ALLCOCK

How to cite:

ALLCOCK, JOSEPH STEVEN (2021) *New Techniques for Coherence Imaging Fusion Plasmas*.
Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/14115/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

New Techniques for Coherence Imaging of Fusion Plasmas

Joseph Steven Allcock

A thesis presented for the degree of
Doctor of Philosophy



Centre for Advanced Instrumentation
Department of Physics
Durham University
UK

July 2021

New Techniques for Coherence Imaging of Fusion Plasmas

Joseph Steven Allcock

Submitted for the degree of Doctor of Philosophy

July 2021

Abstract: Imaging diagnostic techniques are desirable for fusion plasma experiments for their wide coverage and high spatial resolution, which allows for a more complete comparison with the predictions made by plasma physics models than traditional techniques. Benchmarking models against measurements made on current experiments improves our understanding of the physics and reduces the uncertainties involved with designing future experiments and reactors.

This thesis presents new techniques for coherence imaging (CI), an interferometric narrowband spectral imaging technique used to measure the brightness, shift and width of spectral lines emitted by the plasma in the visible range. From these measurements, 2-D maps of emitting species flow velocity, and temperature can be inferred via Doppler shifts and broadening respectively. For passive hydrogen Balmer series emission in the tokamak divertor, Stark broadening is strong enough to provide a 2-D map of electron density n_e .

First, we introduce novel CI instrument designs based on pixelated phase-mask (PPM) interferometry, which improve spatial resolution and robustness over typical linear carrier designs. Secondly, we introduce a new method for absolute calibration of CI flow velocity measurements using emission lines from standard gas-discharge lamps instead of a tuneable laser. This method significantly reduces hardware costs while maintaining high measurement accuracy — ± 1 km/s compared to typical ion flows in the tokamak plasma edge of $\lesssim 30$ km/s. Lastly, we present improved methods for CI measurement of n_e , using modern lineshape models to improve accuracy and using a multi-delay PPM-CI instrument design to minimise errors caused by Doppler broadening, extending the valid measurement range to lower n_e . This is demonstrated with experimental measurements of H_γ and H_δ emission on the Magnum-PSI linear plasma experiment with a direct comparison to Thomson scattering measurements.

Contents

Abstract	ii
List of Tables	v
List of Figures	vi
1 Introduction	1
1.1 Magnetic Confinement Fusion and the Tokamak	2
1.2 The Heat Exhaust Problem	6
1.3 Diagnosing the MAST-U Divertor	9
1.4 Thesis Outline	12
2 Coherence Imaging	14
2.1 Coherence and Two-beam Interferometry	14
2.2 Fourier Transform Spectroscopy	18
2.2.1 Quasi-monochromatic light	19
2.2.2 Accounting for Instrument Dispersion	24
2.2.3 Accounting for an Inhomogeneous Light Source	26
2.3 Polarisation interferometry	27
2.3.1 Stokes vectors	28
2.3.2 Mueller calculus	29
2.3.3 A Simple Polarisation Interferometer	31
2.3.4 Phase Delay due to a Uniaxial Crystal Plate	32
2.3.5 Imaging Polarisation Interferometry	36
2.3.6 Savart plate	41
2.3.7 Instrument function	42
2.4 Previous CI Investigations	44
2.5 CI hardware used in this work	47
2.6 Summary	48

3	Coherence Imaging with a Pixelated Phase-mask	49
3.1	Pixelated Phase-mask Interferometry	50
3.2	Single-delay Coherence Imaging	50
3.2.1	Interferogram Demodulation	55
3.3	Multi-delay Coherence Imaging	57
3.3.1	Multi-delay CI with Pixelated Phase-mask	61
3.3.2	Interferogram Demodulation	63
3.4	Non-ideal Quarter-wave plate Effects	69
3.5	Summary	74
4	Calibration of Flow Measurements using Gas-discharge Lamps	75
4.1	Background	76
4.1.1	Fitting to Wrapped Data	78
4.1.2	Dispersion in Barium Borate	80
4.2	Experimental Testing	85
4.2.1	Tuneable Laser Data	85
4.2.2	Gas-discharge Lamp Data	90
4.2.3	A Note on Group Delay	98
4.2.4	Discussion	98
4.3	Summary	99
5	Coherence Imaging Measurements of Electron Density	100
5.1	Modelling Balmer Series Lineshapes	101
5.1.1	Doppler Broadening	102
5.1.2	Stark Broadening	103
5.1.3	Zeeman Splitting	107
5.2	Measurement Principle	113
5.2.1	Basic Arguments: Choice of Delay and Balmer Line	113
5.2.2	Fitting to Contrast Profiles	117
5.2.3	Systematic Errors	121
5.3	Experimental Testing on Magnum-PSI	123
5.3.1	Magnum-PSI	123
5.3.2	Coherence Imaging Setup	126
5.3.3	Line-integrated Results	131
5.3.4	Non-Gaussian Doppler broadening and Doppler shifts	133
5.3.5	Abel Inversion	133
5.3.6	Inverted Results	137
5.4	Summary	141

6	Summary and Future Work	145
6.1	Pixelated Phase-mask CI	146
6.2	Calibration of Flow Measurements using Gas-discharge Lamps	147
6.3	Coherence Imaging Measurements of Electron Density	148
6.4	Future Work	150
A	Maths	152
A.1	Fourier transforms	152
A.2	Derivation of the Wiener-Khinchin theorem	153
B	Fourier Demodulation of Interference Fringes	155
B.1	Noise Propagation	158
B.1.1	Fourier-domain Filtering	161
C	Bayesian Parameter Estimation	166
C.1	Bayes' Theorem	166
C.2	Parameter Estimation	167
C.3	Markov chain Monte Carlo	168
	Bibliography	169

List of Tables

3.1	Key parameters of the FLIR Blackfly S camera	51
3.2	Pixelated carrier CI interferometer configurations used in this work	63
4.1	Sellmeier coefficients for BBO 1	82
4.2	Sellmeier coefficients for BBO 2	82
4.3	Laser CI calibration data fit results	90
4.4	Cd and Zn calibration line wavelengths	91
4.5	Lamp CI calibration data fit results: extrapolation test	96
4.6	Lamp CI calibration data fit results: interpolation test	97
5.1	Hydrogen Balmer series transition wavelengths	101
5.2	Stark broadening model coefficients	104
5.3	Balmer lineshape model summary	108
5.4	Stokes parameters for Zeeman-split line components	109

List of Figures

1.1	Toroidal and poloidal directions on a torus	3
1.2	Tokamak design	4
1.3	MAST fast camera images	5
1.4	Tokamak divertor design	7
1.5	Modelled MAST-U magnetic equilibria	8
1.6	MAST-U divertor SOLPS modelling	9
2.1	Two classic interferometer designs	15
2.2	Modelled spectra and corresponding interferograms	21
2.3	Wave normal surfaces in a uniaxial crystal	33
2.4	Uniaxial crystal geometry definitions	34
2.5	A waveplate imparts an optical path difference	35
2.6	In general, a uniaxial crystal imparts a lateral displacement	35
2.7	Layout of a simple polarisation interferometer	36
2.8	Projection of ray angles onto the sensor plane	37
2.9	Modelled interferograms due to a simple polarisation interferometer	39
2.10	Raw CI image from HL-2A tokamak	40
2.11	Savart plate composition	41
2.12	Instrument contrast vs. fringe period	44
2.13	Measured carbon impurity flows on MAST	46
2.14	Photo of the MAST CI instrument	48
3.1	Pixelated polariser layout (Sony IMX250MZR sensor)	51
3.2	Single delay pixelated carrier CI instrument layout	54
3.3	Single delay pixelated carrier CI calibration image	55
3.4	Demodulated single delay pixelated carrier CI calibration image	57
3.5	Multi-delay CI instrument layouts	59
3.6	Multi-delay pixelated carrier CI calibration image	64

3.7	Fourier transform of a multi-delay pixelated carrier CI calibration image	66
3.8	Fourier transform ‘synchronous product’ of a multi-delay pixelated carrier CI calibration image	68
3.9	Demodulation phase and contrast of a multi-delay pixelated carrier CI calibration image	70
3.10	Multi-delay PPM CI demodulation algorithm test	71
3.11	Modelling the effect of a non-ideal quarter-wave plate	73
4.1	The von Mises distribution	79
4.2	Simplified CI phase-fitting likelihood functions	80
4.3	Refractive indices and birefringence vs. wavelength for BBO	83
4.4	Variation in phase shift due to variation in Sellmeier coefficients	86
4.5	Experimental setup for CI laser calibration measurements	87
4.6	Measured and modelled phase shift	89
4.7	Cd and Zn calibration line wavelengths, relative intensities and band-pass filter transmission profiles	92
4.8	Measured and modelled phase shift for Cd and Zn lines, linear carrier	95
4.9	Lamp calibration phase residuals, 4.4 nm extrapolation	96
4.10	Lamp calibration phase residuals, 7 nm interpolation	97
5.1	Modelled lineshapes for the hydrogen Balmer series	106
5.2	Origin of Zeeman polarisation	109
5.3	Modelled Zeeman-split spectra	111
5.4	Zeeman contrast	113
5.5	Modelled lineshapes and contrast over a range of electron densities	115
5.6	Measurement sensitivity to changes in electron density	116
5.7	Modelled Doppler and Stark-broadened lineshape and contrast profiles	116
5.8	Measured JET divertor spectrum of contaminated D_δ (410.1 nm)	117
5.9	Contrast fitting tool demo	119
5.10	Inverted logistic curve	121
5.11	Error in inferred electron density due to Doppler broadening	122
5.12	Error in inferred electron density due to lineshape model	124
5.13	Magnum-PSI experiment schematic	125
5.14	Magnum-PSI diagnostic views	127
5.15	Magnum-PSI graphite target and probes	128
5.16	Installed CI instrument on Magnum-PSI	129
5.17	Bandpass filter transmission profiles for Magnum-PSI CI	130

5.18	Raw and demodulated single-delay CI results from Magnum-PSI	132
5.19	Raw and demodulated multi-delay CI results from Magnum-PSI	134
5.20	Predicted error in inferred electron density due to non-Gaussian Doppler broadening	135
5.21	Example CI geometry matrix \mathbf{A} for Magnum-PSI view	136
5.22	Line-integrated and Abel-inverted profiles	138
5.23	CI n_e results (Multi-delay H_γ)	140
5.24	CI n_e results (single-delay)	142
5.25	CI n_e results (Multi-delay H_δ)	143
B.1	Synthetic interferogram with a linear spatial carrier	157
B.2	Synthetic input and demodulated output images	159
B.3	Numerical test for analytical noise model	164

Declaration

The work in this thesis is based on research carried out at the Centre for Advanced Instrumentation, the Department of Physics, Durham University, UK; Culham Centre for Fusion Energy, UK; General Atomics, USA and the Dutch Institute For Fundamental Energy Research, NL. No part of this thesis has been submitted elsewhere for any other degree or qualification. The work within this thesis is that of the author, unless noted otherwise in the text.

Copyright © 2021 Joseph Steven Allcock.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged.”

Acknowledgements

First, I'd like to thank Prof. Ray Sharples for giving me the opportunity to take on this project and for his guidance, support and good humour throughout. I am particularly grateful for Ray's patience when progress was slower than I expected and also for his feedback on early drafts of this thesis.

Next, thanks to my supervisors at Culham for their support and contributions. Neil Conway, for sharing his diagnostic expertise and lightning-fast back-of-the-envelope calculations. James Harrison, of whom I once heard it said: "somebody that smart has no business being that friendly" — I don't think I can improve on that assessment. Scott Silburn, for sharing his deep knowledge of and enthusiasm for all things coherence imaging. I appreciate that Ray, Neil, James and Scott were all able to maintain our weekly meetings — a reliable source of insight and comic relief — throughout the project, despite their many other commitments.

This work has benefited from discussions on statistics with Chris Bowman and also from conversations on coherence imaging with John Howard, Clive Michael, Alex Thorman and Sam Gibson. Also, it has been a pleasure to work with Long Ting, Nie Lin and others at the HL-2A tokamak (work not presented here).

Thanks to Steve Allen at DIII-D for being so generous with his (already limited) time when helping me take the laser data in Chapter 4. Steve really went the extra mile in accommodating me at the lab, even driving me to work after I'd had a nasty bicycle fall! Thanks to Cameron Samuel too, for support and useful discussion. Much of the subsequent work that makes up Chapter 4 would not have been possible without Steve and Cameron first sharing their expertise and hardware in 2017.

Jordy Vernimmen, Serje Brons and Thomas Morgan at DIFFER were key to making the Magnum-PSI measurements in Chapter 5 happen. As were Jack Leland and Fabio Federici of Culham.

Widening the scope, thank you to my friends in the PhD community at Culham. Tom Farley, for effectively running a Python support hotline in our office. Daljeet Gahle for making sure that conversation was never boring. Sam Gibson, Jack Leland

and Charlie Vincent for laughter and warmth. Thanks also to David, Fred, Matt, Tom N., Gale and many others.

Wider still, thank you to the friends that have understood my absences and to my family for a lifetime of love and support. Finally, thanks to Grace for her love and kindness, and who couldn't possibly have known what she was getting herself into.

... suffice it to say that the actual implementation of these ideas in glass and metal is an extended exercise in second order effects.

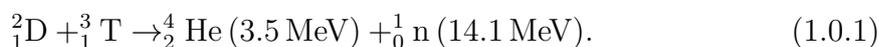
— from *Fourier Transform Spectroscopy* by Davis, Abrams and Brault

Chapter 1

Introduction

Research into controlled nuclear fusion is motivated by its potential as a clean, safe and economic power source[1, 2]. Fuel supplies are expected to be abundant, capable of supporting us for (at least) thousands of years at current levels of global energy consumption[3, 4]. Fusion power would also sidestep many of the disadvantages of fission power, with no high-level radioactive waste, no risk of runaway chain reactions and no risk of nuclear proliferation. However, despite seven decades of research, there are still significant physics and engineering problems to be solved.

In a nuclear fusion reaction, two or more light nuclei combine to form a heavier one. For a reaction to be a viable candidate for controlled fusion, it must meet three conditions. First, the reaction must be exothermic, so the product must be lighter than ^{56}Fe . Secondly, it must occur between two reactants since the reaction cross-section σ becomes vanishingly small for three-body reactions. Thirdly, the reactants must be of low atomic number Z since σ is sensitive to the Coulomb barrier which scales with Z^2 . The fusion of the hydrogen isotopes deuterium (D) and tritium (T) releases a neutron and an alpha particle, with 17.6 MeV of energy shared between them¹:



This is roughly a million times greater than the energy released in a chemical reaction. Since D-T fusion has the highest σ at the lowest reactant particle energies, it will likely be the fuel for at least the first generation of fusion power plants and is therefore the focus of most current research[5]. However, the collision of D and T nuclei almost always results in scattering rather than fusion[6]. Maximising the number of fusion reactions means maximising the average number of collisions each particle

¹The electronvolt (eV) is defined as $1 \text{ eV} = 1.602 \times 10^{-19} \text{ J}$, the kinetic energy gained by an electron after being accelerated through a 1 Volt potential. In plasma physics this is typically used as a unit of temperature too, the Boltzmann constant relating average particle kinetic energy to temperature such that $1 \text{ eV} \approx 11600 \text{ K}$.

makes before escaping the system, precluding a simple beam-into-target approach. The typical approach then is to confine a thermal fuel mixture at a high enough temperature, and for long enough, that a significant number of particles overcome the Coulomb barrier and fuse on collision. It turns out that the temperatures required are ~ 10 keV, ten times hotter than the centre of the Sun and easily hot enough to fully ionise the D-T fuel, creating a plasma. A further complication is that, while D is stable and occurs naturally, T is β -unstable ($t_{1/2} \approx 12$ years) and is therefore scarce. It is hoped that T will be ‘bred’ within the walls of future fusion reactors via neutron bombardment of lithium, but this technology is the subject of current research[7].

A plasma is said to have reached *ignition* when the energy released by fusion reactions can maintain the high temperatures necessary i.e. no external heating is required. Though not strictly necessary for a fusion reactor, ignition is a useful concept and helps formulate a key figure of merit for a reactor: the fusion triple product. It can be shown that for a plasma to ignite, the number density n , temperature T and energy confinement time τ_E must satisfy the ‘fusion triple product’[5]:

$$nT\tau_E > 5 \times 10^{21} \text{ m}^{-3} \text{ keV s.} \quad (1.0.2)$$

Here, τ_E is the characteristic timescale for energy leaving the plasma. Another metric used in fusion research, for non-ignited plasmas, is the fusion gain: the ratio of power produced in the plasma to heating power supplied by external sources $Q \equiv \frac{P_{\text{fus}}}{P_{\text{heat}}}$. Different schemes exist for plasma confinement (see e.g. inertial confinement[8]) but the subject of this work, and the most successful by both metrics just defined, is magnetic confinement fusion (MCF).

1.1 Magnetic Confinement Fusion and the Tokamak

The Lorentz force causes charged particles to move in circular ‘Larmor’ orbits around a magnetic (B-) field line, in a direction determined by the sign of the electric charge. In a uniform B-field, the velocities parallel to the field are unaffected and so a particle traces out a helical path along the field line. MCF uses this constraint on charged particle motion to confine the plasma in the direction perpendicular to the field lines. Electromagnets are arranged so as to produce a B-field geometry which minimises contact between plasma and solid surfaces, which would rapidly cool the plasma. Since the fusion neutrons are not subject to the Lorentz force, they fly out in all directions, depositing their energy in the surrounding walls. In a reactor, this heat will be used to drive a steam turbine and generate electricity. In MCF, toroidal

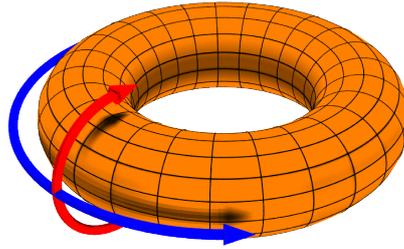


Figure 1.1: Toroidal (blue) and poloidal (red) directions on a torus. The distance from the centre of the hole to the centre of the tube is the major radius R_{maj} , while the radius of the tube is the minor radius R_{min} . Figure created by Dave Burke, licensed under CC BY 2.5.

B-field geometries provide particularly good confinement since the field lines circle back on themselves — there are no ends for the plasma to spill out of. See Figure 1.1 and caption for definitions of the toroidal and poloidal directions on a torus, and major and minor radii R_{maj} and R_{min} . Although a purely toroidal field geometry is ‘closed’, there are still mechanisms for confinement loss. The radial inhomogeneity of the B-field, decreasing with distance from the torus’ central axis, results in a particle drift directed radially outwards across the field lines[5]. A *poloidal* B-field component can be used to stop this leading to a loss of confinement by creating field lines that trace out a helix around the torus, averaging out the drift experienced by a particle to zero.

In a tokamak, shown in Figure 1.2, this poloidal B-field component is created via Ampere’s law by passing an electric current through the plasma in the toroidal direction, of order 1 MA. In today’s experiments, this current is at least partially induced by transformer action, where the primary coil is formed by a solenoid through the torus hole and the secondary coil is the plasma itself. Since plasma has a finite resistance, Ohmic heating of the plasma is a useful side-product². To maintain this induced plasma current, the current in the central solenoid must be continuously ramped, which cannot be done indefinitely. Since pulsed operation of a fusion reactor would be economically undesirable and would also thermally fatigue components, achieving non-inductive current drive in tokamaks is an important research goal. Important methods for this are neutral beam injection (NBI), in which high energy neutral D atoms are fired into the plasma, and electron and ion cyclotron resonant heating (ECRH and ICRH respectively), in which radio waves / microwaves that resonate with the harmonic frequencies of the particle Larmor orbits are fired into the plasma. Both NBI and ECRH/ICRH are also used for plasma heating[5].

Power escapes a tokamak plasma by several routes. Turbulence and collisions

²Although the resistivity of the plasma falls as the electron temperature rises, varying as $T_e^{-3/2}$ [3].

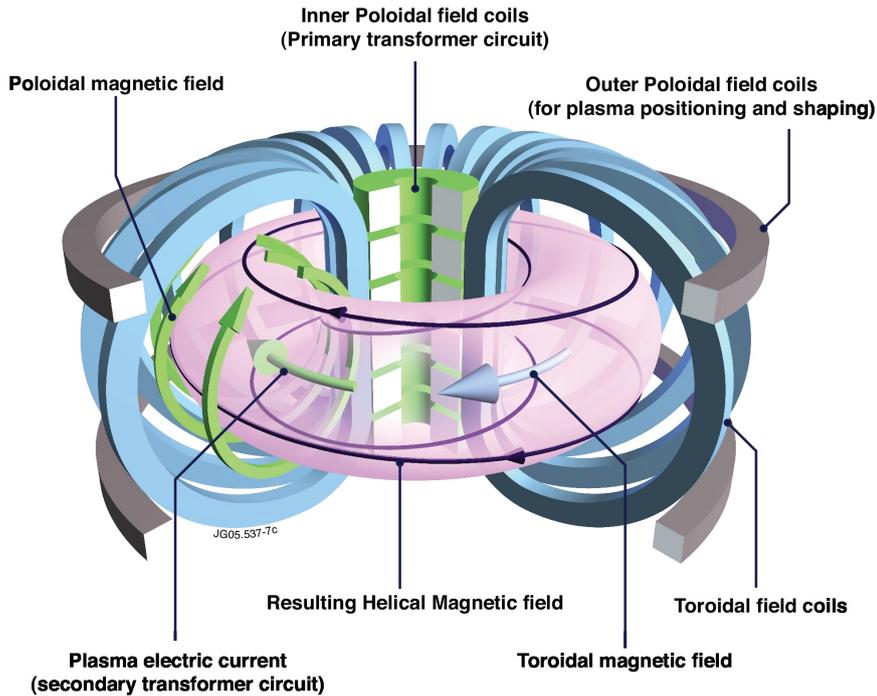


Figure 1.2: A schematic showing the toroidal (blue) and poloidal (green) magnetic field components of a tokamak. The resultant field lines (black) follow a closed helical path on the surface of a torus. Figure courtesy of the EFDA-JET figure database.

between particles ensure that confinement perpendicular to the B-field is not perfect. Particles will eventually be knocked from closed field lines to an outer region of open field lines called the scrape-off layer (SOL). They stream along the SOL until they collide with the surrounding walls. The resulting energy loss is undesirable but the particle loss is necessary to prevent build up of the He fusion reaction product. Radiation is another important mechanism for power loss. An important contribution is bound-bound electron transitions between energy states in atoms, ions and molecules, which emit coherent line radiation. Another is deceleration of charged particles on collision, which emits continuum radiation³. Particle bombardment of the plasma facing components (PFCs) also sputters foreign atoms into the plasma. Any ions or neutrals in the plasma other than the fusion fuel and products are classed as impurities. The heaviest impurities retain some of their bound electrons, even in a 10 keV plasma core, and so can lose a large amount of energy via line radiation. It follows that the requirements for core plasma purity in a fusion reactor are very high, e.g. a tungsten concentration in the core as low as 10^{-3} is enough to make reaching ignition impossible[9].

The tokamak is a Russian design that was demonstrated in the late 1960s to have achieved $T_e \approx 1$ keV, an order of magnitude higher than its contemporaries

³Known as Bremsstrahlung (German for ‘braking radiation’).

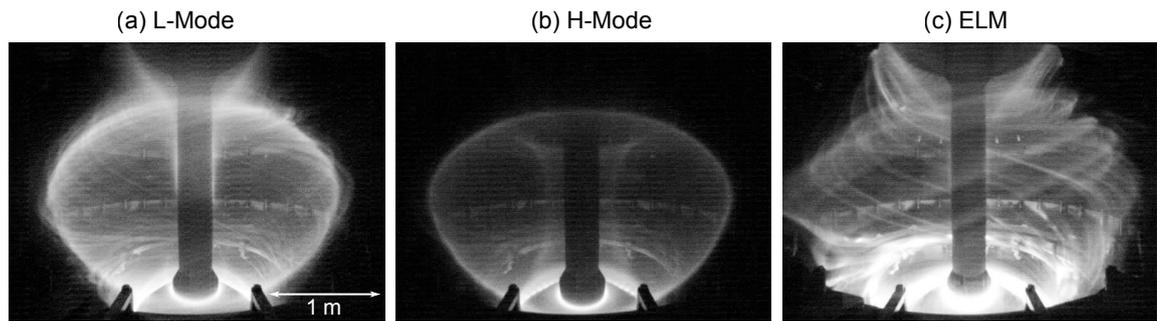


Figure 1.3: Camera images of visible light emitted from a MAST plasma operating in (a) L-mode, (b) H-mode and (c) during an ELM. Figure courtesy of Scott Silburn.

in the west. In the decades that followed, most MCF research switched its focus to tokamaks and substantial progress was made in both theoretical understanding and experimental performance. Study of plasma Magnetohydrodynamics (MHD) was used successfully to predict instabilities, which can lead to loss of confinement and therefore set limits on the available operating space: e.g. the kink limit on plasma current density and the Troyon limit on plasma pressure[3]. Another important development was the discovery in the 1980s of a high confinement operation mode (‘H-mode’): with enough heating power, turbulence in the plasma edge is suppressed (visible in Figure 1.3), leading to a doubling of τ_E [10]. H-mode is not yet fully understood and is prone to periodic bursts of energy directed outwards into the PFCs (called edge-localised modes (ELMs)) which will need to be either suppressed or mitigated to avoid damage in a reactor. The current record for tokamak performance is held by the Joint European Torus (JET), in Oxfordshire, UK. JET is the largest of the current generation of tokamaks ($R_{\text{maj}} = 3 \text{ m}$, $R_{\text{min}} = 1 \text{ m}$) and, unlike most experiments, it is licensed to handle radioactive T and so can produce D-T plasmas (most experiments use only H or D). In 1997, JET produced a record-breaking $P_{\text{fus}} = 16.1 \text{ MW}$ with $P_{\text{heat}} = 25.4 \text{ MW}$ ($Q = 0.62$), sustained for 1 s[11]. A significant next step in MCF research is the enormous multi-national project ITER, an even larger tokamak ($R_{\text{maj}} = 6.2 \text{ m}$, $R_{\text{min}} = 2 \text{ m}$) currently under construction in Cadarache, France. When D-T experiments begin in 2035, ITER aims to produce $P_{\text{fus}} = 500 \text{ MW}$ with $P_{\text{heat}} = 50 \text{ MW}$ ($Q = 10$), sustained for 400 s — demonstrating the physical feasibility of a tokamak fusion reactor[12]. However, ITER is not designed to produce electrical power. It is hoped that the demonstration of electrical power production will come mid-century with a device called DEMO — whose exact design remains an open question.

For JET and ITER (‘conventional’ tokamaks), the torus aspect ratio $A \equiv \frac{R_{\text{maj}}}{R_{\text{min}}} \sim 3$. It was theorised in the 1980s and then demonstrated in the 1990s that reducing A by a factor 2 offers distinct operational advantages, reducing the two MHD limits mentioned

in the previous section, permitting a higher plasma pressure[13]. The plasma volume of these ‘spherical’ tokamaks (STs) better resembles a cored apple than a doughnut, as can be seen in Figure 1.3 which shows the mega-ampere spherical tokamak (MAST) in Oxfordshire, UK. The ST geometry also permits smaller distance between toroidal coils and the plasma, reducing the power needed to achieve an equivalent B-field in the plasma.

Although the tokamak is currently the most popular and successful MCF concept, there are alternatives being researched. The stellarator is a promising toroidal design whose poloidal field is produced by a complicated coil arrangement instead of current in the plasma, allowing it to operate in a steady state. Many excellent technical[3, 5] and non-technical/historical[14, 15] accounts of MCF research can be found in the literature.

1.2 The Heat Exhaust Problem

In modern tokamaks, the divertor controls the interaction between the plasma and material surfaces. Its job is to optimise confinement by screening the plasma core from impurities released from the surfaces, and to prevent impurities building up by continuously pumping them away. This is done by introducing a magnetic X-point at the plasma edge, diverting the field lines away from the core and into two toroidally extended target plates (see Figure 1.4). During operation, these targets can receive extremely high heat fluxes, which must not exceed a limit of $\sim 10 \text{ MW}\cdot\text{m}^{-2}$ set by material erosion[16]. Staying below this power load limit in the divertor whilst producing reactor-relevant levels of steady-state fusion power ($\sim \text{GW}$) in the core is an unsolved problem in fusion research known as the *heat exhaust problem*[17].

It is likely that at least part of the solution to the heat exhaust problem is to ‘detach’ the divertor plasma from the target plates using a blanket of low temperature neutral particles[18]. Power is then lost in this region via ionisation of and volumetric emission from hydrogen and impurities, reducing the heat flux to the targets. The low target plasma temperature ($< 5 \text{ eV}$) and particle flux that characterises detachment also reduces material erosion from physical and chemical sputtering. For ITER, the upper limit on target power load means that 60–70% of the power entering the divertor must be exhausted as radiation in this way[19, 20]. Current tokamaks are investigating the physics of detachment so as to optimise and control the process[21–24]. Tokamaks can reach detachment by increasing the density at the plasma edge, by intense impurity seeding or by increased divertor baffling. However, there is a trade-off since if the low temperature region moves too close to the X-point, it

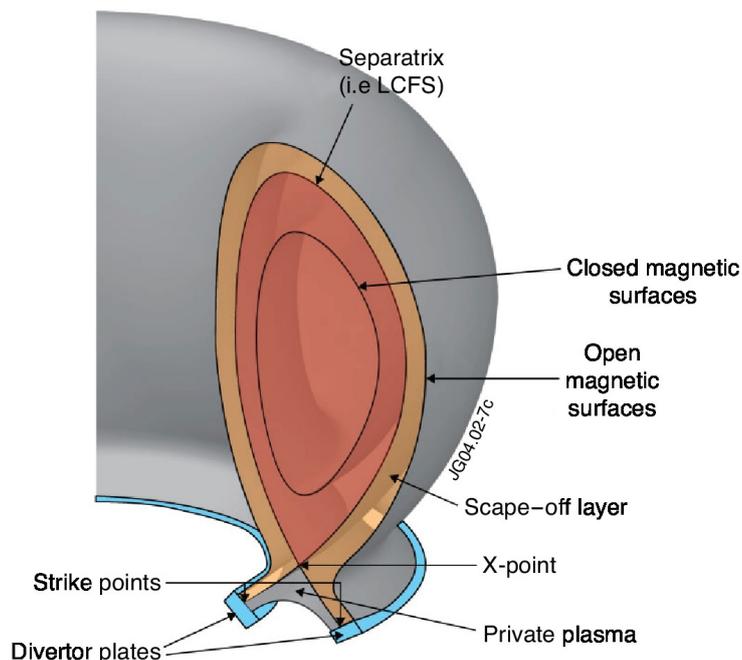


Figure 1.4: A poloidal cross-section showing the magnetic field line geometry of a tokamak with a divertor at its base. Figure courtesy of the EFDA-JET figure database.

can begin to degrade core confinement, and can lead to complete radiative collapse of the plasma. Looking further ahead, the EU roadmap plans for a DEMO reactor producing ~ 500 MW of *electrical* power[17], making the power handling requirements in the divertor even more demanding[25], and it is likely that a controlled seeding of impurities into the main chamber will be required to radiate power from the upstream scrape-off layer (SOL) and further reduce the divertor heat flux[26].

Another way to reduce the heat exhaust is to optimise the magnetic geometry of the divertor. For example, the angle between ITER’s outer divertor target plate and the incident magnetic field lines is minimised ($\sim 3.5^\circ$) so as to spread the heat load over a larger ‘wetted area’[27]. This can give up to a factor ~ 50 reduction in the target heat flux[28]. Further geometry optimisation is being investigated in current experiments to guide DEMO design[17]. This includes the the ‘Snowflake’ divertor, which introduces a second order magnetic null to increase the number of divertor legs, spreading the power load over a larger area. Another advanced design is the ‘Super-X’, which increases magnetic flux expansion and divertor leg length. MAST Upgrade (MAST-U) will experimentally test the predicted benefits of these designs at comparatively low powers when it begins operation in 2020[29]. Figure 1.5 shows the predicted magnetic equilibria for MAST-U in both conventional and Super-X divertor configurations. Modelling suggests that MAST-U’s Super-X configuration will make detachment accessible at lower upstream densities[30]. The introduction

of the divertor baffle prevents neutrals that are produced via recombination at the divertor PFCs from reaching back into the plasma core. This increased neutral pressure allows for more efficient pumping and density control and easier access to detachment. Figure 1.6 shows modelled plasma density and temperature for

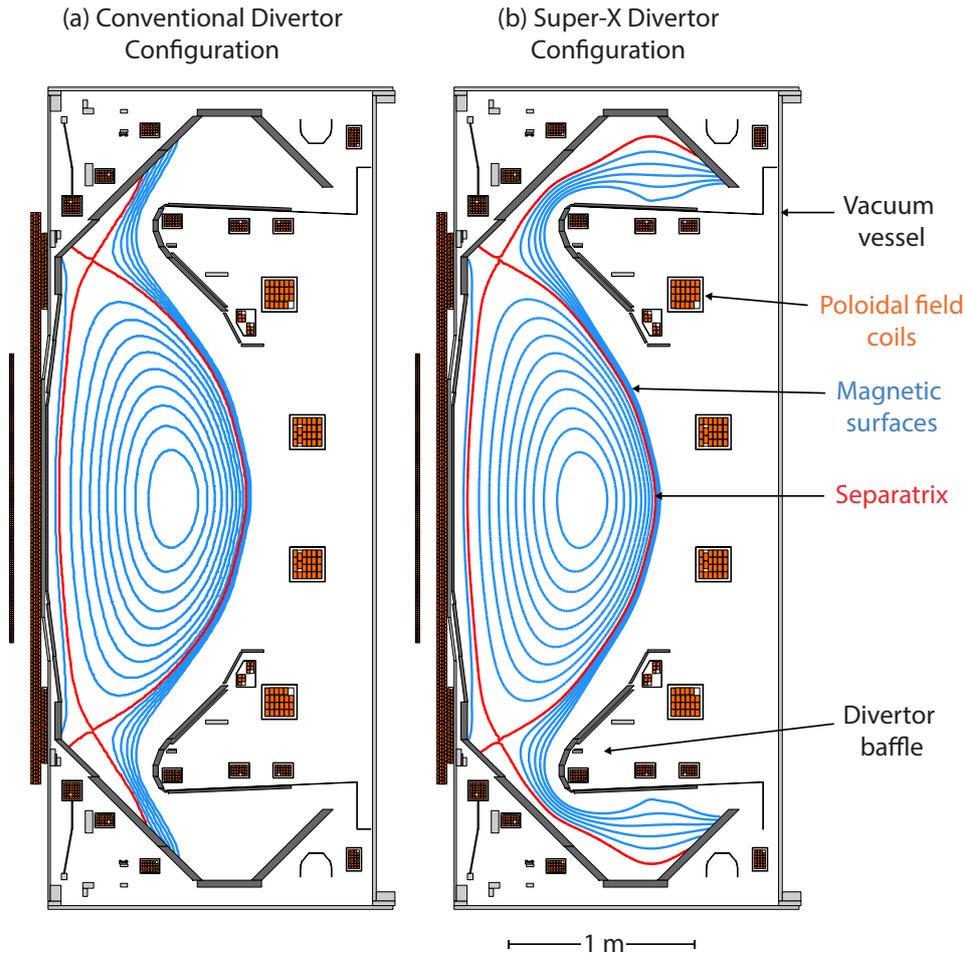


Figure 1.5: Expected magnetic equilibria (poloidal slice) for the MAST-U spherical tokamak, for both (a) conventional and (b) Super-X divertor configurations. Design drawings of MAST-U and information on expected equilibria courtesy of the MAST-U team, as presented in e.g. [31].

the MAST-U divertor in both conventional and Super-X configurations. The scrape-off layer plasma simulation (SOLPS) code used to generate these profiles is widely used in the community to model MCF plasma exhaust physics, using a combination of a fluid model for the plasma and a kinetic model of the neutral particles.

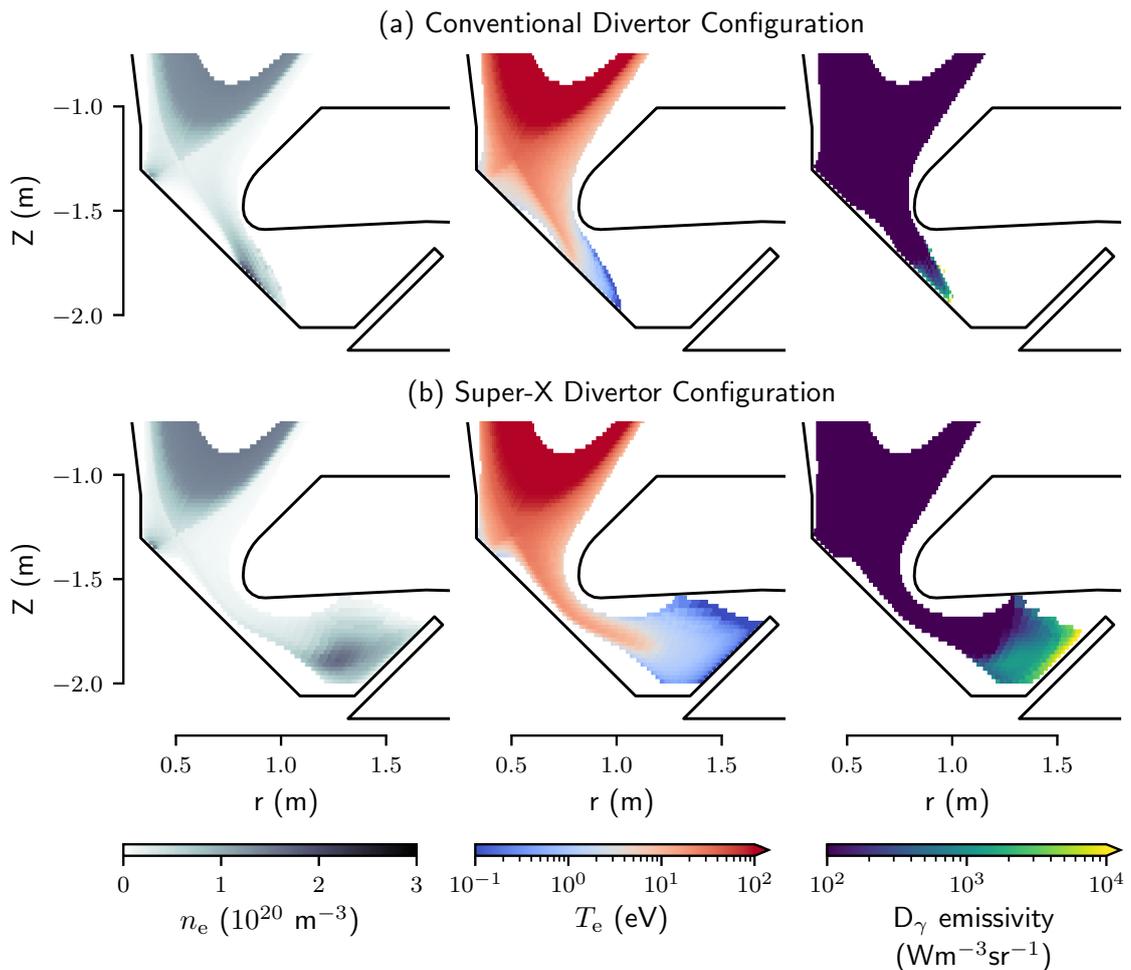


Figure 1.6: SOLPS modelling of the MAST-U divertor in both (a) conventional and (b) Super-X configurations. Shown is the modelled electron density n_e , electron temperature T_e and emissivity of the D_γ emission line (433.9 nm). SOLPS modelling carried out by David Moulton, as presented in e.g. [32].

1.3 Diagnosing the MAST-U Divertor

Every MCF experiment is equipped with a suite of diagnostic instruments for feedback-control of the plasma shape, size and position[33]. High quality measurements are also a prerequisite for any experimental plasma physics study. Tokamak diagnostics are often categorised as being either active or passive, depending on whether they perturb the plasma on measurement. In this thesis, our focus is on developing passive spectroscopic diagnostic techniques for the SOL and divertor of MAST-U. We are motivated, in large part, by the requirements of MAST-U — how can we best investigate the physics and assess the performance of the new divertor configurations? The purpose of the MAST-U divertor diagnostic suite is best explained by the MAST-U team[34]:

“A historical gap in many tokamaks is good information on the plasma parameters at various locations along the SOL and in the divertor; this will be a focus for MAST Upgrade. While the confrontation of the experimental data with SOLPS and other fluid modelling codes will be central to the understanding of the divertor physics, it is hoped to extract the plasma solution directly from an integrated analysis of most, if not all, of the divertor measurements — each measurement corresponds to a location (or chordal integral), local plasma parameters (e.g. n_e and T_e) and thus, with proper implementation can constrain the plasma solution across the divertor region without specifying the physics. The plasma solution derived can then be used to calculate exactly which mechanisms are dominant and where there are additional mechanisms not included (e.g. turbulent-driven cross-field transport). ... The MAST Upgrade diagnostics have been designed for as high space and time resolution as currently feasible.”

We will now briefly cover the relevant MAST-U divertor diagnostic instruments, and the specific diagnostics presented in this thesis.

One of the simplest ways to measure n_e and T_e is by Langmuir probe. By fitting a model to the probe’s current-voltage characteristic curve, n_e and T_e can be inferred. However, to avoid damage, probe measurements are generally limited to the plasma edge and interpretation in low-temperature and/or magnetised plasma conditions is difficult[33]. Despite these limitations, probes are a key part of any edge/divertor diagnostic suite and MAST-U will have over 800[35].

For measuring n_e and T_e further inside the plasma, Thomson scattering (TS) is perhaps the workhorse technique, providing localised measurements over a wide parameter range and with high accuracy[33]. An intense laser beam is directed into the plasma and a spectrometer measures the scattered line radiation. The absolute intensity of the measured line then allows for inference of the local electron density n_e while the shape of the line provides information on the local electron temperature T_e via Doppler broadening. Core and edge TS systems are common on tokamaks (see e.g. [36–39]). A divertor TS system is installed on the DIII-D tokamak[40, 41] and a similar system is planned for MAST-U with a spatial resolution of ~ 1 cm[42]. The main disadvantage of TS is that it is limited to measurements along the laser beam’s 1-D path.

Passive spectroscopy is a valuable tool for diagnosing SOL/divertor plasma conditions. The observed intensity of a spectral line emitted by the plasma is proportional to the density of the corresponding excited atomic state. Relating this to the ground state density of the emitting species requires some assumptions about the thermody-

dynamic equilibrium or else a full collisional-radiative model describing the contributions of all relevant atomic processes[43, 44]. In addition, the spectral lineshape can provide relatively direct measurements of emitter temperature (via Doppler broadening) and/or n_e (via Stark broadening), that are less dependent on assumptions about the plasma equilibrium. Line broadening measurements play an important supplemental role in spectroscopy analyses and are also used as independent diagnostic techniques. Inferring n_e from Stark-broadened hydrogen Balmer emission in the visible light range is a standard diagnostic technique for grating spectrometers[23, 24, 45–47]. As these instruments are fundamentally slit-coupled, they are typically limited to measurements integrated along a 1-D fan of sight lines through the plasma.

It is clear from Figure 1.6 that the divertor plasma has a complex 2-D structure in the poloidal plane with steep gradients in n_e and T_e . To complicate things further, MAST-U’s divertor tiles do not have perfect toroidal symmetry, possibly introducing 3-D effects which must also be accounted for when making inferences about local plasma profiles[48]. In recent years, this complexity has motivated the use of imaging diagnostics to provide 2-D coverage at a high spatial resolution. While the resulting measurements are line-integrated, large number of sight lines make it possible to tomographically invert the raw data and infer a local plasma emission profile. This increase in the quantity of information generally comes at the cost of quality though, as the spectral resolution is inevitably lower than a slit-coupled spectrograph. At one extreme is unfiltered imaging, used to capture the images shown in Figure 1.3. This measurement integrates over the spectral sensitivity of the camera and lens, typically 400–1000 nm for the CCD/CMOS detectors used. The high throughput maximises signal and allows measurements up to \sim MHz sampling rates, fast enough to study the dynamics of turbulence[49]. Filtered imaging uses a narrowband interference filter to isolate a targeted emission line whose intensity distribution, particularly for impurities, can help constrain T_e . MAST-U will also have a multi-wavelength imaging (MWI) diagnostic. Based on a design developed at DIFFER[50], this is a 10-channel filtered imaging system for which each channel has the same view of the plasma but targets a different spectral emission line.

In practice, the ‘integrated analysis’ of MAST-U divertor measurements will use a Bayesian approach to combine information from the different diagnostic instruments in a self-consistent way⁴. A framework has been developed for this purpose and tested using SOLPS simulations of MAST-U similar to those shown in Figure 1.6, incorporating ‘synthetic diagnostic’ measurements made using the TS, LP and MWI capabilities just described. These preliminary tests recover the 2-D SOLPS profiles of

⁴We use Bayesian methods for parameter estimation problems in Chapters 4 and 5. See Appendix C for a brief introduction.

n_e and T_e with $\sim 10\%$ accuracy, but questions remain on how robust the technique will be in practice against molecular emission, toroidal asymmetry and fine spatial structure.

The subject of this thesis is further developments of coherence imaging (CI), a passive, narrowband spectral imaging technique for diagnosing SOL and divertor plasmas. CI uses an imaging interferometric method to produce a fringe pattern which encodes the low-order moments of an isolated spectral emission feature across a wide field of view. This fringe pattern is then recorded by a fast camera and analysed using signal processing techniques to infer the plasma parameters. From Doppler shifts of the emission feature, the (line-integrated) flow of the emitting ion/atomic species can be inferred. This is an established technique for measuring impurity ion velocity in the SOL[51, 52]. From the width and shape of the emission feature, different physics parameters can also be inferred, depending on the plasma regime. For example, emitter temperature can be inferred via Doppler broadening or, when observing line emission from hydrogen in dense divertor plasmas, electron density can be inferred via Stark broadening. A detailed account of CI and its uses is given in Chapter 2. It is planned that CI will contribute to the integrated data analysis of the MAST-U divertor with 2-D measurements of electron density or impurity ion velocity and temperature, depending on the instrument setup.

1.4 Thesis Outline

The original goal of this thesis was to make CI measurements of impurity ion flow and electron density in the MAST-U divertor. However, at the time of writing, MAST-U is still preparing to begin operations, so the project has instead focussed on further developing the CI technique, with experimental testing on a linear plasma device. The rest of this thesis is organised as follows:

- **Chapter 2** covers CI theory: two-beam interferometry, Fourier transform spectroscopy and polarisation interferometry. Following this is an overview of previous CI investigations and a description of the basic CI hardware used in this thesis.
- **Chapter 3** extends the CI technique to incorporate a pixelated polarising filter, resulting in a fringe pattern that is pixelated instead of sinusoidal. We consider the advantages of the technique and also propose a novel multi-delay CI design, which simultaneously encodes the interferogram at four different fixed delays, allowing the instrument to measure more complex spectral lineshapes.

-
- **Chapter 4** introduces a new technique for absolute calibration of CI flow velocity measurements. While existing techniques use expensive tuneable laser and wavemeter equipment, this new technique requires only off-the-shelf gas discharge lamps and interference bandpass filters. This maintains comparable measurement accuracy while significantly reducing hardware costs.
 - **Chapter 5** applies CI to the problem of n_e measurement, building on the limited existing work in a number of ways. Modelling predicts that the simplified analytical Stark-broadening model used in previous work leads to a significant underestimate of n_e when compared to state-of-the-art numerical models. Modelling also predicts that, by using a multi-delay CI instrument, it should be possible to eliminate systematic error caused by Doppler-broadening, extending the lower n_e limit of the measurement's dynamic range. This is then tested experimentally, in tandem with the first tests of the pixelated carrier CI technique introduced in Chapter 3. The linear plasma experiment Magnum-PSI stands in for MAST-U as the test-bed, and we present results over a range of plasma conditions relevant to both MAST-U and future devices like ITER.
 - **Chapter 6** summarises the results and makes recommendations for future work.

Chapter 2

Coherence Imaging

When two beams of light interfere, they produce a fringe pattern that tells us about the coherence of the light — how statistically correlated the electric field is at different points in time and in space. This is a useful thing to know because it is closely related to the light’s frequency spectrum, and therefore also to the physics of the light’s source. coherence imaging (CI) is an application of these ideas to the measurement of light emitted by plasma in nuclear fusion experiments. The clever part is to concentrate on a narrow range of the frequency spectrum, in which any changes can only be caused by a small number of well-understood mechanisms in the plasma. This way, information about the physics can be collected quickly and in 2-D across a wide camera view. Although CI is a relatively new fusion plasma diagnostic technique, the measurement principles date back to the end of the nineteenth century. This chapter begins by introducing the theory of Fourier transform spectroscopy and polarisation interferometry, which underpin CI, and then covers some of the important CI work from the literature.

2.1 Coherence and Two-beam Interferometry

Maxwell’s equations tell us that light is the transverse wave oscillation of electric and magnetic fields travelling through space. An important solution to Maxwell’s equations is a plane wave evolving in space \mathbf{r} and time t , for which the electric field can be written¹:

$$E(\mathbf{r}, t) = |\mathbf{E}| \exp(i[\mathbf{k} \cdot \mathbf{r} - 2\pi\nu t]). \quad (2.1.1)$$

Here, \mathbf{k} is the wave vector and ν is the wave frequency. This scalar treatment ignores polarisation, but is sufficient for introducing coherence effects. The mechanisms that

¹The electric field exerts a much larger force in most light-matter interactions, so is typically the only consideration in optics.

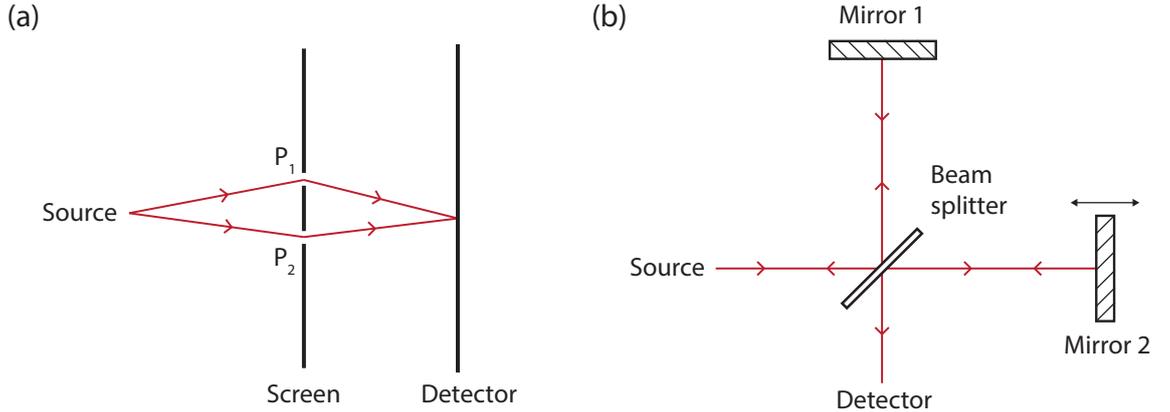


Figure 2.1: (a) Division of wavefront interferometer (Young's double-slit experiment). (b) Division of amplitude interferometer (Michelson's design).

produce light are inherently random and so it follows that all light is the sum of constituent waves whose frequencies and amplitudes are, to some degree, random. Coherence theory provides the required statistical description. In this work, we will assume that the electric field as a function of both space and time can always be considered a *wide-sense stationary* (WSS) random process. This means that the mean and variance of the field do not change with time[53].

Coherence is typically introduced using the example of a two-beam interferometer[54, 55]. Consider an extended light source sitting some distance from an opaque screen. The screen has two pinholes at points P_1 and P_2 and a detector lies behind the screen. This arrangement, shown in Figure 2.1(a), will produce a set of interference fringes at the detector. The electric field at point \mathbf{r} on the detector plane at time t is the sum of the field contributions from the first and second pinhole:

$$E(\mathbf{r}, t) = k_1 E(P_1, t - t_1) + k_2 E(P_2, t - t_2), \quad (2.1.2)$$

where t_1 and t_2 are the times taken for light to traverse the paths between points P_1 and P_2 and point \mathbf{r} respectively. In general $t_1 \neq t_2$ resulting in an optical path difference between the beams. The complex factors k_1 and k_2 account for any absorption and the geometry of the arrangement. We are interested in the irradiance at the detector plane rather than the field amplitude, since this is what our detectors measure². At \mathbf{r} , the irradiance is $I(\mathbf{r}) \equiv \langle E^*(\mathbf{r}, t) E(\mathbf{r}, t) \rangle$, where $\langle \dots \rangle$ is the time-average and $*$ is the complex conjugate. It follows from the field being WSS that the measured irradiance, being a time-average, is independent of the precise acquisition start time.

²For CCD/CMOS sensors, the measured signal is proportional to the number of photons hitting the active sensing area during the exposure time. In this work we measure quasi-monochromatic light, so the number of photons measured is roughly proportional to the irradiance at the pixel.

Using this, we write the irradiance (or interferogram) at \mathbf{r} as³

$$I(\mathbf{r}, \tau) = |k_1|^2 I(P_1) + |k_2|^2 I(P_2) + 2|k_1||k_2|\Re\{\Gamma_{12}(\tau)\}. \quad (2.1.3)$$

The first two terms here are the irradiances that would be observed at \mathbf{r} if each respective pinhole were open in isolation. In the third term, which contains the interference pattern, we have defined the mutual coherence function:

$$\Gamma_{12}(\tau) \equiv \langle E^*(P_1, t)E(P_2, t + \tau) \rangle. \quad (2.1.4)$$

This is the cross-correlation function of the electric field at two points in space P_1 and P_2 at a time delay $\tau \equiv t_2 - t_1$. That $\Gamma_{12}(\tau)$ is a function of τ and not the absolute time is a direct result of the assumed WSS nature of the field. $\Gamma_{12}(\tau)$ has the same units as irradiance.

With the two pinholes separated ($P_1 \neq P_2$), the interfering beams are said to have been created via ‘division of wave-front’ of the incident beam. When the two pinholes coincide ($P_1 = P_2$), it is ‘division of amplitude’. This arrangement requires a beam splitter and beam combiner to separate the beams and introduce time delay τ between them. An example of a division of amplitude interferometer is Michelson’s design, shown in Figure 2.1(b). We will consider these beam splitters to be lossless to and transmit equal power through each port. For this division of amplitude case, the irradiance at the detector plane is:

$$I(\mathbf{r}, \tau) = \frac{1}{2}|k_1|^2 \left(I(P_1) + \Re\{\Gamma_{11}(\tau)\} \right), \quad (2.1.5)$$

calculated as a special case of Equation 2.1.3. Now, the fringes encode the auto-correlation function of the field at point P_1 for time delay τ

$$\Gamma_{11}(\tau) = \langle E^*(P_1, t)E(P_1, t + \tau) \rangle, \quad (2.1.6)$$

which we call the temporal coherence⁴. This is the mutual coherence function but without any information about the spatial coherence of the scene. It can be seen from Equation 2.1.6 that, for $\tau = 0$, the temporal coherence at point P_1 is simply the irradiance at that point: $\Gamma_{11}(0) = I(P_1)$ and similarly $\Gamma_{22}(0) = I(P_2)$. It useful to introduce a normalised version of $\Gamma_{12}(\tau)$ called the complex degree of coherence:

$$\gamma_{12}(\tau) \equiv \frac{\Gamma_{12}(\tau)}{\sqrt{\Gamma_{11}(0)\Gamma_{22}(0)}} \quad (2.1.7)$$

This dimensionless quantity expresses the coherence of the source, independent of

³Using $x + x^* = 2\Re\{x\}$.

⁴Temporal coherence is sometimes called self-coherence.

its total irradiance. When $|\gamma_{12}(\tau)| = 0$, the interference term drops to zero and the source is said to be incoherent. Conversely, $|\gamma_{12}(\tau)| = 1$ represents perfectly coherent light and $0 < |\gamma_{12}(\tau)| < 1$ partially coherent light. It follows that the complex degree of *temporal* coherence reduces to

$$\gamma_{11}(\tau) = \frac{\Gamma_{11}(\tau)}{\Gamma_{11}(0)}. \quad (2.1.8)$$

This allows us to rewrite Equation 2.1.5 for the irradiance at the output of the division of amplitude interferometer as

$$I(\mathbf{r}, \tau) = \frac{I_0(\mathbf{r})}{2} \left(1 + \Re\{\gamma_{11}(\tau)\} \right), \quad (2.1.9)$$

Where we have also introduced the shorthand $I_0(\mathbf{r}) \equiv |k_1|^2 I(P_1)$ for the irradiance at \mathbf{r} observed when $\tau = 0$. This is related to the spectral irradiance $J(\nu)$ at \mathbf{r} by $I_0(\mathbf{r}) = \int_{-\infty}^{\infty} J(\nu) d\nu$.

Temporal coherence is an important observable quantity because the autocorrelation function of a signal can be Fourier transformed to obtain that signal's power spectrum — provided the signal is either deterministic or that the underlying random process is wide-sense stationary (random process) (WSS). This mathematical relationship is called the Wiener-Khinchin theorem (see Appendix A for a derivation). It follows that $\Gamma_{11}(\tau)$ and $J(\nu)$ are a Fourier transform pair:

$$J(\nu) = \int_{-\infty}^{\infty} \Gamma_{11}(\tau) \exp(-2\pi i\nu\tau) d\tau. \quad (2.1.10)$$

Here, the conjugate variable of the light's frequency ν is the interferometer delay τ . It is convenient now to define the area-normalised spectral irradiance $g(\nu) \equiv J(\nu) / \int_{-\infty}^{\infty} J(\nu) d\nu$, which is similarly related to $\gamma_{11}(\tau)$ by Fourier transform:

$$g(\nu) = \int_{-\infty}^{\infty} \gamma_{11}(\tau) \exp(-2\pi i\nu\tau) d\tau. \quad (2.1.11)$$

Since $g(\nu)$ and $\gamma_{11}(\tau)$ are a Fourier transform pair, the characteristic widths of the two distributions, $\Delta\nu$ and $\Delta\tau$ respectively, follow an uncertainty relationship:

$$\Delta\nu\Delta\tau \gtrsim 1. \quad (2.1.12)$$

Interference fringes are only observable when $\tau \sim \Delta\tau$, called the 'coherence time'[55]. Since the spectrum is a real function, and the Fourier transform of a real function has Hermitian symmetry, we can write $\gamma_{11}(\tau) = \gamma_{11}^*(-\tau)$. This means that the observed real part of $\gamma_{11}(\tau)$ is an even function. It follows that $g(\nu)$ can be recovered from an

observation of the real part of $\gamma_{11}(\tau)$ by cosine transform:

$$g(\nu) = 2 \int_0^{\infty} \Re\{\gamma_{11}(\tau)\} \cos(2\pi\nu\tau) d\tau. \quad (2.1.13)$$

So the spectrum can be recovered from a two-beam division of amplitude interferogram (Equation 2.1.9) that is recorded as a function of interferometer path delay τ . This technique is called Fourier transform spectroscopy (FTS). Since FTS is the focus of this work, from this point on we will use ‘coherence’ as a shorthand for the complex degree of temporal coherence, which we will denote $\gamma(\tau)$ (dropping the subscripts). We will also use ‘interferometer’ and ‘interferogram’ as shorthand for ‘two-beam interferometer of the division of amplitude type’ and ‘the interference pattern observed by such an instrument’ (Equation 2.1.9), respectively.

2.2 Fourier Transform Spectroscopy

According to Equation 2.1.13, we measure the interferogram $I(\boldsymbol{r})$ as a function of delay τ , isolate the coherence term then take the cosine transform to recover the spectrum $g(\nu)$. In the example of Michelson’s design in Figure 2.1(b), τ is typically scanned by moving one of the mirrors in the direction parallel to the beam, as shown. In reality, the maximum delay τ_{\max} at which $I(\tau)$ is observed must be finite. If $\gamma(\tau)$ is sampled at equal intervals $\delta\tau$ from 0 to τ_{\max} , then it can be shown that the frequency resolution is $\delta\nu \approx \tau_{\max}^{-1}$ [55, 56]. If the spectrum extends over a range of frequencies $\Delta\nu$, then the number of spectral samples is $\Delta\nu/\delta\nu$. If the spectral range extends from $\nu \approx 0$ then the delay sampling interval required to Nyquist sample the spectrum is $\delta\tau = 1/(2\Delta\nu)$ [55]. FTS instrument designs can be categorised as either *temporal multiplexing* or *spatial multiplexing* designs according to whether the interferogram is sampled in time or space. If the spectrum is encoded in the signal’s time-history then the full spatial resolution of the detector is preserved, at the cost of a reduced time resolution. Instruments which scan τ in space are described as ‘snapshot’, capturing all the spectral information in a single image, but necessarily sacrificing some spatial resolution in the process.

Fourier transform spectroscopy⁵ (FTS) has its roots in the experiments of Fizeau and then later Michelson and others in the late 19th century [54, 57, 58]. This work revealed for the first time the presence of atomic fine structure and hyperfine structure in the line emission from flame tests of different materials. Important observations were also made of pressure broadening, Doppler broadening and Zeeman splitting [57]. Although the accuracy of these measurements was unprecedented (e.g. the observed

⁵FTS is sometimes called ‘interference spectroscopy’ [55].

sodium D-line doublet has a wavelength separation of 0.6 nm), computing the cosine transform in Equation 2.1.11 was difficult, and so dispersive spectroscopy became the standard technique instead. Beginning in the mid 20th century however, FTS has grown in popularity again, due in part to advances in digital computing and the availability of intense coherent laser sources. It is also due to the discovery of two fundamental advantages of FTS over grating spectroscopy[55]. Firstly, the Jacquinot, or throughput advantage: since FTS instruments do not require a slit, a higher throughput can be achieved compared to dispersive (prism or grating) spectrometers of equivalent resolving power⁶. Secondly, the Fellgett or multiplexing advantage: since FTS instruments measure all frequencies simultaneously, a higher signal-to-noise ratio (SNR) is possible compared to dispersive spectrometers of equivalent resolving power — provided the measurement is limited by detector noise and not shot noise. FTS is now a standard technique for spectroscopy, particularly in the infrared where the ability to measure large spectral range at high resolution makes it more practical than dispersive spectrometers[59]. It also allows for compact, low-cost spectral imaging designs, particularly when observing quasi-monochromatic light. This is the subject of this work.

2.2.1 Quasi-monochromatic light

In this section, we will see that the interferogram assumes a simple form when the light being measured is quasi-monochromatic, and that this reduces the delay sampling requirements for FTS. For spectrum $g(\nu)$, the weighted-mean frequency is

$$\nu_c = \int_0^\infty \nu g(\nu) d\nu. \quad (2.2.1)$$

Here we make explicit the presence of a Doppler shift by expressing the mean frequency as $\nu_c = \nu_0 + \Delta\nu_D$, the sum of a fixed stationary term ν_0 and a much smaller shift term $\Delta\nu_D$ which varies with emitter velocity \mathbf{v} :

$$\frac{\Delta\nu_D}{\nu_0} = \frac{\mathbf{v} \cdot \hat{\mathbf{l}}}{c}. \quad (2.2.2)$$

Here, $\hat{\mathbf{l}}$ is the unit vector for the line of sight and c is the speed of light. We next rewrite $g(\nu)$ as the convolution of a Dirac delta function at ν_c and the ‘centred’ spectral distribution $g^c(\nu) \equiv g(\nu + \nu_c)$:

$$g(\nu) = \delta(\nu - \nu_c) * g^c(\nu). \quad (2.2.3)$$

⁶Incidentally, the son of the eponymous Pierre Jacquinot would go on to be director of JET!

Next, by substituting this expression into the Wiener-Khinchin theorem as stated in Equation 2.1.11 and using the convolution theorem (see Appendix A), we can write the coherence as the product of the inverse Fourier transforms of $\delta(\nu - \nu_c)$ and $g^c(\nu)$:

$$\gamma(\tau) = \exp\left(i\left[2\pi\nu_0\tau + 2\pi\Delta\nu_D\tau\right]\right) \int_{-\infty}^{\infty} g^c(\nu) \exp(2\pi i\nu\tau) d\nu. \quad (2.2.4)$$

In general then for quasi-monochromatic light, $\gamma(\tau)$ is the product of a phasor oscillating in τ with period $1/\nu_c$ and the inverse Fourier transform of the centred spectral distribution $g^c(\nu)$. For this second factor we will use the shorthand

$$G^c(\tau) \equiv \int_{-\infty}^{\infty} g^c(\nu) \exp(2\pi i\nu\tau) d\nu. \quad (2.2.5)$$

Our spectrum $g(\nu)$ is quasi-monochromatic if its characteristic spectral width $\Delta\nu$ satisfies $\Delta\nu/\nu_0 \ll 1$. By combining this condition with the uncertainty relation in Equation 2.1.12, it can be seen that $\Delta\tau \gg 1/\nu_c$, meaning $G^c(\tau)$ must vary on τ scales far larger than the phasor period, its magnitude acting as an envelope function. This result can be seen clearly in Figure 2.2, where the interferogram (right column) is modelled for three different spectra (left column) using Equation 2.1.9. In the top row, the interferogram for a broadband spectrum decays quickly to $I_0/2$ with very little structure. The second and third spectra are each quasi-monochromatic and their interferograms show the anticipated behaviour: a rapid oscillation with amplitude slowly decaying to zero as τ increases. We can rewrite Equation 2.2.4 with $G^c(\tau)$ in phasor form too to find

$$\gamma(\tau) = \underbrace{|G^c(\tau)|}_{\text{Contrast}} \exp\left(i\left[\underbrace{2\pi\nu_0\tau}_{\text{Stationary phase}} + \underbrace{2\pi\Delta\nu_D\tau}_{\text{Doppler phase}} + \underbrace{\arg[G^c(\tau)]}_{\text{Asymmetry phase}}\right]\right). \quad (2.2.6)$$

This is the general equation for the coherence of quasi-monochromatic light, as would be measured by an ideal interferometer. It can be used to model the coherence of emission with arbitrary (narrowband) spectral distribution. Small changes must be made to account for instrument dispersion, which will be considered in Section 2.2.2. We will now briefly consider the meaning of each term.

Equation 2.2.6 tells us that $\Re\{\gamma(\tau)\}$ is a rapidly oscillating sinusoid with a slowly varying envelope function and phase perturbation depending on $G^c(\tau)$. The envelope function is called the fringe contrast⁷ and is defined

$$\begin{aligned} \zeta(\tau) &\equiv |\gamma(\tau)| \\ &= |G^c(\tau)| \end{aligned} \quad (2.2.7)$$

⁷This is sometimes called fringe visibility.

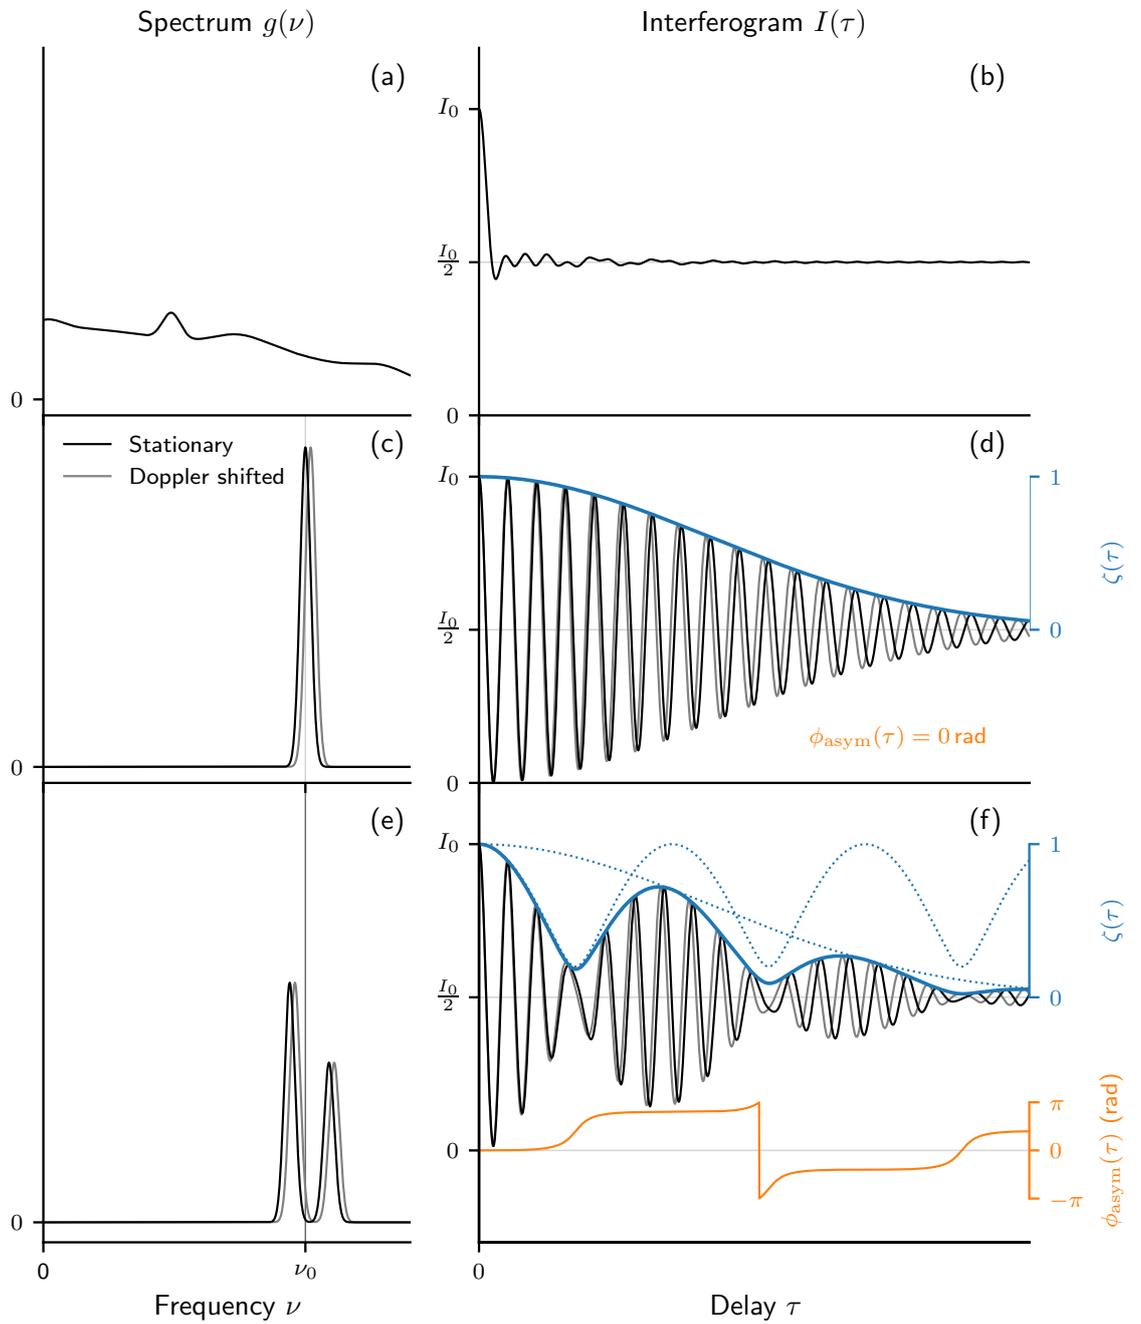


Figure 2.2: Three different modelled frequency spectra are plotted in the left column. The corresponding (ideal) interferograms, calculated using Equation 2.1.9 are plotted in the right column. The first row is a broadband spectral source, the second is a quasi-monochromatic singlet and the third is a quasi-monochromatic multiplet. The effect of a Doppler shift is illustrated for the quasi-monochromatic cases. For clarity, $\Delta\nu/\nu_0$ is significantly larger here than would ever be observed in a plasma.

The contrast is unaffected by shifts in centre frequency of the spectrum, and can be approximated for a given τ by the adjacent maximum and minimum interferogram brightness:

$$\zeta \approx \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}. \quad (2.2.8)$$

The phase of the coherence is defined

$$\begin{aligned} \Phi(\tau) &\equiv \arg[\gamma(\tau)] \\ &= \phi_0(\tau) + \phi_D(\tau) + \phi_{\text{asym}}(\tau) \end{aligned} \quad (2.2.9)$$

where we have introduced a shorthand for each of the three phase terms:

$$\begin{aligned} \phi_0(\tau) &\equiv 2\pi\nu_0\tau \\ \phi_D(\tau) &\equiv 2\pi\Delta\nu_D\tau \\ \phi_{\text{asym}}(\tau) &\equiv \arg[G^c(\tau)]. \end{aligned} \quad (2.2.10)$$

The first phase term $\phi_0(\tau)$ represents a rapidly oscillating phasor, encoding the unshifted centre frequency of the emission feature. We'll call it the 'stationary phase'. The second phase term $\phi_D(\tau)$ is the 'Doppler phase', representing another phasor that oscillates much more slowly than the first (as the Doppler shifts considered in this work are small: $\Delta\nu_D \ll \nu_0$). The increased sensitivity to Doppler shifts at large τ is clearly visible in the Doppler-shifted interferograms of Figures 2.2(d) and 2.2(f). The final phase term $\phi_{\text{asym}}(\tau)$ is related to the transformed centred lineshape $G^c(\tau)$. It is a general property that the Fourier transform of an even function is purely real. It follows that $\phi_{\text{asym}}(\tau) = 0$ rad when $g^c(\nu)$ is even (i.e. when the spectrum is symmetric about ν_c), so we call this third term the 'asymmetry phase'. In the special case of symmetry, the lineshape $g^c(\nu)$ is entirely defined by $\zeta(\tau)$. In plasma experiments, we typically already know ν_0 to a high degree of accuracy, thanks to previous measurements or atomic calculations. So $\phi_0(\tau)$ is a calibration parameter that must be accounted for to infer of the phase terms of interest: $\phi_D(\tau)$ and/or $\phi_{\text{asym}}(\tau)$. However, τ is sensitive to environmental effects like vibration and thermal expansion, regardless of the exact interferometer design, making regular measurement of $\phi_0(\tau)$ necessary to track changes and avoid systematic error. The simplest way to do this is to measure a coherent light source at ν_0 . We will assume here that this calibration procedure can be done perfectly and will discuss the details in Chapter 4.

Returning to the examples in Figure 2.2, we can now say some more about the two quasi-monochromatic spectra and corresponding interferograms. The first (Figure 2.2(c)) is a singlet with a Gaussian profile — as is observed due to Doppler broadening of an emitting gas/plasma with a Maxwellian velocity distribution and negligible absorption[33, 43]. In this case, the Doppler broadened lineshape can be

written as

$$g_D^c(\nu) = \left(\frac{c}{\sqrt{\pi}\nu_0 v_{\text{th}}} \right) \exp \left(- \left[\frac{\nu c}{\nu_0 v_{\text{th}}} \right]^2 \right), \quad (2.2.11)$$

where v_{th} is the characteristic thermal velocity of the emitting species:

$$v_{\text{th}} = \sqrt{\frac{2k_B T_i}{m_i}}. \quad (2.2.12)$$

Here, T_i and m_i are the temperature and mass of the emitting species respectively. Since $g_D^c(\nu)$ is symmetric, it follows that $\phi_{\text{asym}}(\tau) = 0$ rad and the interferogram in Figure 2.2(d) is an unperturbed cosinusoid. The contrast envelope $\zeta(\tau)$ follows a Gaussian decay with increasing τ since the Fourier transform of a (centred) Gaussian profile is another Gaussian profile (see Appendix A). The expression for the contrast envelope is then

$$\zeta_D(\tau) = \exp \left(- \left[\frac{v_{\text{th}} \nu_0 \pi \tau}{c} \right]^2 \right). \quad (2.2.13)$$

So, intuitively, $\zeta(\tau)$ falls with increased T_i (more significant Doppler broadening). It follows that for a Doppler-shifted, Doppler-broadened emission line, a measurement of $\Phi(\tau)$ and $\zeta(\tau)$ at a *single, appropriately chosen delay* is sufficient to recover $\mathbf{v} \cdot \hat{\mathbf{l}}$ and T_i respectively (assuming Maxwellian velocities and perfect calibration of $\phi_0(\tau)$). This massively reduces the τ sampling requirements of FTS to a minimum of three closely spaced (within $\pm\pi$ rad) samples, required to infer the three interferogram unknowns I_0 , $\zeta(\tau)$ and $\Phi(\tau)$. One early and successful application of this idea is the measurement of wind speed and temperature in the earth's upper atmosphere via oxygen line emission [60–62]. This is typically achieved using Michelson's interferometer design, by either temporal[62] or spatial[61] multiplexing, and has been used to measure air flows in the range ± 100 m/s and temperatures ~ 200 K.

The second example of a quasi-monochromatic spectrum (Figure 2.2(e)) has two lines, each subject to the same Doppler broadening / shift as in the previous example. If we consider a general case of N lines, the centred lineshape is now the convolution of the Doppler broadened profile and a sum of weighted delta functions:

$$g^c(\nu) = g_D^c(\nu) * \sum_i^N I_i \delta(\nu - \Delta\nu_i). \quad (2.2.14)$$

Where, if the line frequencies are ν_i , we have defined $\Delta\nu_i \equiv \nu_i - \nu_c$ and the relative intensities I_i are normalised such that $\sum_i^N I_i = 1$. Taking the Fourier transform of Equation 2.2.14 we find

$$G^c(\tau) = \zeta_D(\tau) \sum_i^N I_i \exp(i2\pi\Delta\nu_i\tau). \quad (2.2.15)$$

It follows that $\zeta(\tau) \equiv |G^c(\tau)|$ is the product of $\zeta_D(\tau)$ (given by Equation 2.2.13) and a second term due to the presence of multiple lines — both terms are plotted as dotted blue lines and their product in solid blue in Figure 2.2(f). In our example, $g(\nu)$ is no longer symmetric about ν_c , so $\phi_{\text{asym}}(\tau) \equiv \arg[G^c(\tau)]$ is non-zero and is now plotted in orange below the interferogram (wrapped in the interval $(-\pi, \pi]$). If the observed emission lines are due to fine structure splitting of atomic energy levels, then I_i and ν_i will typically be known and so the contribution to $\zeta(\tau)$ and $\Phi(\tau)$ can be calculated and accounted for in any observations⁸[63]. In this case, the τ sampling requirements for inferring $\mathbf{v} \cdot \hat{\mathbf{l}}$ and T_i are the same as in the previous example of a single line. The only complication is that τ must be chosen so as to avoid minima in $\zeta(\tau)$, to maximise SNR.

If the multiple emission lines are instead due to different ion / atomic species, then the situation is more complicated. I_i will generally not be known a priori and the individual lines may be subject to different degrees of broadening / splitting. In this case, $g(\nu)$ is determined by a larger number of parameters and therefore requires more samples of $\zeta(\tau)$ and $\Phi(\tau)$ to reconstruct. It follows that proper isolation and characterisation of the targeted emission feature is very important for interpreting narrowband FTS measurements of this kind — non-Maxwellian velocity distributions, or contaminant emission can introduce systematic error if not properly accounted for. For more complicated spectral lineshapes, $\zeta(\tau)$ and $\Phi(\tau)$ need to be sampled at multiple delays to unambiguously recover the underlying parameters. We will only consider single-delay CI instrument designs in this chapter, saving the introduction of multi-delay designs for the next chapter.

In the following sections we will consider how this basic theory is extended to account for instrument dispersion and an inhomogeneous light source.

2.2.2 Accounting for Instrument Dispersion

All real interferometers are dispersive. This means that the optical path difference between the two beams depends on the frequency of the light. Since τ is now $\tau(\nu)$, the Fourier transform relationship between $g(\nu)$ and $\gamma(\tau)$ as written in Equation 2.1.11 no longer holds. ‘Group delay’ is a concept in interferometry used to approximate dispersion to first order[55]. We will see in this section that incorporating this approximation requires only small changes to our expression for coherence of quasi-monochromatic light from Equation 2.2.6.

⁸At fusion plasma densities (10^{18} – 10^{21} m⁻³), fine structure levels typically remain collisionally coupled and so are populated according to their known statistical weights[44].

The phase delay between the two interferometer beams can be approximated by a Taylor expansion about the mean stationary frequency ν_0 , truncated at the first term:

$$2\pi\nu\tau(\nu) \approx 2\pi\nu_0\tau_0 + 2\pi(\nu - \nu_0)\hat{\tau}_0. \quad (2.2.16)$$

Here, $\tau_0 \equiv \tau(\nu_0)$ and $\hat{\tau}_0$ is the group delay, defined as the derivative of the phase delay with respect to frequency, at ν_0 :

$$\hat{\tau}_0 \equiv \left[\frac{d}{d\nu}(\nu\tau) \right]_{\nu=\nu_0} = \tau_0 + \nu_0 \left. \frac{d\tau}{d\nu} \right|_{\nu=\nu_0}. \quad (2.2.17)$$

In the literature, $\hat{\tau}_0$ is often expressed as a phase angle $\hat{\phi}_0 \equiv 2\pi\nu_0\hat{\tau}_0$ instead of a time delay. This allows us to rewrite Equation 2.2.16 as

$$\phi(\nu) \approx \phi_0 + \hat{\phi}_0 \left[\frac{\nu - \nu_0}{\nu_0} \right]. \quad (2.2.18)$$

It is also useful to define the ratio between group and phase delays as $\kappa_0 \equiv \hat{\tau}_0/\tau_0$, a dimensionless parameter of order unity that characterises the degree of dispersion. It follows from Equation 2.2.17 that this can be written:

$$\kappa_0 = 1 + \frac{\nu_0}{\tau_0} \left. \frac{d\tau}{d\nu} \right|_{\nu=\nu_0}. \quad (2.2.19)$$

Upon substituting Equation 2.2.16 into the Wiener-Khinchin theorem in Equation 2.1.11 and proceeding as in Section 2.2.1, with the assumption of a quasi-monochromatic spectrum, we obtain:

$$\gamma(\tau_0) \approx \exp\left(i\left[2\pi\nu_0\tau_0 + 2\pi\Delta\nu_D\hat{\tau}_0\right]\right) \int_{-\infty}^{\infty} g^c(\nu) \exp(2\pi i\nu\hat{\tau}_0) d\nu. \quad (2.2.20)$$

Compare this to the equivalent, non-dispersive result from Equation 2.2.4. The rapidly oscillating phasor at ν_0 in the first factor is unchanged, but $\hat{\tau}_0$ has replaced τ as the constant of proportionality in the Doppler phase term and also as the conjugate variable of ν in the Fourier transform of $g^c(\nu)$. We can then rewrite this making the magnitude and phase contributions of $g^c(\nu)$ explicit to find the dispersive version of Equation 2.2.6:

$$\gamma(\tau_0) \approx |G^c(\hat{\tau}_0)| \exp\left(i\left[2\pi\nu_0\tau_0 + 2\pi\Delta\nu_D\hat{\tau}_0 + \arg[G^c(\hat{\tau}_0)]\right]\right). \quad (2.2.21)$$

We can rewrite this using the definitions for the phase terms from Equation 2.2.10 as

$$\gamma(\tau_0) \approx \zeta(\hat{\tau}_0) \exp\left(i\left[\phi_0(\tau_0) + \phi_D(\hat{\tau}_0) + \phi_{\text{asym}}(\hat{\tau}_0)\right]\right). \quad (2.2.22)$$

For the CI instruments considered in this work, κ is in the range 1.05–1.2 across the

visible frequencies. So ignoring dispersion completely — i.e. using $\phi_D(\tau_0)$ instead of $\phi_D(\hat{\tau}_0)$ — leads to a 5–20% underestimate in inferred $\mathbf{v} \cdot \hat{\mathbf{l}}$. Using our best knowledge of CI instrument $\tau(\nu)$, it can be shown that for anticipated flow velocities $\mathbf{v} \cdot \hat{\mathbf{l}} \lesssim 30$ km/s (Doppler shifts of $\Delta\lambda_D \lesssim 0.05$ nm at $\lambda_0 = 500$ nm), the error in inferred $\mathbf{v} \cdot \hat{\mathbf{l}}$ in using the group delay treatment is negligible ($< 0.005\%$). In Chapter 4 we take a more detailed look at dispersion in the context of CI ϕ_0 calibration, where measurements are made of emission lines that are widely separated in wavelength ($\Delta\lambda \lesssim 40$ nm).

2.2.3 Accounting for an Inhomogeneous Light Source

For an extended, inhomogeneous light source, the emission spectrum $g(\nu, \mathbf{r})$ also depends on local position within the source \mathbf{r} . For line emission in the visible range in fusion plasma conditions, it can generally be assumed that the plasma is optically thin (zero absorption). The observed spectrum is therefore integrated along a line of sight through the plasma. Inferring local structure from line-integrated projections is an inverse problem encountered in many fields (e.g. medical imaging and radio astronomy). This section outlines how line-integrated measurements can be handled in narrowband FTS, and follows Section 2.5.1 from Clive Michael’s thesis[64].

The total irradiance measured along line of sight L is

$$\check{I}_0 = \int_L \epsilon_0(\mathbf{r}) dl, \quad (2.2.23)$$

where $\epsilon_0(\mathbf{r})$ is the local volumetric emissivity. The area-normalised spectral irradiance $\check{g}(\nu)$ is then

$$\check{g}(\nu) = \frac{1}{\check{I}_0} \int_L \epsilon_0(\mathbf{r}) g(\nu, \mathbf{r}) dl. \quad (2.2.24)$$

Upon substitution into the Wiener-Khinchin theorem from Equation 2.1.11, we arrive at:

$$\check{\gamma}(\tau) = \frac{1}{\check{I}_0} \int_L \epsilon_0(\mathbf{r}) \gamma(\tau, \mathbf{r}) dl, \quad (2.2.25)$$

where $\check{\gamma}(\tau)$ and $\gamma(\tau, \mathbf{r})$ are the line-integrated and local coherence respectively at interferometer delay τ . We can rewrite each in complex exponential form in terms of contrast and phase to find:

$$\check{\zeta}(\tau) \exp[i\check{\Phi}(\tau)] = \frac{1}{\check{I}_0} \int_L \epsilon_0(\mathbf{r}) \zeta(\tau, \mathbf{r}) \exp[i\Phi(\tau, \mathbf{r})] dl, \quad (2.2.26)$$

So the contrast $\check{\zeta}(\tau)$ and phase $\check{\Phi}(\tau)$ of a measured interferogram observing an extended, inhomogeneous source are each a function of the local contrast $\zeta(\tau, \mathbf{r})$ and local phase $\Phi(\tau, \mathbf{r})$ along line of sight L . To make the inversion of Equation 2.2.26 more

tractable, we express $\Phi(\tau, \mathbf{r})$ as the $\check{\Phi}(\tau)$ plus a local perturbation:

$$\Phi(\tau, \mathbf{r}) = \check{\Phi}(\tau) + \check{\Phi}(\tau, \mathbf{r}). \quad (2.2.27)$$

Substituting the above into Equation 2.2.26 and taking the Maclaurin expansion $\exp[i\check{\Phi}(\tau, \mathbf{r})] \approx 1 + i\check{\Phi}(\tau, \mathbf{r})$, we find:

$$\check{I}_0 \check{\zeta}(\tau) \approx \int_L \epsilon_0(\mathbf{r}) \zeta(\tau, \mathbf{r}) dl \quad (2.2.28)$$

and

$$\check{I}_0 \check{\zeta}(\tau) \check{\Phi}(\tau) \approx \int_L \epsilon_0(\mathbf{r}) \zeta(\tau, \mathbf{r}) \Phi(\tau, \mathbf{r}) dl. \quad (2.2.29)$$

So provided the local phase perturbations are small ($\check{\Phi}(\tau, \mathbf{r}) \ll 1$), then we have three equations (2.2.23, 2.2.28 and 2.2.29) linking the three line-integrated quantities (\check{I}_0 , $\check{\zeta}(\tau)$ and $\check{\Phi}(\tau)$) to the three local quantities ($\epsilon_0(\mathbf{r})$, $\zeta(\tau, \mathbf{r})$ and $\Phi(\tau, \mathbf{r})$) of interest. It is clear that $\epsilon_0(\mathbf{r})$ must be known to find $\zeta(\tau, \mathbf{r})$ and similarly that $\epsilon_0(\mathbf{r})\zeta(\tau, \mathbf{r})$ must be known to find $\Phi(\tau, \mathbf{r})$, so measured irradiance must be inverted first, then contrast and finally phase.

CI implements these general FTS measurement principles using polarisation interferometry, which is the subject of the next section.

2.3 Polarisation interferometry

Michelson's design is an example of a 'double-path' interferometer — its two beams take separate routes before recombining. Coherence imaging uses a 'polarisation interferometer' design instead. In this type of device, the two beams take (roughly) the same route but do so in orthogonal polarisation states. The necessary delay is produced using a birefringent material whose refractive index, determining the phase velocity of the light, depends on polarisation state. A 'common-path' interferometer like this is generally more compact and robust against misalignment than double-path designs[55, 65]. In this section, we will introduce the tools for modelling polarisation interferometers.

Modelling wave polarisation requires a vector field description of the electric field. For now we will only consider light travelling along the optical axis, which we will call z . The field can be written

$$\mathbf{E}(z, t) = \begin{pmatrix} |E_x| \exp(i\psi_x) \\ |E_y| \exp(i\psi_y) \\ 0 \end{pmatrix}. \quad (2.3.1)$$

In general, the amplitude and phase terms of the orthogonal components of the field are modified by an optical system. Two frameworks exist for modelling this action in linear systems: Mueller calculus and Jones calculus[66–68]. In both frameworks light is represented by a column vector but while Mueller calculus deals with irradiance (Stokes vectors), Jones calculus deals with the electric field itself. As such, Mueller calculus discards absolute phase information for the field while Jones calculus retains it. However, Mueller calculus can be used to model mixed and partial polarisation states, while Jones calculus is restricted to pure polarisation states. In the CI literature, Jones calculus is typically used for purely spectroscopic applications[63, 64, 69] while for spectro-polarimetric applications the Mueller treatment is used[70]. To retain the ability to model partial polarisation (needed when dealing with Zeeman splitting in Section 5.1.3), we will use Mueller calculus⁹.

2.3.1 Stokes vectors

Stokes vectors are a convenient way of describing polarised light[66–68], formulated in terms of sums and differences of observable irradiance values. For an operational definition, let I_H , I_V , I_{45} and I_{135} be the irradiance measured behind a linear polariser oriented horizontally, vertically, at 45° and 135° respectively. Also, let I_R and I_L be the irradiance measured behind a right and left-handed circular polariser respectively. The Stokes vector is then the four-element column vector

$$\mathbf{S} = \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix} = \begin{pmatrix} I_H + I_V \\ I_H - I_V \\ I_{45} - I_{135} \\ I_R - I_L \end{pmatrix} = \begin{pmatrix} \langle |E_x|^2 + |E_y|^2 \rangle \\ \langle |E_x|^2 - |E_y|^2 \rangle \\ \langle 2|E_x||E_y| \cos(\psi_x - \psi_y) \rangle \\ \langle 2|E_x||E_y| \sin(\psi_x - \psi_y) \rangle \end{pmatrix}. \quad (2.3.2)$$

The first Stokes parameter S_0 is simply the total irradiance. The second parameter S_1 is positive or negative, depending on whether the polarisation state tends towards linear horizontal or linear vertical respectively, and is zero when there is no preference between the two. Similarly, the third parameter S_2 measures a tendency towards a linear 45° polarisation state for positive values and towards a linear 135° state for negative values. Finally, a positive or negative value for the fourth parameter S_3 indicates a tendency towards right-handed or left handed rotation in the xy -plane. The Stokes vector definition in terms of the amplitude and phase of the x and y field components is also included in Equation 2.3.2. The degree of polarisation (DOP) is the fraction of the light that is polarised. It can be written in terms of the Stokes

⁹Beware: Mueller calculus is not appropriate for modelling *general* interference effects, since it cannot model coherent addition of beams in the same polarisation state. It is only applicable here because the phase delay is between orthogonal polarisation states.

parameters as

$$\text{DOP} \equiv \frac{\sqrt{S_1^2 + S_2^2 + S_3^2}}{S_0}. \quad (2.3.3)$$

For unpolarised light, $\text{DOP} = 0$; for completely polarised light $\text{DOP} = 1$. If $0 < \text{DOP} < 1$ then the light is partially polarised. In this work we are interested in the spectral irradiance of a source, and can therefore work with the ‘Stokes spectrum’

$$\mathbf{S}(\nu) = \begin{pmatrix} S_0(\nu) \\ S_1(\nu) \\ S_2(\nu) \\ S_3(\nu) \end{pmatrix}, \quad (2.3.4)$$

which has an equivalent operational definition to the Stokes vector, but for spectral irradiance. In this work, we will make the frequency dependence explicit when talking about spectral irradiance.

2.3.2 Mueller calculus

In the Mueller matrix formalism, each optical component is represented by a 4×4 matrix \mathbf{M} with real-valued (and generally frequency-dependent) elements. The effect of the component on incident light is calculated by multiplying the associated Stokes vector and Mueller matrix[66, 71]:

$$\mathbf{S}_{\text{out}}(\nu) = \mathbf{M}(\nu)\mathbf{S}_{\text{in}}(\nu) = \begin{pmatrix} m_{00} & m_{01} & m_{02} & m_{03} \\ m_{10} & m_{11} & m_{12} & m_{13} \\ m_{20} & m_{21} & m_{22} & m_{23} \\ m_{30} & m_{31} & m_{32} & m_{33} \end{pmatrix} \mathbf{S}_{\text{in}}(\nu). \quad (2.3.5)$$

If multiple components are encountered in sequence, the Mueller matrix for the total system is found by multiplying the Mueller matrices of the constituent components. For example, a system of three elements, \mathbf{M}_1 , \mathbf{M}_2 and \mathbf{M}_3 has total Mueller matrix $\mathbf{M}_{\text{tot}} = \mathbf{M}_3\mathbf{M}_2\mathbf{M}_1$. To model optical components with arbitrary orientation in the xy -plane, we use the matrix for anti-clockwise frame rotation of angle ρ from the x axis:

$$\mathbf{R}(\rho) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(2\rho) & \sin(2\rho) & 0 \\ 0 & -\sin(2\rho) & \cos(2\rho) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.3.6)$$

It follows that when a component or system of components with Mueller matrix \mathbf{M} is rotated anti-clockwise in xy -plane by angle ρ , its new Mueller matrix is given by $\mathbf{R}(-\rho)\mathbf{M}\mathbf{R}(\rho)$.

In this work, we only need to know the Mueller matrices for two component types: linear polarisers and linear retarders. The Mueller matrix for an ideal linear polariser whose transmission axis makes an angle ρ with the x -axis is

$$\begin{aligned} \mathbf{M}_{\text{pol}}(\rho) &\equiv \mathbf{R}(-\rho) \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{R}(\rho) \\ &= \frac{1}{2} \begin{pmatrix} 1 & \cos(2\rho) & \sin(2\rho) & 0 \\ \cos(2\rho) & \cos^2(2\rho) & \cos(2\rho)\sin(2\rho) & 0 \\ \sin(2\rho) & \cos(2\rho)\sin(2\rho) & \sin^2(2\rho) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \end{aligned} \quad (2.3.7)$$

A linear retarder transmits light in orthogonal polarisation states at different speeds. Incident light is resolved into two components polarised in these two directions, called the fast and slow axes, and a phase delay is introduced between them. The Mueller matrix for a waveplate imparting phase delay $\phi(\nu)$ between components and with fast axis at an angle ρ to the x -axis is

$$\begin{aligned} \mathbf{M}_{\text{LR}}(\rho, \phi) &\equiv \mathbf{R}(-\rho) \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos \phi & \sin \phi \\ 0 & 0 & -\sin \phi & \cos \phi \end{pmatrix} \mathbf{R}(\rho) \\ &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos^2(2\rho) + \sin^2(2\rho)\cos(\phi) & \cos(2\rho)\sin(2\rho)(1 - \cos \phi) & \sin(2\rho)\sin \phi \\ 0 & \cos(2\rho)\sin(2\rho)(1 - \cos \phi) & \cos^2(2\rho)\cos \phi + \sin^2(2\rho) & -\cos(2\rho)\sin \phi \\ 0 & -\sin(2\rho)\sin \phi & \cos(2\rho)\sin \phi & \cos \phi \end{pmatrix} \end{aligned} \quad (2.3.8)$$

The frequency dependence of ϕ is omitted here for clarity but is important to our application. The following identities for two linear retarders in series follow from Equation 2.3.8:

$$\begin{aligned} \mathbf{M}_{\text{LR}}(\rho, \phi_2)\mathbf{M}_{\text{LR}}(\rho, \phi_1) &= \mathbf{M}_{\text{LR}}(\rho, \phi_1 + \phi_2) \\ \mathbf{M}_{\text{LR}}(\rho + 90^\circ, \phi_2)\mathbf{M}_{\text{LR}}(\rho, \phi_1) &= \mathbf{M}_{\text{LR}}(\rho, \phi_1 - \phi_2) \\ \mathbf{M}_{\text{LR}}(\rho, \phi_2)\mathbf{M}_{\text{LR}}(\rho + 90^\circ, \phi_1) &= \mathbf{M}_{\text{LR}}(\rho, \phi_2 - \phi_1). \end{aligned} \quad (2.3.9)$$

So linear retarders can be arranged in series so as to combine their phase delays constructively or destructively.

2.3.3 A Simple Polarisation Interferometer

The simplest layout for a polarisation interferometer is a linear retarder sandwiched between two linear polarisers. The transmission axes of the polarisers are aligned at $\rho = \rho'$ and the orientation of the retarder's fast axis is $\rho' + 45^\circ$. The total Mueller matrix for this layout is:

$$\mathbf{M}_{\text{tot}} = \mathbf{M}_{\text{pol}}(\rho') \mathbf{M}_{\text{LR}}(\rho' + 45^\circ, \phi(\nu)) \mathbf{M}_{\text{pol}}(\rho'), \quad (2.3.10)$$

where $\phi(\nu)$ is the phase delay imparted by the retarder. We will limit our scope for now to measurement of unpolarised light (DOP = 0), so the Stokes vector for incident light is

$$\mathbf{S}_{\text{in}}(\nu) = \begin{pmatrix} S_0(\nu) \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (2.3.11)$$

The Stokes vector for the light at the interferometer output is then

$$\begin{aligned} \mathbf{S}_{\text{out}}(\nu) &= \mathbf{M}_{\text{tot}}(\nu) \mathbf{S}_{\text{in}}(\nu) \\ &= \frac{S_0(\nu)}{4} \left(1 + \cos \phi(\nu)\right) \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}. \end{aligned} \quad (2.3.12)$$

and is purely polarised (DOP = 1) as we would expect. The same result is obtained when the retarder's fast axis is oriented at $\rho' - 45^\circ$. The measured signal at the interferometer output is proportional to the irradiance at the output $I_{0,\text{out}}$, which is the first Stokes parameter of $\mathbf{S}_{\text{out}}(\nu)$, integrated over all frequencies:

$$I_{0,\text{out}} = \int_{-\infty}^{\infty} \frac{S_0(\nu)}{4} \left(1 + \cos \phi(\nu)\right) d\nu \quad (2.3.13)$$

The irradiance at the input is defined $I_{0,\text{in}} \equiv \int_{-\infty}^{\infty} S_0(\nu) d\nu$. The area-normalised spectral irradiance of the input is then $g(\nu) = S_0(\nu)/I_{0,\text{in}}$. The phase delay imparted by the retarder can be written as $\phi(\nu) = 2\pi\nu\tau(\nu)$. Substituting these three relations into Equation 2.3.13 we find

$$I_{0,\text{out}} = \frac{I_{0,\text{in}}}{4} \left(1 + \Re \left\{ \int_{-\infty}^{\infty} g(\nu) \exp [2\pi i \nu \tau(\nu)] d\nu \right\}\right). \quad (2.3.14)$$

The contents of the curly brackets will hopefully be recognisable as the (dispersive) coherence from Equation 2.2.22. So the irradiance at the output of our polarisation

interferometer is:

$$I_{0,\text{out}} = \frac{I_{0,\text{in}}}{4} \left(1 + \Re\{\gamma(\tau_0)\} \right). \quad (2.3.15)$$

That this is a factor two smaller than the irradiance at the output of a Michelson interferometer (Equation 2.1.9) is an unavoidable consequence of the front polariser rejecting half of the unpolarised incident light. However, if a polarising beam splitter is used in place of the front polariser, then this light can be used for other purposes[64]. When the transmission axes of the two polarisers are crossed instead of aligned, the addition in Equation 2.3.15 becomes subtraction. Linear retarders are constructed out of birefringent materials, which we need to understand to model the realistic interferometer response.

2.3.4 Phase Delay due to a Uniaxial Crystal Plate

In the previous section retarders were assumed to impart a fixed (frequency-dependent) delay between orthogonal polarisation states when, in reality, the delay will depend on the path of the incident light ray. Taking advantage of this dependence to scan the delay across the sensor plane is an established technique for snapshot imaging Fourier transform spectroscopy by spatial multiplexing. In this section we will consider the delay path-dependence in uniaxial birefringent crystals, the type of retarder used in this work.

In an anisotropic material, the electron cloud surrounding each atom or molecule distorts under the influence of an applied electric field more freely in one direction than in another. As a result, refractive index n depends on the orientation of the electric field, determined by the light's path and polarisation state. In this work our interest is limited to *uniaxial* crystals, for which a single axis of symmetry determines optical behaviour. If a wave's field vector lies in the plane perpendicular to this axis, the wave will experience the same n . This axis of symmetry is called the crystal optic axis (COA). For a given ray direction in a uniaxial crystal, Maxwell's equations permit two plane wave solutions with different v_{ph} and orthogonal linear polarisation states[54, 66]. One of these solutions is polarised orthogonal to the COA and experiences the same refractive index n_{O} regardless of ray direction — we call this the ordinary (O) ray. The O-ray follows Snell's law at material boundaries. The second solution has a field component parallel to the COA and is called the extraordinary (E) ray. For the E-ray, the refractive index is path-dependent and the ray does not, in general, follow Snell's law. The refractive index n_{E} quoted for a uniaxial crystal is that experienced by an E-ray whose field aligns with the COA, i.e. it's direction of travel is orthogonal to the COA. A special case is when the ray direction aligns with the COA, for which both rays experience refractive index n_{O} and there is no birefringent effect. In general,

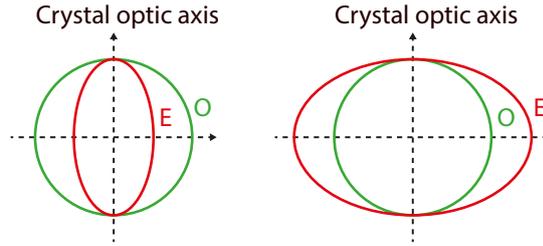


Figure 2.3: Wave normal surfaces for (a) positive and (b) negative uniaxial crystals.

the path of the E-ray can be traced using Huygens’ principle of secondary wavelets [72, 73], only the wavelet for the E-ray is not a sphere but an ellipsoid, as shown in Figure 2.3. A derivation or even a full account of the E-ray path is beyond the scope of this work so we will only quote the key result of the phase difference between the O and E rays here¹⁰.

We start by defining the necessary geometry, with reference to Figure 2.4¹¹. All crystal plates used in this work have two opposite faces that are plane, parallel to one another and perpendicular to the system’s optical axis. Any plane within the crystal that contains both the surface normal and the COA is called a ‘principal section’, as shown in Figure 2.4(a). The angle between the COA and the front face is called the cut angle θ . To generalise for arbitrary crystal orientations in the xy -plane, we define ρ as the angle between the x -axis and the principal section. The plane containing both the incident ray and the surface normal is the plane of incidence, shown in Figure 2.4(b), and the angle of incidence α is the angle between the incident ray and the surface normal. The angle between the x -axis and the plane of incidence is β . Finally, the angle between the plane of incidence and the principal section, shown in Figure 2.4(c), is $\delta \equiv \beta - \rho$.

In general, the phase delay imparted by a uniaxial crystal plate is $\phi = \frac{2\pi}{\lambda} \text{OPD}$ where the optical path difference (OPD) between the two rays is a function of α , θ , δ and the plate thickness L . However, when the light is normally incident ($\alpha = 0^\circ$) and the crystal’s front face is cut parallel to its optic axis ($\theta = 0^\circ$), the O and E rays follow a common path and experience refractive indices n_O and n_E respectively. It follows that phase delay is simply

$$\phi = \frac{2\pi L(n_E - n_O)}{\lambda}. \quad (2.3.16)$$

This situation is shown in Figure 2.5(a). Figure 2.5(b) then shows the path of incident

¹⁰A recent paper states that “the abundance of optical subtleties that can be discovered in anisotropic media seems to be inexhaustible”[74] — probably handy if you work in the field.

¹¹Here, we use the nomenclature of the crystal optics literature[75] and the same coordinate system used in recent CI work[70].

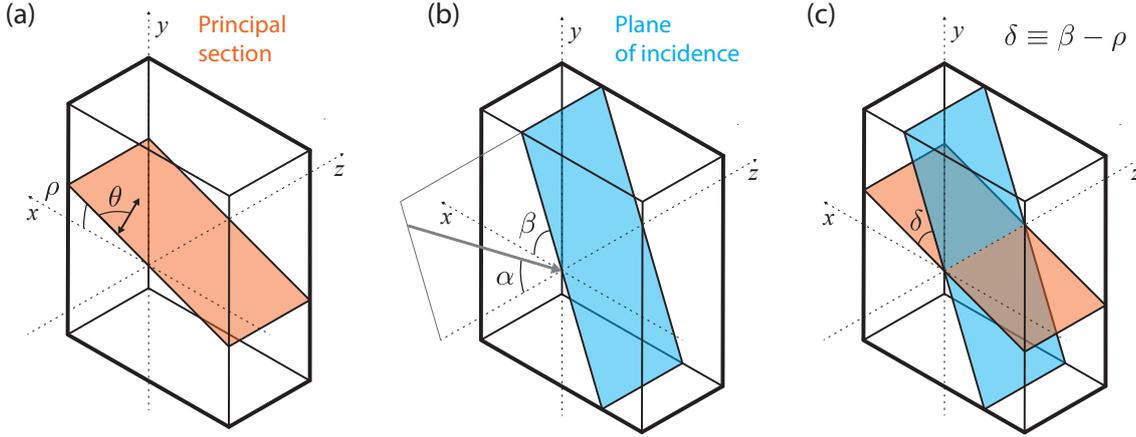


Figure 2.4: Definitions of the angles used to calculate the phase delay imparted by a plane-parallel, uniaxial crystal plate for off-axis rays. The z -axis is the system's optical axis and the crystal optic axis is indicated by the black arrow in (a). See text for details.

light, polarised at 45° . The O-ray is shown in green and the E-ray is shown in red. In the diagram $n_E > n_O$ and so the E-ray lags behind the O-ray. If we maintain normal incidence but consider an arbitrary angle θ between the COA and the crystal front face, we see that there is a lateral displacement between the two rays in the plane of the principal section, as shown in Figure 2.6.

In general, OPD is a function of the angles θ , α and δ , the refractive indices n_O and n_E and the plate thickness L . It can be shown, using either Maxwell's equations or Huygens' principle of secondary wavelets[75], that the phase delay between O-ray and E-ray for light with wavelength λ is:

$$\phi = \frac{2\pi L}{\lambda} \left(\sqrt{n_O^2 - n^2 \sin^2 \alpha} + \frac{n(n_O^2 - n_E^2) \sin \theta \cos \theta \cos \delta \sin \alpha}{n_E^2 \sin^2 \theta + n_O^2 \cos^2 \theta} - \frac{n_O \sqrt{n_E^2 (n_E^2 \sin^2 \theta + n_O^2 \cos^2 \theta) - \{n_E^2 - (n_E^2 - n_O^2) \cos^2 \theta \sin^2 \delta\} n^2 \sin^2 \alpha}}{n_E^2 \sin^2 \theta + n_O^2 \cos^2 \theta} \right), \quad (2.3.17)$$

Here, n is the refractive index of the surrounding medium. The Mueller matrix for an ideal plane-parallel uniaxial crystal plate can be found by combining Equation 2.3.17 with the definition for the Mueller matrix linear retarder in Equation 2.3.8. The meaning of each term in Equation 2.3.17 will be clearer in the next section when we visualise the resulting interferogram.

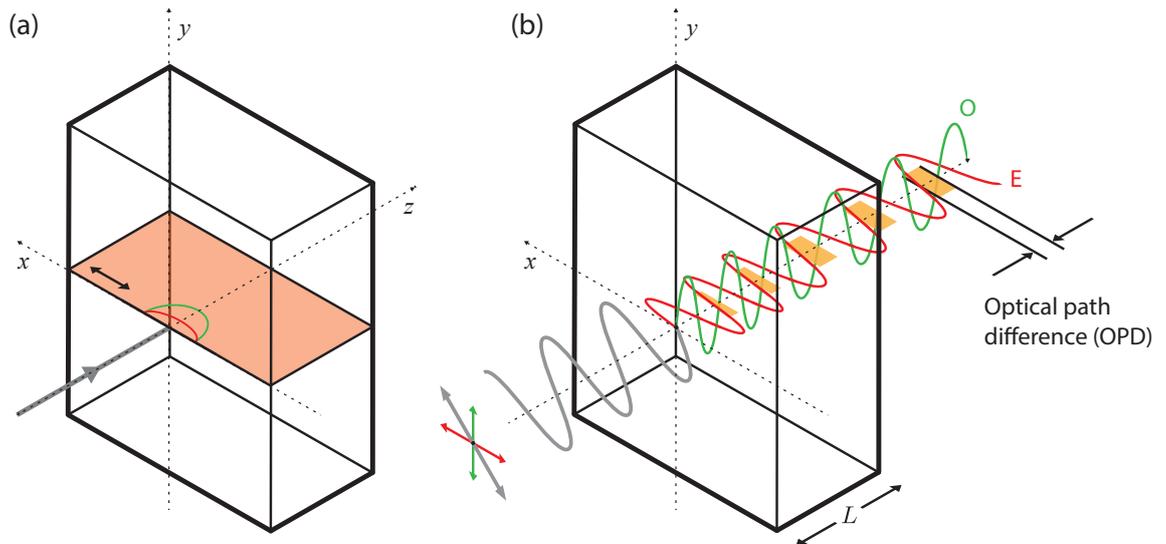


Figure 2.5: (a) Light is normally incident ($\alpha = 0^\circ$) on a plane-parallel uniaxial crystal plate whose optic axis is parallel to the front face ($\theta = 0^\circ$). (b) Considering monochromatic, linearly polarised light (oriented at 45°), the evolution of the electric field through the plate is shown. The ordinary (green) and extraordinary (red) rays are linearly polarised perpendicular to and parallel to the principal section, respectively. The rays travel along a common path at different speeds, resulting in an optical path difference.

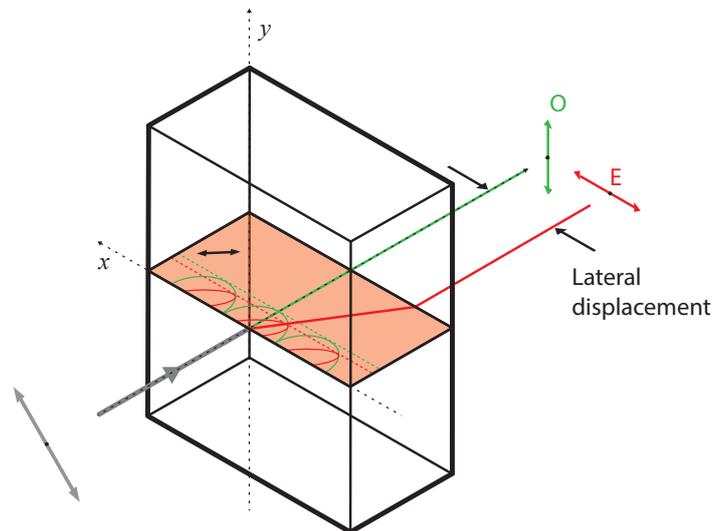


Figure 2.6: Light is normally incident ($\alpha = 0^\circ$) on a plane-parallel uniaxial crystal plate whose optic axis makes an angle θ with the front face. Considering monochromatic, linearly polarised light (oriented at 45°), the paths of the ordinary (green) and extraordinary (red) rays are shown. A lateral displacement is introduced between the rays in the plane of the principal section.

2.3.5 Imaging Polarisation Interferometry

To produce interference fringes with high contrast, light reaching each point on the sensor must experience the same delay. Since the delay imparted by a uniaxial crystal depends on the incident ray's direction, the light passing through the polarisation interferometer must be collimated. CI instruments are imaging polarisation interferometers which typically have the layout shown in Figure 2.7. Light from the source is collected by the objective lens (l_1) and focussed to an intermediate image at the front focal plane of a second lens (l_2). This lens collimates the light and directs it through the interferometer and bandpass filter before it is focussed by a final lens (l_3) onto the sensor plane. The polarisation interferometer shown here — polariser, crystal, analyser — is the simple layout introduced in Section 2.3.3. To model the

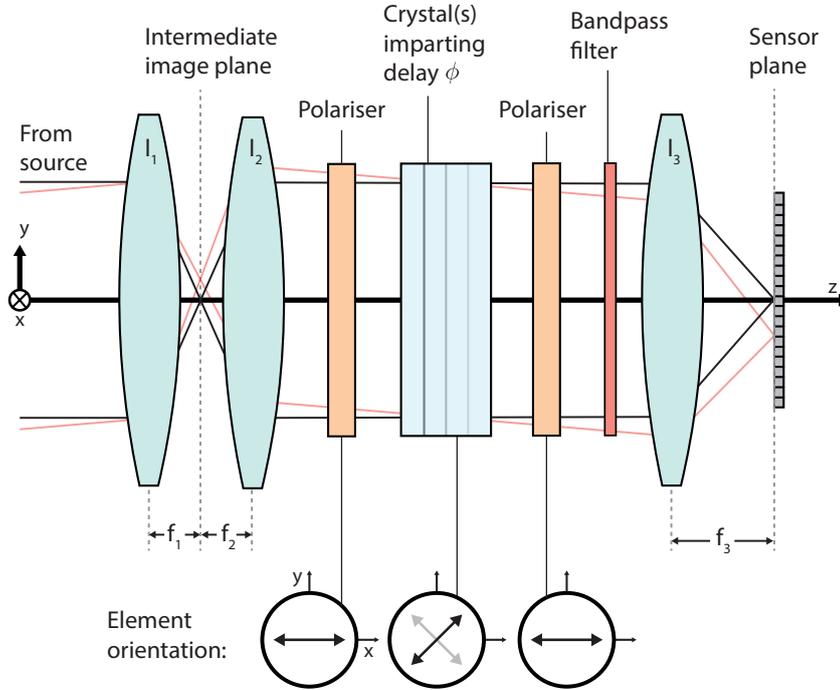


Figure 2.7: The imaging polarisation interferometer optical design used in this work.

interferogram image, delay ϕ must be mapped onto the sensor plane. By treating l_3 as a thin lens, the crystal incidence angle α can be related to position on the sensor plane x, y by

$$\alpha(x, y) = \arctan\left(\frac{\sqrt{x^2 + y^2}}{f_3}\right), \quad (2.3.18)$$

where f_3 is the focal length of l_3 . The angle between the principal section and plane of incidence is $\delta(x, y) \equiv \beta(x, y) - \rho$, where the orientation of the crystal ρ is independent of position on the sensor plane. Therefore, we only need to map the orientation of

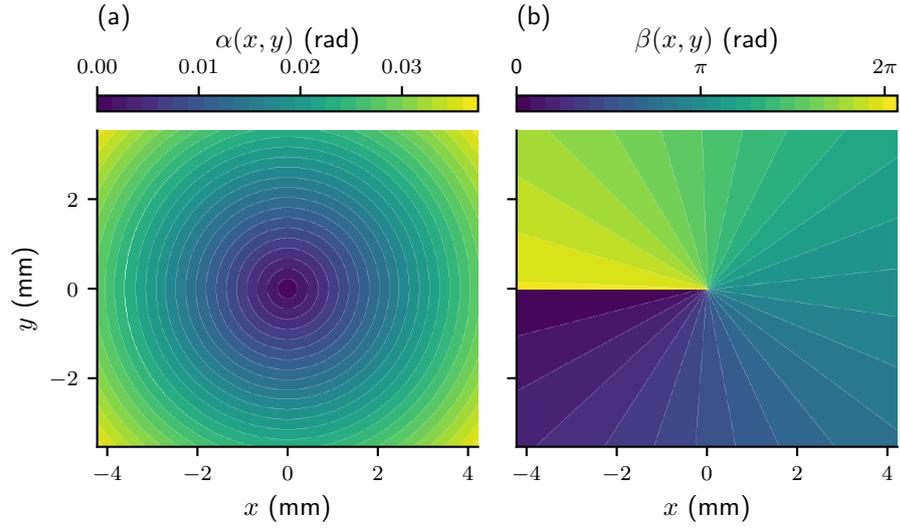


Figure 2.8: Projection onto the sensor plane of the angles defining ray path through the interferometer. (a) Angle of incidence $\alpha(x, y)$ for a lens with focal length $f_3 = 150$ mm. (b) The angle made between the x -axis and the projection of the incident ray onto the front face of the crystal $\beta(x, y)$.

the plane of incidence $\beta(x, y)$ onto the sensor plane:

$$\beta(x, y) = \text{atan2}(y, x) + \pi. \quad (2.3.19)$$

Here, $\text{atan2}(\dots)$ is the ‘two-argument arctangent’ function often used to convert between Cartesian and polar coordinate systems[76] and the extra π term appears since the lens produces an inverted image. Figure 2.8 plots $\alpha(x, y)$ and $\beta(x, y)$ for $f_3 = 150$ mm. By combining Equations 2.3.18, 2.3.19 with Equation 2.3.17, the distribution of phase delay across the sensor plane $\phi(x, y)$ can be calculated and therefore so too can the Mueller matrix for a plane-parallel uniaxial crystal of arbitrary orientation. We can now write a general equation for the irradiance at the sensor plane of an imaging polarisation interferometer with arbitrary polarising and retarding components, observing a scene with arbitrary Stokes spectrum. Let the total Mueller matrix for the interferometer be $\mathbf{M}_{\text{tot}}(x, y, \nu)$ and let $\mathbf{S}_{\text{in}}(x, y, \nu)$ be the Stokes vector for the light that would be arriving at point x, y on the sensor plane *in the absence of the interferometer*. The Stokes vector at x, y is then calculated as

$$\mathbf{S}_{\text{out}}(x, y) = \int_0^\infty \mathbf{M}_{\text{tot}}(x, y, \nu) \mathbf{S}_{\text{in}}(x, y, \nu) d\nu. \quad (2.3.20)$$

The total irradiance at point x, y is then the first Stokes parameter of $\mathbf{S}_{\text{out}}(x, y)$. The top row of Figure 2.9 shows example interferograms modelled using Equation 2.3.20 for an instrument with the configuration shown in Figure 2.7. The modelled source is extended, uniform and monochromatic ($\lambda = 450$ nm) and the delay is produced by a

single uniaxial crystal plate of thickness 10 mm. From Figure 2.9(a)–(e), the cut angle θ of the crystal increases from 0 to 90°. The crystal material modelled here (and the only birefringent material used in this work) is barium borate BaB_2O_4 (BBO), for which $n_E \approx 1.56$, $n_O \approx 1.69$ and $B \equiv n_E - n_O \approx -0.13$ at this wavelength¹². The final lens has focal length 75 mm. It is useful here to consider $\phi(x, y)$ as the sum of two parts: the phase offset ϕ_{offset} , which is the phase delay at normal incidence (i.e. at the nominal image centre $x = y = 0$), and the phase shape $\phi_{\text{shape}}(x, y) \equiv \phi(x, y) - \phi_{\text{offset}}$. The bottom row of Figure 2.9 plots the corresponding $\phi_{\text{shape}}(x, y)$ for the interferograms in the top row and ϕ_{offset} is also stated. An expression for ϕ_{offset} can be found by substituting $\alpha = 0^\circ$ into Equation 2.3.17:

$$\phi_{\text{offset}} = \frac{2\pi L}{\lambda} n_O \left(1 - \frac{n_E}{\sqrt{n_E \sin^2 \theta + n_O^2 \cos^2 \theta}} \right) \quad (2.3.21)$$

At $\theta = 0^\circ$ we see the hyperbolic fringe pattern characteristic of a waveplate and we see that ϕ_{offset} is at its maximum. At the other extreme ($\theta = 90^\circ$), normal incidence corresponds to a ray directed along the COA and so $\phi_{\text{offset}} = 0$. In this case, the size of the component of the field parallel to the COA depends only on the angle of incidence $\alpha(x, y)$ and so the lines of equal phase delay form concentric circles.

For intermediate cut angles $0^\circ < \theta < 90^\circ$, the phase offset ϕ_{offset} takes some intermediate value and, importantly, we see a large phase shear across the image in the direction of the principal section. This shear largely comes about due to lateral displacement between the O and E rays, and reaches a maximum at $\theta \approx 45^\circ$. The minimum fringe period distance P at the sensor plane (indicating maximum phase shear) can be estimated by substituting $\theta = 45^\circ$ and $\delta = 0^\circ$ into Equation 2.3.17, assuming small incidence angles α and combining with Equation 2.3.18:

$$P \approx \frac{f_3 \lambda}{L} \left(\frac{n_O^2 + n_E^2}{n_O^2 - n_E^2} \right). \quad (2.3.22)$$

This phase shear acts as a spatial carrier frequency across the sensor. This carrier is then both amplitude-modulated and phase-modulated by the observed spectral lineshape at each part of the image, according to the FTS theory laid out in Section 2.2. When the phase shear is near-constant across the image, as in Figure 2.9(c), it is called a linear carrier. The (near) constant spatial frequency of the carrier should ensure separation in the spatial frequency domain between the background brightness (I_0) and coherence ($\Re\{\gamma\}$) terms of the interferogram. The brightness, phase and contrast images can then be extracted (demodulated) from the raw image using standard Fourier analysis techniques. See appendix B for an outline of Fourier demodulation of

¹²Dispersion in BBO is treated in detail in Section 4.1.2.

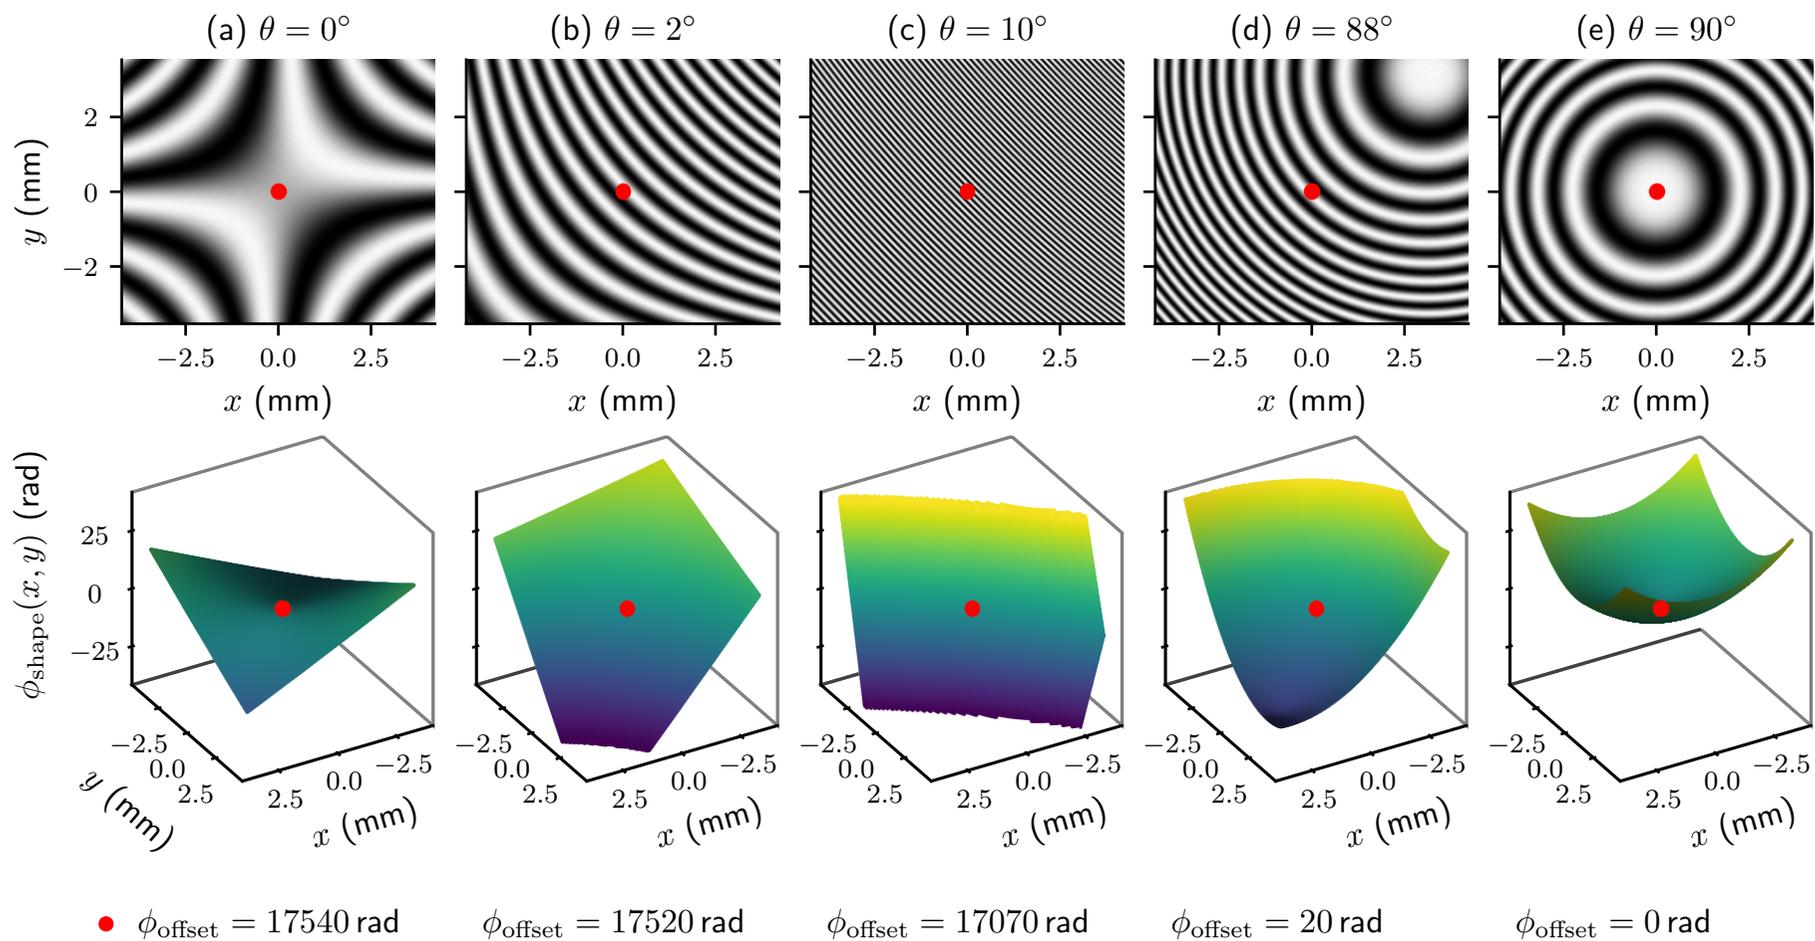


Figure 2.9: Top row: modelled interferograms for a simple imaging polarisation interferometer consisting of a single uniaxial crystal plate sandwiched between two polarisers (component orientations and optical setup as shown in Figure 2.7). Bottom row: corresponding modelled phase shape at the sensor plane. See text for details.

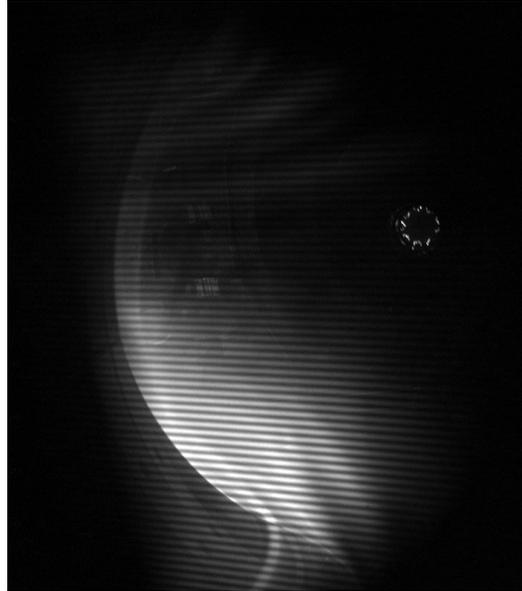


Figure 2.10: A raw CI image observing carbon emission (CIII, 464.9 nm) on the HL-2A tokamak during a disruption, showing the sinusoidal fringe pattern characteristic of a linear spatial carrier. The view is tangential and the curved surface visible on the left is the centre column.

interference fringes with a linear spatial carrier with a treatment of noise propagation from the raw image into the demodulated quantities. In the next chapter we will deal with a pixelated spatial carrier, for which new demodulation techniques need to be introduced. For a linear spatial carrier, the spatial resolution in the direction of the phase shear (orthogonal to the fringes) is roughly one fringe period[63]. The spatial resolution in the orthogonal direction is a single pixel. This is assuming that the resolution is limited by the detector and not the optical components. Figure 2.10 shows an example of a raw CI image taken with a linear spatial carrier (observing C III emission at 464.9 nm on the HL-2A tokamak) showing the characteristic sinusoidal fringe pattern superimposed on top of the observed brightness scene¹³.

As shown in Figure 2.9(c), a single uniaxial crystal plate with intermediate cut angle ($0^\circ < \theta < 90^\circ$) can produce both the phase offset and phase shear required for a single-delay snapshot CI instrument with a linear spatial carrier. This type of crystal is typically called a ‘displacer plate’ in the literature[70, 77, 78]. An alternative is to produce the phase offset and shear using separate birefringent components. Clearly from Figure 2.9(a), a waveplate ($\theta = 0^\circ$) essentially produces a phase offset only. A Savart plate can be used to produce the phase shear only.

¹³CI measurements were made on the HL-2A tokamak by the author and collaborators at the Southwestern Institute of Physics, Chengdu, but this data is not presented in this thesis.

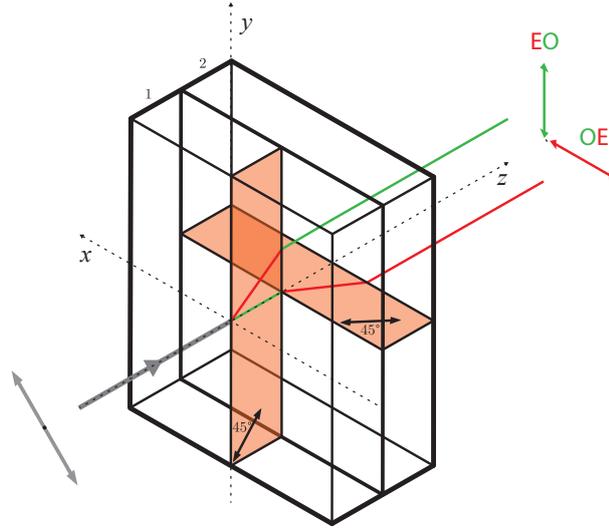


Figure 2.11: The layout of a Savart plate and the O- and E-ray paths for normal incidence.

2.3.6 Savart plate

A Savart plate is a composite component made of two plane-parallel uniaxial crystal plates of equal thickness and cut angle $\theta = 45^\circ$, whose principal sections are oriented at right angles[65]. Figure 2.11 shows a schematic. The result is that the O-ray in the first plate becomes the E-ray in the second (and vice versa), meaning that the ϕ_{offset} imparted by the first crystal is perfectly cancelled out by that of the second—for normal incidence at least. Off-axis, a phase shear is produced that is a factor $\sqrt{2}$ less steep than that produced by a single crystal plate of the same (total) thickness and is oriented at 45° to the principal sections. Approximations for the phase delay ϕ due to a Savart plate can be found in the literature[65, 79], or it can be evaluated by directly modelling the two constituent crystal plates. If the phase delays imparted by the plates are ϕ_1 and ϕ_2 then the Mueller matrices are $\mathbf{M}_{\text{LR}}(\rho, \phi_1)$ and $\mathbf{M}_{\text{LR}}(\rho + 90^\circ, \phi_2)$ respectively. Here, we consider the orientation angle of the Savart plate ρ relative to the principal section of the first plate, so in Figure 2.11 we have $\rho = 90^\circ$. It follows from the relations in Equation 2.3.9 that the total Mueller matrix for a Savart plate with orientation ρ is

$$\begin{aligned} \mathbf{M}_{\text{SP}}(\rho) &= \mathbf{M}_{\text{LR}}(\rho + 90^\circ, \phi_2) \mathbf{M}_{\text{LR}}(\rho, \phi_1) \\ &= \mathbf{M}_{\text{LR}}(\rho, \phi_1 - \phi_2) \end{aligned} \quad (2.3.23)$$

So, ϕ_1 and ϕ_2 are calculated for each crystal using Equation 2.3.17 and then combined subtractively¹⁴. Using a Savart plate/waveplate combination instead of a single displacer plate makes for a more flexible, but less compact instrument design. Generally, it also has an unfavourable effect on the instrument function, which will be briefly discussed in the next section.

2.3.7 Instrument function

For real interferometers, different effects conspire to reduce the observed fringe contrast below the magnitude of the observed coherence – analogous to instrumental broadening in a grating spectrometer. It is important to properly calibrate for these effects to avoid systematic errors in the inferred spectrum. This section explains how some of these effects arise and how they can be calibrated. Maintaining a high fringe contrast is important as the SNR of the measurement scales with both the light intensity and the contrast.

Finite detection area

In section 2.3 we calculated the point-wise irradiance at the detector plane for a simple polarisation interferometer. We have also seen that, for birefringent crystals, the dependence of phase delay on ray path results in a phase distribution as a function of sensor position $\phi(\mathbf{r})$, where we use the shorthand $\mathbf{r} = \{x, y\}$. It follows that the finite detection area of each pixel spans a range of interferometer delays, artificially reducing the contrast of the measured interferogram. This contrast reduction will be greater, the larger the range of delays each pixel's area spans.

This effect will be considered in detail since it is an unavoidable consequence of linear spatial carrier systems, which require a large phase shear across the sensor. We now write the predicted irradiance for a simple, single-delay CI system with the dependence on \mathbf{r} made explicit:

$$I(\mathbf{r}, \tau) = \frac{I_0(\mathbf{r})}{4} \left(1 + \Re\{\gamma(\mathbf{r}, \tau)\} \right). \quad (2.3.24)$$

Now, consider a single pixel on the detector plane \mathcal{P} with dimension d and area A . The measured signal from \mathcal{P} is proportional to the incident number of photons which, in turn, is proportional to the total incident power (for quasi-monochromatic light). To arrive at the total observed power, we take $P = \int_{\mathcal{P}} I(\mathbf{r}, \tau) dA$ where dA is an infinitesimal area element on the detector plane (we will assume that the incident

¹⁴Alex Thorman has shown that Savart plates also introduce unavoidable secondary carrier frequency terms due to off-axis effects, which can be suppressed via Fourier analysis[70].

light's brightness and coherence is constant over the pixel's area) and, similarly for $P_0 = \int_{\mathcal{P}} I_0(\mathbf{r})dA$:

$$P(\tau) = \frac{P_0 A}{4} \left(1 + \Re\{\gamma_m(\tau)\} \right) \quad (2.3.25)$$

where we have used $A \equiv \int_{\mathcal{P}} dA$. This has the same form as Equation 2.3.24 with the coherence encoded in the observed interferogram $\gamma_m(\tau)$ defined as the integral over the point-wise coherence:

$$\gamma_m(\tau) \equiv \frac{1}{A} \int_{\mathcal{P}} \gamma(\mathbf{r}, \tau) dA. \quad (2.3.26)$$

Both observed and point-wise coherence parameters have associated with them contrast and phase:

$$\zeta_m(\tau) \exp[i\Phi_m(\tau)] = \frac{1}{A} \int_{\mathcal{P}} \zeta(\mathbf{r}, \tau) \exp[i\Phi(\mathbf{r}, \tau)] dA, \quad (2.3.27)$$

such that contrast of the observed interferogram will be given by

$$\begin{aligned} \zeta_m(\tau) &= \left| \frac{1}{A} \int_{\mathcal{P}} \zeta(\mathbf{r}, \tau) \exp[i\Phi(\mathbf{r}, \tau)] dA \right| \\ &\approx \zeta(\tau) \underbrace{\left| \frac{1}{A} \int_{\mathcal{P}} \exp[i\Phi(\mathbf{r}, \tau)] dA \right|}_{\equiv \zeta_I}. \end{aligned} \quad (2.3.28)$$

Here, it has been assumed that the contrast of the point-wise interferogram does not vary over the pixel area: $\zeta(\mathbf{r}, \tau) \approx \zeta(\tau)$. The new defined quantity ζ_I is the instrumental contrast, a calibration parameter representing the maximum theoretical fringe contrast that can be measured with the instrument, when the observed light is perfectly coherent ($\zeta(\tau) = 1$). To gauge the size of ζ_I , consider the example of the phase shear required for a linear spatial carrier. The (idealised) fringes are then straight, parallel, cosinusoidal and aligned with the x -axis. The corresponding phase distribution is:

$$\Phi(\mathbf{r}, \tau) = \frac{2\pi}{pd} y, \quad (2.3.29)$$

where p is the number of pixels per fringe, d is the pixel dimension and y is the sensor plane y coordinate. Substituting this into Equation 2.3.28 and evaluating, we find:

$$\zeta_I = |\text{sinc}(\pi/p)| \quad (2.3.30)$$

Figure 2.12 plots ζ_I vs. p . Clearly, choice of fringe period is then a trade-off: maximising spatial resolution by minimising p reduces the observed fringe contrast, lowering the dynamic range and SNR of the measurement. p can be estimated for a displacer plate using Equation 2.3.22 and for a Savart plate by multiplying that expression by $\sqrt{2}$.

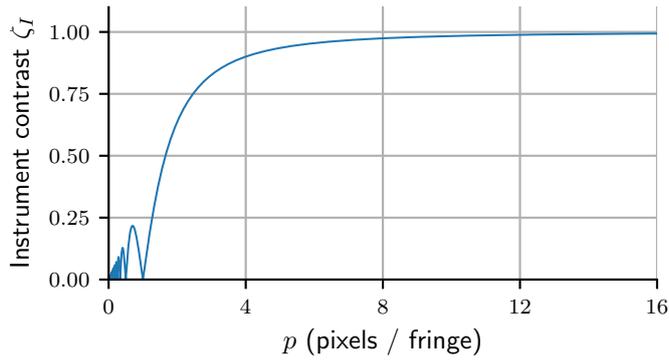


Figure 2.12: Instrument contrast ζ_I as a function of p , the fringe period in pixels for a simplified case of straight, parallel and horizontal fringes.

Linear spatial carrier CI designs typically use around $p = 10$, which gives $\zeta_I = 0.98$ [63]. For these systems, measured ζ_I is typically lower than predicted by Equation 2.3.28. This is due in large part to non-uniformity of the interferometer delay produced across the crystal aperture, which is then integrated over at each point on the sensor. This can be modelled using a similar treatment to Equation 2.3.28, integrating over aperture area instead of detector area.

Since a displacer plate CI system achieves the same phase offset and phase shear as a waveplate + Savart plate system with a lower combined crystal thickness and fewer crystal cut faces, one would expect ζ_I to be higher for this design.

In general, decreased ζ_I can also be caused by component misalignment, imperfect lens focussing and non-ideal behaviour of components. These effects can generally be modelled with the instrument model set out in this chapter.

2.4 Previous CI Investigations

The first application of fixed delay FTS with a polarisation interferometer to measure plasma emission was by John Howard and colleagues at Australian National University, demonstrated on the H1 heliac stellarator. The prototypical CI design was the MOSS (Modulated Optical Solid-State) spectrometer[64, 80, 81]. This was a temporal multiplexing design, rapidly scanning delay over a π rad range by applying an alternating voltage (\sim kV, \sim 10 kHz) across a waveplate. By the electro-optic effect, this changes the birefringence, and so the imparted delay. However, synchronising the delay scanning and camera acquisition is complicated compared to snapshot designs, and the reduced time resolution of the spectral measurement sets a lower limit on the timescales of the dynamic plasma phenomena that can be observed. The first results

from MOSS were captured using a 16 channel photomultiplier tube[64], while a later implementation on the WEGA stellarator used a 164×164 format CCD detector[81]. Another way of scanning the delay in time that has been demonstrated is to use a ferroelectric liquid crystal (FLC), which can be rapidly switched to scan the delay, but with much lower voltages[82].

A static imaging FTS design using a Savart plate to create a linear spatial carrier was used for demonstration purposes as far back as 1970[83], with renewed interest in the design for scientific use from 2002[84]. Since then, CI has broadly tracked developments from the wider optics community, for example the earliest spatial multiplexing designs used Wollaston prisms to produce four separate interferograms in the four quadrants of the sensor. Since 2010, linear spatial carrier CI instrument designs have been more popular.

The most mature CI application at the time of writing is the measurement of impurity flows in the SOL, with the DIII-D, ASDEX-Upgrade (AUG) and Wendelstein 7-X (W7-X) experiments each operating at least one system permanently, and MAST operating a system in 2013. All of these experiments use the spatial heterodyne CI design, with earlier investigations (2010-2015) tending to use Savart plates and more recent investigations (2015-2020) using displacer plates to create the necessary phase shear. Observed SOL flows are consistently of the order ~ 20 km/s. DIII-D was the first to install its system, observing C^{2+} flows in the divertor and finding good agreement with UEDGE fluid modelling[85–87]. More recently, DIII-D has installed a second, wide-field view system covering the main-chamber SOL and the upper and lower divertors. Recent results show good agreement with UEDGE modelling for main-ion flows in helium discharges[52]. Important work on minimising systematic errors due to environmental effects such as thermal drifts, vibration and magnetic fields, as well as direct comparison with high-resolution spectrometer has also been presented[88]. On MAST, CI measurements were made of carbon and helium flows from midplane and divertor views[51, 63]. Measurements of the C^{2+} flows in MAST’s high-field side SOL during neutral deuterium gas puffing (shown in Figure 2.13) were more recently compared to modelling of main ion flows using the EMC3-EIRENE plasma edge fluid and kinetic neutral transport code, finding good qualitative agreement[89]. On AUG, measurements of impurity and neutral deuterium flows around the divertor have been presented[78], while on W7-X, striking images of counter-streaming C^{2+} flows in the island-divertors have been found to agree well with EMC3-EIRENE predictions[90].

An area of recent development for all labs operating Doppler flow CI diagnostics is absolute wavelength calibration. DIII-D, AUG and W7-X have each installed tuneable lasers for the necessary inter-shot phase calibration[77, 91, 92]. We will

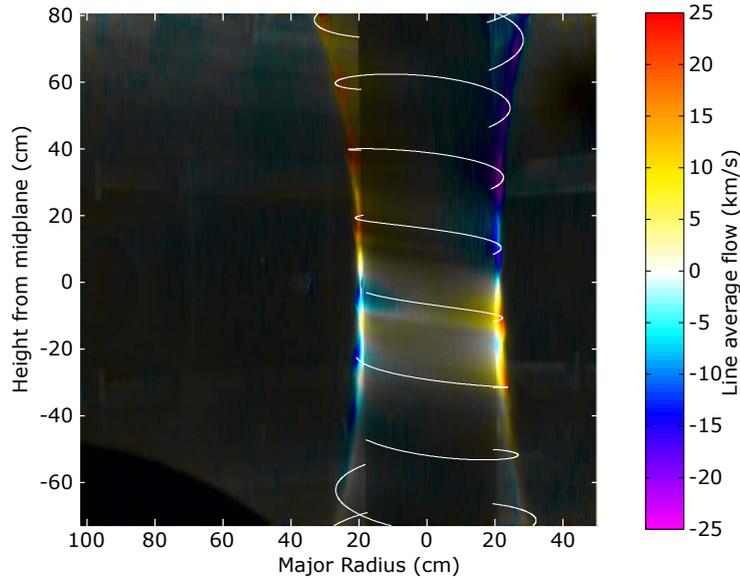


Figure 2.13: Line-integrated carbon impurity flows (C III, $\lambda = 464.9$ nm) around the centre column of the MAST tokamak during midplane fuelling on the high field side. The white line is an estimated magnetic field line. Reproduced from [63] with permission from the author.

save a discussion of this work for Chapter 4, where we introduce an alternative wavelength calibration scheme that uses simpler hardware (gas-discharge lamps) and is significantly cheaper.

A number of proof-of-principle CI investigations have also been carried out without permanent installation or further development. For example, a spatially multiplexed Savart plate design was used to measure impurity temperature and flow via charge exchange recombination emission from a heating neutral beam on the TEXTOR tokamak[93]; see also the (mostly theoretical) work developing CI techniques for measuring electron temperature and density via Thomson scattering[94–96]. Additionally, a demonstration of electron density measurement via Stark broadening of hydrogen Balmer series emission was made on the Pilot-PSI linear plasma experiment[97]. Chapter 5 of this thesis further develops this last application.

Although not the focus of this work, spectropolarimetric CI designs have been developed in parallel to purely spectroscopic designs. Their main application in fusion experiments is the measurement of Balmer- α emission (656 nm) from high-energy neutral hydrogen (or deuterium) beams injected for heating and current control. A strong motional electric field causes Stark splitting and polarisation of the multiplet line which encode the local magnetic field strength and orientation respectively. Since the neutral beam reaches into the plasma core, this technique can measure the local magnetic field along a large radial extent of the plasma, which is important for inferring

the current profile and understanding related instabilities. Spectropolarimetric CI designs typically replace the front polariser of the configuration shown in Figure 2.7 with a delay plate at the same orientation. Each of the two delays needs to be scanned in order to recover both the Doppler shift and the orientation of the linearly polarised light, which both manifest as an interferogram phase shift. Similar to spectroscopic CI designs, the first instruments of this kind used PMT's with a limited number of channels[98] while more recent designs use imaging sensors. Hybrid spatio-temporal or double spatial modulation are now the standard and have been deployed on ASDEX-Upgrade and DIII-D tokamaks.

Finally, a spectro-polarimetric CI design has been proposed for ITER, targeting impurity emission in the SOL[99]. If the strength and orientation of the magnetic field is known, then by comparing the measured and expected polarisation states of a Zeeman-split multiplet, the contribution of depolarised reflections may be unpicked and suppressed, improving the reliability of tomographic reconstructions of impurity flow and temperature. Polarisation due to Zeeman splitting and its effect on conventional CI measurements is discussed for hydrogen emission later on in Section 5.1.3.

2.5 CI hardware used in this work

Although this thesis introduces novel CI interferometer configurations, much of the hardware used here was first used in a previous investigation of impurity ion flow and temperature measurements on MAST, as documented in Scott Silburn's PhD thesis[63]. More complete component specifications can be found in that work, but important information is listed here.

Figure 2.14 shows the MAST CI instrument, whose optical layout matches the schematic from Figure 2.7. The camera and optical components are mounted to a custom anodised aluminium baseplate. Three digital single lens reflex (DSLR) lenses are used as they are optimised for wide angle imaging. The interferometer is enclosed within a temperature controlled cell (Andover Corp. model 101FRDC00-50), with endcaps modified to accommodate more optical components. The cell regulates temperatures in the range 30–60°C with a stated accuracy of $\pm 0.2^\circ\text{C}$. The birefringent crystal components used in this work are three α -BBO waveplates with nominal thicknesses of 4.6, 6.5 and 9.8 mm and measured thicknesses of 4.48, 6.35 and 9.79 mm respectively (± 0.02 mm). The two α -BBO Savart plates used in this work have nominal (total) thicknesses of 4 and 2.2 mm. All crystals have anti-reflective coatings covering the 400–700 nm wavelength range and were purchased from CLaser Photonics. The stand-alone polarisers used are Newport model 20LP-VIS.

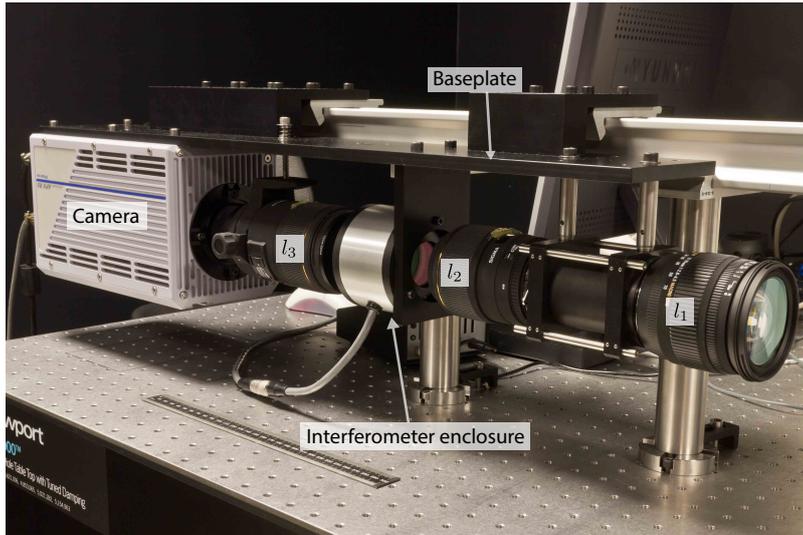


Figure 2.14: The MAST CI instrument[63]. Photo courtesy of Scott Silburn.

For the results in this thesis, we mount to the same style of baseplate and use the same interferometer enclosure (with the exception of Section 4.2.1, where the alternative setup is documented). Multiple different camera models, interferometer configurations and objective lenses (l_3) are used in this work, and are detailed at the appropriate points.

2.6 Summary

In this chapter we have covered the theory that underpins coherence imaging: coherence theory, Fourier transform spectroscopy and polarisation interferometry. By isolating a narrowband emission feature, which depends on a small number of physics parameters and whose functional form is known, the interferometer delay scanning requirements of FTS are drastically reduced. This allows for ‘snapshot’ spectral imaging with a time resolution that is limited only by the brightness of the source. In using the ‘common-path’ design of polarisation interferometers instead of the more traditional double-path design (e.g. Michelson), CI instruments are more robust to vibrations. The result is a compact, rugged and relatively inexpensive spectral imaging diagnostic for fusion plasma experiments. We have given an overview of the development of CI and some of the different applications from previous work. In the next chapter, we will see that a new technology in polarimetric imaging—pixelated polariser arrays—can be used to make CI designs even more compact and robust.

Chapter 3

Coherence Imaging with a Pixelated Phase-mask

In the first chapter, we introduced the motivation behind fusion plasma diagnostics. In the second, we covered the theory and development of CI. In this chapter, we will introduce a novel CI technique that uses a pixelated polarising filter to encode the interference pattern instead of the sinusoidal fringes used in previous investigations. This is an application of an existing technique in optics called pixelated phase mask (PPM) Interferometry, where individual camera pixels sample the coherence at different delays determined by the polariser orientation. The advantages of using this technique over previous snapshot CI designs are that it is more compact, has improved robustness against misalignment and maximises the spatial resolution of the measurement.

After first considering a PPM-CI instrument with a single fixed interferometer delay, we will introduce a multi-delay configuration that combines pixelated and sinusoidal fringe patterns to sample the interferogram at four fixed delays simultaneously. This increases the spectral resolution of the CI measurement which can, depending on the nature of the targeted plasma emission, be used to make inferences about multiple physics parameters at once, and/or to check the validity of the assumed emission model. However, the extra spectral information comes at the price of reduced spatial resolution. Algorithms for interferogram demodulation for single and multi-delay PPM-CI data are also described. Later on, in Chapter 5, these ideas are tested experimentally with measurements of electron density made on the linear plasma device Magnum-PSI.

3.1 Pixelated Phase-mask Interferometry

Information about the polarisation state of light is useful in industrial inspection, remote sensing and many other fields[68, 100]. Traditionally, it is measured by recording a series of images as a polariser rotates in front of the camera sensor. The time taken to capture the required number of images then sets a lower limit on the time resolution of the measurement. Pixelated polarisers allow for ‘snapshot’ encoding of the polarisation state of a scene using an array of polarisers with different orientations over the camera sensor. An example layout for this kind of sensor is shown in Figure 3.1(a). This method of encoding the polarisation state is analogous to how a Bayer filter achieves colour imaging. In the last few years, the key performance indicators for this type of sensor — resolution, pixel size, dynamic range and polariser extinction ratio — have all significantly improved, leading to many commercial products[101]. The use of pixelated polarisers in interferometry was first proposed in 2004 by 4D Technology[102, 103] with a technique called pixelated phase mask (PPM) Interferometry. In the image produced by a PPM instrument, the interferometer phase delay measured by each pixel depends on the orientation of that pixel’s polariser. Since the layout of polariser orientations can be specified at manufacture, arbitrary phase-masks can be created to suit the application. The advantages of this technique over one using a linear phase mask (LPM) (like that produced by a displacer or Savart plate) have already been considered in some detail[102, 104]: for example, the PPM technique is more robust than the linear phase mask (LPM) against systematic effects caused by large variations in the background intensity. PPM has found uses in optical metrology and for imaging both static[105] and dynamic[106] biological specimens. To the author’s knowledge, PPM has not yet been applied to Fourier transform spectroscopy, for measurements of plasma experiments.

3.2 Single-delay Coherence Imaging

The pixelated polariser camera used in this work is the FLIR Blackfly S (model BFS-U3-51S5), shown in Figure 3.1(b), which incorporates the Sony IMX250MZR sensor. This sensor is a complementary metal–oxide–semiconductor (CMOS) type, with $3.45\ \mu\text{m}$ pixel dimension and 2448×2048 format. Table 3.1 lists the key parameters of the camera and sensor. For the pixelated polarisers, wire-grids are formed on the chip at manufacture, with transmission axis orientations arranged in the repeating 2×2 grid pattern shown in Figure 3.1(a). To see how pixelated polarisers can be used in CI, we return to the simple polarisation interferometer shown in Figure 2.7: two polarisers, transmission axes aligned at angle $\rho = 0\ \text{rad}$, sandwiching a single

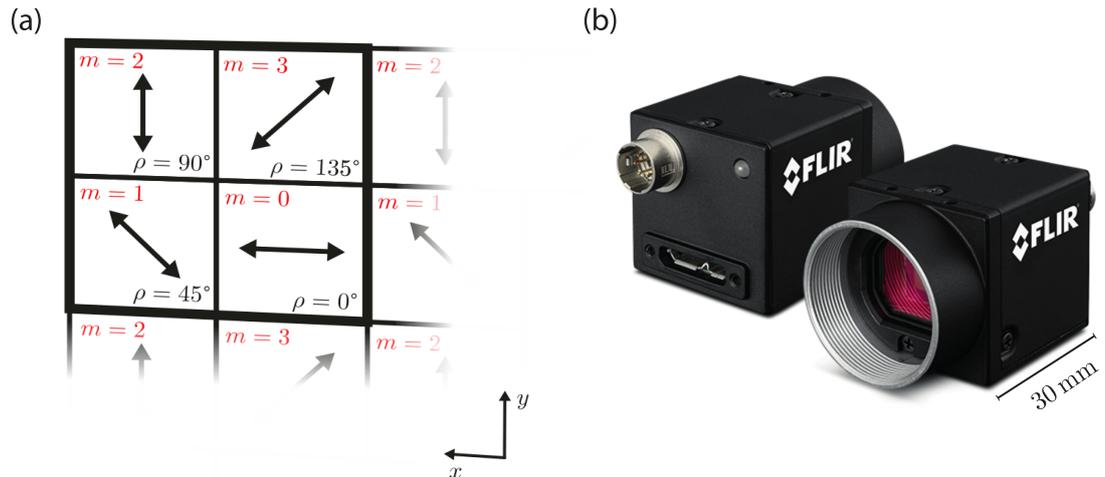


Figure 3.1: (a) Pixelated polarisers arranged in the repeating 2×2 grid pattern of the sensor used in this work (Sony IMX250MZR). The diagram views the sensor from the point of view of a light source (opposite to our convention, hence the flipped x -axis) and a single ‘superpixel’ is bordered in bold. Polariser orientation angle ρ is the orientation of the transmission axis relative to the x -axis. Pixel number m is a convenient shorthand. (b) The FLIR Blackfly S camera, photo from the manufacturer.

waveplate with fast axis oriented at $\rho = \frac{\pi}{4}$ rad. On exiting the waveplate, the light is the sum of two equal-amplitude beams, out of phase and in orthogonal polarisation states that align with the fast and slow axes of the waveplate. The final polariser then transmits equal-sized components of each beam and the interference pattern is produced at the sensor. Instead of introducing a phase shear across the image with a displacer/Savart plate, we now want the phase delay to vary according to the orientation of the final polariser. This can be achieved by introducing a quarter-wave plate (QWP) into the instrument before the final polariser, with fast axis oriented at $\rho = \frac{\pi}{2}$ rad. To calculate the interferogram irradiance we use Mueller calculus. The Mueller matrix for an ideal QWP is just a special case of the Mueller matrix for a

Camera model	FLIR Blackfly S (BFS-U3-51S5)
Sensor model	Sony IMX250MZR CMOS
Format	2448×2048
Pixel size	$3.45 \mu\text{m}$
Quantum efficiency	17% (at 435 nm)
Digitisation	12 bit
Max. framerate (full frame)	75 Hz
Polariser extinction ratio	400 (at 435 nm)

Table 3.1: Key parameters for the pixelated polariser camera used in this chapter and in Chapters 4 and 5.

linear retarder (Equation 2.3.8) with $\phi = \frac{\pi}{2}$ rad. So, for a QWP with its fast axis oriented at an angle ρ to the instrument x -axis:

$$\begin{aligned} \mathbf{M}_{\text{QWP}}(\rho) &\equiv \mathbf{R}(-\rho) \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix} \mathbf{R}(\rho) \\ &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos^2(2\rho) & \cos(2\rho)\sin(2\rho) & -\sin(2\rho) \\ 0 & \cos(2\rho)\sin(2\rho) & \sin^2(2\rho) & \cos(2\rho) \\ 0 & \sin(2\rho) & -\cos(2\rho) & 0 \end{pmatrix}, \end{aligned} \quad (3.2.1)$$

where $\mathbf{R}(\rho)$ is the matrix for anti-clockwise frame rotation by angle ρ , defined in Equation 2.3.6. The orientation of the final polariser is left general and is denoted by ρ_{P} . The total Mueller matrix for the system described is

$$\mathbf{M}_{\text{tot}} = \mathbf{M}_{\text{P}}(\rho_{\text{P}})\mathbf{M}_{\text{QWP}}\left(\frac{\pi}{2}\right)\mathbf{M}_{\text{LR}}\left(\frac{\pi}{4}, \phi(\nu)\right)\mathbf{M}_{\text{P}}(0). \quad (3.2.2)$$

Here, the waveplate imparts phase delay $\phi(\nu)$. Next, we proceed as in Section 2.3.3. If the incident light is unpolarised, its Stokes vector is

$$\mathbf{S}_{\text{in}}(\nu) = \begin{pmatrix} S_0(\nu) \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (3.2.3)$$

The Stokes vector for light at the output of the interferometer (the sensor plane) is then calculated as

$$\begin{aligned} \mathbf{S}_{\text{out}}(\nu) &= \mathbf{M}_{\text{tot}}(\nu)\mathbf{S}_{\text{in}}(\nu) \\ &= \frac{S_0(\nu)}{4} \left(1 + \cos[\phi(\nu) + 2\rho_{\text{P}}]\right) \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \end{aligned} \quad (3.2.4)$$

where the dependence of $\mathbf{M}_{\text{tot}}(\nu)$ on frequency has now been written explicitly. Assuming our sensor is insensitive to polarisation state, the measured signal is proportional to the total irradiance at the output $I_{0,\text{out}}$, which is the first Stokes parameter of $\mathbf{S}_{\text{out}}(\nu)$, integrated over all frequencies. It follows that

$$I_{0,\text{out}} = \int_{-\infty}^{\infty} \frac{S_0(\nu)}{4} \left(1 + \cos[\phi(\nu) + 2\rho_{\text{P}}]\right) d\nu$$

$$\begin{aligned}
&= \frac{I_{0,\text{in}}}{4} \left(1 + \Re \left\{ \int_{-\infty}^{\infty} g(\nu) \exp(i[\phi(\nu) + 2\rho_{\text{P}}]) d\nu \right\} \right) \\
&= \frac{I_{0,\text{in}}}{4} \left(1 + \Re \left\{ \gamma(\tau_0) \exp(2i\rho_{\text{P}}) \right\} \right).
\end{aligned} \tag{3.2.5}$$

Here, we have used $g(\nu) \equiv S_0(\nu)/I_{0,\text{in}}$ as the area-normalised spectral irradiance. Also, $\gamma(\tau_0)$ is the degree of (temporal) coherence as a function of time delay τ_0 , where $\tau_0 = \frac{\phi(\nu_0)}{2\pi\nu_0}$. Comparing this to Equation 2.3.15 we see that the introduction of the QWP results in an extra phase term $\exp(2i\rho_{\text{P}})$ in the encoded coherence. Equation 3.2.5 matches previous results from the literature[103]. We can then write

$$I_{0,\text{out}} = \frac{I_{0,\text{in}}}{4} \left(1 + \zeta(\hat{\tau}_0) \cos[\Phi(\tau_0)] \right) \tag{3.2.6}$$

Where the observed interferogram phase $\Phi(\tau_0)$ is given by

$$\Phi(\tau_0) = \phi_0(\tau_0) + \phi_{\text{D}}(\hat{\tau}_0) + \phi_{\text{asym}}(\hat{\tau}_0) + 2\rho_{\text{P}} \tag{3.2.7}$$

Here, the first three terms were defined in Section 2.2.1 and the fourth term is new. It follows that the final polariser can be rotated through π rad to scan $\Phi(\tau_0)$ over a complete 2π rad cycle. So if the final polariser is actually a grid of pixelated polarisers, whose orientations span a range of π rad in N equal steps, we sample one full cycle of the interferogram in $2\pi/N$ rad increments. Figure 3.2 shows a schematic of this instrument layout. We can now apply Equation 3.2.6 to the 2×2 repeating grid of pixelated polarisers from the Sony IMX250MZR sensor shown in Figure 3.1(a). With the pixel numbers m as shown in red, we write the orientation of each polariser's transmission axis as $\rho_{\text{P}} = m\frac{\pi}{4}$ rad. On substituting this into Equation 3.2.7 we find the observed phase is

$$\Phi(\tau_0) = \phi_0(\tau_0) + \phi_{\text{D}}(\hat{\tau}_0) + \phi_{\text{asym}}(\hat{\tau}_0) + m\frac{\pi}{2} \tag{3.2.8}$$

Therefore, each repeating 2×2 polariser grid samples the interferogram at $\frac{\pi}{2}$ rad intervals and the measured image has a pixelated fringe pattern. It can be shown that, if the waveplate orientation is $\rho = -\frac{\pi}{4}$ rad instead of $\rho = \frac{\pi}{4}$ rad, the new phase term at the end of Equation 3.2.8 switches sign.

To realise this CI design experimentally requires only minor changes to the MAST CI hardware described in Section 2.5. The QWP used here is a zero-order polymer film WP140HE based on polycarbonate (Edmund Optics model #88-252) that provides

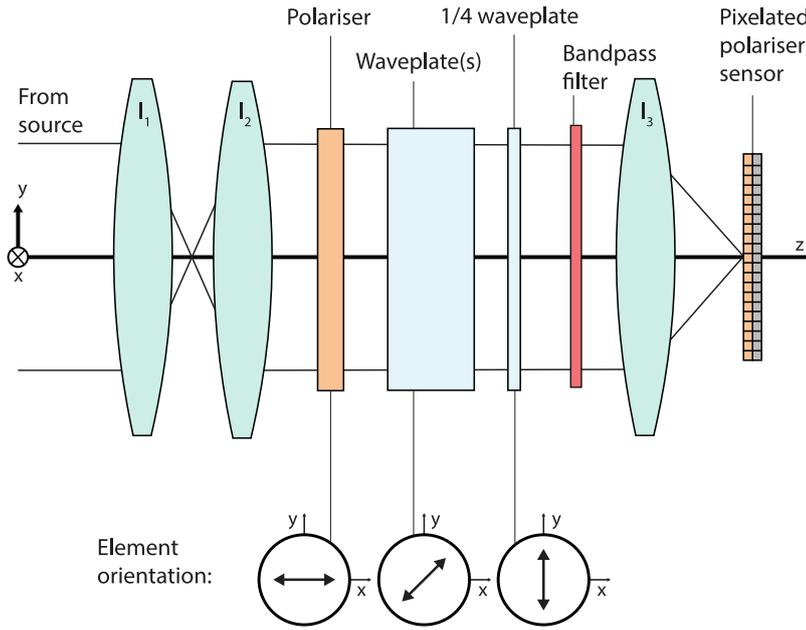


Figure 3.2: The optical set-up for a PPM-CI instrument with a single fixed delay.

a retardation of 142 ± 5 nm across the visible range¹. This film is secured to the back of the interferometer enclosure by a 3D-printed plastic ring that screws into the back of the enclosure endcap. The FLIR camera is mounted to the baseplate using a second 3D-printed component (This setup is visible in Figure 5.16(a) in Chapter 5). Figure 3.3(a) shows a raw CI image measured by an instrument in this configuration. This is a calibration image measuring an integrating sphere illuminated by a Cd discharge lamp and with a bandpass filter isolating the Cd I emission line at 467.8 nm. The delay $\phi(\nu)$ for this instrument was produced by three α -BBO waveplates (measured thicknesses 4.48, 6.35 and 9.79 mm, ± 0.02 mm) with fast axes aligned such that the constituent phase profiles sum together. Figure 3.3(c) shows four images corresponding to the four polariser orientations (and so the four $\pi/2$ rad shifts) that are interleaved to form the single raw image. Figure 3.3(b) then shows an illustration of how the interferogram is sampled at the four marked delays (on-axis). The hyperbolic fringe pattern seen in each of the four constituent images is caused by off-axis rays in a waveplate (for example, compare this to the modelled fringes in Figure 2.9(a)). The maximum (combined) waveplate thickness available was chosen here to accentuate this hyperbolic fringe pattern.

¹This non-ideal behaviour results in deviations from the modelled results at the wavelengths measured in this work ($410 \text{ nm} < \lambda < 510 \text{ nm}$) and can introduce systematic error into the demodulated phase and contrast images. Testing is currently under way to gauge and factor out the resulting systematic error from the results in Chapters 4 and 5, but unfortunately the error was caught too late for this to be included in this thesis.

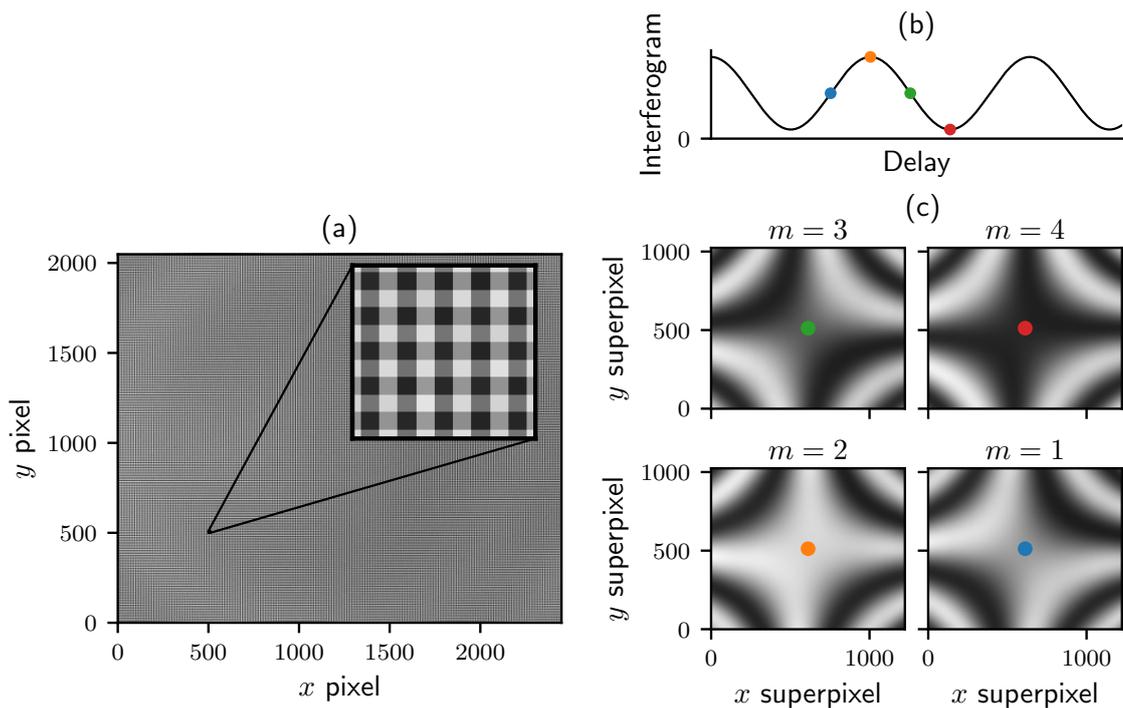


Figure 3.3: (a) A raw calibration image captured using the pixelated phase-mask CI instrument configuration (single delay). The light source is Cd I (467.8 nm) and the delay is produced by three α -BBO waveplates (4.48, 6.35 and 9.79 mm) with fast axes aligned. (b) The four orientations of the pixelated polarisers sample the interferogram at $\pi/2$ rad intervals. (c) The four interleaved images are shown here separately.

3.2.1 Interferogram Demodulation

We will now consider how the background brightness, phase and contrast can be extracted from the PPM-CI interferogram in Equation 3.2.6. If the variation in Φ is small over each 2×2 superpixel, then the four interferogram samples are $\pi/2$ rad spaced over a single interferogram cycle and there exist simple expressions for recovering the desired quantities[102]. It can be shown using Equations 3.2.6 and 3.2.8 that the input irradiance can be recovered as

$$I_{0,\text{in}} \approx \sum_{m=0}^3 I_{0,\text{out}}(m). \quad (3.2.9)$$

Similarly, the phase can be recovered as

$$\Phi^* \approx \arctan \left(\frac{I_{0,\text{out}}(m=3) - I_{0,\text{out}}(m=1)}{I_{0,\text{out}}(m=0) - I_{0,\text{out}}(m=2)} \right). \quad (3.2.10)$$

The asterisk here indicates that this recovered phase image does not include the pixelated carrier term, i.e. $\Phi = \Phi^* + m\frac{\pi}{2}$. The contrast ζ is estimated as

$$\zeta \approx \frac{1}{I_{0,\text{in}}} \sqrt{8 \sum_{m=0}^3 \left(I_{0,\text{out}}(m) - \frac{I_{0,\text{in}}}{4} \right)^2}. \quad (3.2.11)$$

These equations return demodulated profiles at superpixel resolution, a factor of two lower resolution in each dimension than the raw image. The top row of Figure 3.4 plots the demodulated phase and contrast images for the Cd I calibration PPM-CI image from Figure 3.3. The bottom row of Figure 3.4 then plots the corresponding modelled phase and contrast images, generated using Equation 2.3.20 and the nominal values for all model parameters. The mean on-axis value the modelled phase and contrast images is set to match that of the corresponding measured image for ease of comparison. The discrepancy is caused by the crystal thickness values not being accurate enough to reproduce the phase offset and the full range of contrast degradation effects not being modelled. The observed phase shape matches the model but is displaced vertically by ~ 20 superpixels. The observed contrast is roughly $\zeta \approx 0.75$ on-axis and exhibits a non-ideal hyperbolic pattern. This pattern is partially recovered in the Mueller matrix model by accounting for the non-ideal behaviour of the polymer film QWP used in this work — the delay is 142 ± 5 nm for all wavelengths. This suggests the effect would disappear were a more standard achromatic waveplate used.

The error in using demodulation Equations 3.2.9–3.2.11 due to phase variation across a superpixel has been considered in some detail elsewhere[107, 108]. Similar to the problem of reduced contrast due to finite detection area discussed in Section 2.3.7, this error can be mitigated by using field-widened birefringent components. Alternatively, Fourier demodulation[109] or synchronous demodulation[108] techniques can be used to demodulate the images instead of the simple equations listed above. For the waveplates used for single-delay PPM-CI in this work, modelling suggests that the error introduced in using Equations 3.2.9–3.2.11 to demodulate is not significant. This effect can be mitigated entirely for the single-delay PPM configuration by using the ‘synchronous demodulation’ technique described later, in Section 3.3.2.

The pixelated carrier design avoids the steep phase shear across the sensor that is fundamental to the LPM discussed in the previous chapter. Since the carrier is produced at the sensor instead of by a displacer / Savart plate, it is fixed, making the instrument more robust against crystal misalignment in the interferometer. The spatial resolution at which ζ and Φ are recovered from an LPM interferogram is anisotropic. Perpendicular to the direction of phase shear, it is set by the pixel width (assuming detector-limited operation). Parallel to the direction of phase shear, the spatial resolution is worse, depending on the fringe frequency and the width of the

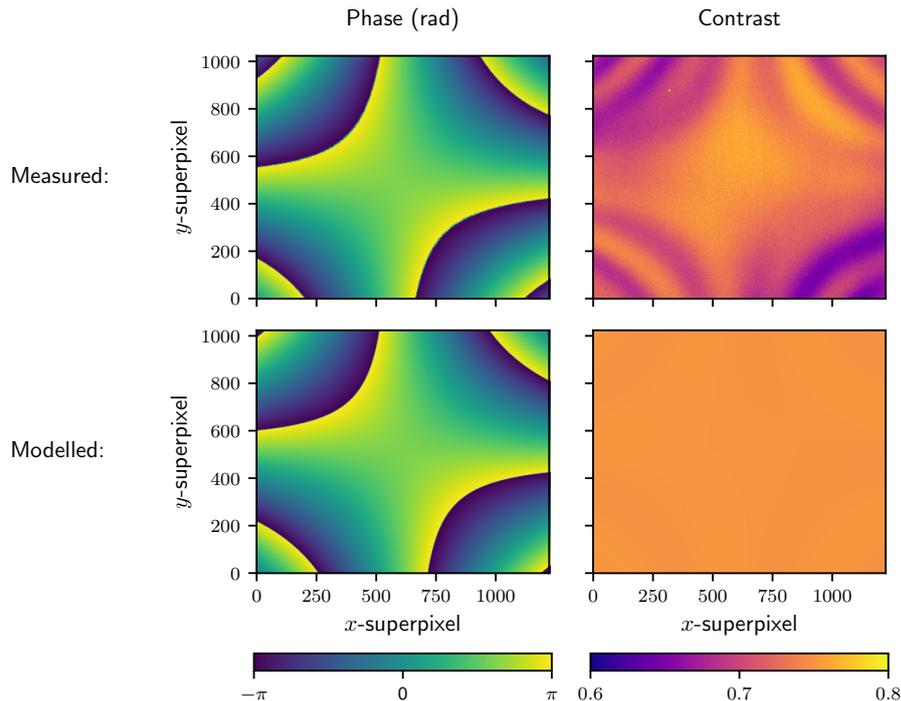


Figure 3.4: Top row: demodulated phase (wrapped) and contrast images from the raw calibration image in Figure 3.3. Bottom row: the corresponding modelled images.

Fourier-domain filter used in the demodulation process. Modelling carried out in previous work[63] of MAST CI images suggests 2 fringe periods as a typical spatial resolution in this direction. The LPM fringe period is typically chosen to be > 6 pixels, to avoid contrast degradation due to phase shear across each pixel's collection area. The PPM spatial resolution is isotropic and is 2 pixels or better in both dimensions, depending on the demodulation algorithm used[104].

3.3 Multi-delay Coherence Imaging

Sampling the coherence at a single delay is not sufficient to resolve more complicated spectral shapes, for example when several line-broadening mechanisms are in competition or multiple lines are being observed. In these situations, a multi-delay measurement can be desirable. Measuring plasma emission lines using Fourier transform spectroscopy at multiple fixed delays was first carried out over 50 years ago, when a fast-scanning Michelson interferometer with six fixed delays was used to estimate plasma temperature in Cd gas discharge tube[110]. Similarly, but much later, some of the early CI investigations utilised coherence information at multiple delays. For example, the Modulated Optical Solid-State (MOSS) spectrometer was used to diagnose asymmetric non-Maxwellian Doppler broadening in argon plasmas in the

H-1NF stellarator[64]. However in this instance, multiple single-delay measurements were taken over a series of identical plasma discharges using different delays and then combined for analysis. More recently, Peter Urlings' masters thesis tested snapshot CI designs for simultaneously encoding multiple spatial carriers in a single image, each corresponding to a different fixed delay[111]. These delays were optimised for measurement of charge exchange recombination emission, which typically contains active and passive components with different Doppler widths and Doppler shifts corresponding to the plasma conditions of the different emission regions. While the results of multi-delay CI measurements taken on the KSTAR tokamak were inconclusive, the work presented several novel instrument configurations and analysis tools. In this section we introduce a snapshot PPM-CI instrument design which simultaneously encodes four fixed delays via a mixture of linear and pixelated spatial carriers.

Perhaps the simplest way to measure multiple fixed interferometer delays is to substitute a polarising beam splitter for the front polariser in the single delay design presented in the previous chapter. This way, the light that is currently being discarded can be diverted to a second CI instrument with a different delay and a different camera. It is cheaper, and often more practical, to encode multiple fixed delays onto a single camera sensor. However, this necessarily reduces the signal to noise ratio for each delay since the maximum total irradiance at the sensor cannot exceed $I_{0,\text{in}}/2$. A multi-delay polarisation interferometer can be made by combining two single-delay interferometers in series, as in the schematic in Figure 3.5(a). Here, the analyser of the first instrument (in the red box) serves as the front polariser of the second (in the green box). If the fixed delays imparted by the two systems are $\phi_1(\nu)$ and $\phi_2(\nu)$, then the Mueller matrix for the total instrument is given by

$$\mathbf{M}_{\text{tot}} = \mathbf{M}_{\text{P}}(0)\mathbf{M}_{\text{LR}}\left(\frac{\pi}{4}, \phi_2(\nu)\right)\mathbf{M}_{\text{P}}(0)\mathbf{M}_{\text{LR}}\left(\frac{\pi}{4}, \phi_1(\nu)\right)\mathbf{M}_{\text{P}}(0), \quad (3.3.1)$$

To derive an expression for the measured interferogram, we will again proceed as in Sections 2.3.3 and 3.2. If the Stokes vector for the incident light $\mathbf{S}_{\text{in}}(\nu)$ is unpolarised, then it is given by Equation 3.2.3. We then find the Stokes vector for light at the output of the interferometer (the sensor plane) to be

$$\begin{aligned} \mathbf{S}_{\text{out}}(\nu) &= \mathbf{M}_{\text{tot}}(\nu)\mathbf{S}_{\text{in}}(\nu) \\ &= \frac{S_0(\nu)}{8} \left(1 + \cos[\phi_1(\nu)]\right) \left(1 + \cos[\phi_2(\nu)]\right) \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}. \end{aligned} \quad (3.3.2)$$

Then we take only the S_0 component and integrate over all frequencies to find the

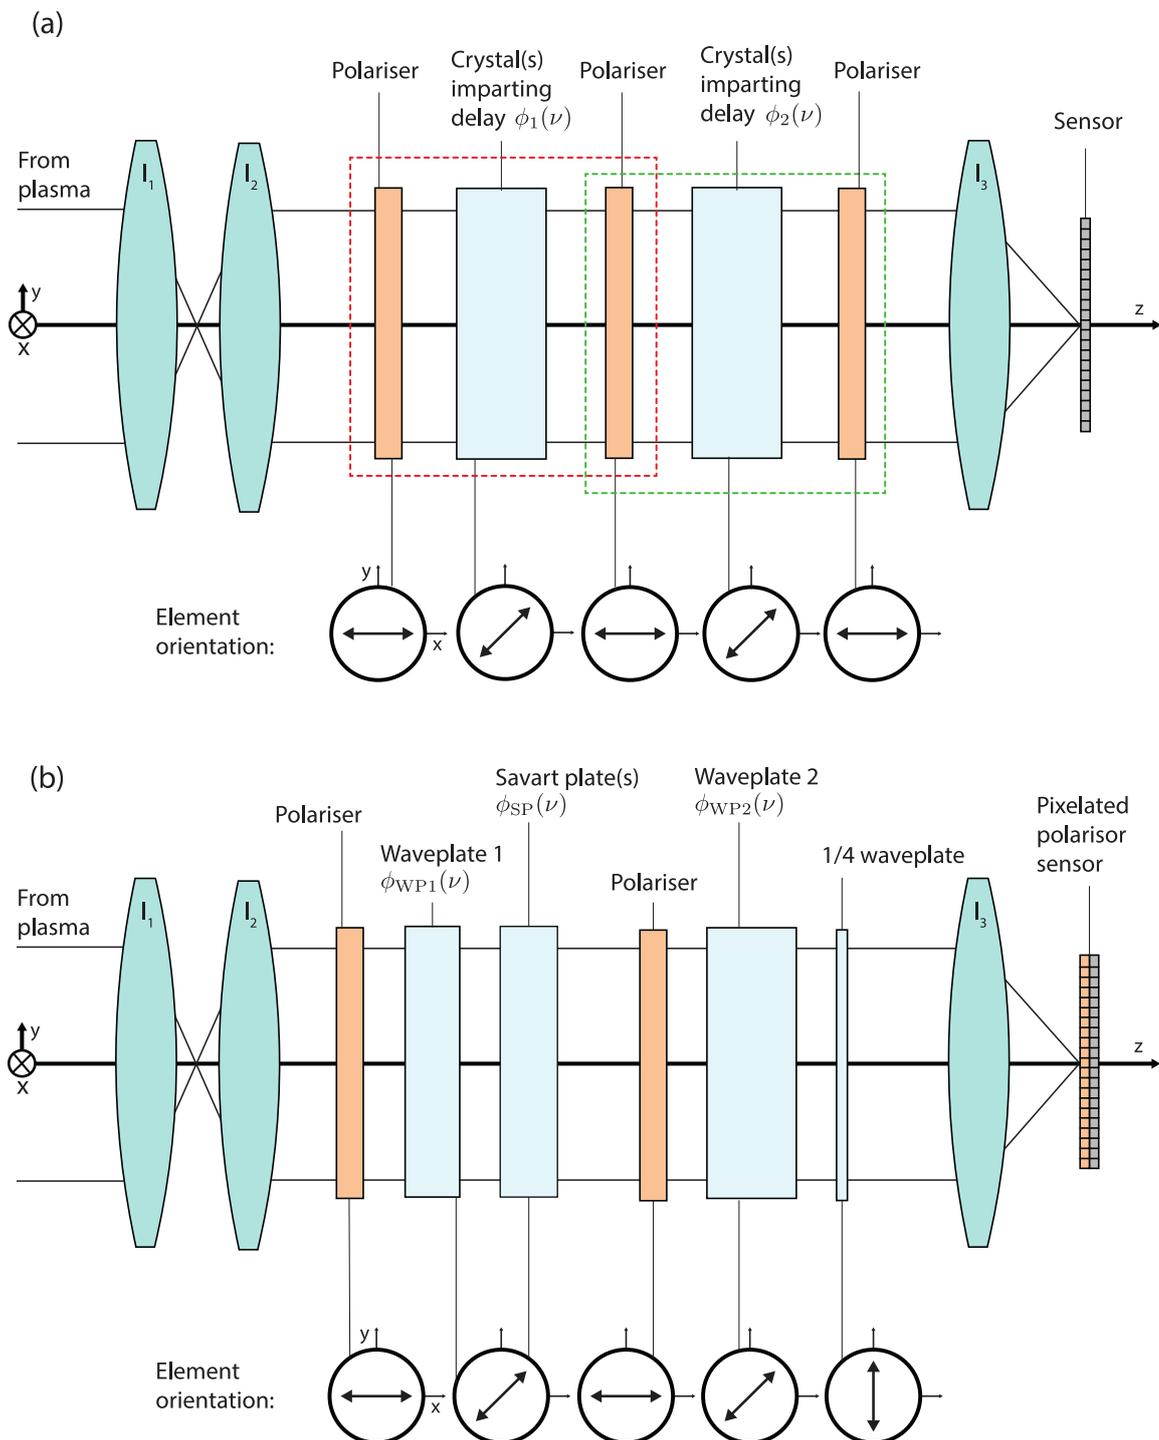


Figure 3.5: (a) Schematic diagram showing the general layout of a multi-delay CI instrument. (b) Diagram showing the layout of the pixelated/linear carrier hybrid multi-delay CI instrument introduced in this section and tested experimentally in Chapter 5. For clarity, the bandpass filter is not shown in either diagram.

total irradiance:

$$\begin{aligned}
I_{0,\text{out}} = \frac{I_{0,\text{in}}}{8} & \left(1 + \Re \left\{ \int_{-\infty}^{\infty} g(\nu) \exp[2\pi i \nu \tau_1(\nu)] d\nu \right\} \right. \\
& + \Re \left\{ \int_{-\infty}^{\infty} g(\nu) \exp[2\pi i \nu \tau_2(\nu)] d\nu \right\} \\
& + \frac{1}{2} \Re \left\{ \int_{-\infty}^{\infty} g(\nu) \exp(2\pi i \nu [\tau_2(\nu) + \tau_1(\nu)]) d\nu \right\} \\
& \left. + \frac{1}{2} \Re \left\{ \int_{-\infty}^{\infty} g(\nu) \exp(2\pi i \nu [\tau_2(\nu) - \tau_1(\nu)]) d\nu \right\} \right). \quad (3.3.3)
\end{aligned}$$

Here, we have written out the phase in full: $\phi_1(\nu) = 2\pi\nu\tau_1(\nu)$ and $\phi_2(\nu) = 2\pi\nu\tau_2(\nu)$. Applying the group delay approximation from Equation 2.2.20 to each integral, we can write $I_{0,\text{out}}$ in terms of the coherence at *four* different interferometer delays:

$$\begin{aligned}
I_{0,\text{out}} \approx \frac{I_{0,\text{in}}}{8} & \left(1 + \Re \{ \gamma(\tau_{0,1}) \} \right. \\
& + \Re \{ \gamma(\tau_{0,2}) \} \\
& + \frac{1}{2} \Re \{ \gamma(\tau_{0,2+1}) \} \\
& \left. + \frac{1}{2} \Re \{ \gamma(\tau_{0,2-1}) \} \right). \quad (3.3.4)
\end{aligned}$$

Here we have defined a shorthand for the time delay at the rest frequency imparted by each of the two systems: $\tau_{0,1} \equiv \tau_1(\nu_0) = \frac{\phi_1(\nu_0)}{2\pi\nu_0}$ and similarly $\tau_{0,2} \equiv \tau_2(\nu_0) = \frac{\phi_2(\nu_0)}{2\pi\nu_0}$. We have also defined shorthand for the sum and difference of the two delays: $\tau_{0,2+1} \equiv \tau_{0,2} + \tau_{0,1}$ and similarly $\tau_{0,2-1} \equiv \tau_{0,2} - \tau_{0,1}$. The four delays simultaneously encoded into the image, arise as follows. The polarised light exiting the second polariser in the system is the sum of two beams with delay $\tau_1(\nu)$ between them. The second waveplate then splits each of these beams again into two, introducing a further delay $\tau_2(\nu)$ between each pair. The end result after the final analyser is therefore four beams delayed by 0, $\tau_1(\nu)$, $\tau_2(\nu)$ and $\tau_1(\nu) + \tau_2(\nu)$. Combining these four beams, as in Section 2.1, encodes the coherence at the *relative* delay between each pair of beams, giving the four delays indicated in Equation 3.3.4: $\tau_1(\nu)$, $\tau_2(\nu)$, $\tau_2(\nu) + \tau_1(\nu)$ and $\tau_2(\nu) - \tau_1(\nu)$.

We can write $I_{0,\text{out}}$ in terms of the phase and contrast at each of the four delays:

$$\begin{aligned}
I_{0,\text{out}} \approx \frac{I_{0,\text{in}}}{8} & \left(1 + \zeta_1 \cos \Phi_1 \right. \\
& + \zeta_2 \cos \Phi_2 \\
& \left. + \frac{1}{2} \zeta_{2+1} \cos \Phi_{2+1} \right)
\end{aligned}$$

$$+ \frac{1}{2} \zeta_{2-1} \cos \Phi_{2-1} \Big). \quad (3.3.5)$$

Here, we have introduced a further shorthand: $\zeta_1 \equiv \zeta(\tau_{0,1})$, $\zeta_2 \equiv \zeta(\tau_{0,2})$, $\zeta_{2+1} \equiv \zeta(\tau_{0,2+1})$ and $\zeta_{2+1} \equiv \zeta(\tau_{0,2+1})$. And similarly for the phase: $\Phi_1 \equiv \Phi(\tau_{0,1})$, $\Phi_2 \equiv \Phi(\tau_{0,2})$, $\Phi_{2+1} \equiv \Phi(\tau_{0,2+1})$ and $\Phi_{2+1} \equiv \Phi(\tau_{0,2+1})$. Similar to the single-delay design, the spatial carriers must be separable from the DC component in frequency space (and also separable from one another) in order for them to be demodulated without systematic error. As for single-delay designs, the choice of fixed delay(s) here strongly depends on the application. Where do the changes in the spectral feature that relate to the physics of interest manifest most strongly in the coherence profile? Generally, we want to spread out the delays so as to maximise the amount of unique information provided by each sample of the coherence profile, which for our configuration means avoiding for example $\tau_{0,1} \approx |\tau_{0,2} - \tau_{0,1}|$.

We will introduce and discuss a four-delay instrument design of this kind in the next section and will test it experimentally in Chapter 5. Combining two polarisation interferometers in series increases the sampling of the interferogram fourfold. Using the same principle, it can be shown that combining *three* polarisation interferometers in series samples the coherence at 13 fixed delays. However, the necessary reduction in signal and spatial resolution for each carrier that this entails could limit the usefulness of this or any further extension of this kind. The instruments in [111] achieve three fixed delays by incorporating a half-wave plate.

3.3.1 Multi-delay CI with Pixelated Phase-mask

Multi-delay CI can be achieved using two linear carriers, as long as the spatial frequencies are sufficiently well separated. This kind of design has been applied in previous work to spectropolarimetric CI, specifically the ‘amplitude double spatial heterodyne’ IMSE instrument from [70]. In this work we instead focus on a multi-delay design combining in series the linear and pixelated spatial carrier single-delay CI designs from Figures 2.7 and 3.2 respectively. A schematic for this instrument configuration is shown in Figure 3.5(b). We will again use the repeating 2×2 pixelated polariser layout of the Sony IMX250MZR sensor shown in Figure 3.1(a), with the same convention for pixel number m . The total Mueller matrix for the system as shown is

$$\mathbf{M}_{\text{tot}} = \mathbf{M}_{\text{P}}(\rho_{\text{P}}) \mathbf{M}_{\text{QWP}}\left(\frac{\pi}{2}\right) \mathbf{M}_{\text{LR}}\left(\frac{\pi}{4}, \phi_2(\nu)\right) \mathbf{M}_{\text{P}}(0) \mathbf{M}_{\text{LR}}\left(\frac{\pi}{4}, \phi_1(\nu)\right) \mathbf{M}_{\text{P}}(0), \quad (3.3.6)$$

Where $\rho_{\text{P}} = m \frac{\pi}{4}$ rad is the pixel-dependent relative orientation of the final polariser, as it was in the case of the single-delay PPM configuration. Here, we maintain the 1 and 2

designations for the first and second instruments, but now $\phi_1(\nu) = \phi_{\text{WP1}}(\nu) + \phi_{\text{SP}}(\nu)$ is the phase contribution of the first waveplate and the Savart plate and $\phi_2(\nu) = \phi_{\text{WP2}}(\nu)$ is the phase contribution of the second waveplate. Of course, these phase terms also depend on the position on the sensor plane but this is left out for clarity. Proceeding as above to calculate the interferogram irradiance, if the incident light is unpolarised then the irradiance at the sensor plane can be written as a function of m as

$$I_{0,\text{out}} = \frac{I_{0,\text{in}}}{8} \left(1 + \Re\{\gamma(\tau_{0,1})\} + \Re\{\gamma(\tau_{0,2}) \exp(im\frac{\pi}{2})\} + \frac{1}{2}\Re\{\gamma(\tau_{0,2+1}) \exp(im\frac{\pi}{2})\} + \frac{1}{2}\Re\{\gamma(\tau_{0,2-1}) \exp(im\frac{\pi}{2})\} \right). \quad (3.3.7)$$

$$= \frac{I_{0,\text{in}}}{8} \left(1 + \zeta_1 \cos \Phi_1 + \zeta_2 \cos \Phi_2 + \frac{1}{2}\zeta_{2+1} \cos \Phi_{2+1} + \frac{1}{2}\zeta_{2-1} \cos \Phi_{2-1} \right). \quad (3.3.8)$$

Where the shorthand time delay symbols τ and contrast symbols ζ have the same definitions as in the previous section, but now a pixelwise term must be included in our expressions for the observed phase for three of the four carriers:

$$\begin{aligned} \Phi_1 &\equiv \Phi(\tau_{0,1}) = \phi_0(\tau_{0,1}) + \phi_{\text{D}}(\hat{\tau}_{0,1}) + \phi_{\text{asym}}(\hat{\tau}_{0,1}) \\ \Phi_2 &\equiv \Phi(\tau_{0,2}) = \phi_0(\tau_{0,2}) + \phi_{\text{D}}(\hat{\tau}_{0,2}) + \phi_{\text{asym}}(\hat{\tau}_{0,2}) + m\frac{\pi}{2} \\ \Phi_{2+1} &\equiv \Phi(\tau_{0,2+1}) = \phi_0(\tau_{0,2+1}) + \phi_{\text{D}}(\hat{\tau}_{0,2+1}) + \phi_{\text{asym}}(\hat{\tau}_{0,2+1}) + m\frac{\pi}{2} \\ \Phi_{2-1} &\equiv \Phi(\tau_{0,2-1}) = \phi_0(\tau_{0,2-1}) + \phi_{\text{D}}(\hat{\tau}_{0,2-1}) + \phi_{\text{asym}}(\hat{\tau}_{0,2-1}) + m\frac{\pi}{2} \end{aligned} \quad (3.3.9)$$

So the first delay is encoded using the linear carrier provided by the Savart plate. The second is encoded in the pixelated carrier provided by the QWP and pixelated polarisers. The third and fourth carriers, which correspond to the sum and difference delays, are encoded in the sum and difference respectively of the linear and pixelated carriers.

This type of multi-delay instrument has been realised experimentally using same hardware as the single-delay PPM-CI instrument from Section 3.2. Although, with more crystals, all of the room in the temperature-controlled cell is taken up. Table 3.2 lists for reference two different multi-delay PPM-CI delay configurations that are

	Single-delay	Multi-delay 1	Multi-delay 2
CI format	pixelated carrier	pixelated / linear carrier hybrid	pixelated / linear carrier hybrid
Waveplate 1 (mm)	4.48	6.35	4.48
Waveplate 2 (mm)	-	4.48	9.79
Savart plate(s) (mm)	-	4 + 2.2	4 + 2.2
Delay(s) at 410.2 nm (rad)	8830	3680, 8830, 12 500, 21 300	8830, 10 500, 19 300, 28 100
Delay(s) at 434.0 nm (rad)	8260	3450, 8260, 11 700, 20 000	8260, 9790, 18 100, 26 300

Table 3.2: Three configurations of pixelated phase-mask coherence imaging (PPM-CI) instruments used in this work. All waveplate and Savart plate crystals used in this work are α -BBO. The delays listed are calculated for normal incidence at the wavelengths of the two hydrogen Balmer transitions that are measured experimentally in Chapter 5.

used in the remainder of this work, with the four delays listed for each configuration at two wavelengths corresponding to hydrogen Balmer series transitions that will be the subject of investigation in Chapter 5. Figure 3.6(a) shows a raw experimental calibration image captured using a CI instrument in the ‘Multi-delay 1’ configuration. The calibration source is the same as that used in the previous section, a Cd gas discharge lamp illuminating an integrating sphere and the bandpass filter isolates Cd I emission at 467.8 nm. In Figure 3.6(b), the four interleaved images representing the different polariser orientations are shown separately. The LPM created by the Savart plate(s) is clearly present in each constituent image, as is the hyperbolic fringe pattern characteristic of a waveplate.

3.3.2 Interferogram Demodulation

The parameters of interest to be extracted from a multi-delay image are the background brightness, the four contrast terms: ζ_1 , ζ_2 , ζ_{2+1} , ζ_{2-1} , and the four phase terms: Φ_1 , Φ_2 , Φ_{2+1} , Φ_{2-1} corresponding to the four fixed interferometer delays. Unfortunately, the simple arithmetic approach for demodulating the single-delay images from Equations 3.2.9–3.2.11 is no longer valid for a multi-delay image of the kind

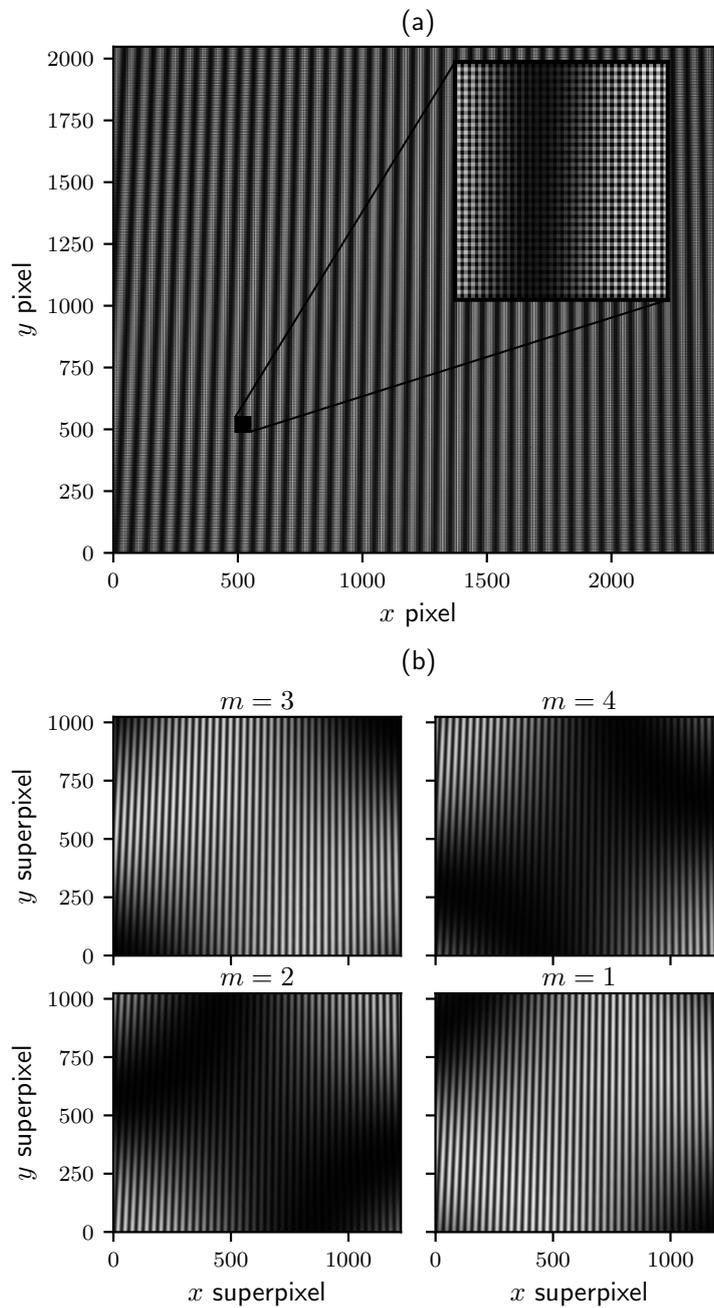


Figure 3.6: (a) A raw experimental CI calibration image captured using an instrument in the multi-delay pixelated phase-mask configuration. (b) The four interleaved images shown separately, each corresponding to a different polariser orientation on the sensor. See text for more details of the instrument setup.

shown in Figure 3.6. Instead, we will need to use Fourier techniques, which have been discussed briefly in the previous chapter and in more detail for the case of a single linear carrier in Appendix B.

We will start by looking at the Fourier transform of the pixelated spatial carrier for the more simple single-delay image. Figure 3.7(a) shows the Fourier transform of the measured single-delay CI calibration image from Figure 3.3(a). The background brightness term at the centre is labelled ‘DC’. Given the image is roughly uniform in brightness, this term is localised to the lower spatial frequencies. It can be seen that the pixelated spatial carrier, labelled Φ here, lies at *both* the horizontal and vertical Nyquist frequencies (for this layout of pixelated polariser orientations at least). The phase and contrast modulation is convolved with (and so mostly localised around) the carrier terms. This means the modulation, in places, extends beyond the Nyquist frequency (0.5 cycles per pixel) and so appears ‘folded’ back onto itself about the Nyquist frequency in frequency space. This aliasing obscures the Φ and ζ that would be recovered using the standard Fourier demodulation approach described in Appendix B: Fourier-domain windowing the carrier and then calculating the analytic signal. A different approach is needed to demodulate pixelated carrier.

‘Synchronous demodulation’² is a standard technique for demodulation of spatial carrier interference fringe patterns [112], relying on the frequency shifting property of the Fourier transform (see Appendix A). It can be stated: multiply the raw interferogram image by the known spatial carrier in complex exponential form, shifting carrier to the zero frequency and allowing the analytic signal to be extracted by an appropriate low-pass filter. Synchronous demodulation of a pixelated spatial carrier is considered in more detail and for a range of micro-polariser layouts in J. M. Padilla’s thesis[113]. To demonstrate, consider an interferogram whose phase $\Phi = \phi_c + \phi_u$ is the sum of a known spatial carrier ϕ_c and an unknown term ϕ_u that is of interest. Assume also that the brightness $I_{0,\text{in}}$ and contrast ζ are unknown and of interest. We will assume that the spatial frequency (or frequencies) of the spatial carrier are significantly higher than the spectral composition of $I_{0,\text{in}}$, ϕ_u and ζ , ensuring the carrier is isolated from the background brightness image in the frequency domain. The measured interferogram signal is then

$$I_{0,\text{out}} \propto I_{0,\text{in}} \left(1 + \zeta \cos[\phi_c + \phi_u] \right), \quad (3.3.10)$$

where the constant of proportionality depends on interferometer type. We now multiply $I_{0,\text{out}}$ by $\exp(-i\phi_c)$, which is the known spatial carrier in (negative) complex

²The name comes from the original and more popular application: electronic demodulation of incoming time signals.

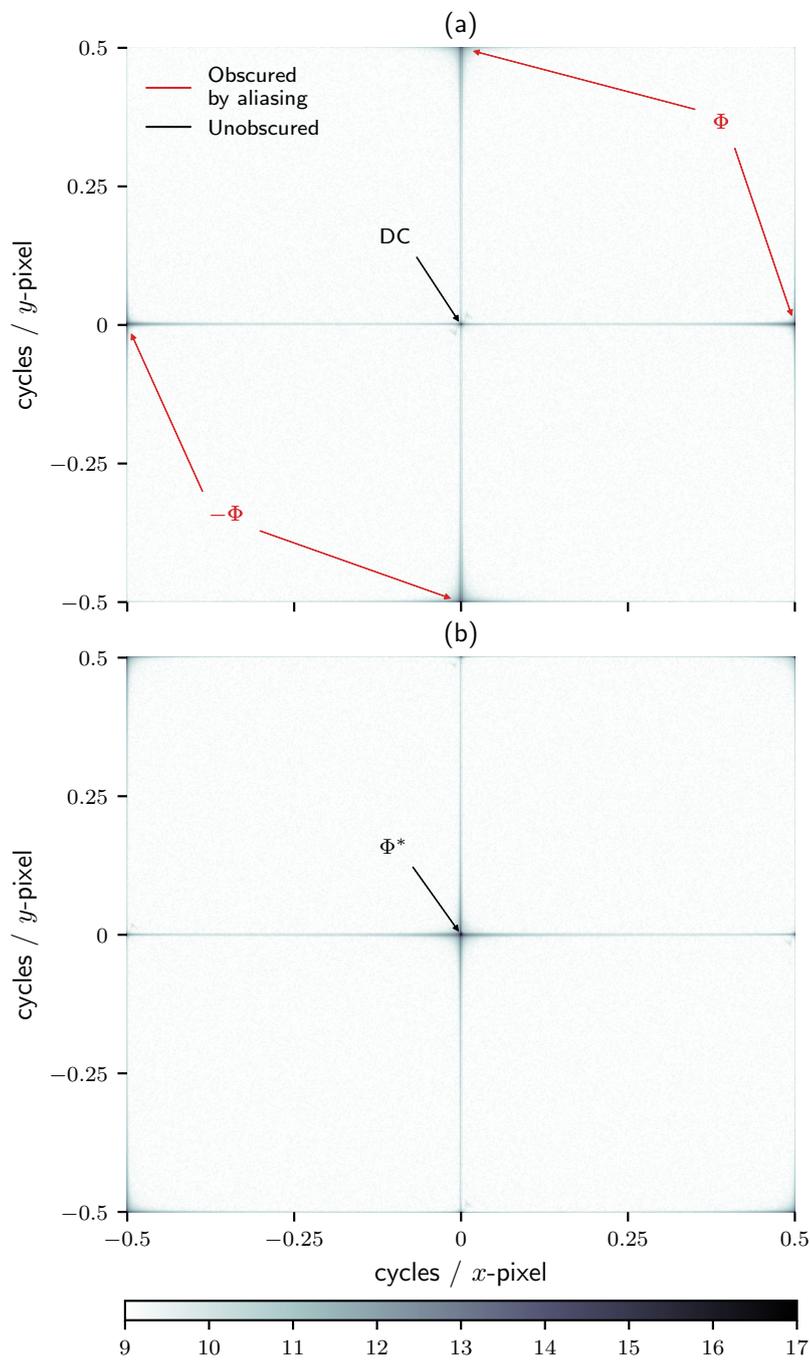


Figure 3.7: (a) Fourier transform of the single-delay CI calibration image shown in Figure 3.3. (b) Fourier transform of the same image's 'synchronous product', see text for details. Both images plot the power spectral density on a log scale and are smoothed for clarity.

exponential form. This gives us

$$I_{0,\text{out}} \exp(-i\phi_c) \propto I_{0,\text{in}} \left(\underbrace{\frac{1}{2}\zeta \exp(i\phi_u)}_{\text{Low spatial frequency}} + \underbrace{\exp(-i\phi_c) + \frac{1}{2}\zeta \exp(-i[\phi_u + 2\phi_c])}_{\text{High spatial frequency}} \right). \quad (3.3.11)$$

This complex image is called the ‘synchronous product’. By the Fourier transform shifting property, the ϕ_u term is now shifted in frequency space down to near zero. The Fourier transform of the synchronous product for the single-delay pixelated calibration image from Figure 3.3 is shown in Figure 3.7(b). It can be seen that the term at the zero frequency is $\propto I_{0,\text{in}}\zeta \exp(i\phi_u)$, labelled with the shorthand Φ . This term, now free of any aliasing effects and already in complex form, can be filtered and processed in the standard way to obtain ζ and ϕ_u .

Next, we show how this technique can be used to demodulate multi-delay PPM-CI images, using the measured calibration image from Figure 3.6(a) as an example. The Fourier transform of this calibration image is shown in Figure 3.8(a). The DC term is labelled and the pixelated carrier terms, which are common to both single and multi-delay PPM images are here labelled Φ_2 . The other three carrier terms are present too: Φ_1 due to the linear carrier only and the sum and difference terms Φ_{2+1} and Φ_{2-1} , also labelled. As a first step, the DC and linear carrier term Φ_1 can be extracted using conventional Fourier demodulation. The other three carriers are obscured by aliasing in the same way as the pixelated carrier in the single-delay example, and so synchronous demodulation is required. Figure 3.8(b) shows the Fourier transform of the synchronous product for the multi-delay calibration image, showing that the three carriers have been shifted to the lower spatial frequencies and are no longer obscured by aliasing. Standard Fourier demodulation techniques can then be used to extract the phase and contrast from the Φ_2 , Φ_{2+1} and Φ_{2-1} carriers.

This can be seen by evaluating the synchronous product for the irradiance at the output of the multi-delay instrument from Equation 3.3.8. Written out in full, this is

$$\begin{aligned} I_{0,\text{out}} \exp(-im\frac{\pi}{2}) = \frac{I_{0,\text{in}}}{8} & \left(\exp(-im\frac{\pi}{2}) \right. \\ & + \frac{\zeta_1}{2} \left[\exp(i[\Phi_1 - m\frac{\pi}{2}]) + \exp(-i[\Phi_1 + m\frac{\pi}{2}]) \right] \\ & + \frac{\zeta_2}{2} \left[\exp(i\Phi_2^*) + \exp(-i[\Phi_2^* + m\pi]) \right] \\ & + \frac{\zeta_{2+1}}{4} \left[\exp(i\Phi_{2+1}^*) + \exp(-i[\Phi_{2+1}^* + m\pi]) \right] \\ & \left. + \frac{\zeta_{2-1}}{4} \left[\exp(i\Phi_{2-1}^*) + \exp(-i[\Phi_{2-1}^* + m\pi]) \right] \right). \end{aligned}$$

Here, we use of the same convention as in Section 3.2.1 whereby an asterisk indicates that the phase does not include the pixelated term, i.e. $\Phi_{2+1} = \Phi_{2+1}^* + m\frac{\pi}{2}$. Figure 3.8

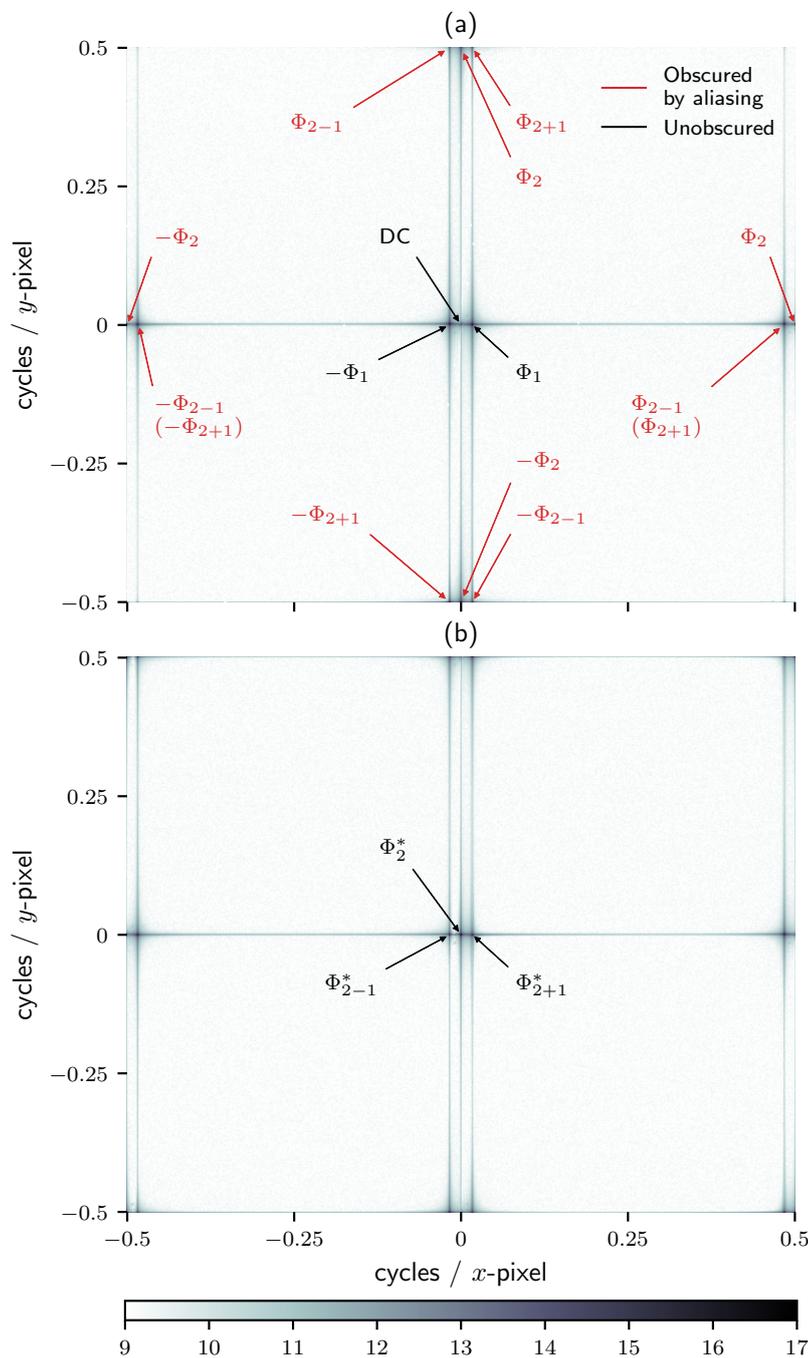


Figure 3.8: (a) Fourier transform of the multi-delay calibration image shown in Figure 3.6. (b) Fourier transform of the same image's 'synchronous product', see text for details. Both images plot the power spectral density on a log scale and are smoothed for clarity.

shows the phase and contrast profiles for each of the four carriers, extracted from the multi-delay calibration image shown in Figure 3.9(a) using the demodulation scheme described.

Rewriting Equation 3.3.12, listing only the terms with low spatial frequency (i.e. the terms that are not a function of m) gives us:

$$I_{0,\text{out}} \exp(-im\frac{\pi}{2}) = \frac{I_{0,\text{in}}}{8} \left(\frac{\zeta_2}{2} \exp(i\Phi_2^*) + \frac{\zeta_{2+1}}{4} \exp(i\Phi_{2+1}^*) + \frac{\zeta_{2-1}}{4} \exp(i\Phi_{2-1}^*) + \text{High spatial frequency terms} \right). \quad (3.3.12)$$

The contrast and phase images for these three terms can be extracted from the synchronous image using Fourier-domain filtering.

Finally, a numerical test of the synchronous demodulation procedure for multi-delay CI images was carried out. First, synthetic spectra corresponding to realistic plasma conditions were modelled for three different values temperature (0.1, 1 and 5 eV) and flow speeds (0, 15 and 30 km/s). Doppler broadening is assumed to be the only broadening mechanism and the emission line is hydrogen Balmer- γ at 434.0 nm. These spectra are plotted in Figure 3.10(a). The forward model described in Section 2.3.5 is then used to generate synthetic multi-delay calibration CI images at full resolution for each of the three modelled spectra, for the multi-delay instrument configuration shown schematically in Figure 3.5(b). The instrument's fixed delays correspond to those of the 'Multi-delay 1' configuration from Table 3.2. A cropped portion of each synthetic CI image is shown in Figure 3.10(b). The phase shift between images is visible by eye for the linear carrier and the contrast drop is visible for both the linear and pixelated carriers. The demodulation procedure is then used to extract phase and contrast images for each of the four fixed delays and the mean values of all the demodulated parameters are taken within a 200×200 central region of the images. Figure 3.10(c) then compares the demodulated contrast values with the analytical contrast profiles. Similarly, Figure 3.10(d) compares the demodulated and analytical phase shifts, relative to the stationary spectrum. The demodulated points track the analytical profiles well.

3.4 Non-ideal Quarter-wave plate Effects

In this final section we will consider how non-ideal quarter-wave plate performance affects the observed interferogram. Figure 3.11(a) plots the retardance of the zero-

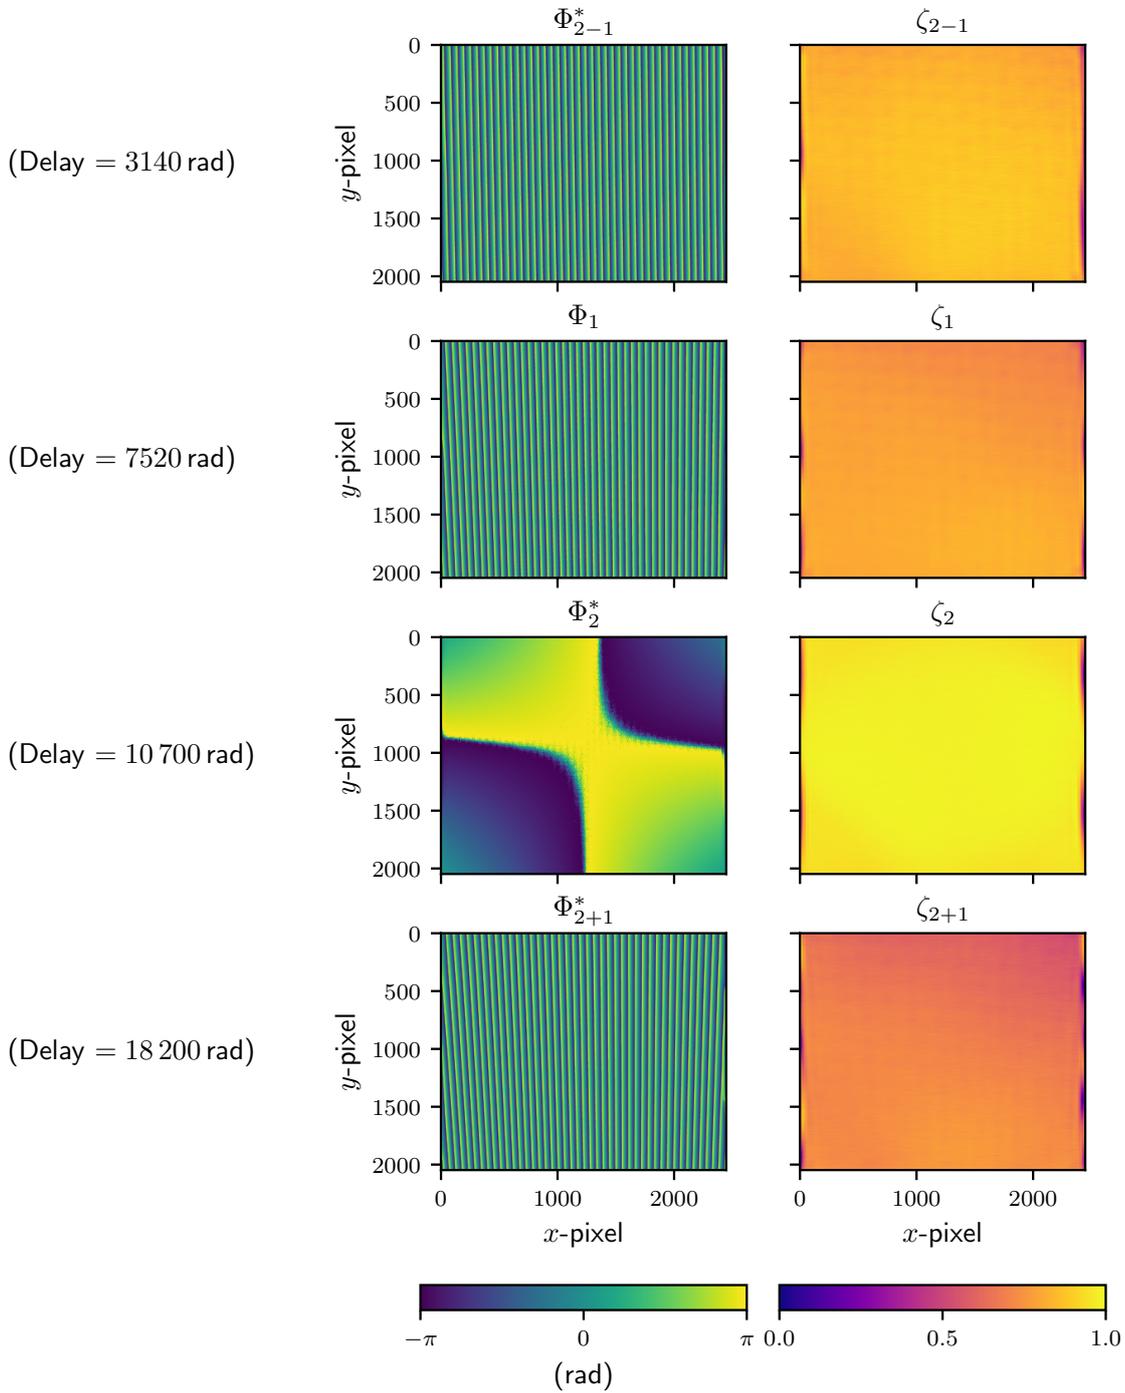


Figure 3.9: Demodulated phase (modulo 2π) and contrast images from the raw calibration image shown in Figure 3.6, captured using an instrument in the multi-delay pixelated phase-mask configuration..

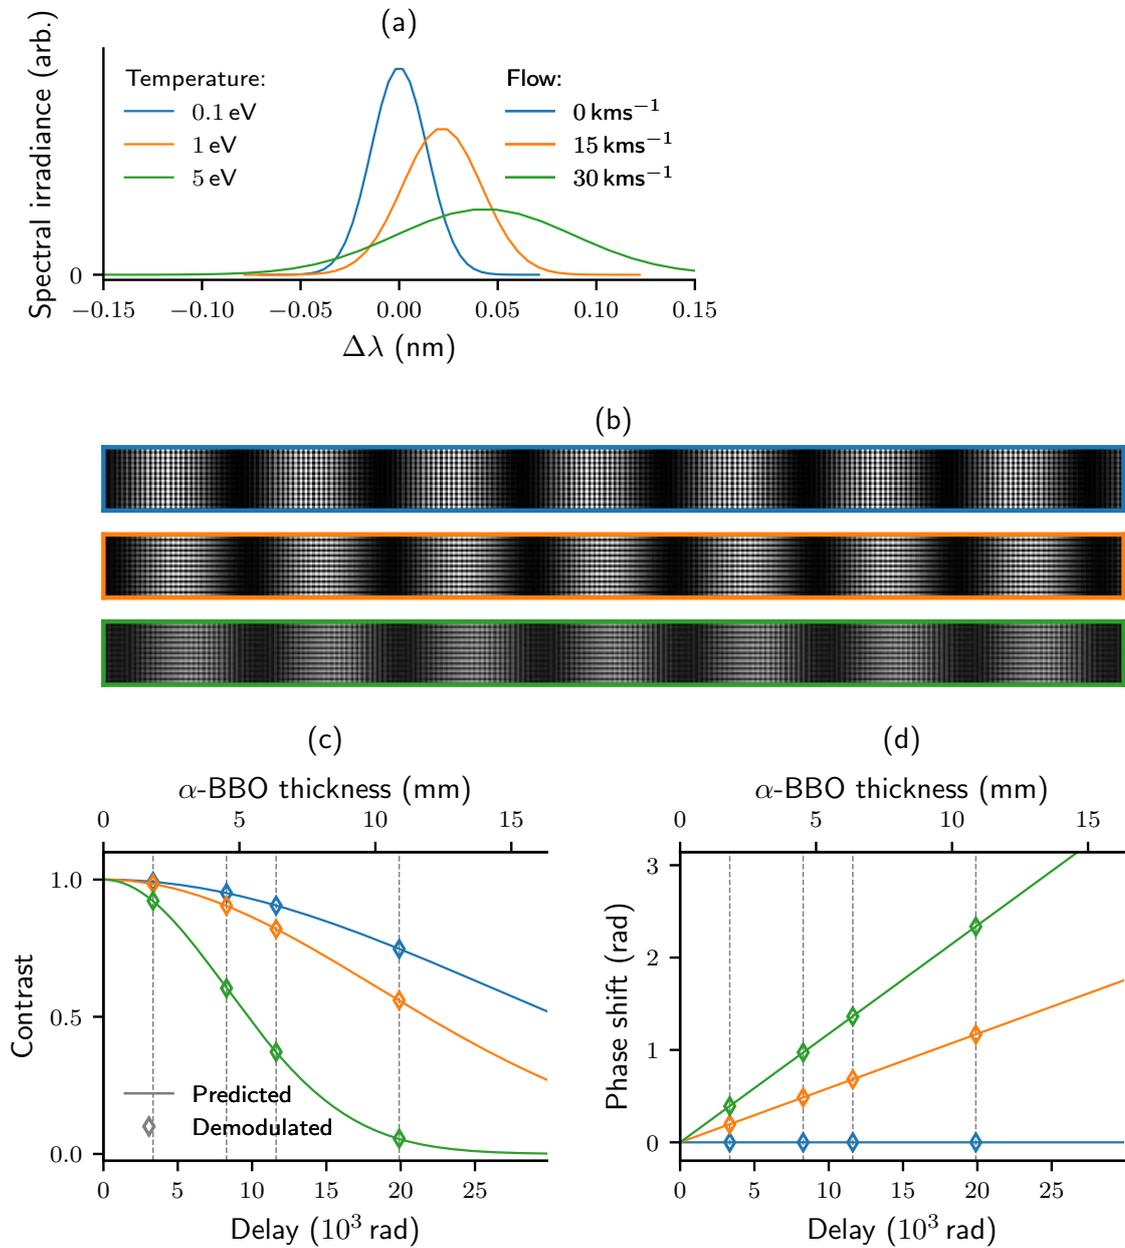


Figure 3.10: Testing the demodulation procedure for a multi-delay CI instrument of the pixelated / linear Carrier hybrid type introduced in this section. (a) Modelled lineshapes for H_γ emission (434.0 nm) with different emitter temperatures (widths) and flows (shifts). (b) Synthetic CI images (modelled at full resolution and realistic image noise, cropped here) for an instrument in the PPM multi-delay 1 configuration (see Table 3.2) observing the modelled emission with a uniform brightness. Predicted and demodulated profiles are plotted for (c) contrast and (d) phase shift.

order polymer QWP used in this work as a function of wavelength, according to the manufacturer. The plot also indicates some key wavelengths used in this work. The $5 \rightarrow 2$ and $6 \rightarrow 2$ hydrogen Balmer transitions are at 434.0 nm and 410.2 nm respectively, and are measured in Chapter 5. The remaining three lines are the Cd lines measured in Chapter 4. By combining this data with the Mueller matrix for a linear retarder, we can model PPM-CI images and gauge the resulting non-ideal effects.

The single-delay PPM-CI configuration described in Section 3.2 is used as the basis for this modelling, with the delay provided by an α -BBO waveplate with thickness 4.48 mm. Full resolution CI images were generated numerically for each of the five wavelengths marked in Figure 3.11, using both an ideal and non-ideal QWP. No noise was added to these images, so as to highlight the systematic effects. The demodulation algorithm described in Section 3.2.1 was then used to extract the contrast and phase from each image. The deviations between non-ideal and ideal behaviour are plotted for phase (additive contribution) and contrast (multiplicative contribution) in Figures 3.11(b) and 3.11(c) respectively. Only the profiles corresponding to the central column slice through the images are plotted here. The contribution of the non-ideal QWP is small for the phase profiles. For reference, a phase shift of 0.005 rad corresponds to a flow equivalent of < 200 m/s for this delay and at a wavelength of 500 nm. This is smaller than we are currently able to measure (see Chapter 4). For the contrast, our typical measurement uncertainty in Chapter 5 is ~ 0.02 , so the non-ideal effect is only significant for the $7 \rightarrow 2$ hydrogen Balmer line at 410.2 nm. Repeating this modelling with the multi-delay PPM-CI configuration described in Section 3.3, we find the same results for the three delays with the pixelated encoding and, as might be expected, no effect on the LPM term.

The non-ideal effects described above can be accounted for if the retardance vs. wavelength curve is well known. Start by defining the Mueller matrix for the non-ideal QWP:

$$\mathbf{M}_{\text{QWP}}^*(\rho) \equiv \mathbf{M}_{\text{LR}}\left(\rho, \frac{\pi}{2}x\right), \quad (3.4.1)$$

where x is an arbitrary multiplicative deviation of the delay from its ideal. By substituting $\mathbf{M}_{\text{QWP}}^*(\rho)$ in place of $\mathbf{M}_{\text{QWP}}(\rho)$ in Equation 3.2.2, and proceeding as in Section 3.2, we can write the irradiance at the sensor plane as

$$I_{0,\text{out}} = \frac{I_{0,\text{in}}}{4} \left(1 + \zeta_m \zeta(\hat{\tau}) \cos[\Phi(\tau)]\right). \quad (3.4.2)$$

Here, we have defined a new m -dependent contrast term that arises due to the non-

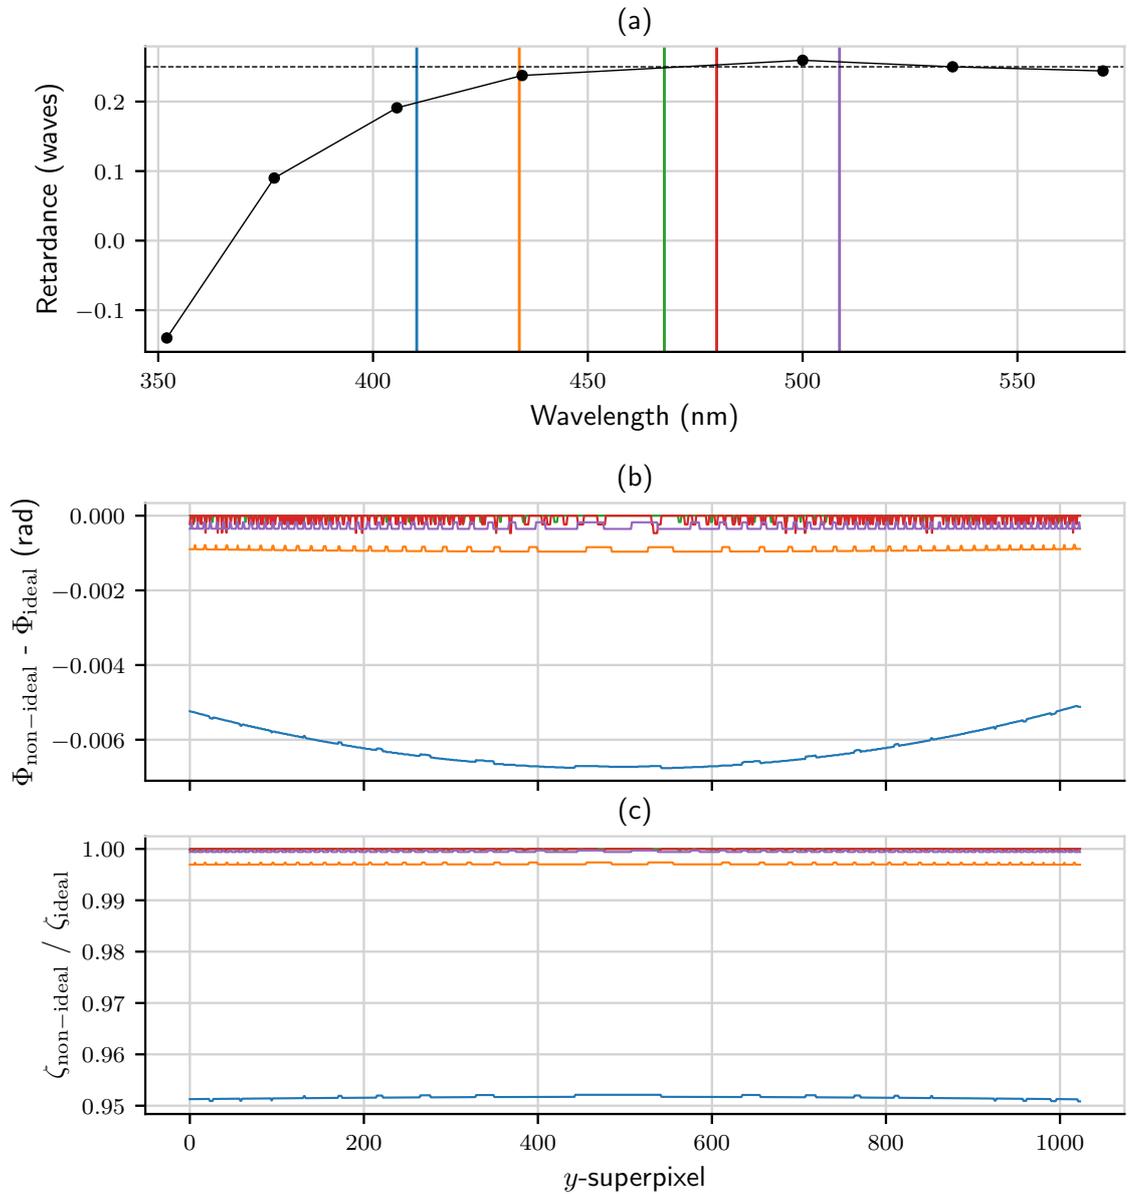


Figure 3.11: (a): Retardance as a function of wavelength for the quarter-wave plate used in this work, as supplied by the manufacturer. Vertical lines mark some of the key wavelengths used in this work. (b) and (c): predicted error introduced into the phase and contrast images respectively, where line colour corresponds to the wavelengths marked in (a). The profiles correspond to a single column through the image centre. See text for details of the modelled CI instrument configuration.

ideal behaviour of the QWP:

$$\zeta_m = \begin{cases} 1 & \text{if } m = 0, 2 \\ \sin\left(\frac{\pi}{2}x\right) & \text{if } m = 1, 3 \end{cases} \quad (3.4.3)$$

Only a small change needs to be made to the synchronous demodulation algorithm described in Section 3.3.2 to eliminate this non-ideal effect. When multiplying the image by the known phase-mask in exponential form, the reciprocal of ζ_m must also be included. So the multiplication factor becomes:

$$\frac{1}{\zeta_m} \exp\left(im\frac{\pi}{2}\right). \quad (3.4.4)$$

This correction has not been applied to any of the results in this thesis. The modelling presented in Figure 3.11 suggests that the only results that will be meaningfully affected by the non-ideal QWP behaviour is the measurement of the $7 \rightarrow 2$ hydrogen Balmer line contrast in Chapter 5.

3.5 Summary

In this Chapter we have introduced coherence imaging using a pixelated phase-mask. We have seen that this instrument design has a number of advantages over the existing linear carrier design: improved robustness, compactness and spatial resolution. We have also considered a hybrid multi-delay CI instrument design that combines pixelated and linear spatial carriers to encode the coherence profile into the interferogram at four fixed delays. The demodulation procedure for extracting the phase and contrast images from the multi-delay CI interferograms is slightly more involved than for a standard single linear carrier. A routine is introduced, demonstrated on an experimental calibration image and also tested numerically using synthetic images generated using an instrument model.

Chapter 4

Calibration of Flow Measurements using Gas-discharge Lamps

Phase shifts in the measured CI fringe pattern encode the flow velocity of the emitting ion species. Calibration of the zero-point of this flow measurement requires a set of unshifted reference fringes measured at the rest-frame wavelength of the targeted spectral feature. When discussing calibration in Chapter 2, we assumed that we had to hand a coherent laboratory source at the correct wavelength, but this is often not the case. Tuneable lasers are fit for this purpose but can be expensive, particularly if the flexibility of a large wavelength range is required. In this chapter, we show that the phase calibration necessary for flow CI measurements can be obtained using standard gas-discharge lamps, analogous to the wavelength calibration of a slit-coupled spectrograph. The method works by measuring separate phase images at a small number of spectral lines and fitting a multi-parameter instrument model to the data. This model then permits extrapolation of the measured calibration phase over small wavelength ranges ($\Delta\lambda \sim 5$ nm) to the required wavelength while maintaining a high calibration accuracy (± 1 km/s compared to anticipated flows of up to 30 km/s). This method significantly reduces hardware costs by forgoing the need for a tuneable laser. Constraining the instrument model at just a few, sparse wavelengths is made possible using a wrapped distribution to account for the ‘ 2π ambiguity’ in the measured phase data, which is inherently wrapped in the interval $(-\pi, \pi]$ rad[114]. This method will be used to calibrate CI measurements of C and He impurity ion flows on MAST-U when it begins operations later this year.

This chapter begins with an overview of CI phase calibration methods used in previous work. We then introduce the tools necessary to fit to wrapped data and to model CI instrument dispersion, using a Bayesian framework. Next, we measure CI phase images over a $\Delta\lambda \sim 3$ nm wavelength range using a tuneable laser and

show that accurate phase extrapolation over this range is possible by constraining an instrument model at just two closely-spaced ($\Delta\lambda = 0.5$ nm) wavelengths. Finally, we measure CI phase images over a $\Delta\lambda \sim 40$ nm wavelength range using cadmium and zinc gas-discharge lamps and bandpass interference filters. We show that by fitting a model to phase measurements made at four wavelengths over this range, we can interpolate phase inside or extrapolate phase outside of this range by $\Delta\lambda \sim 5$ nm with high accuracy. It was necessary to adjust the available Sellmeier coefficients describing dispersion in the birefringent crystal material by up to 4% to fit the observed measurements over $\Delta\lambda \sim 40$ nm.

4.1 Background

The impurity emission lines previously targeted for CI flow measurement on MAST[51, 63] were C III ($\lambda_0 = 464.8810$ nm), He II ($\lambda_0 = 468.5701$ nm) and C II ($\lambda_0 = 514.1842$ nm)¹. Carbon is naturally present in the plasma due to the graphite plasma-facing components while helium is naturally present due to the glow discharge wall conditioning that takes place between plasma experiments[51, 63]. Other, more recent, CI investigations have targeted just the C III and He II lines[52, 78, 90]. These two lines will be the focus of this chapter. We saw in Section 2.2.3 that, when observing an inhomogeneous source, the observed CI phase can be written

$$\begin{aligned}\check{\Phi} &= \phi_0 + \check{\phi}_D \\ &= \phi_0 + \hat{\phi}_0 \frac{1}{I_{0C}} \int_L I_0(\mathbf{r}) \mathbf{v}(\mathbf{r}) \cdot d\mathbf{l},\end{aligned}\quad (4.1.1)$$

provided the interferometer delay is chosen appropriately. The first term in Equation 4.1.1 is the fringe phase at the unshifted wavelength λ_0 and the second term arises due to Doppler shifts along the line of sight L . A third phase term due to asymmetry in the observed spectral line or feature was introduced in Section 2.2 but it is ignored here for simplicity. The two calibration parameters in flow CI are then the phase delay ϕ_0 and the group delay $\hat{\phi}_0$, which calibrate the flow measurement's zero point and magnitude respectively. From its definition in Equation 2.2.17, $\hat{\phi}_0$ can be rewritten in terms of wavelength:

$$\hat{\phi}_0 \equiv \nu_0 \left. \frac{\partial \phi}{\partial \nu} \right|_{\nu_0} = -\lambda_0 \left. \frac{\partial \phi}{\partial \lambda} \right|_{\lambda_0} \quad (4.1.2)$$

¹Quoted wavelength values are the weighted mean over the fine structure components, calculated using component wavelength, transition probabilities and statistical weights from NIST[115]. See Section 3.2 of [69] or Section 3.1 of [63] for more detail.

so that, for small wavelength shifts $\Delta\lambda \equiv \lambda - \lambda_0$, we can write²:

$$\phi(\lambda) \approx \phi_0 - \hat{\phi}_0 \frac{\Delta\lambda}{\lambda}. \quad (4.1.3)$$

Both ϕ_0 and $\hat{\phi}_0$ vary across the field of view and so both represent calibration *images*.

The most direct way of measuring ϕ_0 is to illuminate the CI instrument with an extended, quasi-monochromatic source at λ_0 . This modulo 2π measurement of ϕ_0 is sufficient since the flow information is encoded in phase shifts about ϕ_0 . To measure $\hat{\phi}_0$, the source is tuned about λ_0 to mimic the anticipated Doppler shifts and the phase gradient with respect to wavelength is inferred. Both of these measurement procedures have been realised using a tunable laser source and high-resolution wavemeter[63, 77, 91, 92]. The sensitivity of the calibration presents difficulties in practice: small changes in crystal temperature or optical alignment that occur between measuring the calibration images and measuring the plasma can introduce large systematic errors. Inference of $\mathbf{v}(\mathbf{r})$ is far more sensitive to absolute error in the calibration parameter ϕ_0 than it is to error in $\hat{\phi}_0$. This is because $\mathbf{v}(\mathbf{r}) \ll c$, meaning that the Doppler phase shifts are orders of magnitude smaller than $\hat{\phi}_0$. It follows that regular calibration of ϕ_0 is required to track environmental changes while a single measurement of $\hat{\phi}_0$ is sufficiently accurate[63]. For this reason, our focus here is on ϕ_0 first and we will return to $\hat{\phi}_0$ later in the chapter. Previous work[63] has found that the sensitivity of the ϕ_0 calibration to crystal temperature is (12 ± 1) km/s/°C at 468 nm. This makes regular calibration and/or interferometer temperature control very important.

The ideal calibration source for measuring ϕ_0 is a bright, coherent emitter at λ_0 . Unfortunately, the standard gas-discharge lamps used to calibrate slit-coupled spectrographs are not hot enough to emit intensely at the wavelengths of interest in the SOL. In the absence of an appropriate lab source, previous work has estimated ϕ_0 using emission from the tokamak plasma itself, assuming that Doppler shifts are negligible either viewed along a radial sight line[63] or during the breakdown phase of the discharge[82]. Samuëll et al.[91, 92] at DIII-D were the first to demonstrate an absolute phase calibration that is independent of the plasma being measured. They did this with a diode laser tunable to the targeted C III and He II wavelengths, and over the few nanometres in between. More recently, Gradic et al.[77] have demonstrated that a widely tunable, continuous-wave laser can calibrate almost any targeted plasma line in the 450–650 nm range. Despite these developments, calibration of flow CI using gas-discharge lamps remains desirable to some labs, since it significantly reduces hardware complexity and cost. To this end, Samuëll et al. demonstrated that a calibration phase image can be measured at a nearby lamp line and then extrapolated

²In this chapter we will work in terms of wavelength instead of frequency.

in wavelength over 3 nm to the targeted plasma line — while maintaining accuracy better than ± 1 km/s. However, this work relied on tunable laser measurements, using ~ 50 phase images made over the 3 nm wavelength range to constrain an instrument model [91]. In this chapter, we will build upon this existing work, modifying the fitting tools such that emission lines from gas-discharge sources can also be used to constrain the instrument model itself.

4.1.1 Fitting to Wrapped Data

Measurements of interferogram phase are inherently wrapped into the interval $(-\pi, \pi]$ rad, introducing ambiguity and discontinuities. Unwrapping to recover an unambiguous, continuous signal is a common inverse problem in interferometry[55]. For a set of wrapped samples from a 1-D function $\{f_x\}$ to be unwrapped naively, the inequality $-\pi \leq (f_x - f_{x-1}) < \pi$ must be satisfied for all x ('Itoh's smoothness condition'[116]). This is violated when the sampling rate is too low or when measurements are too noisy. Since flow CI instruments are designed to measure Doppler shifts ($\lesssim 0.05$ nm) that are far smaller than the typical separation of calibration lines from gas-discharge lamps ($\gtrsim 1$ nm), separate phase measurements made at these lines inevitably violate the smoothness condition and cannot be naively unwrapped. One way of moving forward is to fit the instrument model directly to the wrapped data points, accounting for the '2 π ambiguity' by using a wrapped distribution for each data point's likelihood function³.

For a data set \mathcal{D} that is well described by a model $f(\boldsymbol{\theta})$ with parameters $\boldsymbol{\theta}$, the likelihood function $\mathcal{L}(\mathcal{D} | \boldsymbol{\theta})$ describes the plausibility of a particular set of parameter values. In the absence of any prior knowledge, the most plausible values are those which maximize $\mathcal{L}(\mathcal{D} | \boldsymbol{\theta})$. Consider a single wrapped data point \mathcal{D}_i whose measurement error is normally distributed with zero mean and variance σ_i^2 . The 2π ambiguity in \mathcal{D}_i is formally accounted for in the likelihood function using a *wrapped* normal distribution:

$$\mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta}) = \frac{1}{\sigma_i \sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{1}{2} \left[\frac{\mathcal{D}_i - f_i(\boldsymbol{\theta}) + 2\pi k}{\sigma_i}\right]^2\right). \quad (4.1.4)$$

The sum to infinity here makes the function difficult to work with numerically, so a closed-form approximation called the von Mises distribution is typically used instead[114]:

$$\mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta}) \approx \frac{\exp(s_i \cos[\mathcal{D}_i - f_i(\boldsymbol{\theta})])}{2\pi J_0^m(s_i)}. \quad (4.1.5)$$

³Model fitting in this chapter uses Bayesian parameter estimation, see Appendix C for a brief introduction.

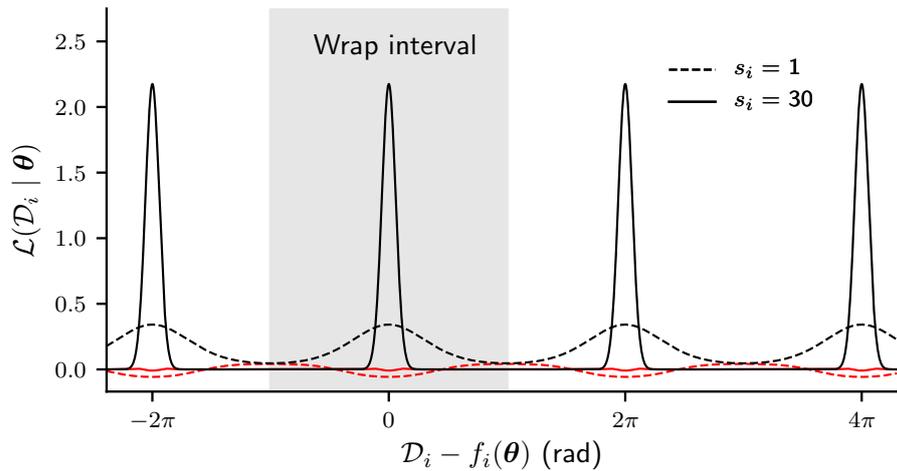


Figure 4.1: The von Mises distribution approximates the likelihood function $\mathcal{L}(\mathcal{D}_i | \theta)$ of a single, wrapped data point \mathcal{D}_i drawn from a normal distribution. The black lines show the von Mises distributions for two different scale parameters s_i while the red lines show the corresponding approximation error when compared to the wrapped normal distribution.

Here, $s_i = 1/\sigma_i^2$ is the scale parameter and $J_0^m(s_i)$ is the modified Bessel function of order 0. Figure 4.1 plots the von Mises distribution for two different values of s_i . Also plotted is the corresponding deviation from the wrapped normal distribution, showing that the approximation is more accurate for larger s_i . The infinitely multimodal distribution for a single data point demonstrates the 2π ambiguity. When more (statistically independent) data points are measured, the total likelihood function is the product of each point's individual likelihood function. So, generally, by making measurements at different values of some independent variable, the multiple modes of $\mathcal{L}(\mathcal{D} | \theta)$ become more and more separated in parameter space. In a well-designed experiment, enough data is collected such that the modes corresponding to incorrect parameter values can be completely dismissed by reasonable assumptions or available prior information.

We will now consider a simple example which highlights one of the difficulties of fitting to CI phase data. Let's assume we have two coherent light sources at wavelengths $\lambda_0 = 465$ nm and $\lambda_1 = 465.3$ nm. We want to fit a model to the CI phase measured at these wavelengths in order to predict the phase at some third wavelength, for which we have no light source. To keep things simple, we will consider a single delay CI instrument with a waveplate (uniaxial crystal with $\theta = 0^\circ$) only, and we will only consider normal incidence ($\alpha = 0^\circ$), so we are dealing with a single data point for each wavelength instead of an image. The model for the phase is then

$$\phi(\lambda) = 2\pi LB(\lambda)/\lambda, \quad (4.1.6)$$

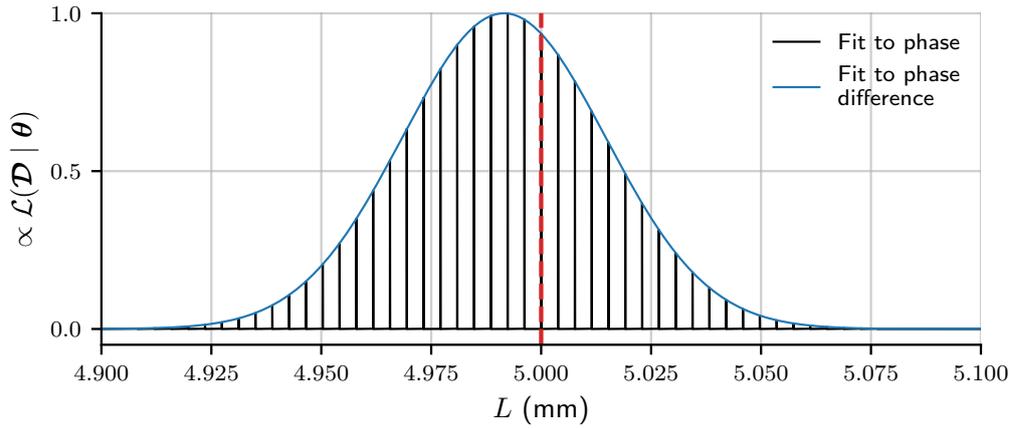


Figure 4.2: Likelihood functions for a simplified fitting problem: fitting to CI phase measurements made at two wavelengths: 465 and 465.3 nm. The instrument’s fixed delay is due to a single waveplate of thickness 5 mm.

where L is the waveplate thickness and $B(\lambda)$ is the birefringence. For now we will assume that $B(\lambda)$ and the wavelengths are known perfectly well, so the only unknown is L . We also assume zero-mean Gaussian noise with variance 10^{-4} rad. Figure 4.2 plots total likelihood $\mathcal{L}(\mathcal{D} | \theta)$ as a function of L . While the maximum likelihood approximates the correct value of L (5 mm), the distribution is highly multi-modal. This can become a problem when the number of fit parameters $\gtrsim 3$ since the necessary optimisation algorithms struggle to find the global maximum in $\mathcal{L}(\mathcal{D} | \theta)$. Also plotted in Figure 4.2 is the likelihood when fitting a model to the *difference between the two phase data points*. This distribution is much smoother and easier for an optimisation algorithm to navigate. It is still multi-modal, but here the other modes are at $L \approx 10$ mm and $L \approx 0$ mm, so are easy to dismiss with any reasonable prior distribution on L . It is obvious from Figure 4.2 that fitting to the absolute phase can result in a significantly more accurate measurement of L and that fitting to phase difference does not use all the information we have. It is almost certainly possible to find the global maximum when fitting to the the absolute phase images in all of the fitting problems discussed in the rest of this chapter by e.g. using a smarter optimisation algorithm, or by using the results of a fit to the phase difference as a starting point. However, we find that fitting only to the phase difference images is numerically simple and is sufficiently accurate for our purposes.

4.1.2 Dispersion in Barium Borate

In Section 2.2.2 we considered the effect of instrument dispersion on CI measurements. In the context of the observed Doppler shifts, a first-order treatment of the dispersion is sufficient. In this chapter we will be fitting a model to phase shifts over ~ 40 nm,

so we need to consider a higher-order treatment. Dispersion is typically described using the Sellmeier equation[117], which can take the form:

$$n_{E,O}(\lambda_\mu) = \left(A_{E,O} + \frac{B_{E,O}}{\lambda_\mu^2 + C_{E,O}} + D_{E,O}\lambda_\mu^2 \right)^{1/2}, \quad (4.1.7)$$

or sometimes

$$n_{E,O}(\lambda_\mu) = \left(A_{E,O} + \frac{B_{E,O}}{\lambda_\mu^2 + C_{E,O}} + \frac{D_{E,O}}{\lambda_\mu^2 + E_{E,O}} \right)^{1/2}. \quad (4.1.8)$$

Here, the subscripts (O) and (E) refer to the ordinary and extraordinary refractive indices, λ_μ is the wavelength of the light (in microns) and $A_{E,O}$, $B_{E,O}$, $C_{E,O}$, $D_{E,O}$ and $E_{E,O}$ are empirically determined coefficients. The forms of Equations 4.1.7 and 4.1.8 can be derived by considering the bound electrons in the crystal as classical, damped oscillators, forced by the electric field of the incident light. Despite the simplicity of the model, “the refractive index of most materials with good homogeneity can be modelled to a few parts in 10^5 over their entire transparent region with a Sellmeier fit of a few terms”[117]. The birefringent material used in this work is barium borate BaB_2O_4 (BBO) in its α phase. There exist multiple sets of α -BBO Sellmeier coefficients in the literature and on the websites of crystal manufacturers. We will also consider sets of coefficients for BBO’s β phase, since these have been used in previous CI investigations and are very similar[63]. Tables 4.1 and 4.2 list all sets of α and β -BBO Sellmeier coefficients that could be found at the time of writing. These coefficients are typically inferred from measurements of the refractive indices made over the visible wavelength range and beyond, made using the standard ‘method of minimum deviation’[118] or else by the method of ‘phase-matching’, which makes use of the material’s non-linear optical properties[119, 120]. Figure 4.3 plots $n_O(\lambda)$, $n_E(\lambda)$ and $B(\lambda) \equiv n_E(\lambda) - n_O(\lambda)$ over the visible wavelength range for each set of listed Sellmeier coefficients. It seems there is a broad consensus that the magnitude of $B(\lambda)$ for α -BBO is $\sim 3\%$ larger than for β -BBO.

It should be noted that the sets of Sellmeier coefficients listed are not all entirely independent of one another. Specifically, $\alpha/1$ and $\alpha/2$ share seven and six of their respective eight coefficients with $\beta/2$. This suggests that these two sets left only one or two coefficients free when fitting to measurements, borrowing the values for the remaining coefficients from the existing source $\beta/2$. Although $\beta/3$ has a very similar $n_O(\lambda)$ to $\beta/2$ (the two cannot be distinguished in Figure 4.3(a)), their coefficients are entirely distinct.

BBO phase/ identifier	Index	$A_{E,O}$	$B_{E,O}$ (μm^2)	$C_{E,O}$ (μm^2)	$D_{E,O}$ (μm^{-2})	Source
$\alpha/1$	E	2.3753	0.01224	-0.01667	-0.01516	agoptics.com
	O	2.7471	0.01878	-0.01822	-0.01354	
$\alpha/2$	E	2.37153	0.01224	-0.01667	-0.01516	Kim et al.[121]
	O	2.7471	0.01878	-0.01822	-0.01354	
$\alpha/3$	E	2.31197	0.01184	-0.01607	-0.00400	newlightphotonics.com
	O	2.67579	0.02099	-0.00470	-0.00528	
$\beta/1$	E	2.3730	0.0128	-0.0156	-0.0044	Eimerl et al.[118]
	O	2.7405	0.0184	-0.0179	-0.0155	
$\beta/2$	E	2.3753	0.01224	-0.01667	-0.01516	Kato (1986)[119]
	O	2.7359	0.01878	-0.01822	-0.01354	

Table 4.1: Sellmeier coefficients describing dispersion of the ordinary (O) and extraordinary (E) refractive indices in barium borate (BBO). The coefficients $A_{E,O}$, $B_{E,O}$, $C_{E,O}$ and $D_{E,O}$ reference Equation 4.1.7.

BBO phase/ identifier	Index	$A_{E,O}$	$B_{E,O}$ (μm^2)	$C_{E,O}$ (μm^2)	$D_{E,O}$ (μm^2)	$E_{E,O}$ (μm^2)	Source
$\beta/3$	E	3.33469	0.01237	-0.01647	79.0672	-82.2919	Kato et al. (2010)[120]
	O	3.63357	0.018778	-0.01822	60.9129	-67.8505	

Table 4.2: Sellmeier coefficients describing dispersion of the ordinary (O) and extraordinary (E) refractive indices in barium borate (BBO). The coefficients $A_{E,O}$, $B_{E,O}$, $C_{E,O}$, $D_{E,O}$ and $E_{E,O}$ reference Equation 4.1.8.

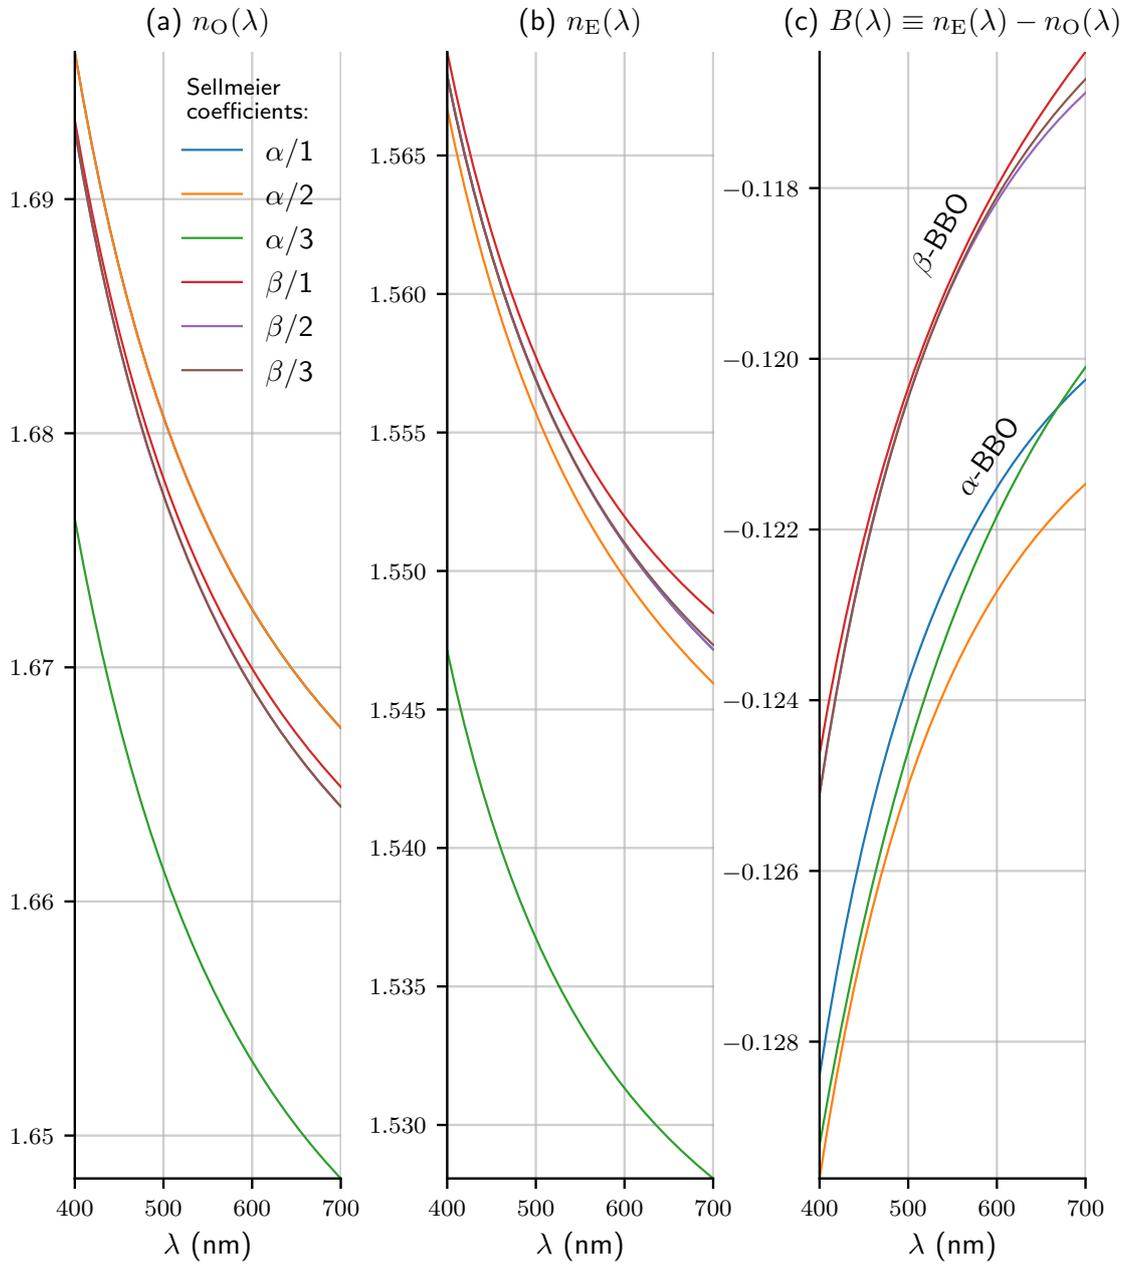


Figure 4.3: (a) Ordinary refractive index, (b) extraordinary refractive index and (c) birefringence of barium borate modelled using the available Sellmeier coefficients from Tables 4.1 and 4.2, plotted over the wavelength range of visible light. See the main text for comments.

Are these Sellmeier Coefficients Accurate Enough?

We have no information about the accuracy of the different dispersion curves for the ordinary and extraordinary indices of refraction in BBO. How do we know which of these curves — if any — can accurately model the measured calibration phase data? A good place to start with this question is to take the spread of the dispersion curves in Figure 4.3 as representative of our uncertainty. Then we can propagate the uncertainty into our model for phase shift and gauge its effect.

To keep things simple, we will revisit the case of a single waveplate at normal incidence that we used in the previous section, with phase delay defined in 4.1.6. If we define a phase shift about wavelength λ_0 as $\Delta\phi \equiv \phi - \phi_0$, then we have:

$$\Delta\phi = 2\pi L \left(\frac{B(\lambda)}{\lambda} - \frac{B_0}{\lambda_0} \right). \quad (4.1.9)$$

Writing $B(\lambda)$ as a Taylor series about λ_0 gives us:

$$\begin{aligned} B(\lambda) &= \sum_{p=0}^{\infty} \frac{\Delta\lambda^p}{p!} \left. \frac{\partial^p B}{\partial \lambda^p} \right|_{\lambda_0} \\ &= B_0 + \Delta\lambda \left. \frac{\partial B}{\partial \lambda} \right|_{\lambda_0} + \sum_{p=2}^{\infty} \frac{\Delta\lambda^p}{p!} \left. \frac{\partial^p B}{\partial \lambda^p} \right|_{\lambda_0}, \end{aligned} \quad (4.1.10)$$

where $\Delta\lambda \equiv \lambda - \lambda_0$ is the wavelength shift about λ_0 and we have explicitly written out the zeroth and first order terms in the sum to infinity. Substituting this expression into Equation 4.1.9 and using the definition of $\hat{\phi}_0$ from Equation 4.1.2 we find

$$\Delta\phi = \underbrace{-\hat{\phi}_0 \frac{\Delta\lambda}{\lambda}}_{\text{Linear term}} + \underbrace{\frac{2\pi L}{\lambda} \sum_{p=2}^{\infty} \frac{\Delta\lambda^p}{p!} \left. \frac{\partial^p B}{\partial \lambda^p} \right|_{\lambda_0}}_{\text{Higher-order terms}}. \quad (4.1.11)$$

Here, we have deliberately separated the linear term from the higher order ($p \geq 2$) terms. When $\Delta\phi$ is written in this form, it is easy to see that for small enough $\Delta\lambda$, the higher order terms will vanish to zero, leaving Equation 4.1.3. Figure 4.4(a) plots $\Delta\phi$ predicted by each available set of BBO Sellmeier coefficients for the case of $L = 4.5$ mm over a 20 nm wavelength range about $\lambda_0 = 460$ nm. Figure 4.4(b) then plots the higher-order phase terms from Equation 4.1.11 only, which we will call $\Delta\phi_{\text{H.O.}}$. The equivalent flow velocity on the second y-axis is calculated using

$$\mathbf{v} \approx c \frac{\Delta\phi_{\text{H.O.}}}{\hat{\phi}_0}. \quad (4.1.12)$$

It can be seen that even though the linear term in Equation 4.1.11 is $\gtrsim 2$ orders of magnitude larger than the higher order terms for the wavelength range shown, the higher order terms become important (i.e. contribute $\gtrsim 1$ km/s flow-equivalent) for

$\Delta\lambda \gtrsim 1$ nm. This has been noted in previous work[91].

By combining Equations 4.1.6 and 4.1.2 we can write group delay as

$$\hat{\phi}_0 = 2\pi L \left(\frac{B_0}{\lambda_0} - \frac{\partial B}{\partial \lambda} \Big|_{\lambda_0} \right). \quad (4.1.13)$$

Suppose we have measured $\Delta\phi$ over a small enough $\Delta\lambda$ that only the linear term in Equation 4.1.11 need be considered. When fitting the linear model to the data, the parameter L is perfectly (negatively) correlated with a function of B_0 and $\frac{\partial B}{\partial \lambda} \Big|_{\lambda_0}$ (in the brackets in Equation 4.1.13). It follows that if we try to fit the parameter L to any $\Delta\phi$ data measured over a sufficiently small $\Delta\lambda$, choice of Sellmeier coefficients will affect the optimised value of fit parameter L but will not affect the quality of the fit. If we consider measurements over a larger $\Delta\lambda$, where the higher order terms *do* need to be accounted for, then the optimised value of L remains almost entirely determined by B_0 and $\frac{\partial B}{\partial \lambda} \Big|_{\lambda_0}$. This means that, in the context of accurately modelling phase shifts, we are not concerned about discrepancies in $B(\lambda)$ and $\frac{\partial B}{\partial \lambda}$ between the different Sellmeier coefficient sets. It is the uncertainty in the higher order ($p \geq 2$) terms in Equation 4.1.11 that will affect our ability of the model the data, so our concern is the discrepancy in $\Delta\phi_{\text{H.O.}}$ between coefficient sets. Figure 4.4(c) plots $\Delta\phi_{\text{H.O.}}$, calculated for each set of Sellmeier coefficients, with the mean over all sets subtracted so as to highlight the spread. This rough treatment suggests that, for $\Delta\lambda \gtrsim 5$ nm, uncertainty in the Sellmeier coefficients is likely too large to model phase shifts to the goal accuracy of ± 1 km/s.

4.2 Experimental Testing

4.2.1 Tuneable Laser Data

To test some of these ideas experimentally, we measured CI calibration data using a tuneable diode laser⁴. The laser was tuneable over a 464–468 nm wavelength range and was fibre-coupled to a 6 inch diameter integrating sphere to provide a uniform, extended, quasi-monochromatic calibration source. The laser wavelength was measured using a High Finesse WS7 wavemeter with a reported accuracy of ± 0.03 pm. The CI instrument was of the single-delay layout shown in Figure 2.7, using three commercial DSLR lenses of focal length of 85 mm. The camera used was a PCO.edge 5.5 with

⁴Experimental data in this section was collected at DIII-D in Summer 2017 with the help of Steve Allen of Lawrence Livermore National Lab (LLNL). The tuneable laser, wavemeter, camera, opto-mechanical components and temperature control components used here were provided by LLNL. In addition, the Labview software for acquisition of the raw camera images was written by Cameron Samuell (LLNL).

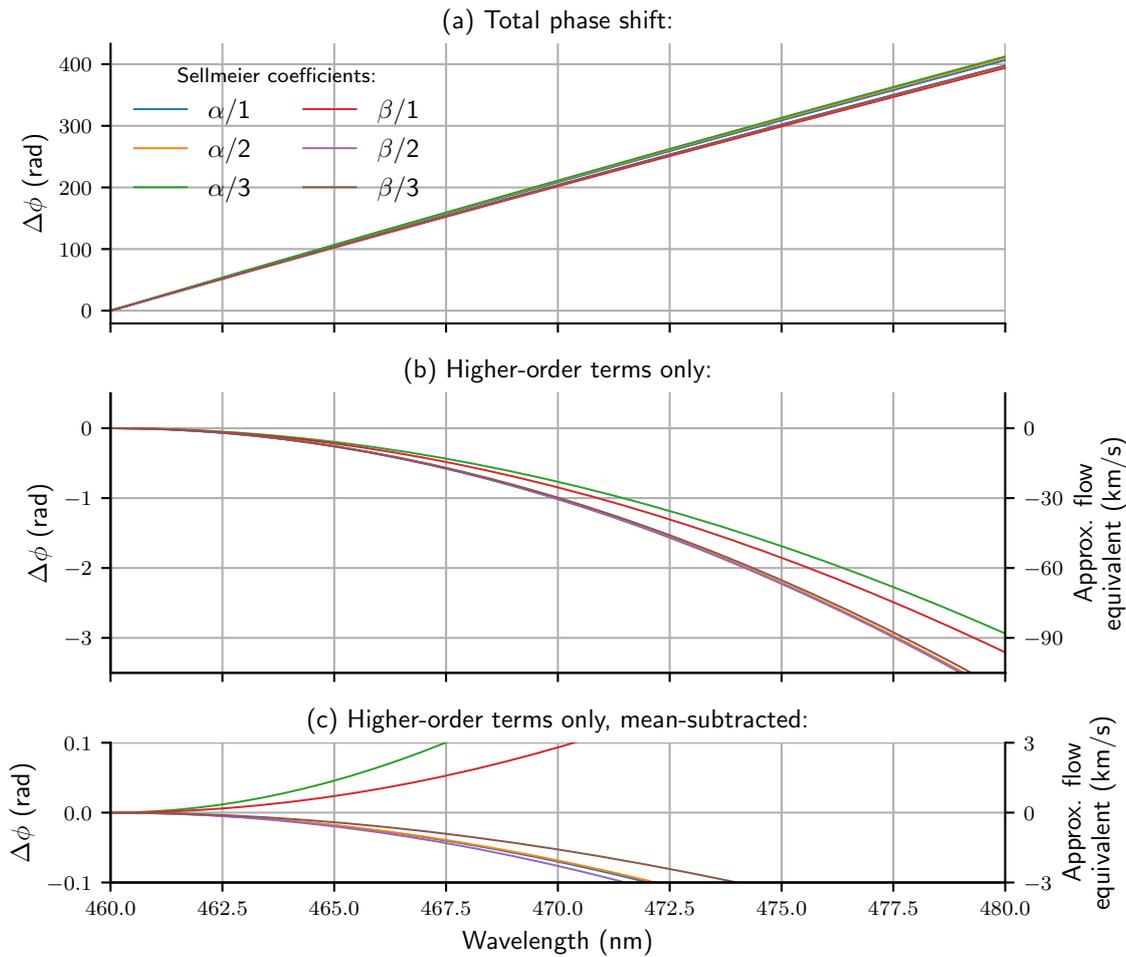


Figure 4.4: (a) Modelled phase shift vs. wavelength for a simplified CI instrument (with a single waveplate of thickness 4.5 mm, and considering only light at normal incidence) for all available sets of BBO Sellmeier coefficients. (b) The higher-order phase shift terms only (see Equation 4.1.11). (c) As in (b) but with the mean value across all sets of Sellmeier coefficients subtracted.

a 2560×2160 format complimentary metal-oxide-semiconductor (CMOS) detector. Figure 4.5(a) shows the CI instrument and mount. The interferometer was of a single delay, linear carrier type, consisting of a waveplate and a Savart plate with nominal thicknesses of $L_{wp} = 4.6$ mm and $L_{sp} = 4.0$ mm respectively. Temperature stability of the interferometer over the measurement period was achieved using the same three levels of control described in Section III of [92]. Firstly, the experiment took place in a lab that was temperature-controlled to 18.65 ± 0.25 °C. Secondly, the CI instrument was installed in a large insulated box with a 63 W Peltier cooling assembly and a proportional integral differential (PID) feedback-control system maintaining 27 ± 0.1 °C. Finally, the interferometer was surrounded with resistive heaters (Figure 4.5(b)) and a thermistor was secured next to the waveplate crystal using epoxy (Figure 4.5(c)).

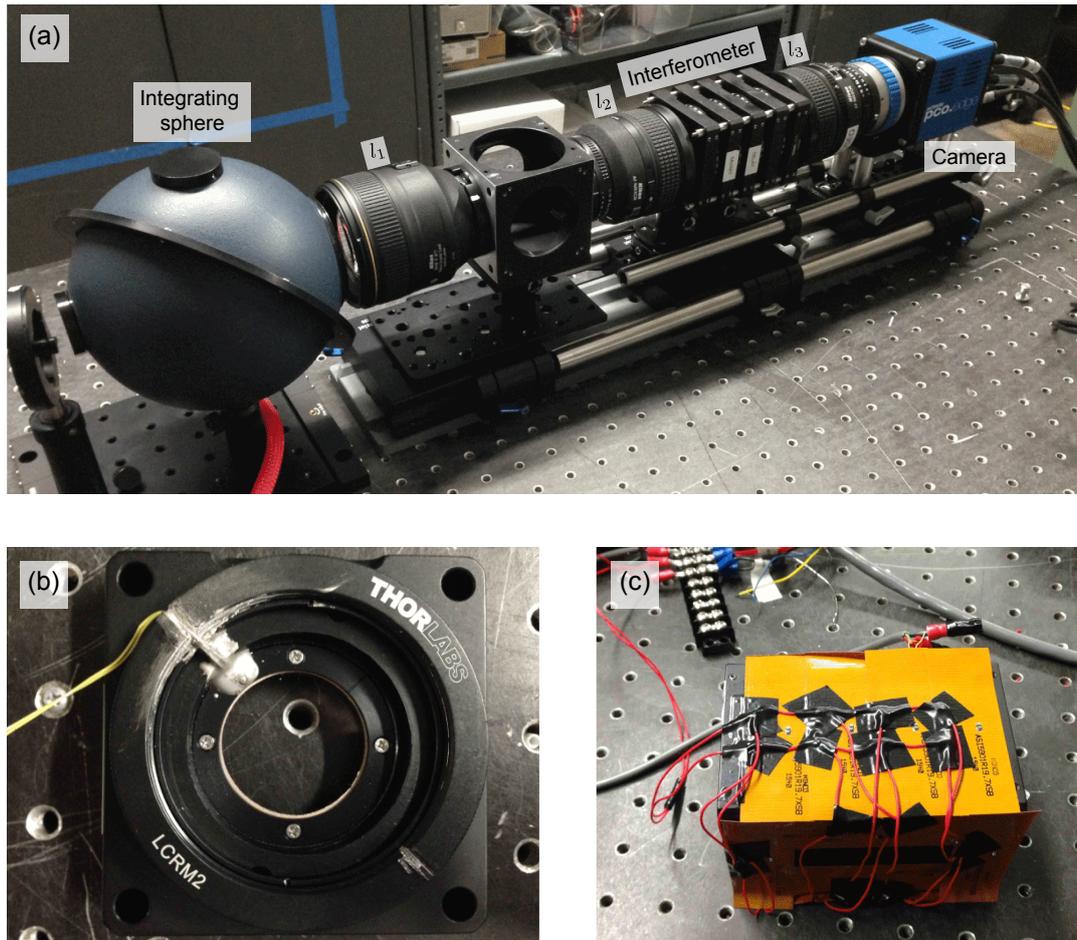


Figure 4.5: (a) The coherence imaging instrument used to make calibration measurements in this section (prior to wrapping interferometer with resistive heaters). (b) A $100\text{ k}\Omega$ thermistor is secured to the waveplate mount for crystal temperature feedback. (c) The interferometer is surrounded by $16 \times 20\ \Omega$ resistive heaters.

A second PID feedback-control system used this to stabilise the crystal at the desired temperature $33 \pm 0.01\text{ }^\circ\text{C}$. The temperature stability of the waveplate corresponds to a flow stability of roughly $\pm 0.1\text{ km/s}$. In total, 68 calibration images were captured over the available wavelength range. Figure 4.6(a) shows one raw image, measured at $\lambda = 464.336\text{ nm}$. Figure 4.6(c) then plots $\Delta\phi \equiv \phi - \phi_0$ about $\lambda_0 = 436.336\text{ nm}$ as a function of laser wavelength for three different points in the image — labelled A, B and C in Figure 4.6(a). Note that $\frac{\partial\phi}{\partial\lambda} \propto \hat{\phi}$ clearly varies between chosen image points.

Fourier demodulation is used to extract the wrapped phase from the captured CI calibration images⁵. To model these phase images, we use the instrument model detailed in Section 2.3.5. Directly modelling phase images using Equations 2.3.17, 2.3.18 and 2.3.19, instead of using the more general Mueller matrix treatment to model

⁵See Appendix B for details.

the raw interferograms (Equation 2.3.20) minimises model evaluation time for efficient fitting. However, the Mueller matrix treatment would be required to realistically model crystal misalignment or partial polarisation effects. Figure 4.6(b) shows a modelled interferogram image for illustrative purposes, generated as $\propto (1 + \cos \phi)$ where ϕ is the modelled phase image and using the nominal values for the model parameters. When constructing the likelihood function for the fit, we use the von Mises distribution (Equation 4.1.5) for each $\Delta\phi$ data point, with each data point assumed to be statistically independent. The four model parameters that are optimised are the crystal thicknesses L_{wp} and L_{sp} , the focal length of lens 3 f_3 and the orientation of the interferometer relative to the sensor ρ_{int} . Based on the discussion in Section 4.1.1 we fit to the difference between measured phase images instead of the phase images themselves. Based on the discussion in Section 4.1.2, uncertainty in Sellmeier coefficient values ought not to (and does not) significantly affect the goodness of fit over the wavelength range measured ($\Delta\lambda < 5 \text{ nm}$) so any of the coefficient sets from Tables 4.1 and 4.2 can be used. The $\Delta\phi$ images are downsampled by a factor 50 in both linear dimensions to reduce model evaluation time. The $\Delta\phi$ data plotted in Figure 4.6(c) clearly satisfies the smoothness condition introduced in Section 4.1.1 and so could be trivially unwrapped before fitting. Previous work[91] has already shown that an instrument model can be fit to 50+ images over this wavelength range. With a view to extending the model-fitting technique to gas-discharge lamps, whose bright lines are (relatively) few and sparse such that the smoothness condition is not met, we concentrate here on the smallest number of the laser images that can be used to constrain the fit.

It was found that fitting to just two wavelengths is sufficient to constrain the model and extrapolate over the $\Delta\lambda \sim 3 \text{ nm}$ wavelength range with better than $\pm 1 \text{ km/s}$ accuracy. This corresponds to a fit to a single $\Delta\phi$ image. Figure 4.6(c) shows the fit to the calibration images at wavelengths 464.336 nm and 464.784 nm (highlighted) for the three image points A, B and C. Figure 4.6(d) then plots the corresponding fit residuals in units of radians and equivalent flow in km/s. The residuals are consistently less than $\pm 0.5 \text{ km/s}$ flow equivalent when using the fit to extrapolate across the 3 nm wavelength range, significantly better than our goal accuracy of $\pm 1 \text{ km/s}$. The fit shown assumes that the $\alpha/3$ set of Sellmeier coefficients is correct, though the fit residuals are similarly small regardless of which set of coefficients is used. Table 4.3 lists the maximum a posteriori (MAP) estimates and uncertainty of the fit parameters for fits using each of the available Sellmeier coefficient sets in turn. As expected, choice of coefficient set changes the inferred L_{wp} but does not significantly change the inferred group delay $\hat{\phi}_0$. Also as expected, the measured value of $L_{\text{wp}} = 4.48 \pm 0.05 \text{ mm}$ is closer to the MAP values inferred using the α -BBO Sellmeier coefficient sets.

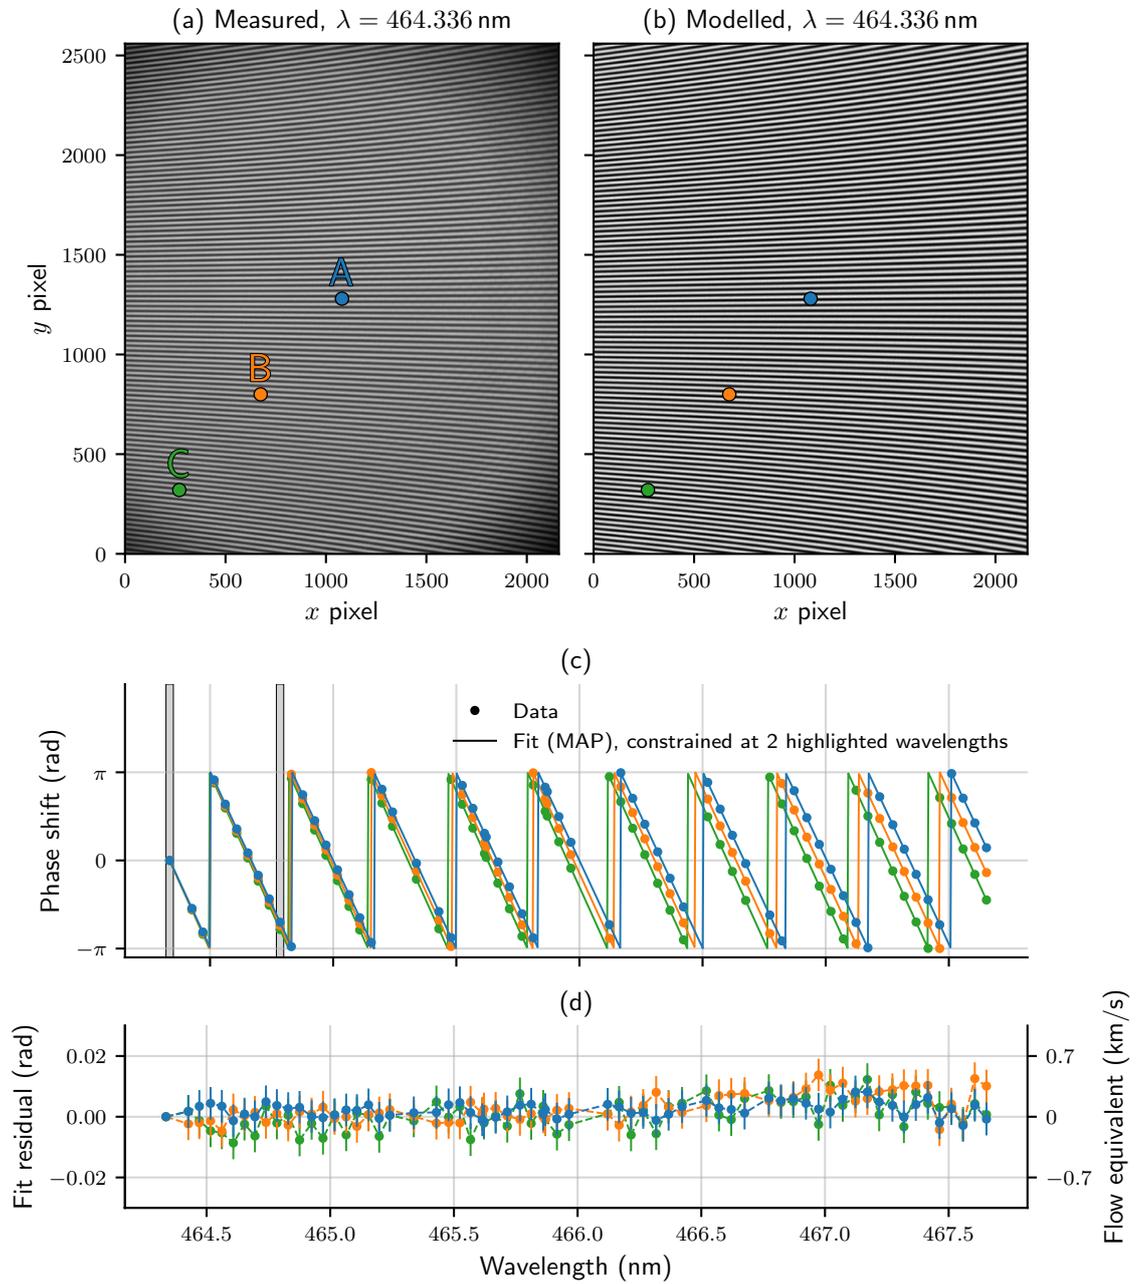


Figure 4.6: (a) Measured CI calibration image observing coherent laser light at 464.362 nm. (b) The corresponding modelled image. (c) Measured and fit phase shift plotted against laser wavelength over the 3.5 nm range for the three image points marked in (a). Fit is constrained by highlighted wavelengths only. (d) The corresponding residuals between measured and fit phase shift.

Sellmeier coefficient set	Fit parameter value (MAP)			Group delay $\hat{\phi}_0$ (waves)
	L_{wp} (mm)	L_{sp} (mm)	f_3 (mm)	
$\alpha/1$	4.532 ± 0.001	4.04 ± 0.01	85.0 ± 0.2	1401.7 ± 0.3
$\alpha/2$	4.495 ± 0.001	4.00 ± 0.02	85.0 ± 0.3	1401.6 ± 0.3
$\alpha/3$	4.472 ± 0.001	3.93 ± 0.01	85.0 ± 0.2	1401.5 ± 0.3
$\beta/1$	4.680 ± 0.002	4.16 ± 0.02	84.9 ± 0.2	1401.5 ± 0.6
$\beta/2$	4.635 ± 0.002	4.13 ± 0.01	85.0 ± 0.2	1401.6 ± 0.6
$\beta/3$	4.632 ± 0.002	4.12 ± 0.01	85.0 ± 0.1	1401.5 ± 0.6

Table 4.3: Optimised model parameter values, fitting to the two wavelengths highlighted in Figure 4.6, for each available set of Sellmeier coefficients. Value of interferometer orientation was found to be the same for each fit $\rho_{\text{int}} = -0.1 \pm 0.1^\circ$. The measured value for L_{wp} is 4.48 mm and the nominal values for the other three parameters are $L_{\text{sp}} = 4$ mm, $f_3 = 85$ mm and $\rho_{\text{int}} = 0^\circ$. Also listed is the inferred group delay $\hat{\phi}_0 = \frac{L_{\text{wp}} B_0 \kappa_0}{\lambda_0}$ in units of waves at $\lambda_0 = 464.336$ nm.

4.2.2 Gas-discharge Lamp Data

Results from the previous section show that phase calibration images measured at two narrowly separated wavelengths ($\Delta\lambda \sim 0.5$ nm) can be used to constrain a model that can accurately predict the phase at a third wavelength nearby ($\Delta\lambda \sim 3$ nm). We found that the choice between available sets of Sellmeier coefficients describing instrument dispersion does not affect the quality of the fit over this wavelength range. However, bright spectral lines from a gas-discharge lamp will typically cover a larger wavelength range and so fitting to phase data measured using such a source means fitting the dispersion model to the calibration measurements as well as the CI instrument model used in the previous section. Measuring lamp lines rather than laser light also requires optical bandpass filters with appropriate transmission curves for isolating each calibration line from its neighbouring lines. These complicating factors will be considered in this section using Cd and Zn gas-discharge lamps.

Table 4.4 lists the wavelengths (in air) of the brightest emission lines observed from Cd and Zn gas-discharge lamps within the 460–520 nm range. Also listed are the relative intensities of the lines, measured using a slit-coupled grating spectrometer. This wavelength range is relevant as it roughly spans the three plasma impurity emission lines that will be targeted for CI flow measurements on MAST-U (precise wavelengths listed in Section 4.1). Figure 4.7(a) plots the Cd and Zn calibration lines relative to these three targeted plasma impurity lines. Also plotted are the transmission curves of the four bandpass interference filters used to isolate the emission lines for CI measurement, measured using a stabilised continuum source and a slit-coupled spectrometer. The filters used have a nominal transmission of 10^{-4} at

wavelengths far away from the bandpass, and block up to $\lambda = 1000$ nm. Table 4.4 lists the wavelengths of all the Cd and Zn calibration lines considered. Two ‘nuisance’ Cd lines are indicated with dashed lines in Figure 4.7(a). These lines are not targeted for the calibration measurements, but contribute a small offset to the measured phase of their neighbouring lines. The size of the phase contamination due to improperly isolated lines can be calculated using the FTS theory from Section 2.2 and the observed relative line intensities and filter transmission profiles in Figure 4.7. It can be shown that the nuisance Cd I line at 515.5 nm, contributes the largest phase contamination, to the measurement of the Cd I line at 508.6 nm. The size of this contamination is equivalent to < 0.5 km/s and is an order of magnitude larger than the predicted contamination on any other line measured here. Since this is smaller than our goal accuracy (± 1 km/s), phase contamination is ignored for the rest of this section.

Ion	Wavelength (nm)	Observed intensity (norm.)	Filter
Cd I	466.23520	0.0035	-
Cd I	467.81493	0.4980	1
Zn I	468.013590	0.4243	1
Zn I	472.215690	0.8614	2
Cd I	479.99123	0.8872	3
Zn I	481.053210	1	3
Cd I	508.58217	1	4
Cd I	515.46605	0.0039	-

Table 4.4: Wavelengths (in air) of the Cd and Zn calibration lines used in this work (from NIST[115]). Also listed are the intensities of the lines, as observed in the lab using a grating spectrometer. These values are normalised to the most intense line of each lamp. The number identifier for the filter used to isolate each line refers to the black, circled numbers in Figure 4.7(a).

For the measurements in this section, the PPM CI instrument configuration introduced in Chapter 3 was used. Specifically, we used the ‘Multi-delay 1’ configuration from Table 3.2. Although the measured interferogram samples the coherence at four fixed delays simultaneously (the demodulation procedure is outlined in Section 3.3.2), *we will only consider data from two of these delays here for simplicity*. The two delays chosen correspond to the two waveplates with measured thicknesses 4.48 ± 0.02 mm and 6.35 ± 0.02 mm (i.e. it is the sum and difference terms that we are ignoring). These two delays are referred to here as the ‘linear carrier’ and ‘pixelated carrier’ respectively, since the first is encoded using the linear phase shear produced by the Savart plates and the second is encoded in the pixelated pattern of the micro-polariser array. See Figure 3.6 in Chapter 3 for an example of a raw calibration image taken with an instrument in this configuration. For the calibration measurements, the

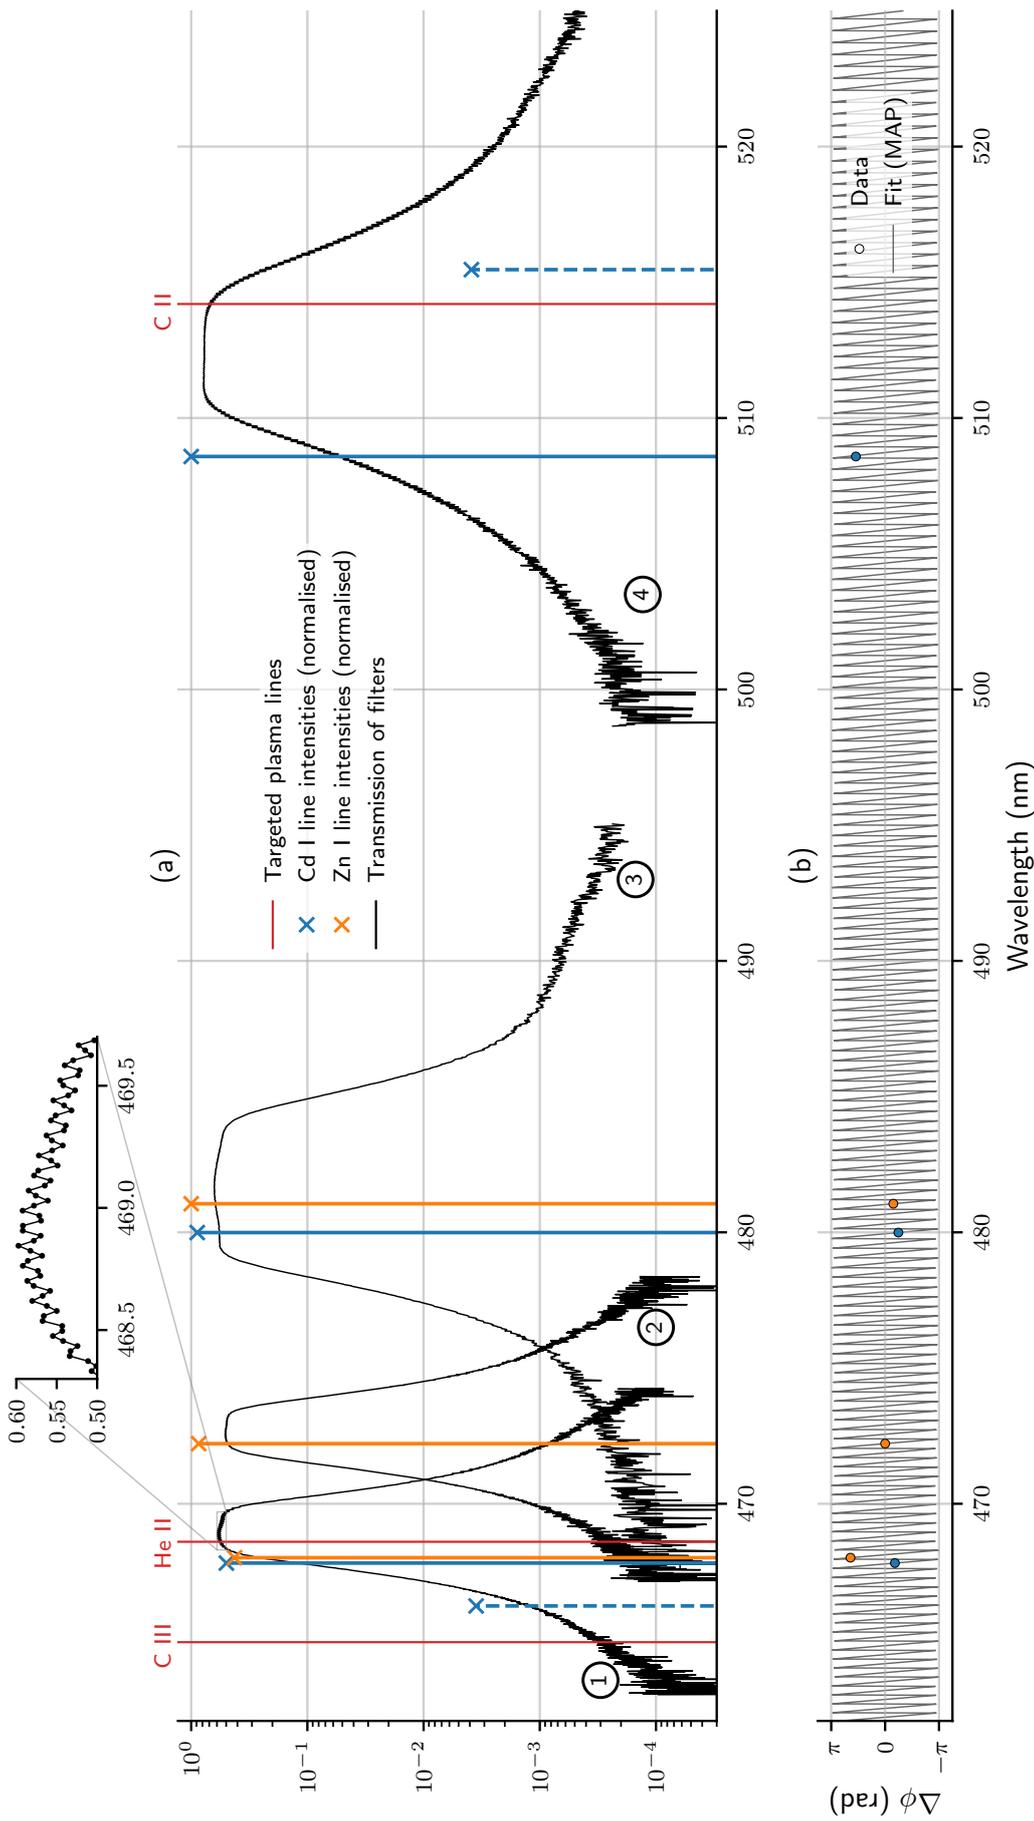


Figure 4.7: (a) Measured intensities of the Cd (blue) and Zn (orange) lamp calibration lines, normalised to the most intense line of each lamp. Dashed lines indicate ‘nuisance’ emission lines. Plasma impurity line wavelengths are shown in red. Also shown are the measured transmission profiles of the bandpass filters. (b) Measured and fit (wrapped) phase shift $\Delta\phi$ relative to $\lambda_0 = 472.2$ nm for the image centre. See text for details.

interferometer temperature controller was set to a nominal temperature and stability of $34 \pm 0.25^\circ\text{C}$ and the instrument was left alone for four hours to equilibrate prior to measurements being made. Systematic error in the measured phase images due to ambient temperature drift was estimated by comparing the phase offset between calibration images measured at the same wavelength (Cd I at 467.8 nm) at the start and end of the measurement period. For the linear carrier the mean of this error was found to be 0.003 rad while for the pixelated carrier it was found to be 0.001 rad. Converted to effective flow velocity, these errors are small: 0.1 km/s and 0.03 km/s respectively (calculated using the on-axis delay due to a 4.48 mm α -BBO waveplate). Calibration images were captured at each of the six brightest Cd and Zn wavelengths (467.8, 468.0, 472.2, 480.0, 481.1 and 508.6 nm) in turn. The filters were installed and uninstalled manually, though mounted independently of the CI instrument using a clamp stand so as to avoid mechanical disturbance.

Test 1: Extrapolation of Phase Beyond Measured Wavelengths

Figure 4.7(b) shows the measured phase shift $\Delta\phi$ calculated relative to the Zn line at $\lambda_0 = 472.2$ nm for the image centre (data only shown for the linear carrier here). For the first test, we include in the fit to the data the four highest wavelength lines only (472.2, 480.0, 481.1 and 508.6 nm), reserving the Cd I line at 467.8 nm line for testing the accuracy of our phase extrapolation outside of the wavelength range used to constrain the fit. In Figure 4.8(a), the left column plots the measured $\Delta\phi$ images (linear carrier) relative to $\lambda_0 = 472.2$ nm for the three other lines included in the fit. The left column of Figure 4.8(b) plots the corresponding $\Delta\phi$ images for the pixelated carrier. The phase shear characteristic of the linear carrier causes a shear in the $\Delta\phi$ images too, roughly proportional to the wavelength shift. The pixelated carrier $\Delta\phi$ images show the hyperbolic phase pattern characteristic of a waveplate. In fitting a model to this data, it was found that the available sets of BBO Sellmeier coefficients result in large residuals (> 1 km/s). Based on the discussion in Section 4.1.2, this is consistent with the estimated uncertainty in the available coefficients. Instead then, it is necessary to fit the dispersion curves to the measured data. The most obvious way to do this is to leave one or more Sellmeier coefficients free to fit the data. Here, we find that by fitting to just two of the coefficients for the extraordinary refractive index (A_E and B_E), and leaving the remaining extraordinary and ordinary coefficients fixed to one set of available values, we are able to accurately fit the measured $\Delta\phi$ data. Here, we use the $\alpha/3$ coefficient set for the static coefficient values since it gave the closest match to measured waveplate thickness in the previous section.

The other fit parameters optimised in this section are f_3 and ρ_{int} , the lens 3 focal length (see Figure 2.7) and the interferometer orientation, as in the previous section.

For the linear carrier the Savart plate thickness L_{sp} is also optimised, but there is no Savart plate contribution to the phase of the pixelated carrier. In contrast to the previous section, we do not fit the waveplate thicknesses L_{wp} here, since this parameter will be extremely highly correlated with A_{E} . Instead, we simply use the measured values $L_{\text{wp}} = 4.48 \text{ mm}$ and $L_{\text{wp}} = 6.35 \text{ mm}$ for the linear and pixelated carriers respectively. Assuming these measured values to be perfectly correct will likely cause a discrepancy between the optimised Sellmeier coefficient values for the linear and pixelated carriers. This shouldn't worry us though as our immediate concern here is the predictive power of the optimised instrument model, not the universal applicability of our optimised Sellmeier coefficients.

The optimised values of the fit parameters for both linear and pixelated carriers for the extrapolation test are listed in Table 4.5 along with the nominal values. Figures 4.8(a) and (b) plot in the centre column the modelled $\Delta\phi$ images for the linear and pixelated carriers respectively, modelled using these optimised fit parameters. The final column of Figures 4.8(a) and (b) then plots the residual between fit and measured $\Delta\phi$ images for each wavelength. The pattern of concentric circles in the residual images (most prominently at $\lambda = 508.6 \text{ nm}$) is likely caused by an etalon effect in the bandpass filters, due to imperfect anti-reflective coating between the layers. Further evidence for this effect can be seen in the 'ringing' in the measured filter bandpass profiles, highlighted in the inset of Figure 4.7(a). As a test of the predictive power of the model outside of the wavelength range used to constrain the fit, we can compare the measured $\Delta\phi$ for the Cd line at 467.8 nm (relative to 472.2 nm) with the $\Delta\phi$ modelled using the optimised fit parameters. This corresponds to an extrapolation of the CI phase over $\Delta\lambda \approx 4.4 \text{ nm}$ from the nearest wavelength used to constrain the fit. Figure 4.9 plots this residual for both linear and pixelated carriers as images (left column) and sliced profiles (right column). It can be seen that the residual in the extrapolated phase calibration corresponds to an effective flow velocity that is mostly smaller than the $\pm 1 \text{ km/s}$ accuracy aim for both carriers.

Test 2: Interpolation of Phase Between Measured Wavelengths

For a second test, we will look at how accurately we can interpolate in wavelength between measured phase images. To do this, we will fit to a different subset of the six available Cd/Zn lines: the four lines at 467.8 nm , 468.0 nm , 472.2 nm and 508.7 nm . This leaves the Cd line at 480.0 nm available to test the phase interpolation accuracy over the $\Delta\lambda \approx 7.8 \text{ nm}$ range from the nearest measurement included in the fit (at 472.2 nm). The instrument model and fit procedure is as in the previous section. The residuals with respect to the $\Delta\phi$ measurements used to constrain the fit are not shown, but are comparable to those of the previous section, shown in Figure 4.8.

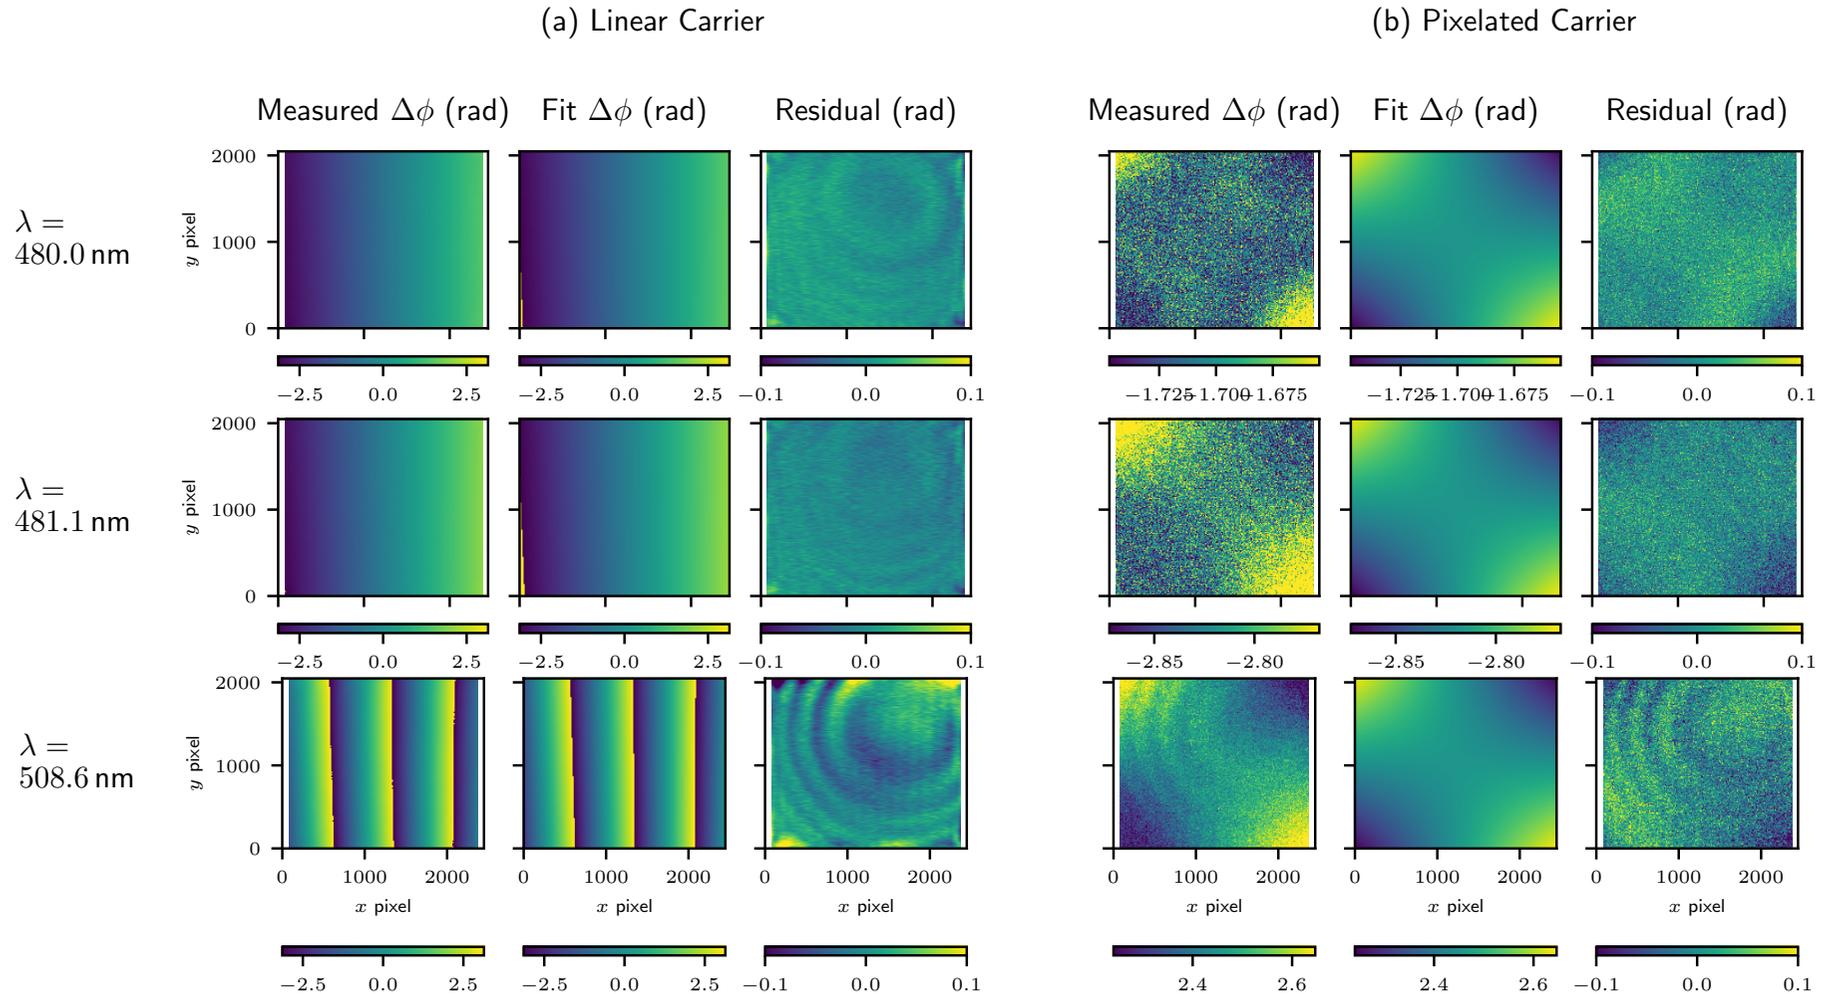


Figure 4.8: The left columns of (a) and (b) show measured phase shift $\Delta\phi$ images at three wavelengths for the linear and pixelated carriers respectively of a multi-delay CI instrument. Phase shift here is relative to the phase measured at 472.2 nm. The centre columns of (a) and (b) show the corresponding modelled phase shift images, fit to the data in the left columns. The right columns of (a) and (b) show the corresponding fit residuals. See text for details.

Parameter	Nominal value	Fit value (MAP)	
		Linear carrier	Pixelated carrier
L_{sp} (mm)	6.2 (4 + 2.2)	6.35 ± 0.01	-
f_3 (mm)	150	149.8 ± 0.1	149.8 ± 0.1
ρ_{int} ($^\circ$)	0	0.76 ± 0.01	-0.4 ± 0.5
A_E	2.31197	2.3203 ± 0.0006	2.3218 ± 0.0003
B_E (μm^2)	0.01184	0.011372 ± 0.000004	0.011352 ± 0.000002

Table 4.5: Optimised model parameter values for the extrapolation test fit to $\Delta\phi$ measured at 472.2 nm, 480.0 nm, 481.1 nm and 508.7 nm.

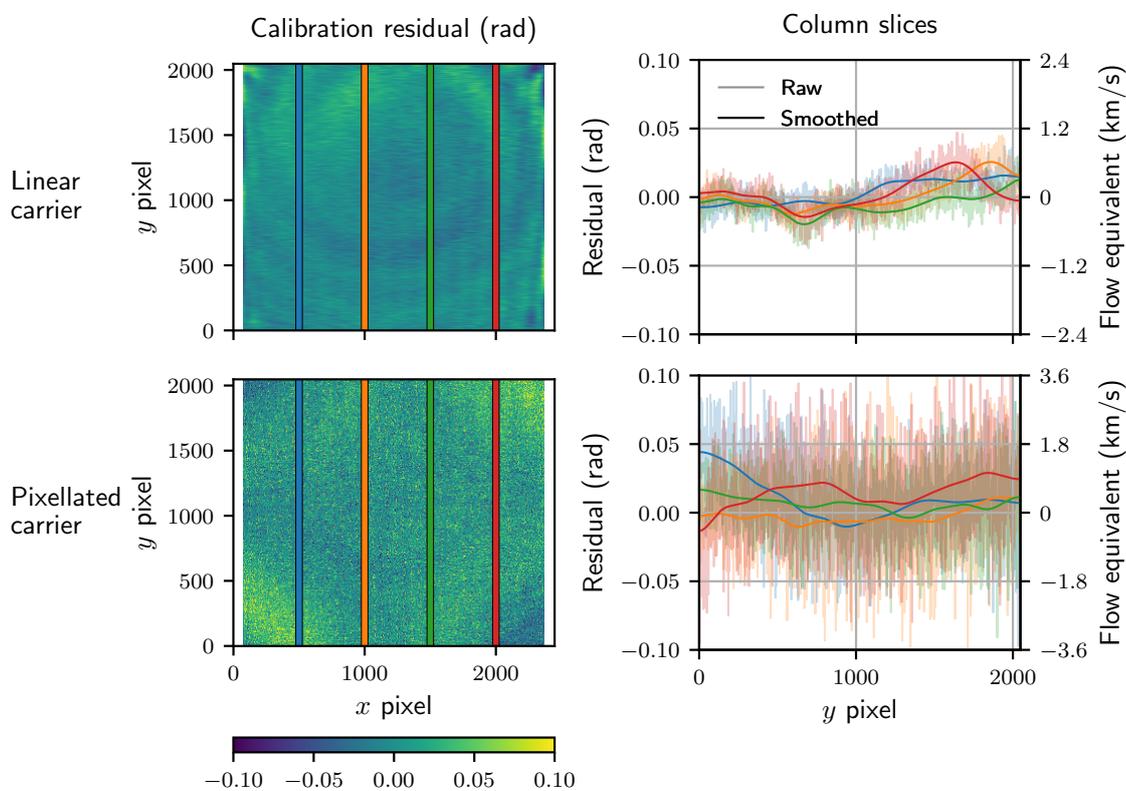


Figure 4.9: Residuals between the predicted and measured calibration phase images at 467.8 nm. Predicted phase at this wavelength is *extrapolated* over $\Delta\lambda \approx 4.4$ nm from the measurement at 472.2 nm. The model used for extrapolation is constrained using measurements at four wavelengths (472.2 nm, 480.0 nm, 481.1 nm and 508.7 nm). Data is shown for two carriers of a multi-delay CI instrument, see text for details.

Table 4.6 lists the optimised values of the fit parameters for both linear and pixelated carriers for the interpolation test, along with the nominal values. Figure 4.10 then plots the residuals between the measured and interpolated phase images at 480.0 nm. The residuals are again mostly smaller than our goal accuracy of ± 1 km/s. With the filter ringing effect less prominent, the residuals for the linear carrier are smaller than

in the extrapolation test (Figure 4.9). However, the residuals for the pixelated carrier are larger, and show the same hyperbolic pattern characteristic of waveplates.

Parameter	Nominal value	Fit value (MAP)	
		Linear carrier	Pixelated carrier
L_{sp} (mm)	6.2 (4 + 2.2)	6.35 ± 0.01	-
f_3 (mm)	150	149.8 ± 0.1	149.9 ± 0.1
ρ_{int} ($^\circ$)	0	0.80 ± 0.01	-0.4 ± 0.6
A_E	2.31197	2.3203 ± 0.0004	2.3215 ± 0.0004
B_E (μm^2)	0.01184	0.011373 ± 0.000004	0.011375 ± 0.000004

Table 4.6: Optimised model parameter values for the interpolation test fit to $\Delta\phi$ measured at 467.8 nm, 468.0 nm, 472.2 nm and 508.7 nm.

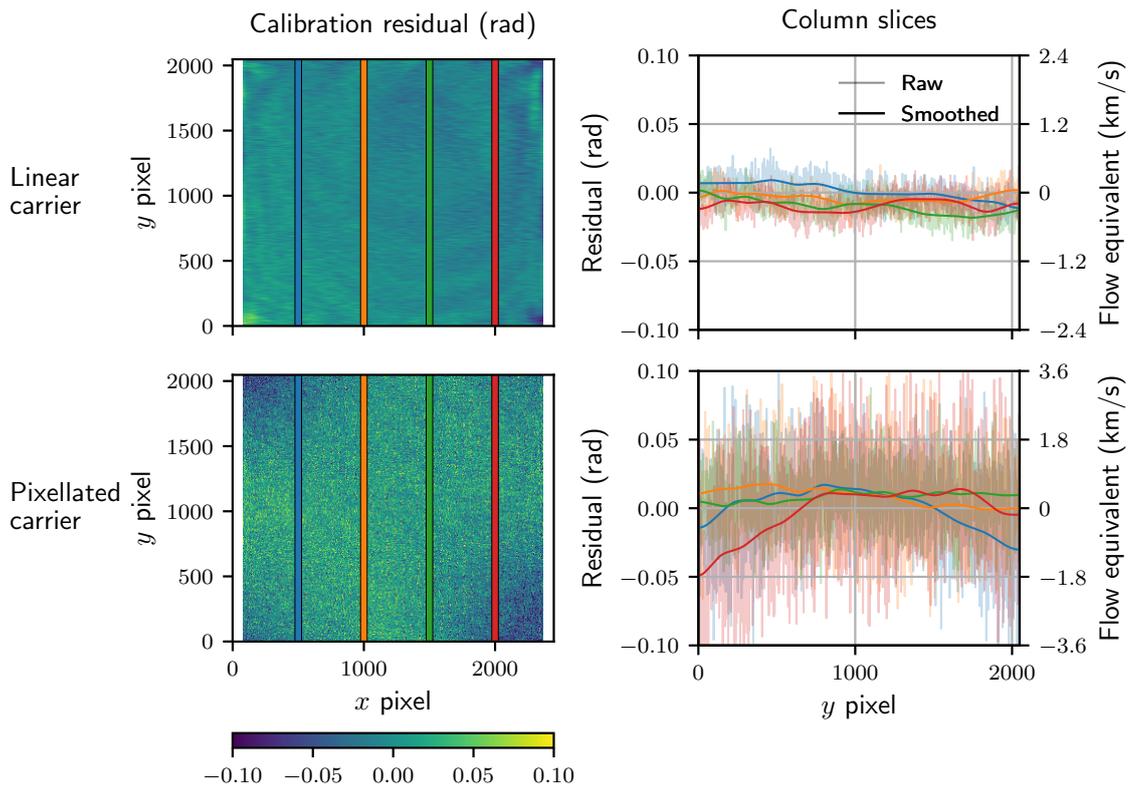


Figure 4.10: Residuals between the predicted and measured calibration phase images at 480.0 nm. Predicted phase at this wavelength is *interpolated* over $\Delta\lambda \approx 7.8$ nm from the measurement at 472.2 nm. The model used for interpolation is constrained using measurements at four wavelengths (467.8 nm, 468.0 nm, 472.2 nm and 508.7 nm). Data is shown for two carriers of a multi-delay CI instrument, see text for details.

4.2.3 A Note on Group Delay

We will now briefly return to the second interferometric calibration parameter: group delay $\hat{\phi}_0$. We can combine the optimised model parameters in Table 4.5 with the equation for on-axis phase delay due to a waveplate from Equation 4.1.6 and the definition of $\hat{\phi}_0$ from Equation 4.1.2 to calculate the (on-axis) $\hat{\phi}_0$ that best matches the measured calibration lamp data, for both linear and pixelated carriers. We evaluate $\hat{\phi}_0$ at $\lambda_0 = 464.336$ nm for a direct comparison to the $\hat{\phi}_0$ values inferred from the tuneable laser data in the previous section, listed in Table 4.3. The linear carrier lamp data corresponds to the same waveplate used in the laser measurements (with measured $L_{\text{wp}} = 4.48 \pm 0.02$ mm). The value inferred from the lamp measurements is $\hat{\phi}_0 = 1401 \pm 2$ waves, in agreement with the laser measurement of $\hat{\phi}_0 = 1401.5 \pm 0.3$ waves. This waveplate's $\hat{\phi}_0$ was also measured in previous work[51, 63] using a tuneable laser, for calibration of the MAST CI measurements. Since this measurement was made at $\lambda_0 = 460.85$ nm, it needs to be extrapolated to $\lambda_0 = 464.336$ nm for comparison to the other two values. This extrapolation is done using Equation 5.3.6 from [63] and the $\alpha/3$ set of Sellmeier coefficients, giving a value of $\hat{\phi}_0 = 1387 \pm 10$ waves. This discrepancy of 1% in the measured $\hat{\phi}_0$ corresponds to a change of the same proportion in the calibration of the flow *magnitude*. Even for the largest flows this is smaller than the expected uncertainty in the flow *zero point*. Broad agreement between measurements of the waveplate $\hat{\phi}_0$ made over a period of > 5 years using different cameras and, for the laser data from Section 4.2.1, different opto-mechanical components, shows that regular $\hat{\phi}_0$ calibration is unnecessary. That said, if multiple lamp lines, or multiple tuned laser wavelengths, are being measured for ϕ_0 calibration anyway, then $\hat{\phi}_0$ can be regularly monitored with little extra effort.

4.2.4 Discussion

The C III, He II and C II impurity ion emission lines that will be targeted for CI measurements on MAST-U all lie within ~ 5 nm of at least one of the calibration lines used in this section (see Figure 4.7(a)). It follows from the results of this section that measurements of these Cd and Zn lines can be used to calibrate each targeted line to a high accuracy (± 1 km/s). Though we have used the multi-delay PPM CI instrument design in this demonstration, there is no reason at all why the method described cannot be applied to the more simple and popular single-delay linear carrier design. The method should extend to arbitrary targeted plasma lines, provided a set of bright calibration lines can be measured with at least one at a nearby ($\Delta\lambda \sim 5$ nm) wavelength. Generally, interpolation is preferable to extrapolation and so calibration lines surrounding the targeted plasma line should be sought. Also, care should be

taken to ensure the bandpass filters used are properly suppressing neighbouring lines to avoid significant contamination. From the residuals shown in Figures 4.8, 4.9 and 4.10, it looks like the largest current source of systematic error for this calibration method, at least for the linear carrier, is the ringing pattern caused by the filters. Switching to higher quality filters with better anti-reflective coatings is the obvious way to reduce this error. The residuals for the pixelated carrier data show a different systematic error with what looks like a hyperbolic pattern. Preliminary Mueller matrix modelling suggests this could be caused by small rotational misalignment ($\sim 1^\circ$) of one or more of the crystals in the xy -plane. If this is the cause, improved mechanical mounting of the interferometer components would reduce the effect.

This work improves our understanding of how much information can be extracted from just a few CI calibration images measured over a relatively large wavelength range ($\Delta\lambda \sim 40$ nm). Existing work[91] used 50+ phase images measured using a laser tuned over $\Delta\lambda \sim 3$ nm. That said, in this work we are clearly discarding information by fitting to the phase *difference* images instead of to the absolute phase. With a smarter fitting algorithm, it is feasible that fewer lamp lines could be used to perform the calibration than the four we have used here.

4.3 Summary

In this Chapter we have introduced a new method for absolute calibration of the zero point of CI measurements of C and He impurity flow in the tokamak plasma edge. The new method does not require a tunable laser source or wavemeter, and instead uses standard Cd and Zn gas-discharge lamps and bandpass interference filters. This represents a significant reduction in hardware costs and can be easily implemented on-machine. The accuracy of the calibration demonstrated is better than ± 1 km/s, comparable to existing methods. This method will be used to calibrate MAST-U CI measurements of impurity flow when operations begin later this year.

Chapter 5

Coherence Imaging Measurements of Electron Density

In this chapter we apply CI to the measurement of electron density n_e in the relatively cool ($T_e \sim 1$ eV) and dense ($n_e \sim 10^{20} \text{ m}^{-3}$) conditions representative of the tokamak divertor. A rise in n_e manifests as a drop in the interference fringe contrast due to Stark broadening of the observed spectral lineshape. State-of-the-art models for Stark-broadened lineshapes are outlined and a fitting tool for inferring n_e from measured CI contrast profiles is introduced. We show that using multi-delay CI instead of single delay makes the n_e measurement more robust against systematic error due to Doppler broadening and can also provide a simultaneous estimate of the emitting species' temperature. Finally, we present experimental CI measurements of n_e made on the Magnum-PSI linear plasma experiment over a range of conditions ($10^{19} < n_e < 10^{21} \text{ m}^{-3}$) matching the divertor plasma conditions of current and next generation tokamaks. Also being tested is the PPM CI instrument design introduced in Chapter 3.

CI measurement of n_e has been demonstrated at a proof-of-principle level by Lischtschenko et al.[97] on the Pilot-PSI linear plasma experiment. A match between the CI and TS n_e profiles was found for a discharge at $T_e \sim 1$ eV and $n_e \sim 10^{20} \text{ m}^{-3}$. A similar (unpublished) investigation on the MAGPIE linear plasma experiment attempted to benchmark CI measurements of n_e using established techniques — interferometry, slit-coupled spectroscopy and probes in lower density conditions $n_e \sim 10^{19} \text{ m}^{-3}$, but the results were inconclusive. The work presented in this chapter develops the technique in two ways. First, we use state-of-the-art numerical calculations for Stark-broadened lineshape instead of the analytical approximation used previously. Second, we use a multi-delay CI instrument design to improve the diagnostic power of the measurement. We will discuss each of these in turn.

5.1 Modelling Balmer Series Lineshapes

Balmer emission spectroscopy is a useful diagnostic tool for the exhaust of fusion experiments because it emits intensely in the visible part of the spectrum under the conditions of the SOL and the divertor. Table 5.1 lists the wavelengths of the first six lines in the Balmer series for H and for D. The isotope shift is comparable to the linewidths observed in divertor plasma conditions, so would have to be accounted for in any CI measurements of isotope mixtures. However, we only consider pure H plasma in this work. Fine structure splitting, due to spin and relativistic effects, is much smaller for the Balmer series at $\Delta\lambda \sim 0.005$ nm and is left unresolved in Table 5.1. The CI contrast profile ζ is most sensitive to this splitting at larger delays but the contribution to the contrast profile is small for all results in this chapter: $\zeta_{\text{fs}} > 0.985$. The main contributors to the Balmer lineshapes observed in this work

Name	n_{upper}	λ_{H} (nm)	λ_{D} (nm)	$\lambda_{\text{H}} - \lambda_{\text{D}}$ (nm)
α	3	656.279	656.1012	0.1778
β	4	486.135	486.00013	0.13487
γ	5	434.0472	433.92833	0.11887
δ	6	410.1734	410.06186	0.11154
ϵ	7	397.0075	396.89923	0.10827
ζ	8	388.9064	388.79902	0.10738

Table 5.1: Balmer series transition wavelengths ($n_{\text{upper}} \rightarrow 2$) for hydrogen and deuterium (in air, with fine-structure unresolved)[115].

are Doppler broadening, Stark broadening and Zeeman splitting. If each of these phenomena is statistically independent, then the total lineshape is the convolution of the constituent lineshape contributions:

$$g^c(\nu) = g_{\text{D}}^c(\nu) * g_{\text{S}}^c(\nu) * g_{\text{Z}}^c(\nu). \quad (5.1.1)$$

Since n_{e} affects only the lineshape, and does not shift the central frequency ν_{c} , we have considered only the centred lineshape $g^c(\nu)$ here, related to the total lineshape by $g(\nu)$ by Equation 2.2.3. In Chapter 2 we saw that it is the magnitude of the Fourier transform of $g^c(\nu)$ that gives us the contrast ζ as a function of interferometer delay τ . If Equation 5.1.1 holds then we can write ζ as the product of the constituent contrast contributions:

$$\zeta(\tau) = \zeta_{\text{D}}(\tau)\zeta_{\text{S}}(\tau)\zeta_{\text{Z}}(\tau), \quad (5.1.2)$$

In this section we will introduce the model for each contribution in turn. Other broadening mechanisms can influence plasma emission (e.g. Natural, resonance, Van der Waals and opacity broadening), but these are negligible for Balmer lines in scrape-

off layer/divertor conditions and so will not be considered any further here (see Section 6.2 of [23] and also [44]).

5.1.1 Doppler Broadening

In general, Doppler broadening encodes the velocity distribution function (VDF) of the emitting species into the observed lineshape. For a Maxwellian VDF, it can be shown that the local Doppler-broadened lineshape $g_D(\nu)$, centred at frequency ν_0 , has a Gaussian profile[43]:

$$g_D(\nu) = \frac{1}{\nu_0} \left(\frac{\pi v_{\text{th}}^2}{c^2} \right)^{-1/2} \exp \left(- \frac{(\nu - \nu_0)^2 c^2}{\nu_0^2 v_{\text{th}}^2} \right). \quad (5.1.3)$$

Here, v_{th} is the thermal velocity:

$$v_{\text{th}} = \sqrt{\frac{2k_B T_i}{m_i}}, \quad (5.1.4)$$

where T_i and m_i are the temperature and mass respectively of the emitting species¹. It is noted that the Doppler full width at half maximum (FWHM) of the line is given by

$$\Gamma_{\text{FWHM},\nu} = \sqrt{\frac{8 \ln(2) \nu_0^2 k_B T_i}{c^2 m_i}}, \quad (5.1.5)$$

and so the width of the Doppler-broadened line for hydrogen will be a factor $\sqrt{2}$ larger than for deuterium.

The corresponding degree of coherence due to Doppler broadening γ_D can be found using Equation 2.2.21 and the standard result that the Fourier transform of a Gaussian function is another Gaussian (see Appendix A). Since the lineshape in Equation 5.1.3 is symmetric about ν_0 , γ_D is real for the case of a Maxwellian VDF and Doppler broadening contributes a contrast term ζ_D only. It is simple to show that, as a function of interferometer group delay $\hat{\phi}_0$, this is given by[63]:

$$\zeta_D = \exp \left(- \frac{k_B T_i}{2m_i c^2} \hat{\phi}_0^2 \right). \quad (5.1.6)$$

Of course, a non-Maxwellian VDF invalidates Equation 5.1.6 and, in general, CI is not well-suited to characterising an unknown VDF. That said, early CI work demonstrated that a multi-delay instrument² could be used to infer the presence of a VDF with an asymmetric tail on the H1 stellarator[64].

¹We use T_i to generally refer to the kinetic temperature of species i . In general this may be the temperature of an ion species, but for H and D it is the temperature of the neutral atoms.

²Actually, a single delay CI instrument repeating its measurements at different delays.

5.1.2 Stark Broadening

The presence of an external electric field perturbs the electron energy states in an atom, splitting and shifting the observed line emission—this is known as the Stark effect[122]. For hydrogen, the line splitting is proportional to the external field strength and is symmetric about the initially degenerate line, so there is no net shift in line frequency. When the external field is due to nearby charged particles in a plasma, the observed line profile is smoothly broadened with a width that increases with plasma density and whose shape has a weak dependence on temperature. Modelling the Stark-broadened line profile is complicated by the long-range Coulomb interactions of the many particles that must be considered. However, the field is mature and simulation results for hydrogen lines are in very good agreement with experiments across a wide range of temperatures and densities[43, 123, 124]. A detailed review of the physics models for Stark broadening is beyond the scope of this work so we will discuss broad arguments and practical considerations only. A useful and accessible review that focusses on plasma diagnostics can be found by Gigosos[125].

The earliest models for Stark broadening are based on two extreme approximations: the *impact* approach and the *quasi-static* approach[33, 44]. The impact approach assumes that the perturbing particles move quickly compared to the time over which the atoms are actually radiating photons, occasionally interrupting the process. The resulting discontinuities in the emitted wave train necessarily broaden the frequency spectrum. If these interruptions arrive according to Poisson statistics then the line profile has a Lorentzian form:

$$g_S(\nu) = \frac{1}{\pi} \left(\frac{\Gamma_{\text{HWHM},\nu}}{(\nu - \nu_0)^2 + \Gamma_{\text{HWHM},\nu}^2} \right). \quad (5.1.7)$$

Here, the profile is area-normalised and $\Gamma_{\text{HWHM},\nu}$ is the half width half maximum in units of frequency. Since $\Gamma_{\text{HWHM},\nu}$ is proportional to the collision frequency, the profile width scales with perturber density. It is often appropriate to treat electron perturbers using the impact approach due to their relatively high mobility[33].

Conversely, the quasi-static approach assumes that the perturbing particles move slowly compared to the time over which the atoms are radiating—their electric field is quasi-static. The resulting lineshape can then be found via superposition of all the possible Stark-split spectra, weighted by the statistical distribution of electric fields experienced by the emitting atom. This is analogous to how a Doppler broadened profile is the superposition of many Doppler-shifted lines, weighted by the emitter VDF. Since this treatment was first proposed, more and more accurate electric field distributions have been derived, accounting for Debye shielding, ion-ion correlations and other effects[44]. It can be shown that the quasi-static approach predicts lineshape

n_{upper}	line name	λ_0 (nm)	Model 1	Model 2		
			$\alpha_{1/2}$	C_i	a_i	b_i
4	β	486.1	0.08	8.425e-18	0.7803	0.050
5	γ	434.0	0.09	1.310e-15	0.6796	0.030
6	δ	410.2	0.17	3.954e-16	0.7149	0.028

Table 5.2: Coefficients for calculating the Stark-broadened profile for three Balmer lines for two of the models described in this section. For Model 1, the values for Ba- β and Ba- δ come from tabulated calculations[123] while the value for Ba- γ is empirical and was used in a previous Stark CI investigation[97]. For Model 2, the values are from [47].

wings which fall off with $g_S(\nu) \propto (\nu - \nu_0)^{-5/2}$ and with a line width that is proportional to $n^{2/3}$ where n is the density of the perturber[33, 47].

It is often appropriate to treat the ion perturbers using this quasi-static approach. The ‘standard theory’ of Stark broadening combines an impact treatment of the electron perturbers with a quasi-static treatment of the ion perturbers. To a very good approximation, the Stark-broadened lineshapes of one-electron systems are symmetric[123] and so affect the CI contrast only. In this section we will introduce three existing lineshape models of varying complexity and compare the corresponding CI contrast profiles.

Model 1: Lorentzian with Width $\propto n_e^{2/3}$

The simplest lineshape model considered here assumes that the lineshape has a Lorentzian profile with a width that scales according to a combination of ion quasi-static broadening and electron impact broadening. This gives a simple relationship between the full width half maximum $\Gamma_{\text{FWHM},\lambda}$ (in units of m) of the lineshape and the plasma density n_e (in units of m^{-3})[33]:

$$\Gamma_{\text{FWHM},\lambda} \approx k_S \alpha_{1/2} n_e^{2/3}. \tag{5.1.8}$$

Here, $k_S = 2.50 \times 10^{-23}$ is a constant and $\alpha_{1/2}$ is the key result of calculations detailed in [123]. Representative values of $\alpha_{1/2}$ for three Balmer lines considered for this work are listed in Table 5.2. For H_β and H_δ these values can be used to infer density from measured $\Gamma_{\text{FWHM},\lambda}$ to within $\pm 15\%$ [33]. Converting Equation 5.1.8 from wavelength to frequency we can write

$$\Gamma_{\text{FWHM},\nu} \approx \frac{c}{\lambda_0^2} k_S \alpha_{1/2} n_e^{2/3}. \tag{5.1.9}$$

This model's assumption of a Lorentzian lineshape leads to a simple expression for the Stark contribution to the CI contrast profile ζ_S . Using Equations 5.1.9 and 5.1.7, with the definition for degree of coherence from 2.2.2 and the standard result that the Fourier transform of Lorentzian is an exponential decay (see Appendix A), we can write

$$\zeta_S \approx \exp\left(-\frac{k_S \alpha_{1/2} n_e^{2/3}}{2\lambda_0} |\hat{\phi}_0|\right). \quad (5.1.10)$$

Figure 5.1 then plots the lineshapes and corresponding contrast profiles for the H_γ and H_δ Balmer lines for $T_i = T_e = 1$ eV and three different densities for each of the three lineshape models introduced in this section.

While the width scaling with $n_e^{2/3}$ is likely to be accurate for Stark lineshape Model 1, the assumed Lorentzian form of the lineshape is a first order approximation and so will introduce error at the line centre and wings. There are two reasons for considering this model here despite its shortcomings. First, it has been used in the previous density CI investigations[97, 126] and the associated CI error when compared to more accurate models has not been considered in this context. Secondly, the simple analytical form of ζ_S makes it easy to express the basic arguments of density CI in Section 5.2.1.

Model 2: Model Microfield Method (MMM)

Modern calculations typically consider the intermediate broadening regime: when it is not valid to consider either the electron or ion perturbations as entirely static or instantaneous. One such method is the model microfield method (MMM), where the dynamics of the statistical distribution describing the ion field are now accounted for to some extent[127]. An extensive database of MMM simulation results for hydrogen lines exists[124] and has been extensively used in tokamak spectroscopic studies, for example on MAST[24], JET[47], NSTX[45], Asdex Upgrade[21] and Alcator C-Mod[128]. These tables are intended for use in astrophysics and so cover an extremely wide range of densities: $10^{16} < n_e < 10^{25} \text{ m}^{-3}$ and temperatures: $0.2 < T_e < 112$ eV.

Due to the coarseness of the frequency grid of these tabulated profiles, subsequent work has been done to parametrise the results for accurate and fast interpolation at any frequency[47]. This parametrisation assumes a modified Lorentzian lineshape profile of the form

$$g_S(\nu) \propto \frac{1}{(\nu - \nu_0)^{5/2} + \Gamma_{\text{FWHM},\nu}^{5/2}}. \quad (5.1.11)$$

The FWHM (in wavelength space) is expressed as a function of n_e and T_e as

$$\Gamma_{\text{FWHM},\lambda} = C_i \frac{n_e^{a_i}}{T_e^{b_i}} \quad (5.1.12)$$

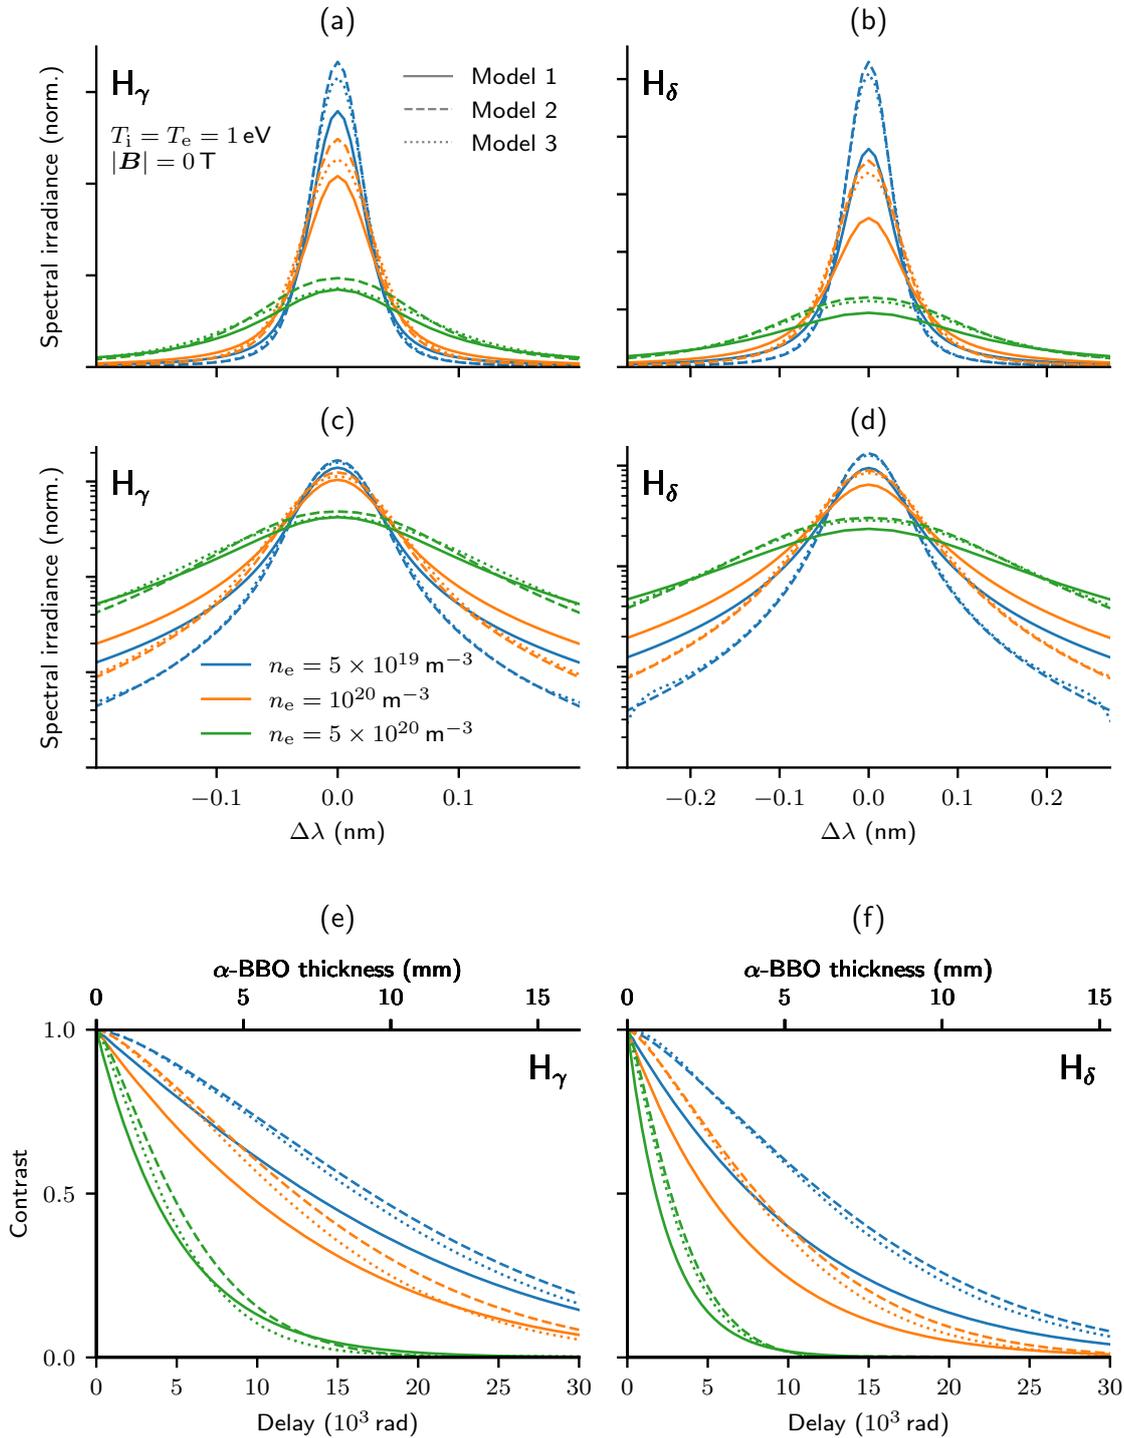


Figure 5.1: Balmer lineshapes modelled using the three different models introduced in this section for a range of densities at $T = 1$ eV. (a) and (b) show H $_{\gamma}$ (434.0 nm) and H $_{\delta}$ (410.2 nm) respectively plotted on a linear scale. (c) and (d) are the same lineshapes plotted on a logarithmic scale. (e) and (f) are the corresponding contrast profiles as a function of interferometer delay.

where C_i , a_i and b_i are coefficients depending on upper principal quantum number (shorthand i). Table 5.2 lists the coefficient values for the H_β , H_γ and H_δ lines. The range over which the coefficients are accurate is $10^{19} \leq n_e \leq 10^{21} \text{ m}^{-3}$. This parametrisation recovers the original tabulated widths to better than $\pm 10\%$ and maintains a good fit in the wings[47]. There is a loss of resolution at the line centre, but this part of the profile is typically obscured by Doppler broadening. For this model there is no simple analytical expression for the contrast contribution ζ_S .

Model 3: Numerical Simulation

State-of-the-art models for Stark broadening use numerical simulations to predict the time-dependent electric microfield experienced by an emitter and then numerically integrate over the Schrödinger equation to produce the lineshape. Modelling from fundamental principles in this way allows the presence of an external magnetic field (the Zeeman effect, see Section 5.1.3) to be modelled with the Stark broadening in a self-consistent way. This type of simulation is computationally intensive. It is not suitable for routine use in plasma diagnostics and so we again must use a database of tabulated lineshapes as a lookup table. For the third Stark broadening model considered here, we use tabulated lineshapes of this kind that are explicitly intended for use in divertor plasma spectroscopy[129, 130]. These lineshapes have a finer frequency grid than model 2, but interpolate over a coarser grid in temperature and density parameter space, which can result in artefacts.

Figure 5.1 compares the lineshapes and calculated contrast profiles for the three Stark broadening models introduced in this section. It can be seen that the Lorentzian profile ($g_S(\nu) \propto (\nu - \nu_0)^{-2}$) of model 1 results in broader wings than the more realistic MMM profile ($g_S(\nu) \propto (\nu - \nu_0)^{-5/2}$) of model 2 and the numerical simulation of model 3. The contrast profiles show that Model 1's broad wings result in a significantly lower contrast at smaller delays than models 2 and 3. Since models 2 and 3 will produce more accurate lineshapes, it follows that using the Lorentzian approximation results in an underestimate of n_e . For this model, as with Model 2, there is no simple analytical expression for the contrast contribution ζ_S .

5.1.3 Zeeman Splitting

An external magnetic field will also perturb the electron energy states in an atom, splitting the observed line emission[122]. In general, the degree of splitting depends on the field strength while the relative intensities and polarisation states of the line components depend on the field direction relative to the observer. Measurements

Model	Doppler profile	Stark profile	Zeeman splitting	Density range (m^{-3})	\mathbf{B} -field range (T)
1	Gaussian	Lorentzian	Independent, strong-field	-	$ \mathbf{B} \gtrsim 1$
2	Gaussian	MMM	Independent, strong-field	$10^{19} \lesssim n \lesssim 10^{21}$	$ \mathbf{B} \gtrsim 1$
3	Gaussian	Computer simulation, Stark-Zeeman		$10^{19} < n < 10^{22}$	$0 \leq \mathbf{B} \leq 5$

Table 5.3: A summary of the three Balmer lineshape models used in this chapter.

of Zeeman splitting have been used to infer the poloidal magnetic field in tokamaks[131, 132] and in the ITER CI design it is hoped that the difference in polarisation state between the Zeeman-split emission and reflections from the vessel wall will help reduce tomography artefacts[99]. In this work we consider Zeeman splitting only as a ‘nuisance parameter’ in the measured CI contrast. It must be factored out to infer an accurate n_e . Calculating the Zeeman contrast is complicated by the polarised nature of the emission, as the CI instrument is polarisation sensitive via its front polariser. These effects have been discussed briefly in relation to CI measurements of impurity emission in previous work[64, 69, 133]. It is noted that, since Zeeman splitting is symmetric about the central frequency ν_0 , it contributes only to the CI contrast.

The external magnetic field \mathbf{B} interacts with the orbital and spin magnetic moments of the electron, breaking the degeneracy of Δm_l transitions. Selection rules dictate that, in any transition between states, the change in the magnetic quantum number must satisfy $\Delta m_l \in [-1, 0, 1]$. We call the line corresponding to $\Delta m_l = 0$ the π component and its frequency is unchanged at $\nu_\pi = \nu_0$. The $\Delta m_l = \pm 1$ transitions are called the σ^\pm components and have shifted frequencies $\nu_{\sigma^\pm} = \nu_0 \pm \Delta\nu_Z$, where[122, 134]

$$\Delta\nu_Z \approx \frac{e}{4\pi m_e} |\mathbf{B}|. \quad (5.1.13)$$

Here, e is the elementary charge and m_e is the electron mass. However, the orbital and spin magnetic moments also self-interact (spin-orbit coupling) to produce the fine structure splitting observed in the absence of \mathbf{B} . The treatment of the Zeeman effect that leads to Equation 5.1.13 assumes that the interaction with \mathbf{B} dominates this self-interaction[122], and is called the ‘strong-field’ Zeeman effect³. The validity of the strong-field treatment can be checked by comparing the predicted Zeeman splitting to the observed fine structure splitting. For example, for H_γ emission (434.0 nm), the separation of the fine structure components from the line centre is ~ 0.003 nm[115],

³Also known as the Paschen-Back effect.

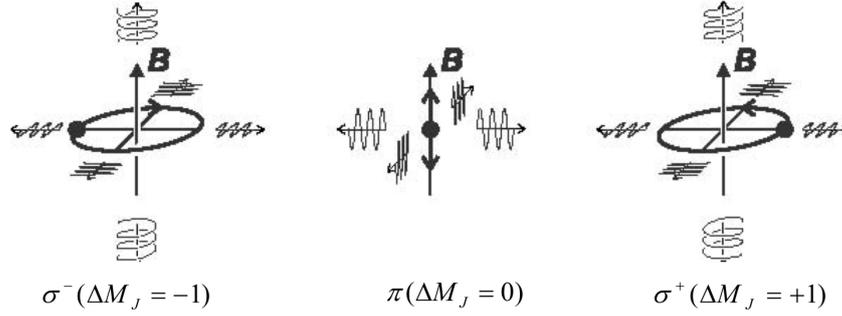


Figure 5.2: The polarisation states of the π and σ^\pm Zeeman components for longitudinal and transverse views.

while a 1 T field will cause a split from the line centre of ~ 0.009 nm. The strong-field treatment will be used as a starting point in this analysis. Codes exist for calculating the Zeeman multiplet feature in the more complicated case when the spin-orbit coupling cannot be ignored (e.g. ADAS603[135]). It will also be assumed, unless stated otherwise, that the Zeeman splitting is independent of the Doppler and Stark broadening mechanisms.

For a transverse view of the magnetic field—when the angle between the line of sight \hat{l} and field \mathbf{B} is $\chi = 90^\circ$ —the π component is linearly polarised parallel to \mathbf{B} and the σ^\pm components are linearly polarised perpendicular to \mathbf{B} . For a longitudinal view—when $\chi = 0^\circ$ —the π component is not observed and the σ^\pm components are right and left-hand circularly polarised respectively. This is shown schematically in Figure 5.2. For intermediate views of the field ($0^\circ < \chi < 90^\circ$) the π and σ^\pm components' intensity ratio and polarisation states are more complicated[134]. The Stokes parameters for each component are listed in Table 5.4. Here, the Stokes parameters are defined such that a linear horizontal polarisation state (positive S_1) oscillates parallel to the the projection of \mathbf{B} onto the CI instrument's x - y imaging plane. So the π component is linearly polarised parallel to the projection of \mathbf{B} and

Stokes parameter	$\pi (\Delta m_l = 0)$	$\sigma^\pm (\Delta m_l = \pm 1)$
S_0	$\frac{1}{2} \sin^2 \chi$	$\frac{1}{4}(1 + \cos^2 \chi)$
S_1	$\frac{1}{2} \sin^2 \chi$	$-\frac{1}{4} \sin^2 \chi$
S_2	0	0
S_3	0	$\pm \frac{1}{2} \cos \chi$

Table 5.4: Stokes parameters for the three Zeeman split line components, assuming the strong-field Zeeman effect is a valid treatment[134]. Here, χ is the angle between the line of sight and the magnetic field vector \mathbf{B} . The total irradiance of the three line components is independent of χ and has been normalised to 1.

the σ^\pm components are each a mix of circularly polarised light and linearly polarised

light that is oriented perpendicular to the projection of \mathbf{B} . While the relative strength of the π and σ^\pm components depends on χ , the total intensity of all components is independent of χ and in Table 5.2 is normalised to 1.

To model the transmission of the Zeeman-split line through the CI instrument, we have to express the line as a (centred) Stokes spectrum. We do this by multiplying the π and σ^\pm Stokes parameters from Table 5.4 by delta functions $\delta(\nu)$ and $\delta(\nu \mp \Delta\nu_Z)$ respectively and then summing the three line components. The centred, Zeeman-split Stokes spectrum is then

$$\mathbf{S}_Z^c(\nu) = \begin{pmatrix} \frac{1}{2} \sin^2 \chi \delta(\nu) + \frac{1}{4} (1 + \cos^2 \chi) [\delta(\nu - \Delta\nu_Z) + \delta(\nu + \Delta\nu_Z)] \\ \frac{1}{2} \sin^2 \chi \delta(\nu) - \frac{1}{4} \sin^2 \chi [\delta(\nu - \Delta\nu_Z) + \delta(\nu + \Delta\nu_Z)] \\ 0 \\ \frac{1}{2} \cos \chi [\delta(\nu - \Delta\nu_Z) - \delta(\nu + \Delta\nu_Z)] \end{pmatrix}. \quad (5.1.14)$$

The top row of Figure 5.3 plots example Zeeman-split H_γ Stokes spectra for different $|\mathbf{B}|$ and χ . We now have to rewrite the expression for the centred lineshape from Equation 5.1.1 as the centred Stokes spectrum:

$$\mathbf{S}^c(\nu) = g_D^c(\nu) * g_S^c(\nu) * \mathbf{S}_Z^c(\nu) \quad (5.1.15)$$

Where the convolution operation is element-wise over the four Stokes parameters. This result is normalised such that the total irradiance $S_0 \equiv \int S_0(\nu) d\nu = 1$. The bottom row of Figure 5.3 plots $\mathbf{S}^c(\nu)$ for the same $|\mathbf{B}|$ and χ values as the top row and with realistic Doppler and Stark broadening. It is noted that the Zeeman-split line is always net unpolarised ($S_1 = S_2 = S_3 = 0$) when integrating over the full multiplet feature.

Zeeman Contrast

We will now use Equation 5.1.15 to find the Zeeman contribution ζ_Z to the measured CI contrast. First, we need to generalise Equation 5.1.14 for $\mathbf{S}_Z(\nu)$ to cases where the projection of \mathbf{B} onto the CI x - y imaging plane does not align with the x -axis but instead makes some angle ρ_B with it. We can do this by multiplying $\mathbf{S}_Z(\nu)$ by the Mueller matrix for anti-clockwise frame rotation (Equation 2.3.6) of angle $-\rho_B$: $\mathbf{R}(-\rho_B)\mathbf{S}(\nu, \chi)$. An important additional effect to consider is that the Zeeman-split Stokes spectrum emitted by the plasma is modified by the front polariser of the CI instrument. We account for this in the standard way by multiplying the incident Stokes spectrum by the Mueller matrix for a linear polariser whose transmission axis is at some angle ρ_{pol} to the x -axis. Putting this together, we can write the Zeeman-split

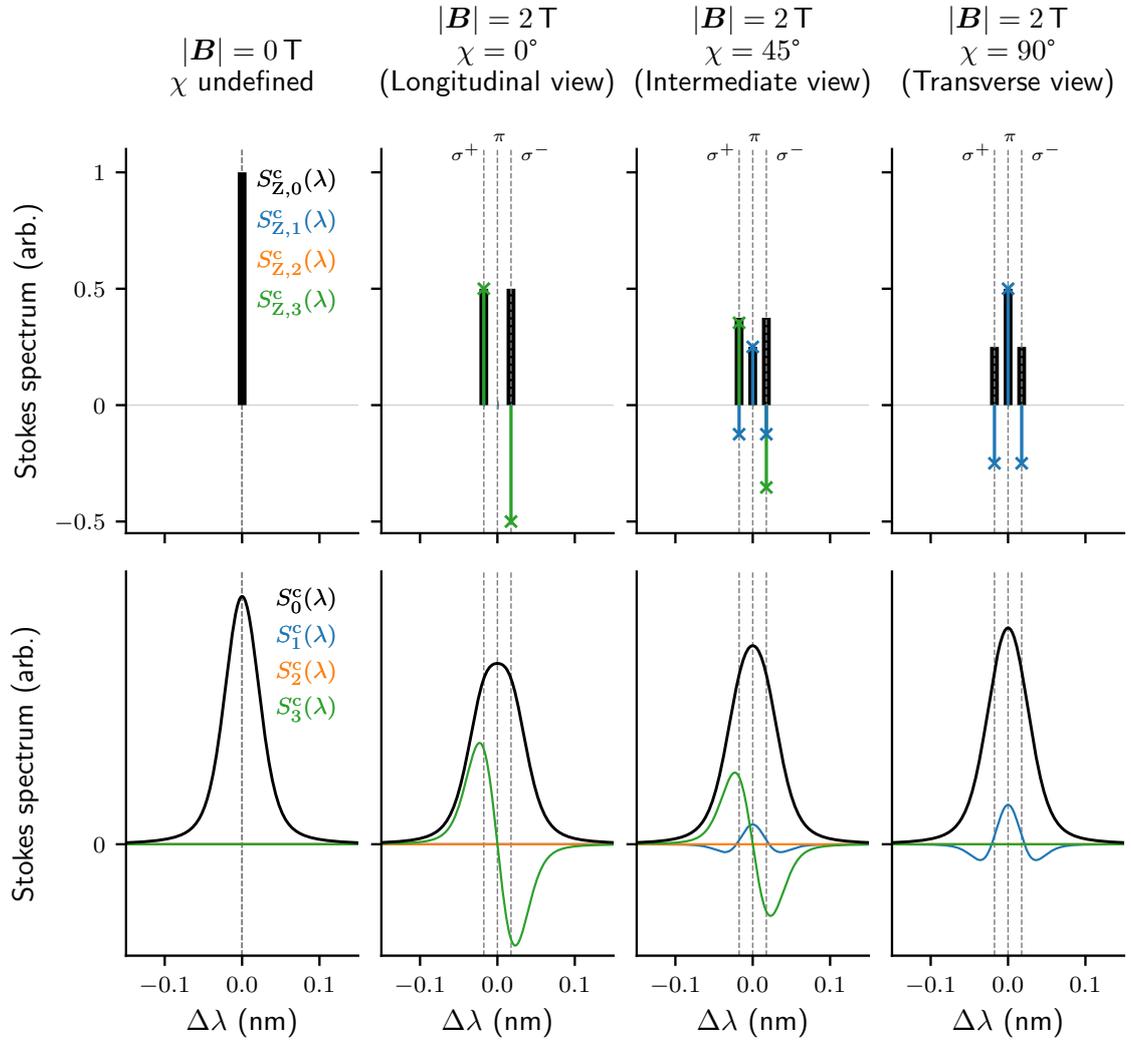


Figure 5.3: Modelled (centred) Stokes spectra for H_γ emission in the presence of a strong external magnetic field for various field strengths $|B|$ and view angles χ . The top row shows the Zeeman-split line components only while the bottom row convolves these components with modelled Doppler-broadened and Stark-broadened line profiles for divertor-relevant conditions: $n_e = 5 \times 10^{19} \text{ m}^{-3}$ and $T_e = T_i = 1 \text{ eV}$.

Stokes spectrum, modified by the front polariser, as:

$$\mathbf{S}_{Z,\text{mod}}^c(\nu) = \mathbf{M}_{\text{pol}}(\rho_{\text{pol}})\mathbf{R}(-\rho_B)\mathbf{S}_Z^c(\nu). \quad (5.1.16)$$

Since $\mathbf{S}_{Z,\text{mod}}^c(\nu)$ is purely linearly polarised, it is only the total spectral irradiance of (i.e. the S_0 component) that is encoded in the coherence measured by CI. We will call this component $g_{Z,\text{mod}}^c(\nu)$. An expression for $g_{Z,\text{mod}}^c(\nu)$ can be found by evaluating 5.1.16 and taking only the S_0 component, leaving

$$g_{Z,\text{mod}}^c(\nu) \approx \sin^2 \chi \cos^2(\rho_\Delta) \delta(\nu) + \frac{1}{4} \left(1 + \cos^2 \chi - \sin^2 \chi \cos[2\rho_\Delta]\right) \left[\delta(\nu - \Delta\nu_Z) + \delta(\nu + \Delta\nu_Z)\right]. \quad (5.1.17)$$

Here, we have defined $\rho_\Delta \equiv \rho_B - \rho_{\text{pol}}$. So the spectrum encoded in the CI fringes, modified by the front polariser, depends only on the difference between the orientation of the magnetic field projection and the orientation of the front polariser. The modified spectrum, including all broadening effects can be written

$$g_{\text{mod}}^c(\nu) = g_{\text{D}}^c(\nu) * g_{\text{S}}^c(\nu) * g_{Z,\text{mod}}^c(\nu). \quad (5.1.18)$$

This means that the expression for total contrast as the product of the constituent contributions from Equation 5.1.2 holds and the Zeeman contrast can be calculated by taking the magnitude of the Fourier transform of $g_{Z,\text{mod}}^c(\nu)$. The result, using the definition of $\Delta\nu_Z$ from Equation 5.1.13, can be written

$$\zeta_Z \approx \left| \sin^2 \chi \cos^2(\rho_\Delta) + \frac{1}{2} \left(1 + \cos^2 \chi - \sin^2 \chi \cos[2\rho_\Delta]\right) \cos\left(\frac{e|\mathbf{B}|\hat{\phi}_0}{4\pi m_e \nu_0}\right) \right|. \quad (5.1.19)$$

This is plotted as a function of delay ϕ_0 and field strength $|\mathbf{B}|$ for a range of viewing geometries in Figure 5.4. For a longitudinal view ($\chi = 0^\circ$), Equation 5.1.19 reduces to

$$\zeta_Z \approx \left| \cos\left(\frac{e|\mathbf{B}|\hat{\phi}_0}{4\pi m_e \nu_0}\right) \right|. \quad (5.1.20)$$

Similarly, for a transverse view ($\chi = 90^\circ$) we find that

$$\zeta_Z \approx \left| \cos^2(\rho_\Delta) + \sin^2(\rho_\Delta) \cos\left(\frac{e|\mathbf{B}|\hat{\phi}_0}{4\pi m_e \nu_0}\right) \right|. \quad (5.1.21)$$

So for a longitudinal view there is no dependence on front polariser orientation. This makes sense as only the σ^\pm components will be observed from this view and they will be circularly polarised so the normalised spectrum will be unaffected by the polariser. For a transverse view, front polariser orientation can be chosen such that

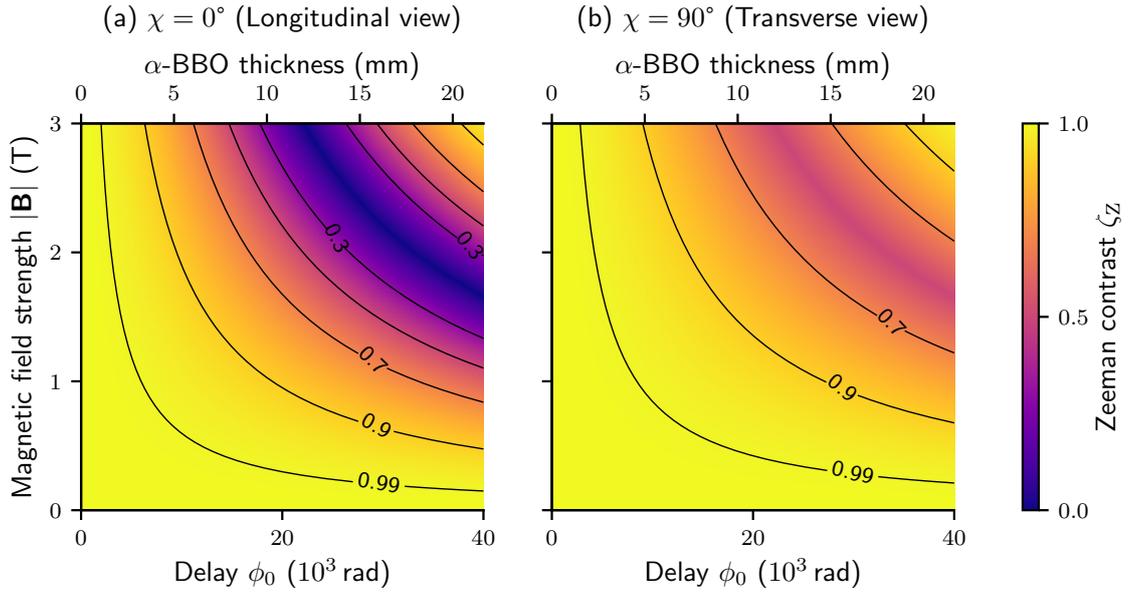


Figure 5.4: Modelled drop in contrast due to Zeeman splitting of H_γ in the presence of a strong external magnetic field, plotted as a function of field strength and interferometer delay. (a) For a longitudinal view of the field or, equivalently, a transverse view with $\rho_\Delta = 90^\circ$. (b) For a transverse view with $\rho_\Delta = 45^\circ$.

$\rho_\Delta = 0^\circ$, extinguishing the now linearly polarised σ^\pm components and leaving $\zeta_Z = 1$. As shown in Figure 5.4, the Zeeman contrast drop can be unavoidable and significant for longitudinal / intermediate views, large ϕ_0 and/or large $|\mathbf{B}|$. If the magnetic geometry of an experiment is simple and a transverse view is practical, then a drop in ζ_Z can be avoided, regardless of field strength. For other views, where the effect cannot be avoided outright, it can be mitigated by choice of front polariser orientation and accounted for via modelling if the field strength and geometry is known.

5.2 Measurement Principle

Before diving into how the accurate numerical lineshape models can be used to infer n_e from coherence imaging contrast measurements, we will briefly cover the basics of density CI using a simplified analytical model.

5.2.1 Basic Arguments: Choice of Delay and Balmer Line

For a CI contrast measurement made at a single delay, the competing Stark and Doppler broadening mechanisms cannot be unpicked (let's assume here that the Zeeman contribution can be modelled and factored out). We must therefore make a

careful choice of delay and targeted Balmer line to minimise systematic error in the inferred n_e due to Doppler broadening while also maximising sensitivity to changes in n_e . We will assume, for now, that we have no prior knowledge of the plasma state. By combining the Lorentzian approximation for the Stark contrast (Model 1, Equation 5.1.10) with a Gaussian profile for the Doppler contrast (Equation 5.1.6), we can write the total contrast as

$$\zeta = \zeta_D \zeta_S \approx \exp\left(-\left[\frac{n_e}{n_c}\right]^{2/3} - \frac{T_i}{T_c}\right). \quad (5.2.1)$$

Here, we follow existing work[64, 97] and introduce n_c and T_c , the ‘characteristic’ density and temperature of the instrument respectively. Both characteristic terms are functions of ϕ_0 . The definition of n_c can be inferred from Equation 5.1.10:

$$n_c = \left(\frac{2\lambda_0}{k_S \alpha_{1/2} |\hat{\phi}_0|}\right)^{3/2}. \quad (5.2.2)$$

Similarly, the definition of T_c can be inferred from Equation 5.1.6:

$$T_c = \frac{2m_i c^2}{k_B \hat{\phi}_0^2}. \quad (5.2.3)$$

Figure 5.5 shows the modelled line spectra and contrast profiles for H_γ emission for relatively high-density ($\sim 3 \times 10^{20} \text{ m}^{-3}$) and low temperature ($T_i = 0.5 \text{ eV}$) conditions. In this regime the Doppler contrast (shown) can be ignored ($\zeta_D \approx 1$) for the smaller delays where the contrast is most sensitive to changes in n_e . In this regime, n_e can be inferred from ζ as $n_e \approx n_c (-\ln \zeta)^{3/2}$.

The first consideration for choice of fixed delay is maximising the sensitivity of the measured contrast ζ to changes in n_e . From Equation 5.2.1, we calculate this sensitivity as

$$\frac{\partial \zeta}{\partial n_e} \propto \frac{1}{n_e} \left(\frac{n_e}{n_c}\right)^{2/3} \exp\left(-\left[\frac{n_e}{n_c}\right]^{2/3}\right) \quad (5.2.4)$$

In Figure 5.6, this is plotted as a function of n_e/n_c and normalised. It can be seen that $\frac{\partial \zeta}{\partial n_e}$ reaches a maximum when the delay is chosen such that $n_c = n_e$. This (group) delay can then be written as a function of targeted n_e as

$$|\hat{\phi}_{0,\text{sens}}| \approx \frac{2\lambda_0}{k_S \alpha_{1/2} n_e^{2/3}}. \quad (5.2.5)$$

So larger/smaller delays are better tuned to measure lower/higher n_e respectively. It also follows that, for equivalent delays, observing higher n_{upper} Balmer lines tunes the sensitivity to lower densities. In figure 5.5, $\hat{\phi}_{0,\text{sens}}$ is indicated for H_γ at $n_e = 3 \times 10^{20} \text{ m}^{-3}$, corresponding to an α -BBO waveplate of thickness $\approx 4 \text{ mm}$ (viewed at normal incidence). When $\zeta_D \not\approx 1$, $\hat{\phi}_{0,\text{sens}}$ will be lower than predicted by Equation

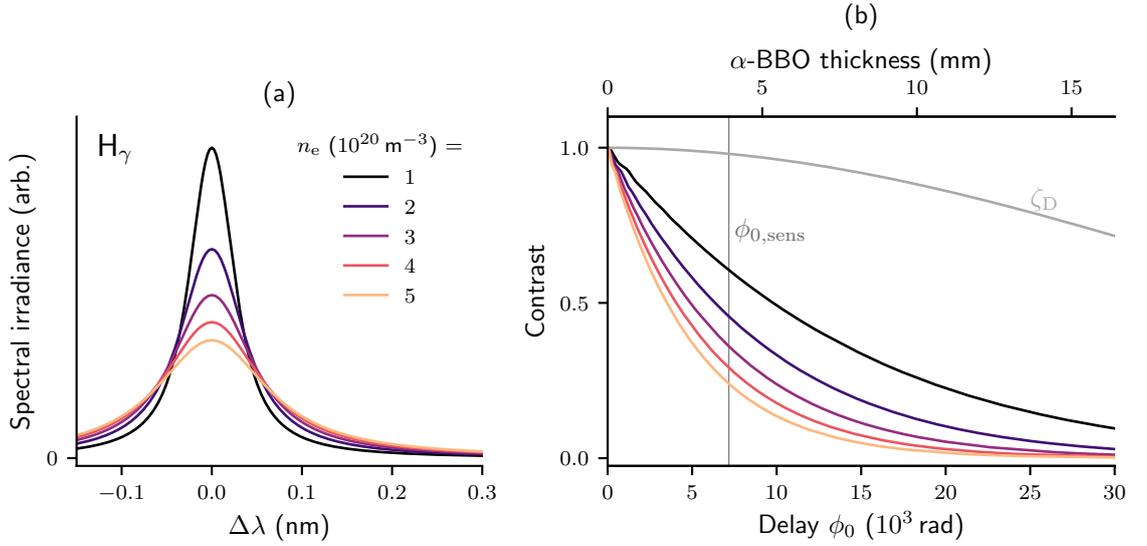


Figure 5.5: (a) Modelled lineshapes for H_γ emission over a range of electron densities for $T_i = 0.5$ eV. (b) The corresponding contrast profiles as a function of delay. Also shown in (b) is the Doppler contrast ζ_D and the optimum delay $\phi_{0,\text{sens}}$ for maximising sensitivity to changes in n_e about $3 \times 10^{20} \text{ m}^{-3}$.

5.2.5 but this is a second order effect compared to the systematic error introduced, which we will consider next.

A second constraint on the choice of delay is that the systematic error due to Doppler broadening δn_e must be small. In this simplified treatment, δn_e is the difference between the inferred density $n_c(-\ln \zeta)^{3/2}$ and the true density n_e . Expressing this error as a fraction of n_e we can write

$$\frac{\delta n_e}{n_e} \approx \left(1 + \frac{k_B}{c^2 k_S} \left[\frac{T_i \lambda_0 |\hat{\phi}_0|}{n_e^{2/3} m_i \alpha_{1/2}} \right] \right)^{3/2} - 1, \quad (5.2.6)$$

where the constants have been grouped outside of the square brackets. Of course, $\delta n_e/n_e$ is larger for higher T_i and lower n_e . As for the instrumental parameters, $\delta n_e/n_e$ increases with delay. This can be seen in Figure 5.7, where the constituent Doppler and Stark components of the spectrum and contrast profile are plotted for H_γ at $n_e = 10^{20} \text{ m}^{-3}$ and $T_i = 1$ eV. Also plotted is $\delta n_e/n_e$, calculated using Equation 5.2.6 and $\hat{\phi}_{0,\text{sens}}$. At $\hat{\phi}_{0,\text{sens}}$ the systematic error is significant (20%), showing that when Stark and Doppler effects are comparable, sensitivity may have to be sacrificed by choosing a smaller delay to reduce systematic error to an acceptable level. Finite SNR sets a limit on the highest contrast that can be resolved and so then also on the smallest delay that can be used.

A consideration in the choice of Balmer line for Stark CI is the presence of contaminating emission. Figure 5.8 shows measured spectra from the JET tokamak

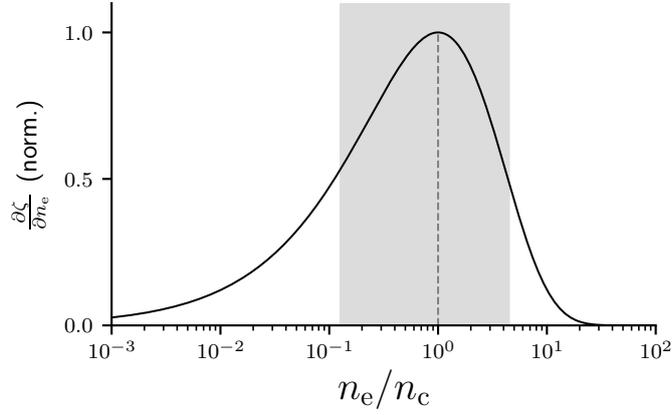


Figure 5.6: The sensitivity of the measured contrast to the ratio of the electron density to the characteristic instrument density. Shaded region highlights approximate dynamic range for n_e measurement.

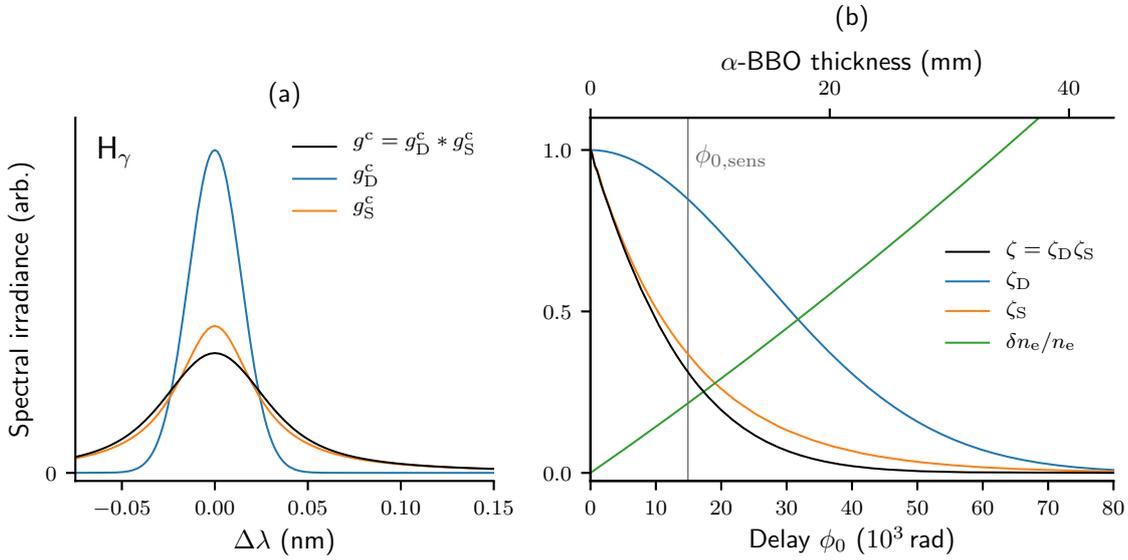


Figure 5.7: (a) Modelled lineshape (blue) for H_γ emission for $n_e = 10^{20} \text{ m}^{-3}$ and $T_i = 1 \text{ eV}$. The constituent Doppler and Stark components are shown in blue and orange respectively. (b) The corresponding contrast profiles as a function of delay. The vertical dashed line indicates the optimum delay for sensitivity, while the green line is the fractional systematic error in the inferred n_e due to Doppler broadening.

divertor (from the KT3 spectrometer diagnostic[136]) in the region of D_δ (410.1 nm) during and after nitrogen-seeding experiments. The N III emission at the wings preclude density CI of the line during nitrogen seeding experiments. Even in the case of trace nitrogen presence (blue), the presence of the lines can lead to unacceptable systematic error in the inferred density if not properly accounted for.

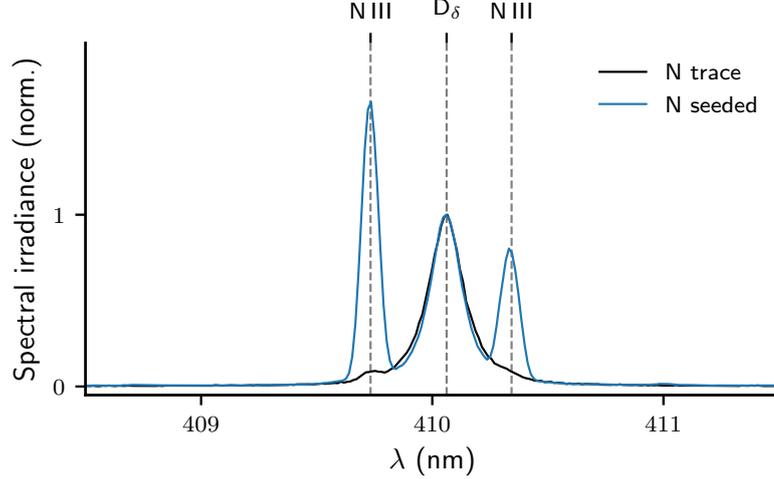


Figure 5.8: Measured spectra from the JET tokamak divertor during L-mode nitrogen seeding experiments (blue, #90423, $n_e \approx 2.9 \times 10^{20} \text{ m}^{-3}$) and in L-mode without N seeding (black, #90531, $n_e \approx 3.4 \times 10^{20} \text{ m}^{-3}$). Measured using the KT3 spectrometer diagnostic[136]. Data courtesy of Bart Lomanowski.

5.2.2 Fitting to Contrast Profiles

We will now consider the more general case of inferring n_e from CI measurements of Balmer emission made at one or more interferometer delays and with a lineshape whose modelled contrast cannot be expressed analytically. For this curve fitting problem we will again use the framework of Bayesian parameter estimation (see Appendix C) as it allows for easy incorporation of prior information on the model parameter values. Measured contrast $\zeta(\tau_0; \boldsymbol{\theta})$ is a function of interferometer delay τ_0 at the central frequency ν_0 of the targeted Balmer line and the set of parameters which determine the observed lineshape $\boldsymbol{\theta} = \{n_e, T_i, T_e, |\mathbf{B}|, \chi, \rho_\Delta\}$. The contrast observed at τ_0 for lineshape $g^c(\nu; \boldsymbol{\theta})$ is calculated as

$$\zeta(\tau_0; \boldsymbol{\theta}) \approx \left| \int_{-\infty}^{\infty} g^c(\nu; \boldsymbol{\theta}) \exp(2\pi i \nu \hat{\tau}_0) d\nu \right|, \quad (5.2.7)$$

Where $\hat{\tau}_0 = \kappa_0 \tau_0$ is the group delay, in units of time. As discussed, this equation must be evaluated numerically for the most accurate lineshape models (models 2 and

3 in Section 5.1.2). An example modelled contrast profile is plotted in 5.9(a) for H_γ emission from a plasma with $n_e = 10^{20} \text{ m}^{-3}$, $T_i = T_e = 1 \text{ eV}$ and $|\mathbf{B}| = 0 \text{ T}$. Lineshape model 3 is used here. Also plotted in Figure 5.9(a) are simulated measurements of the contrast profile made at four delays. These measurements are modelled with zero-mean additive Gaussian noise with standard deviation $\sigma_\zeta = 0.02$ and the four delays are chosen to match a CI instrument in the Multi-delay 1 configuration (see Table 3.2). We will use this example profile and simulated measurements to demonstrate a contrast profile fitting tool developed for inferring n_e .

When curve fitting, quick evaluation of the model is important for repeated comparison with the measured data. To quickly map given values of τ_0 and $\boldsymbol{\theta}$ to the modelled $\zeta(\tau_0; \boldsymbol{\theta})$ without repeatedly evaluating Equation 5.2.7, a set of lookup tables (LUTs) has been pre-calculated. Linear interpolation over these LUTs can then be used to quickly find $\zeta(\tau_0; \boldsymbol{\theta})$ for comparison to data. The LUTs in this work are generated on log-spaced grids with 400 elements in both n_e and T_i over the ranges $10^{19} \leq n_e \leq 5 \times 10^{21} \text{ m}^{-3}$ and $0.32 \leq T_i \leq 15 \text{ eV}$. It is assumed here and for the rest of this thesis that $T_i = T_e$. Figures 5.9(b)–(e) are contour plots showing 2-D slices through the LUT at four different delays, plotted as a function of n_e and T_i . The LUT shown is for H_γ emission generated using lineshape model 3 and the four delays correspond to the simulated measurements in Figure 5.9(a).

We will now consider how the LUTs can be used for rapid inference of plasma parameters from contrast measurements. Let \mathcal{D} be a set of contrast measurements $\{\mathcal{D}_1, \mathcal{D}_2, \dots\}$ made at N fixed interferometer delays $\boldsymbol{\tau} = \{\tau_1, \tau_2, \dots\}$ of the lineshape $g(\nu; \boldsymbol{\theta})$. If we assume that the random error on each measurement is normally distributed with standard deviation $\boldsymbol{\sigma} = \{\sigma_1, \sigma_2, \dots\}$ and that the measurements are uncorrelated, then the likelihood function of the i^{th} measurement can be written

$$\mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\left[\frac{\mathcal{D}_i - \zeta(\tau_i, \boldsymbol{\theta})}{\sigma_i}\right]^2\right). \quad (5.2.8)$$

Remember that the likelihood function describes the plausibility of a particular set of model parameter values, given the measured data. Figures 5.9(b)–(e) show coloured regions that indicate the 68% confidence interval for the likelihood functions of each simulated contrast measurement. It has been assumed here that $\boldsymbol{\tau}$, $|\mathbf{B}|$, χ and ρ_Δ are all perfectly well known. This reduces $\mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta})$ to a function of two parameters: n_e and T_i . It can be seen that $\mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta})$ corresponds to an inverted L-shaped band in density-temperature log-space whose width depends on σ_i and whose path depends on τ_i . If each measurement is statistically independent then the total likelihood function is the product over all N delays of the individual likelihood functions of each

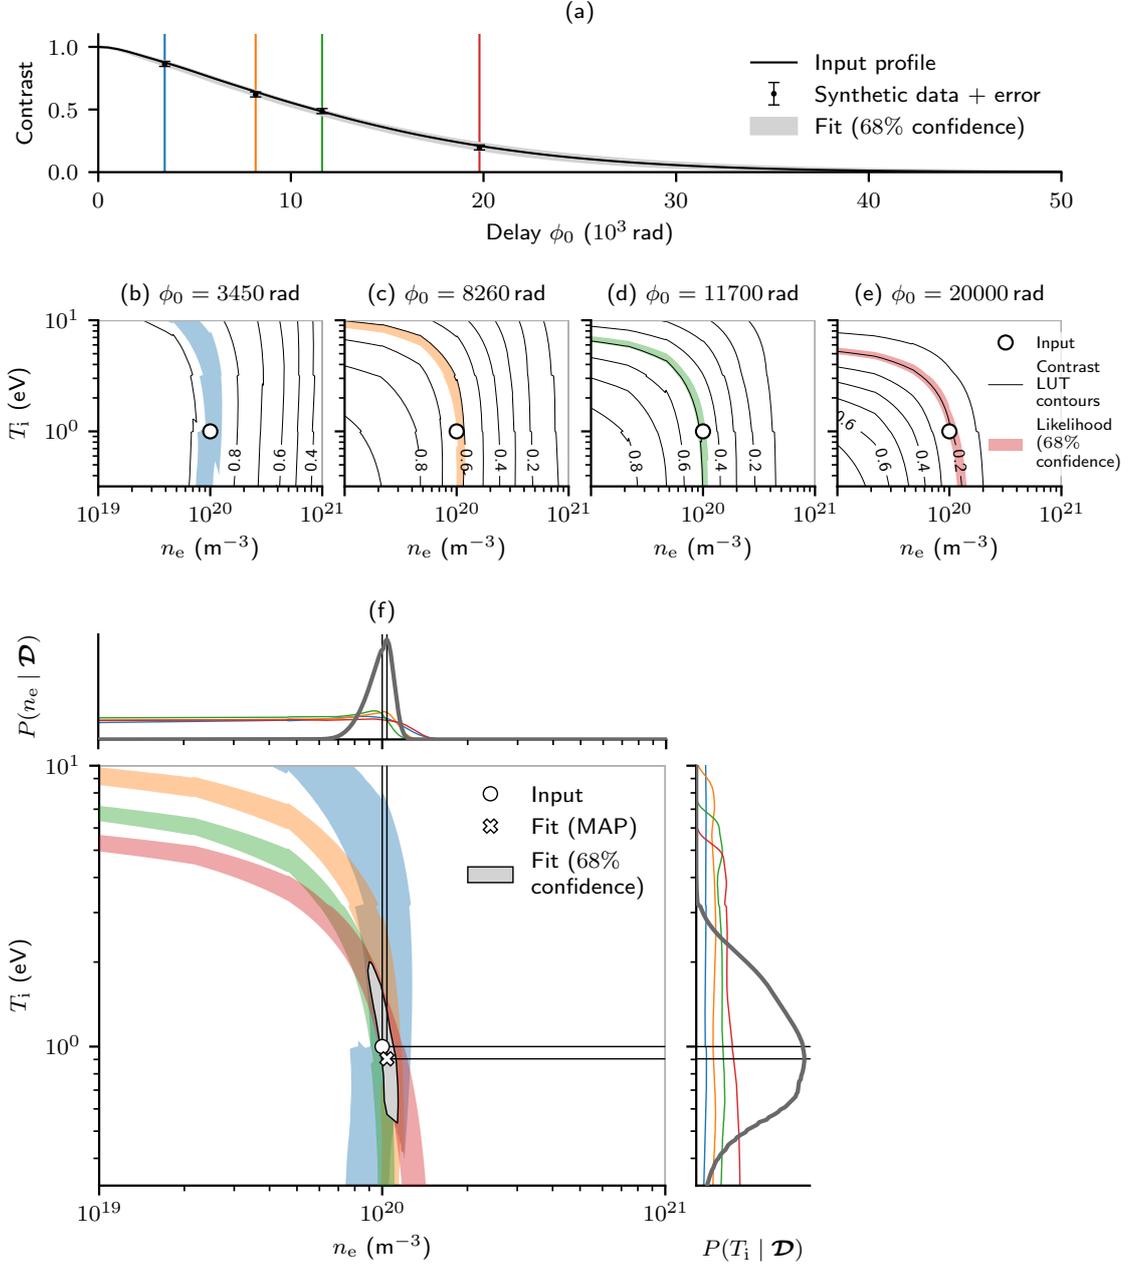


Figure 5.9: (a) Simulated contrast profile and measurements for H $_{\gamma}$ emission (point source) at $n_e = 10^{20} \text{ m}^{-3}$ and $T_i = T_e = 1 \text{ eV}$, observed with a CI instrument in the multi-delay 1 configuration (see Table 3.2). (b)–(e) Slices of the look-up table between measured contrast and n_e and T_i for the four interferometer delays listed. (f) Joint and marginal PDFs for each of the four delays individually (colours) and the combined delays (grey).

measurement

$$\mathcal{L}(\mathcal{D} | \boldsymbol{\theta}) \propto \prod_{i=1}^N \mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta}). \quad (5.2.9)$$

Figure 5.9(f) overlays the 68% confidence interval for $\mathcal{L}(\mathcal{D}_i | \boldsymbol{\theta})$ at each of the four delays in colour and then plots the same interval for $\mathcal{L}(\mathcal{D} | \boldsymbol{\theta})$ in grey. The total likelihood function is localised in parameter space around the input values of n_e and T_i . At this point, any prior information on the model parameter values $P(\boldsymbol{\theta})$ can be factored in using Bayes' theorem to find the posterior probability density function (PDF):

$$P(\boldsymbol{\theta} | \mathcal{D}) \propto \mathcal{L}(\mathcal{D} | \boldsymbol{\theta})P(\boldsymbol{\theta}). \quad (5.2.10)$$

In our multi-delay example here we will assume that we have no prior knowledge about the parameters and so the posterior PDF is $P(\boldsymbol{\theta} | \mathcal{D}) \propto \mathcal{L}(\mathcal{D} | \boldsymbol{\theta})$. We now take as the best estimate of $\boldsymbol{\theta}$ the parameter values that maximise $P(\boldsymbol{\theta} | \mathcal{D})$: the maximum a posteriori (MAP) estimate (defined in Equation C.2.1). The MAP is indicated in 5.9(f). The marginal posterior PDF for n_e is calculated by taking

$$P(n_e | \mathcal{D}) \propto \int_{-\infty}^{\infty} P(\boldsymbol{\theta} | \mathcal{D})dT_i, \quad (5.2.11)$$

and similarly for T_i by taking

$$P(T_i | \mathcal{D}) \propto \int_{-\infty}^{\infty} P(\boldsymbol{\theta} | \mathcal{D})dn_e. \quad (5.2.12)$$

Figure 5.9(f) plots the marginal distributions for both n_e and T_i for the total likelihood and also for the likelihood of each of the four separate delays. For the total likelihood, the MAP estimate of n_e and (approximate) uncertainty is $n_e = (1.0 \pm 0.1) \times 10^{20} \text{ m}^{-3}$ matching the input value of 10^{20} m^{-3} . In the example shown in Figure 5.9, the reduction of the posterior PDF to 2-D allowed us to evaluate it directly on a grid of n_e and T_i . If uncertainty in the other model parameters must be accounted for, then this increases the dimensionality, which quickly makes grid evaluation impossible, or at least impractical. In these cases, Monte-Carlo sampling must be used to evaluate the joint posterior PDF[137].

For single-delay CI measurements, the most that can be inferred about n_e without making any assumptions about T_i is an upper limit. This can be seen clearly in Figure 5.9(f) in the marginal PDFs corresponding to the individual measurements (the coloured lines). Previous work on CI measurement of n_e [97] has ignored the Doppler contribution, effectively assuming $T_i \ll 1 \text{ eV}$. In this work, we will make a more conservative assumption when analysing single-delay data, instead assuming a (soft) upper limit on T_i . This limit is introduced as a prior PDF on T_i with the form

of an inverted logistic curve:

$$P(T_i) \propto \frac{1}{1 + \exp(k[T_i - T_u])} \quad (5.2.13)$$

The choice of constants values ($T_u = 6 \text{ eV}$ and $k = 4 \text{ eV}^{-1}$) is informed by the SOLPS modelling of the MAST-U divertor plasma conditions (see Figure 1.6), which predict that the vast majority of the Balmer series emission comes from regions below 5 eV . Figure 5.10 plots this prior PDF. We will find that this upper limit is also applicable to the experimental measurements presented later in Section 5.3.

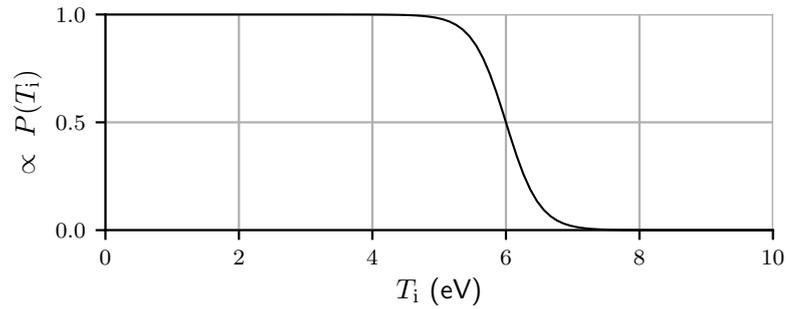


Figure 5.10: An inverted logistic curve (Equation 5.2.13) is used as the prior PDF for emitter temperature when interpreting single-delay CI results in this work.

5.2.3 Systematic Errors

Next, we will test the contrast fitting tool with synthetic data, using the lineshape models introduced in Section 5.1, to quantify the systematic errors. The testing in this section follows the general pattern: model the lineshape for a certain set of plasma parameters and then model the corresponding contrast measurement(s) (assuming zero-mean, independent Gaussian noise with a realistic standard deviation $\sigma_\zeta = 0.02$) for a certain instrument configuration. Then, use the fitting tool on the modelled contrast to get the MAP estimate of n_e . Do this for 30 independent samples of the measurement noise and take the average inferred n_e over the samples. This ensures that any discrepancy between the modelled and final inferred n_e is a systematic effect.

Error due to Doppler Broadening

First, consider the systematic error due to Doppler broadening. Figure 5.11(a) plots inferred n_e versus modelled n_e for a single-delay CI instrument observing H_γ (solid line) and H_δ (dashed line) emission over a range of n_e and T_i values and at $|\mathbf{B}| = 0 \text{ T}$.

Lineshape model 3 is used both to model and to fit to the data. The instrument is in the single-delay configuration from Table 3.2 with delay $\phi_0 = 8260$ rad on-axis for H_γ . In fitting to the modelled data we apply the soft upper limit prior of $T_i \lesssim 5$ eV discussed in the previous section. As expected, the systematic error is larger for larger modelled T_i and is consistently smaller for H_δ than it is for H_γ (Stark-broadened lineshape width scales with n_{upper}). The overestimate of n_e can be significant even for moderate conditions — e.g. $> 25\%$ for H_γ at $n_e = 7 \times 10^{19} \text{ m}^{-3}$ and $T_i = 3$ eV — and becomes very large for low n_e and moderate T_i . Figure 5.11(b) then plots the same test, but for an instrument in the Multi-delay 1 configuration from Table 3.2 with four delays at $\phi_0 = 3450, 8260, 11700$ and 20000 rad (for H_γ emission). This shows a significantly reduced systematic error, extending the dynamic range over which n_e can be measured down to low n_e , moderate T_i conditions.

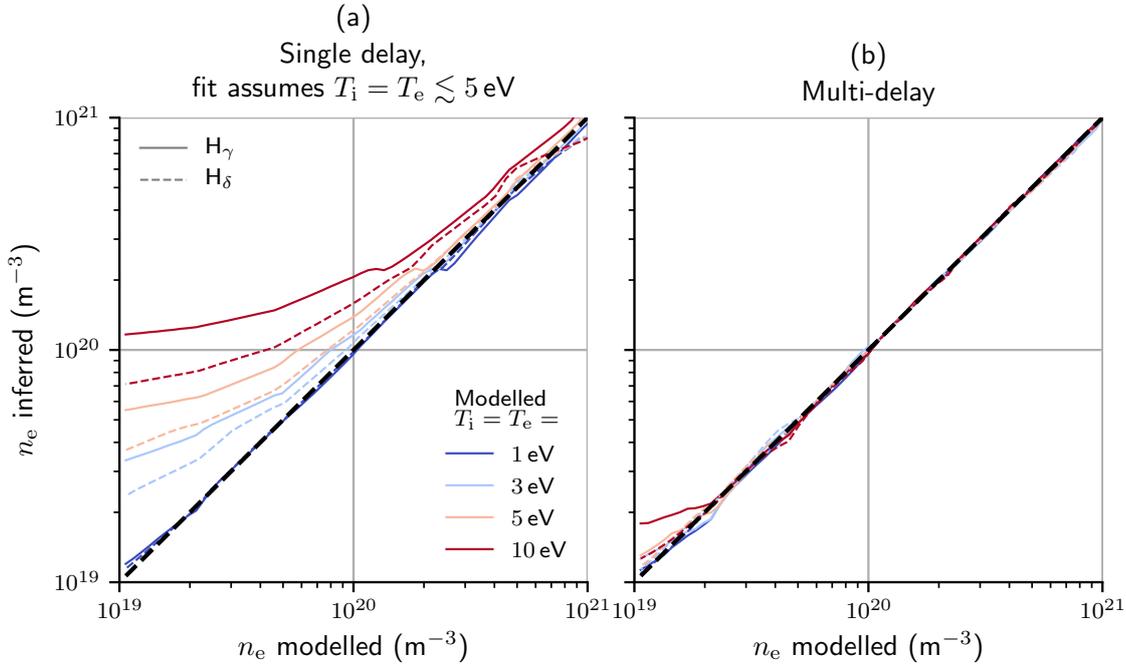


Figure 5.11: Predicted systematic error in inferred electron density n_e due to Doppler broadening, plotted over a range of plasma conditions for (a) single-delay and (b) multi-delay CI instrument configurations. See text for details.

Error due to Lineshape Model

Next, consider systematic error due to the choice of lineshape model used to infer n_e . Figures 5.12(a) and 5.12(b) plot inferred n_e versus modelled n_e for single-delay and multi-delay CI instruments respectively, observing H_γ emission over a range of n_e and T_i values at $|\mathbf{B}| = 0$ T. The single-delay and multi-delay instrument configurations

are the same as in the previous section. The modelled lineshapes are generated using lineshape model 3 (numerical simulation of Stark-Zeeman profile, see Table 5.3) but the contrast fitting tool used to infer n_e assumes models 1, 2 or 3 to be accurate (plotted in different line styles). The ‘pointy’ nature of the inferred n_e lines for models 1 and 2 is a result of linear interpolation of the modelled lineshape between model 3’s coarse n_e grid points (see Section 5.1). Since the Lorentzian lineshape profile assumed in model 1 falls off less steeply in the wings than the other two models (see Figure 5.1), we would expect a fit assuming this model to underestimate n_e for synthetic measurements produced using model 3. This is born out for both the single-delay and multi-delay tests. For the single-delay case, systematic error in n_e relative to model 3 is relatively small ($< 25\%$) for $n_e > 10^{20} \text{ m}^{-3}$ but is significant at lower n_e . For the multi-delay case, assuming model 1 to be valid leads to a very large systematic error relative to model 3 for the whole range of n_e and T_i considered. Assuming Model 2 is valid gives a closer match to the spectra generated using model 3, which is to be expected as both of these models account for more physics and are regularly used in spectroscopic diagnostics. It can be seen in Figure 5.1 that model 2 predicts a slightly narrower lineshape than model 3. This leads to the overestimate seen in model 2’s inferred n_e from model 3’s lineshape. This could possibly be caused by the fine structure splitting not being accounted for in model 2. Since we cannot be certain which of these two models is more accurate, the fractional density error between in these tests ($\lesssim 20\%$ for low temperatures) can be taken as characteristic of the uncertainty in the lineshape model. In the next section, when we apply these methods to experimental CI measurements, we use lineshape model 3 and bear in mind this lineshape model uncertainty.

Figures 5.12(c) and 5.12(d) show the inferred n_e versus the modelled n_e for the same instrument configurations and plasma conditions as 5.12(a) and 5.12(b) respectively, but for $|\mathbf{B}| = 1 \text{ T}$. It can be seen that the discrepancy between models 2 and 3 is similar under the influence of a magnetic field, indicating that the strong-field Zeeman effect assumed in model 2 is valid.

5.3 Experimental Testing on Magnum-PSI

5.3.1 Magnum-PSI

Magnum-PSI is a linear plasma experiment in Eindhoven, NL, at the Dutch Institute for Fundamental Energy Research (DIFFER)[138–140] whose main purpose is to study plasma-surface interactions (PSI) and to test technologies and components for fusion reactors. Magnum-PSI is unique in its ability to produce the particle flux

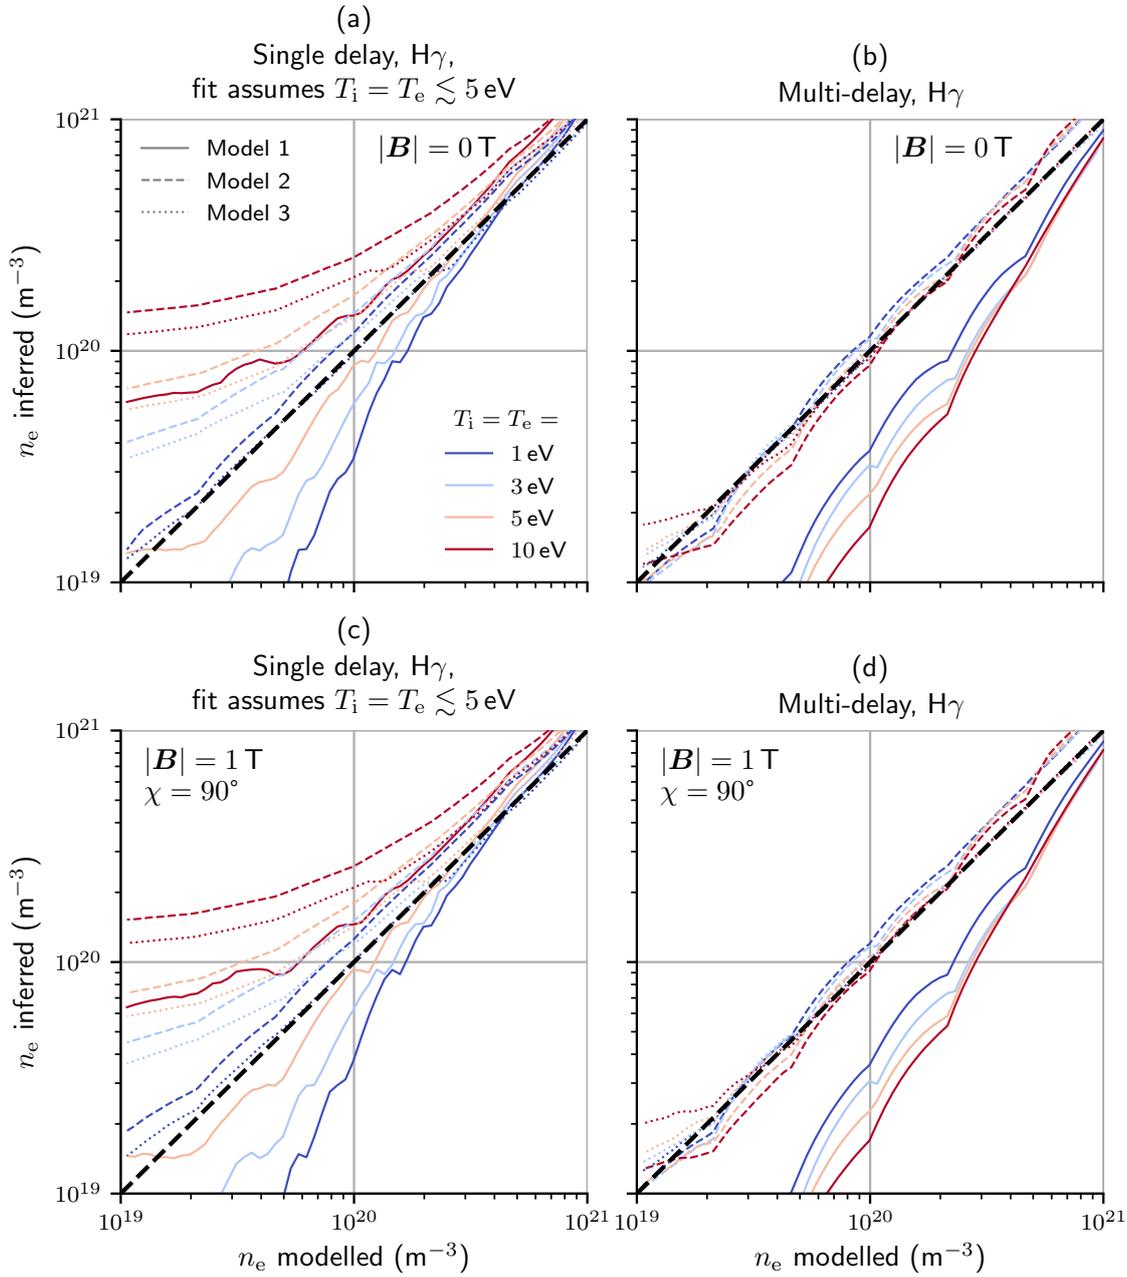


Figure 5.12: Predicted systematic error in inferred electron density n_e due to choice of lineshape model, plotted over a range of plasma conditions for (a) single-delay and (b) multi-delay CI instrument configurations for magnetic field strength $|\mathbf{B}| = 0$ T. (c) and (d) plot the same test for single-delay and multi-delay instruments respectively, but for a transverse view ($\chi = 90^\circ$) of a $|\mathbf{B}| = 1$ T field. See text for details.

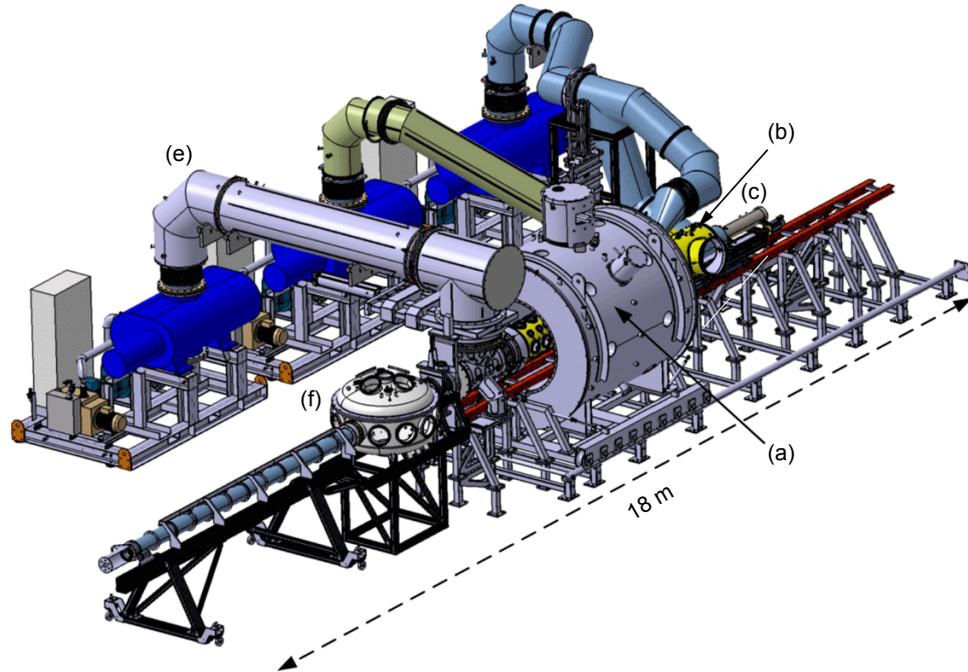


Figure 5.13: The linear plasma experiment Magnum-PSI. (a) The superconducting magnet. (b) The vacuum vessel (in yellow). (c) The moveable plasma source. (d) Three pump sets. (e) The target exchange and analysis chamber, into which the target can be moved for surface analysis. Reproduced from [140] with modified labels.

($10^{23} - 10^{25} \text{ m}^{-2}\text{s}^{-1}$) and heat flux ($\geq 10 \text{ MWm}^{-2}$) conditions (both transient and steady-state) that are expected in the ITER divertor. The device has been used to test the performance of tungsten divertor components[141] as well as novel concepts like liquid metal[142] and heat pipe plasma-facing components[143]. Magnum-PSI is flexible in its range of plasma temperature (0.1 – 5 eV), density ($10^{19} - 10^{21} \text{ m}^{-3}$) and magnetic field strength (0 – 2.5 T), covering the expected MAST-U divertor conditions (See Figure 1.6). This, and the simple magnetic field geometry make it appropriate for testing the CI techniques developed in this work. Figure 5.13 shows a diagram of Magnum-PSI. The moveable plasma source is of the cascaded arc type[144], and the confining field is produced by five superconducting niobium-tin solenoid magnets[145]. The plasma is channelled into a beam of diameter $\sim 10 \text{ mm}$ in a vacuum vessel of diameter 500 mm where it meets a moveable and replaceable target. Plasma conditions are controlled by changing the source current, the source gas flow and the magnetic field strength. Though Magnum-PSI is capable of producing hydrogen, deuterium and helium plasmas of high isotopic purity, only measurements of H discharges are reported in this work.

Magnum-PSI Diagnostics

For study of PSI, Magnum-PSI has a suite of diagnostics measuring conditions in the plasma beam and also at the target surface. Since we are interested here in benchmarking density CI against existing plasma diagnostics, the surface diagnostics (e.g. infrared cameras, laser-induced ablation spectroscopy) will not be discussed here. While Magnum-PSI normally operates a high-speed camera for filtered imaging of the target, this camera's mount and port were occupied instead by CI for this work. The relevant plasma diagnostics for our investigation are:

Thomson scattering (TS) measures n_e and T_e profiles across the beam immediately in front of the target at a spatial resolution of 1.6 mm[146]. The accuracy of the inferred n_e and T_e is reportedly 3% and 6% respectively at $n_e = 9.4 \times 10^{18} \text{ m}^{-3}$ when summing 30 laser pulses at a 10 Hz repetition rate. The same accuracy can be achieved with fewer summed pulses at higher n_e . The minimum n_e and T_e values measurable with this system are $1 \times 10^{17} \text{ m}^{-3}$ and 0.07 eV respectively. Figure 5.14 shows the TS laser path and collection view relative to the other diagnostics used in this investigation. The TS n_e profile will be the benchmark for the CI measurements.

Magnum-PSI does not have Langmuir probes permanently installed, but the experiments for this work were in tandem with tests of a set of MAST-U probes carried out by another PhD student (Jack Leland, University of Liverpool). The probes were flush-mounted on a curved graphite target which, in turn, was mounted to a water-cooled plate and then to Magnum-PSI's target holder. Figures 5.15(a) and (b) show photos of the probe array during and after operation, respectively. Measurements of the Balmer series lines targeted by CI were also made using a multi-chord grating spectrometer whose 40 sight-lines span the width of the plasma beam profile in increments of 0.8 mm. However, due to time constraints, data from these diagnostics is not presented in this thesis.

5.3.2 Coherence Imaging Setup

The CI instrument configurations tested in this chapter are all of the PPM type introduced in Chapter 3, using the same FLIR camera and interferometer hardware as discussed there. The three instrument configurations for which data is presented here are listed in Table 3.2: one single-delay and two multi-delay configurations.

The main change to the CI instrument for the Magnum-PSI experiments is in the imaging optics. The stray magnetic field from Magnum-PSI's coils can cause camera electronics to malfunction. This means that the pre-existing mount for the high-speed camera sits 5.5 m from the plasma beam axis, on the railing of a mezzanine area next

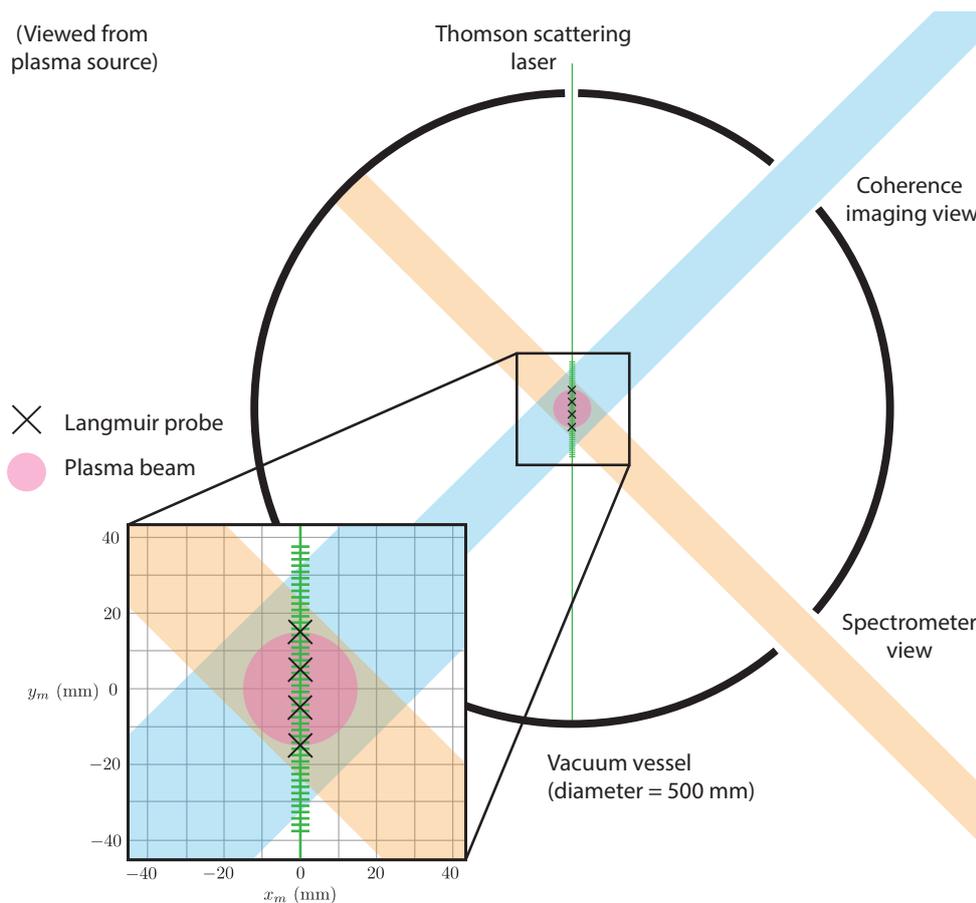


Figure 5.14: The views / locations of the Magnum-PSI diagnostics considered in this chapter, viewed from the plasma source to the target. The CI and spectrometer sight lines and the Thomson scattering laser all lie in the same plane while the probes are set back, on the target itself.

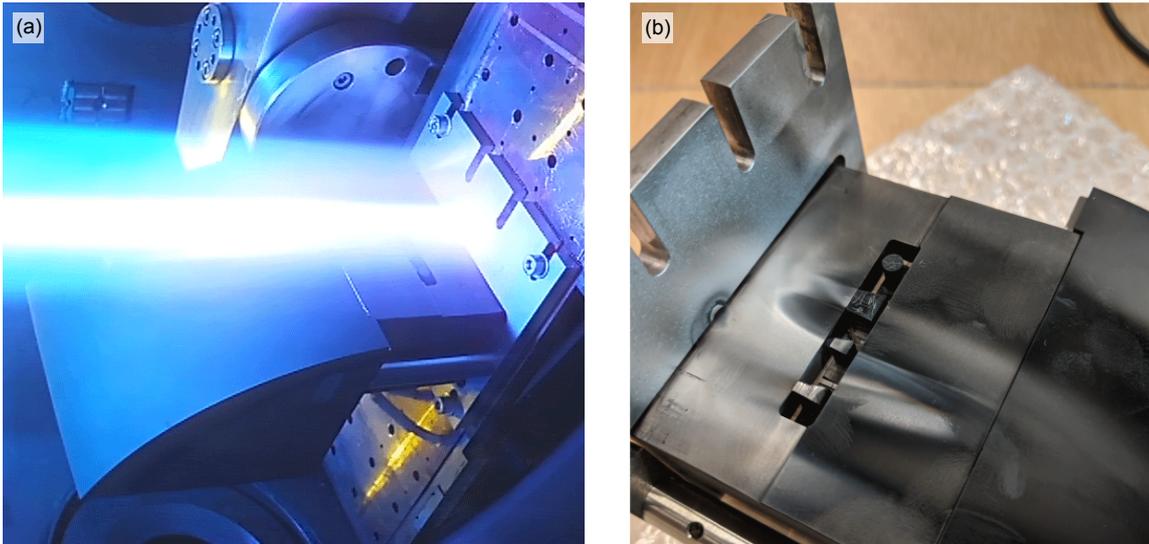


Figure 5.15: (a) The plasma target used in this experiment was a curved graphite tile mounted to a water-cooled molybdenum target holder. Photo taken during a He Magnum-PSI discharge using a monitor camera. (b) Set into the target, orthogonal to the plasma beam axis, is a row of four Langmuir probes.

to the device. This arrangement can be seen in the photo in Figure 5.16(b). From this mount, the camera views the plasma through a 70 mm diameter port in the vacuum vessel, as seen in Figure 5.16(c). Since the plasma occupies a small vertical angle when viewed from this distance (0.23° for a 20 mm beam diameter), a narrow field of view (FOV) is required. With the sensor area acting as the field stop, the FOV in our three lens CI setup (Figure 2.7 is determined by the focal length of the front lens, the ratio of the focal lengths of the second and third lenses (magnification) and the sensor size, which is fixed by the available camera. For the second and third lenses, we opted to keep the Sigma DSLR lenses that were used in the MAST instrument, with focal lengths $f_2 = 105$ mm and $f_3 = 150$ mm respectively, and both with f-number $f/2.8$. For the front lens, we used a Sigma telephoto DSLR lens with focal length 150–600 mm and f-number $f/5$ – 6.3 . At $f_1 = 600$ mm, the vertical FOV is 0.47° , small enough that the plasma beam fills a significant portion of the image vertically. Figure 5.14 shows the CI view relative to the other Magnum-PSI diagnostics used in this investigation. The CI port lies in the same plane as the TS and OES ports. Figure 5.16(a) shows a photo of the CI instrument installed on the mezzanine while Figure 5.16(b) is a wider shot showing the instrument in relation to the Magnum-PSI experiment. The camera was connected to a computer nearby on the mezzanine and its settings were remotely controlled from the experiment control room.

The hydrogen Balmer lines targeted in this investigation were H_γ and H_δ . These were chosen as they both result in $n_c \sim 10^{20} \text{ m}^{-3}$ for the available α -BBO waveplate

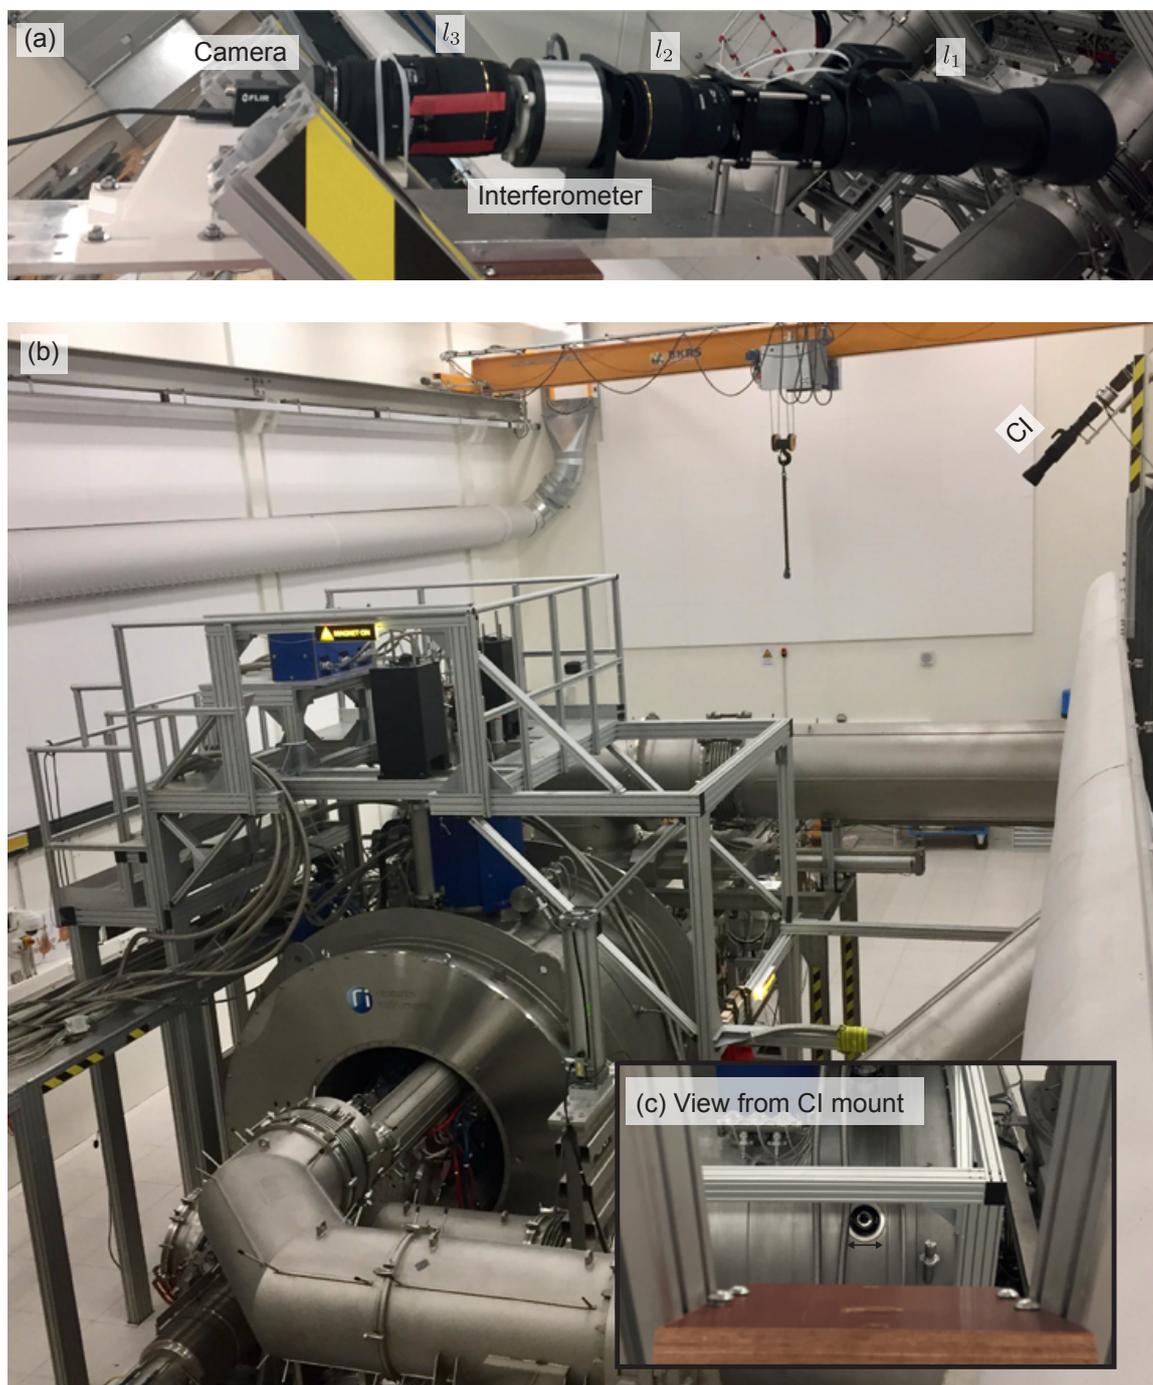


Figure 5.16: (a) The Magnum-PSI machine hall. The coherence imaging instrument is mounted to the railing of a mezzanine area, and can be seen in the upper right corner. (b) The CI instrument.

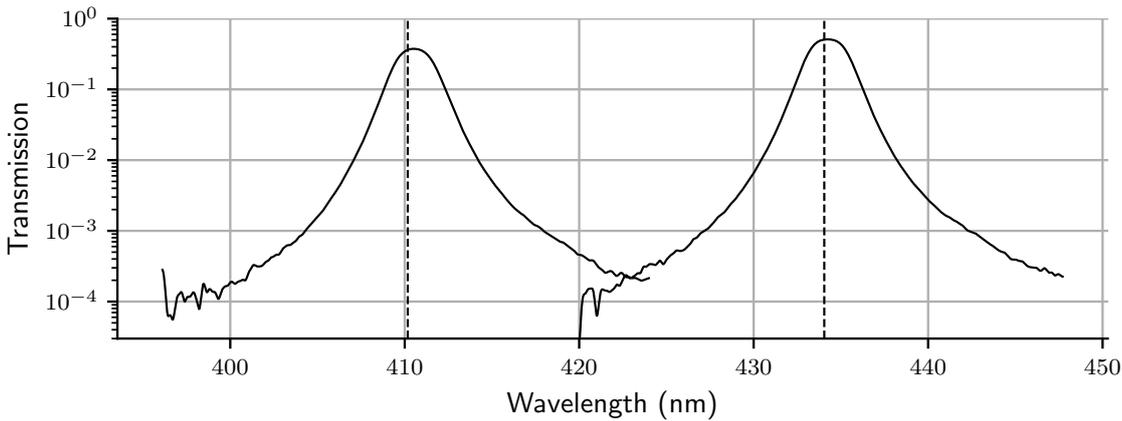


Figure 5.17: Measured transmission profiles for isolating the targeted H_γ and H_δ emission lines.

with the smallest thickness ($L_{wp} = 4.48$ mm). The bandpass filters used to isolate these lines were each two-cavity interference filters supplied by Andover Corporation. The measured transmission profiles of the filters are plotted in Figure 5.17. For H_γ and H_δ filter profiles the peak transmission is 0.50 and 0.44 respectively and the FWHM is 2.67 nm and 2.63 nm respectively.

Calibration

Dark frame subtraction is important when measuring interferogram contrast since a positive offset applied to the image, if unaccounted for, will artificially lower the inferred contrast. Since the dark current depends on the temperature of the camera sensor, which in the case of the FLIR Blackfly is not actively cooled, dark frames are captured after each plasma measurement using the same exposure and gain settings. These dark frames are then subtracted before analysis of the images.

As discussed in Chapters 2 and 4, the CI instrument function is sensitive to changes in the alignment and temperature of the birefringent optical components, making regular calibration necessary. However, since instrument contrast ζ_I does not depend directly on crystal temperature via thermal expansion or thermo-optic effects, the temperature dependence is weaker than for instrument phase[63]. Instead, ζ_I is sensitive to variations in phase delay across the crystal body. These could be caused by variation in thickness or refractive index across the crystal. Previous work has measured significant variation in phase delay (up to $\sim \pm 2$ rad) across the surface of each of the crystals used in this work[63]. Since ζ_I is found by integrating over the phase variations across the aperture (similar to the treatment in Section 2.3.7), it follows that if the calibration source under-fills the crystal aperture relative

to the plasma source, ζ_I will be different for the two cases, introducing systematic error into the inferred contrast. In this investigation, CI calibration images were taken using Cd I line emission (467.8 nm) from a gas discharge lamp, isolated from the surrounding emission lines using Filter 1 from Figure 4.7 and illuminating an integrating sphere. To ensure the calibration source filled the entrance pupil of the instrument, and therefore the crystal body, in an identical way to the plasma emission, the Cd source is situated 5 m from the instrument for each measurement, which took place on the mezzanine area in the Magnum-PSI hall. By monitoring the instrument contrast values ζ_I throughout the day of each measurement set, we can estimate the contribution to the measurement uncertainty due to systematic shifts in the instrument function. Measurements from each of the three instrument configurations were each taken over a ~ 2 hr period. The temperature controller maintains the temperature of the birefringent components at 35 ± 0.25 °C.

We are interested here in systematic shifts in ζ_I so, for each calibration image, we average over a 200×200 pixel central region. Our confidence in this value is then expressed by the standard error on the mean. The standard deviation of the seven ζ_I measurements is 0.003. This is significantly larger than the standard error on the mean, so can be taken as representative of the uncertainty in ζ_I due to systematic calibration error for the single delay measurements presented in this section. For the multi-delay instrument configurations, we use the same method. It was found that the uncertainty due to systematic calibration error was almost an order of magnitude larger than for the single delay instrument at 0.02. The reason for this is not clear. It could be because this interferometer configuration has six birefringent crystals (each Savart plate counting for two) instead of one.

5.3.3 Line-integrated Results

Figure 5.18 shows raw CI images and demodulated contrast images for the single-delay instrument observing H_γ and H_δ emission. Parts (a) and (b) correspond to (on-axis) plasma conditions of $n_e = 1.5 \times 10^{20} \text{ m}^{-3}$, $T_e = 2.8 \text{ eV}$ and $n_e = 4.8 \times 10^{20} \text{ m}^{-3}$, $T_e = 1.8 \text{ eV}$ respectively, as reported by the TS diagnostic. Details of the plasma control parameters are provided in the Figure and caption. The higher brightness of the H_γ emission over H_δ is apparent in the raw images, as is the higher brightness of the higher n_e discharge in part (b) of the Figure.

Figure 5.19 shows (a) a raw CI image, (b) the corresponding brightness image and (c)–(f) the corresponding demodulated contrast images for an instrument in the Multi-delay 1 configuration observing H_γ emission. This discharge had source current 160 A, gas flow 5 SLM and $|\mathbf{B}| = 1.2 \text{ T}$. The TS diagnostic reported the on-axis

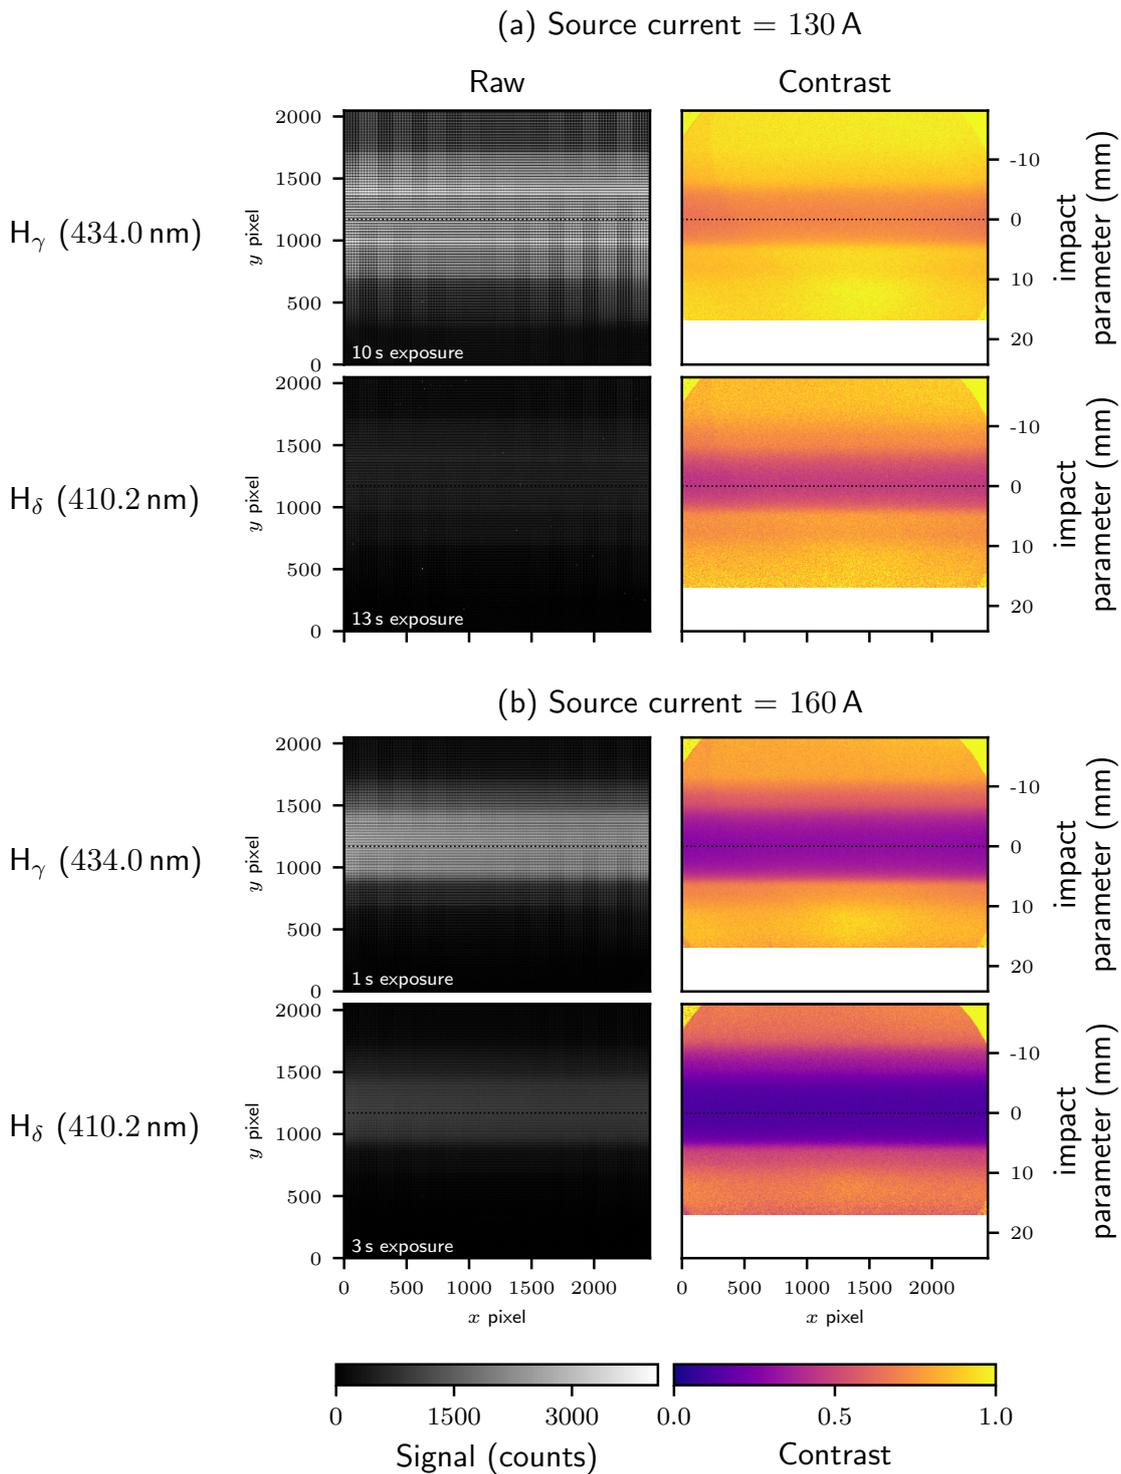


Figure 5.18: Raw CI Magnum-PSI results for a single-delay instrument observing H_γ and H_δ emission from a discharge with field strength $|\mathbf{B}| = 1.5$ T, source gas flow $Q_s = 4.5$ SLM and source current (a) $I_s = 130$ A and (b) $I_s = 160$ A. Raw interferograms are shown in the left column while calibrated (line-integrated) contrast is shown in the right. Regions with $\text{SNR} < 5$ in either the raw image or the calibration image are masked in the contrast images.

plasma conditions to be $n_e = 4.2 \times 10^{20} \text{ m}^{-3}$ and $T_e = 1.8 \text{ eV}$. That the contrast drops monotonically with increasing delay for a given image region, suggests that the observed spectral feature is unimodal, as expected. For a direct comparison to the TS values, the line-integrated CI contrast must be tomographically inverted.

5.3.4 Non-Gaussian Doppler broadening and Doppler shifts

The presence of non-Gaussian Doppler broadening or significant Doppler shifts could complicate the interpretation of the CI data. Previous work by Shumack et al.[147] used a grating spectrometer to view H_β (486.1 nm) at the plasma source of the Pilot-PSI machine (a smaller, non-superconducting forerunner to Magnum-PSI that used the same cascaded arc plasma source). From the observed lineshape, the presence of two atomic hydrogen populations was inferred: one coupled to the ions via charge exchange with T_n at a few eV and the other one cold at 0.1–0.5 eV. The intensity ratio between the two populations was constant across the beam profile at roughly 2 to 1 (hot to cold). In addition, the hot population was observed to be rotating around the beam axis, at velocities up to 10 km/s, due to an $\mathbf{E} \times \mathbf{B}$ drift. Since this rotation was observed to decrease with axial distance from the source, with a characteristic decay length of 0.5 m, it is not observed in our measurements at the Magnum-PSI target, which sits 1.5 m from the source.

The presence of two emitting populations at different temperatures results in a non-Gaussian Doppler broadening contribution to the lineshape, which could introduce systematic error into the inferred n_e if not accounted for. The size of this error was modeled by the same fitting procedure used in Section 5.2.3. Representative T_n chosen for the hot and cold populations in this test are 3 eV and 0.3 eV. Figure 5.20(a) shows the predicted error for the single-delay CI configuration and Figure 5.20(b) shows the error for the multi-delay CI configuration. As might be expected, the presence of the cold population reduces the error due to Doppler broadening in the single-delay case when compared to the Gaussian broadening model with $T_n = 3 \text{ eV}$. In the case of the multi-delay CI, the error in inferred n_e is only larger than 5% for $n_e < 5 \times 10^{19} \text{ m}^{-3}$. Since the predicted effect is small, we do not consider non-Gaussian Doppler broadening in the analysis of the Magnum-PSI CI results.

5.3.5 Abel Inversion

To infer the local n_e profiles, we need the local CI fringe contrast. Obtaining the local contrast from the observed, line-integrated contrast $\check{\zeta}$ requires tomographic inversion. The CI inversion problem was introduced in Section 2.2.3. Since the line-integrated

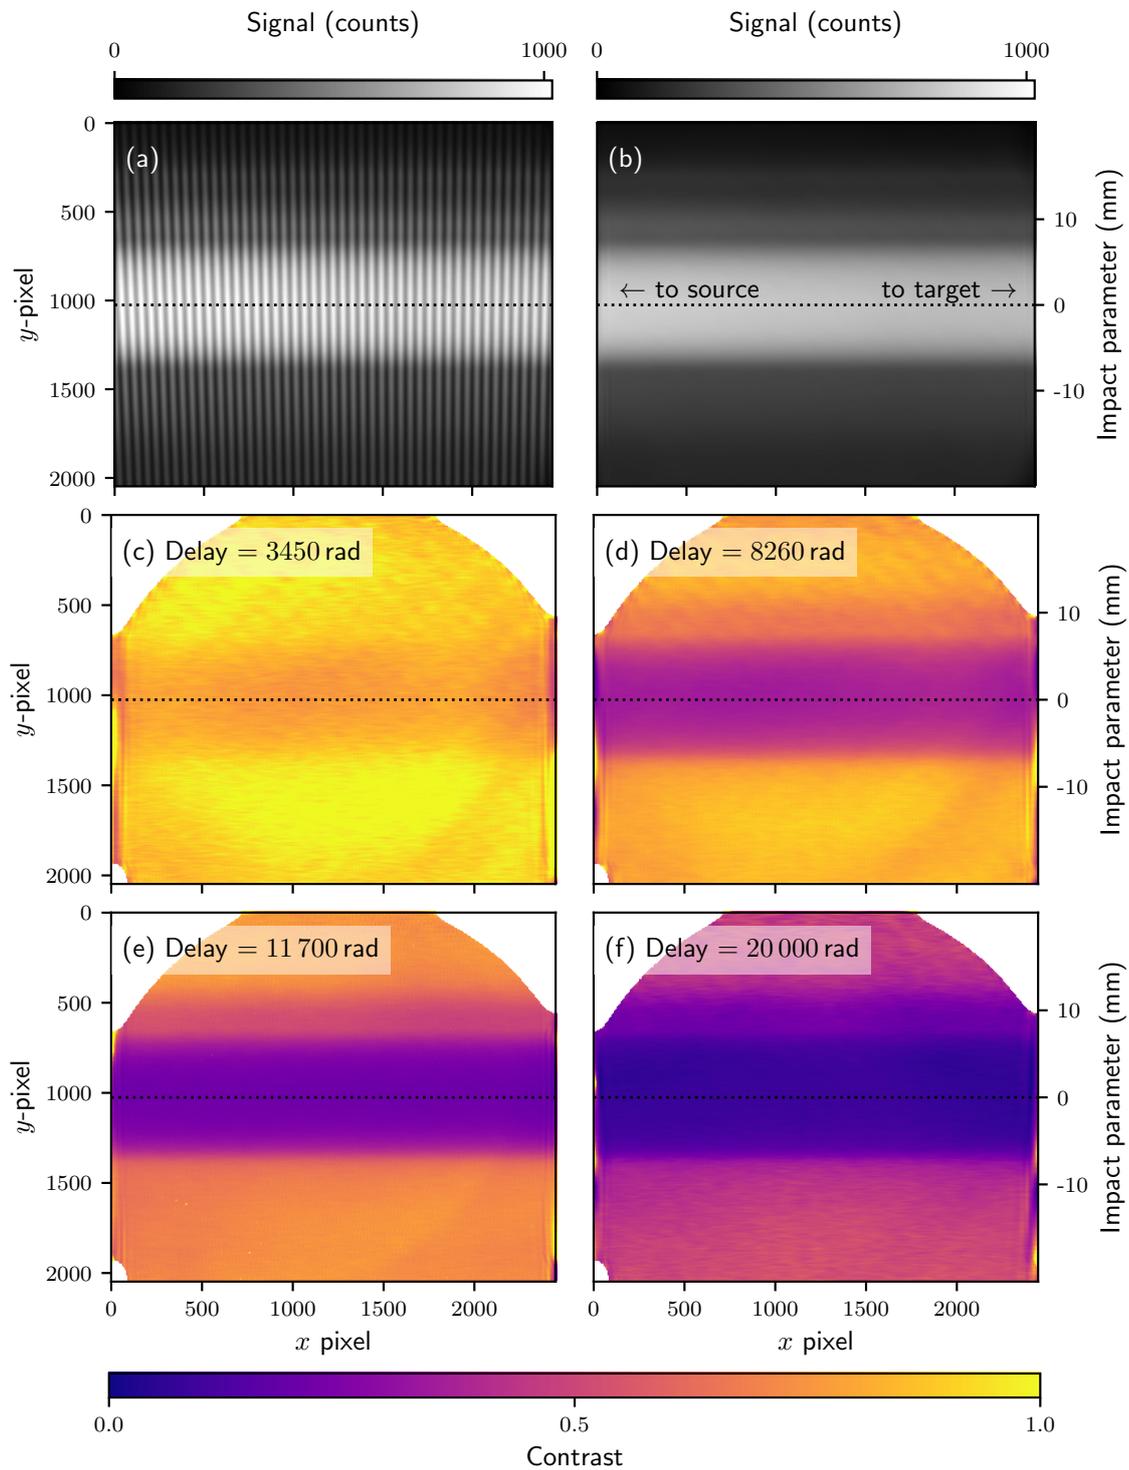


Figure 5.19: Raw CI Magnum-PSI results for a multi-delay instrument observing H_γ emission from a discharge with $|\mathbf{b}| = 1.2$ T, $Q_s = 5$ SLM and $I_s = 160$ A. The images shown are (a) the raw interferogram, (b) the background brightness (c)–(f) the calibrated contrast profiles. Regions with $\text{SNR} < 5$ are masked in (c)–(f).

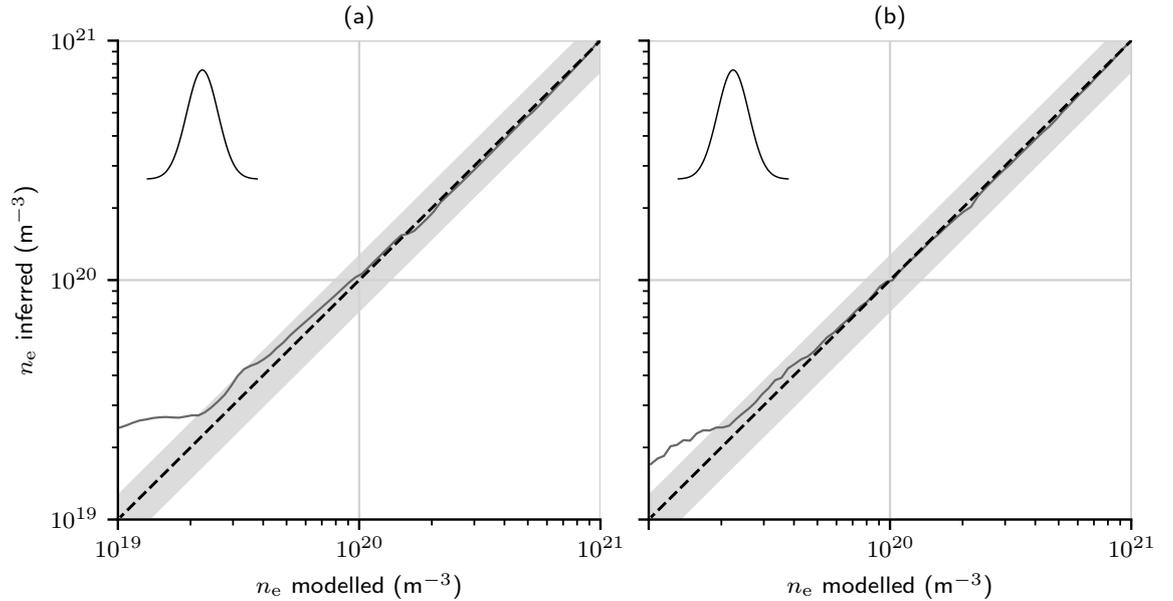


Figure 5.20: Predicted systematic error in inferred electron density n_e due to non-Gaussian Doppler broadening for (a) single-delay and (b) multi-delay CI instrument configurations. The grey shaded region indicates $\pm 25\%$ error. The Doppler broadened lineshape contribution is plotted inset, and corresponds to two emitting hydrogen populations with temperatures 3 eV and 0.3 eV and intensity ratio 2 : 1.

contrast is weighted by the local emission brightness, Equation 2.2.23 must first be inverted to recover the local brightness $I_0(\mathbf{r})$ before Equation 2.2.28 can be inverted to find the local contrast. Since the CI lines of sight all essentially view the plasma beam from the same direction, we must rely on assumptions about the beam profile in order to obtain a unique inversion solution. We will assume that the Magnum-PSI plasma beam is cylindrically symmetric about the central axis for both brightness and contrast. Since the line-integrated CI brightness and contrast measurements show no significant variation along the central axis of the beam (see results in Figures 5.18 and 5.19), for the remainder of this section *we will present results from column slices instead of the full resolution images*. The assumption of cylindrical symmetry about the beam axis is not perfect. To make a first-order account for the cylindrical asymmetry, while maintaining a unique inversion solution, we split the line-integrated measurements in to two halves along the beam axis and invert each side separately. This then assumes that we have two half-cylinders, each with its own cylindrical symmetry, joined along the plane containing the beam axis and the lines of sight.

The mathematical operation that relates a cylindrically symmetric 3-D object to its 2-D projection is the Abel transform[148]. Of interest here is the inverse Abel transform, which reconstructs the original 3-D object from measurements of its 2-D projection. In the absence of an obvious simple analytical form for the measured

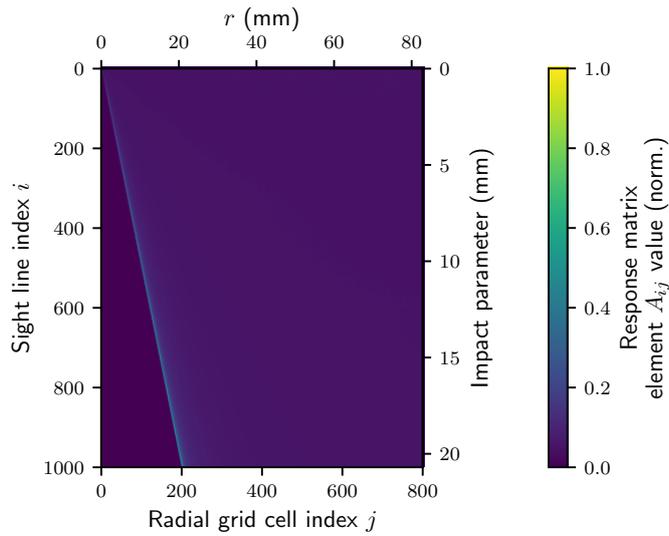


Figure 5.21: Example CI geometry matrix for Magnum-PSI view.

projections, we use numerical methods to solve this problem, of which there are many[148]. We opt for a ‘pixel-based’ approach that casts the problem as a system of linear equations. The transformation from local values \mathbf{f} to projected values \mathbf{F} is treated as a multiplication by a ‘response matrix’ \mathbf{A} :

$$\mathbf{A}\mathbf{f} = \mathbf{F}. \quad (5.3.1)$$

Here, \mathbf{F} is an m -length column vector describing the observed projection along each of the m lines of sight, and \mathbf{f} is an n -length column vector describing the local function at each of the n radial grid cells. \mathbf{A} is therefore a $m \times n$ matrix. The first step is to define the radial grid. Each element A_{ij} of the response matrix is then assigned a value in proportion to the length of sight line i falling within grid cell j . This is calculated by incrementally ray tracing along each sight line in steps much smaller than the size of the radial grid cells. Here, each pixel’s line of sight is treated as being narrow and having equal collection power along its length — i.e. depth of field effects are ignored. As our problem has a simple radial grid and near-parallel lines of sight from a single direction, \mathbf{A} takes a particularly simple form, with an example plotted in Figure 5.21. The matrix indices are shown, as are the conversions to radius of grid cell r and sight line impact parameter. Intuitively, the elements of \mathbf{A} for which $r \approx$ impact parameter have the highest values, as these elements correspond to sight lines passing tangentially through the annular grid cells.

The system of linear equations in Equation 5.3.1 is then solved for \mathbf{f} by finding \mathbf{A}^{-1} . Many numerical techniques exist for solving this problem, but here we use simultaneous algebraic reconstruction technique (SART), a technique popular in

computerised tomography. This finds the solution for \mathbf{f} by iteratively minimising the residuals $\mathbf{F} - \mathbf{A}\mathbf{f}$. The SART implementation used here was written in Python by James Harrison and Scott Silburn, and based on an existing Matlab implementation used in tomographic inversions of flow CI data on MAST[63]. Figure 5.22 shows the inversion process for a single column of the raw multi-delay CI data shown in Figure 5.19. The left figure column shows the measured projections of (a) brightness \check{I}_0 , (c) contrast $\check{\zeta}$ and (e) the product of the two $\check{I}_0\check{\zeta}$. Prior to inversion these measurements are smoothed and, where necessary, the brightness is brought to zero outside the measurement range. The first step is the inversion of \check{I}_0 , which results in the local profile in 5.22(b). The projection is split in half at the beam axis, which is assumed to be at the point of maximum \check{I}_0 . This is indicated by the black, vertical dotted line in the left-hand column plots of Figure 5.22. Each side of the projection is then inverted independently using the SART algorithm and then stitched back together to make up the inverted profile. As the assumed cylindrical symmetry of the two beam halves is only approximate, the two halves of the inverted profiles will not, in general, agree on the $r = 0$ value of the local profile. The discrepancy between the two halves, visible in the inverted I_0 profile in Figure 5.22(b) then provides some indication of the uncertainty associated with the assumption of (half) cylindrical symmetry. According to Equation 2.2.28, to find the local contrast profile, the product of the projected brightness and projected contrast $\check{I}_0\check{\zeta}$, shown in Figure 5.22(e) must be inverted. This results in the product of the local brightness and local contrast $I_0\zeta$, shown in Figure 5.22(f), which is divided through by the inverted brightness \check{I}_0 from Figure 5.22(b) to find the local contrast, shown in Figure 5.22(d). For a multi-delay CI system, each delay's $\check{I}_0\check{\zeta}$ profile must be inverted independently. This is shown in the bottom row of Figure 5.22.

5.3.6 Inverted Results

Once the inverted contrast profile(s) have been calculated, the fitting tool described in Section 5.2.2 can be used to infer the local n_e profile. All the results shown in this section fit to the data using the contrast LUT generated using lineshape model 3 — as this is the most sophisticated of the available models. Zeeman splitting is accounted for the fitting process by assuming the Magnum-PSI control value to be completely correct.

For the rest of this section, we will directly compare the local n_e profiles inferred from CI measurements to the corresponding TS measurements. We will also plot the TS measurements of T_e for context. For the multi-delay CI results, we can compare this T_e profile to the T_i inferred from the CI measurements. The CI-TS comparisons will

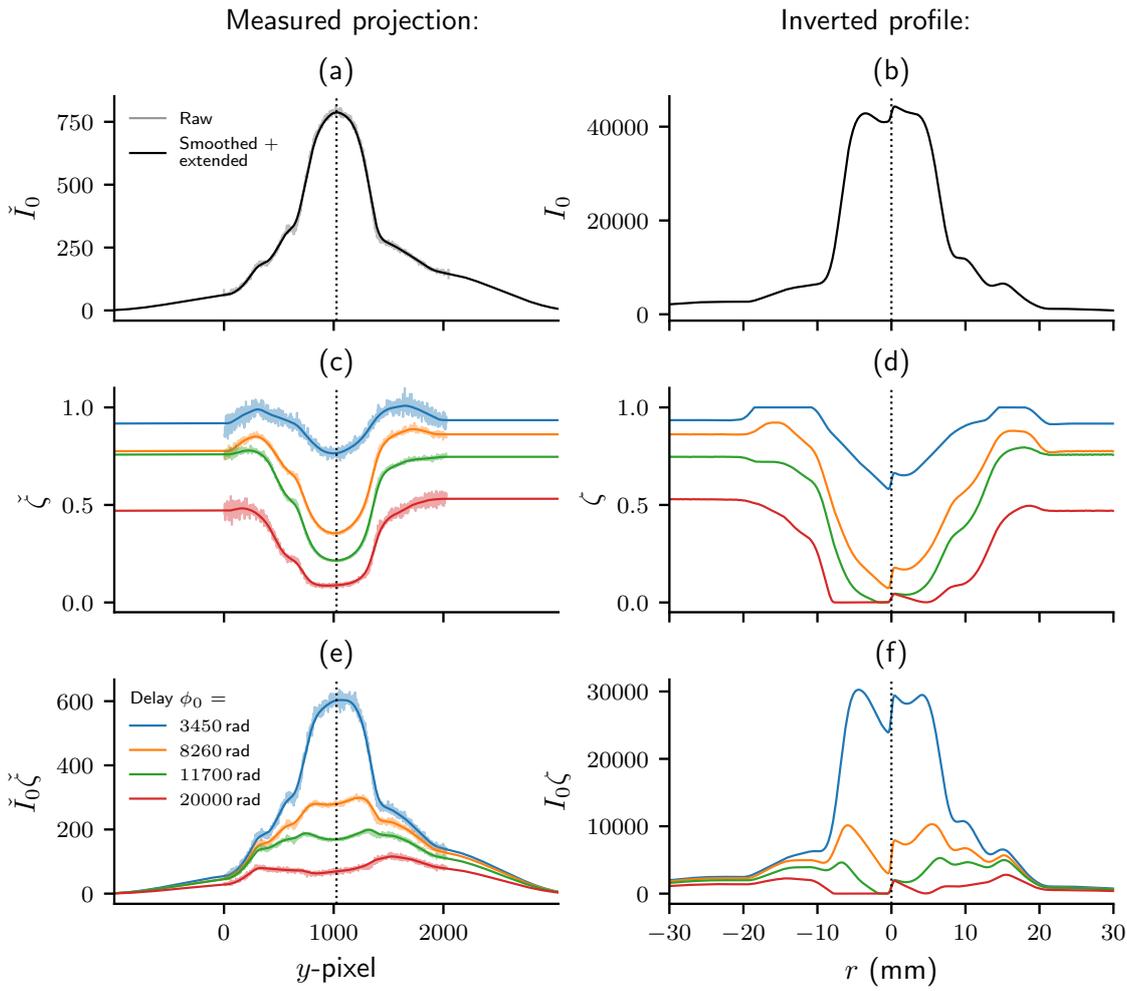


Figure 5.22: Example tomographic inversion for the Multi-delay CI H_γ images shown in Figure 5.19.

be presented in a consistent format for each CI instrument configuration, for example see the multi-delay delay H_γ CI results in Figure 5.23. Each column corresponds to different plasma conditions, as specified in the figure caption. The top row shows 3-D plots of the inverted contrast profile(s) in black as a function of beam radius r and interferometer delay ϕ_0 . For the multi-delay results there are contrast profiles at four delays. The grey mesh surface is the MAP value for the fit to the contrast profile at each r , plotted over a range of delays. Three slices through this ‘fit surface’ are made at radial points $r = 3, 9$ and 15 mm, plotted in blue, orange and green respectively. In the second row, each of these three slices is plotted in as a function of ϕ_0 only, along with the measured contrast data at the same r . This allows for a visual inspection of the goodness of fit — particularly useful for the multi-delay CI data. The third row then plots the radial n_e profiles as inferred from the CI and TS diagnostics. For the CI results, the MAP values are plotted as a solid line and

the 68% confidence interval is shaded. The r values of the three slices are indicated here with vertical lines of the appropriate colour. Finally, the fourth row plots the corresponding radial temperature profiles: T_e for the TS diagnostic and T_i for the multi-delay CI diagnostic. For the CI data, only the 68% confidence interval is shown, as the posterior PDFs are relatively wide.

Figures 5.23(a)–(e) correspond to a scan in n_e from low to high values. Part (d) of this figure corresponds to the same CI data presented earlier in Figures 5.19 and 5.22. The lineshape model appears to fit the contrast data well across these measurements, particularly closer to the beam axis (see $r = 3$ mm slice plotted in blue). Agreement between CI and TS n_e profiles is mostly good across this range. It is important to note that the results in Figures 5.23(a) and (b) were taken with the CI instrument in the Multi-delay 2 configuration while the results in Figures 5.23(c)–(e) were taken in the Multi-delay 1 configuration (See Table 3.2 for reference). The larger maximum delay of the Multi-delay 2 configuration ought to make it more sensitive at lower n_e . While the Multi-delay 2 configuration does provide four fixed interferometer delays, two of these delays are similar in value. This represents a sub-optimal sampling in delay space, but this is not a significant problem. In fact, the close proximity of the two delay samples here provides a welcome consistency check. Figures 5.23(a) and (b) were measured at a lower magnetic field strength: $|\mathbf{B}| = 0.8$ T compared to $|\mathbf{B}| = 1.2$ T for parts (c)–(e). While there is agreement between the two diagnostics for n_e , the inferred CI emitter temperature T_i is consistently higher than the T_e profiles reported by the TS. Whether this is accurate or is a systematic error in the CI results is not clear.

Figure 5.24 plots H_γ CI results measured using an instrument in the single-delay configuration (see Table 3.2). Results are presented in the same format as Figure 5.23. These single-delay results correspond to a scan in plasma source current I_s with a gas flow rate of $Q_s = 4.5$ SLM and a magnetic field strength of $|\mathbf{B}| = 1.5$ T. For these single-delay results, the soft upper limit of $T_i < 5$ eV discussed in Section 5.2.2 is used as a prior PDF. This is justified for the Magnum-PSI plasma based on the TS measurements of T_e , which are < 5 eV for all discharges reported. Additional justification comes from previous work modelling discharges on Pilot-PSI, a smaller linear plasma experiment that served as a prototype for Magnum-PSI[149]. This work combined fluid plasma and kinetic neutral models, and predicted a relatively low neutral temperature ($T \lesssim 2$ eV). While the single-delay n_e results show good agreement with the TS results at higher densities ($n_e \gtrsim 10^{20}$ m $^{-3}$), Figure 5.24(a) shows that at lower densities ($n_e \approx 5 \times 10^{19}$ m $^{-3}$) we see a considerable overestimate ($\sim 300\%$) in the CI n_e . Based on the discussion in Section 5.2.2, we expect a degree of n_e overestimation for the single-delay instrument. However, the observed

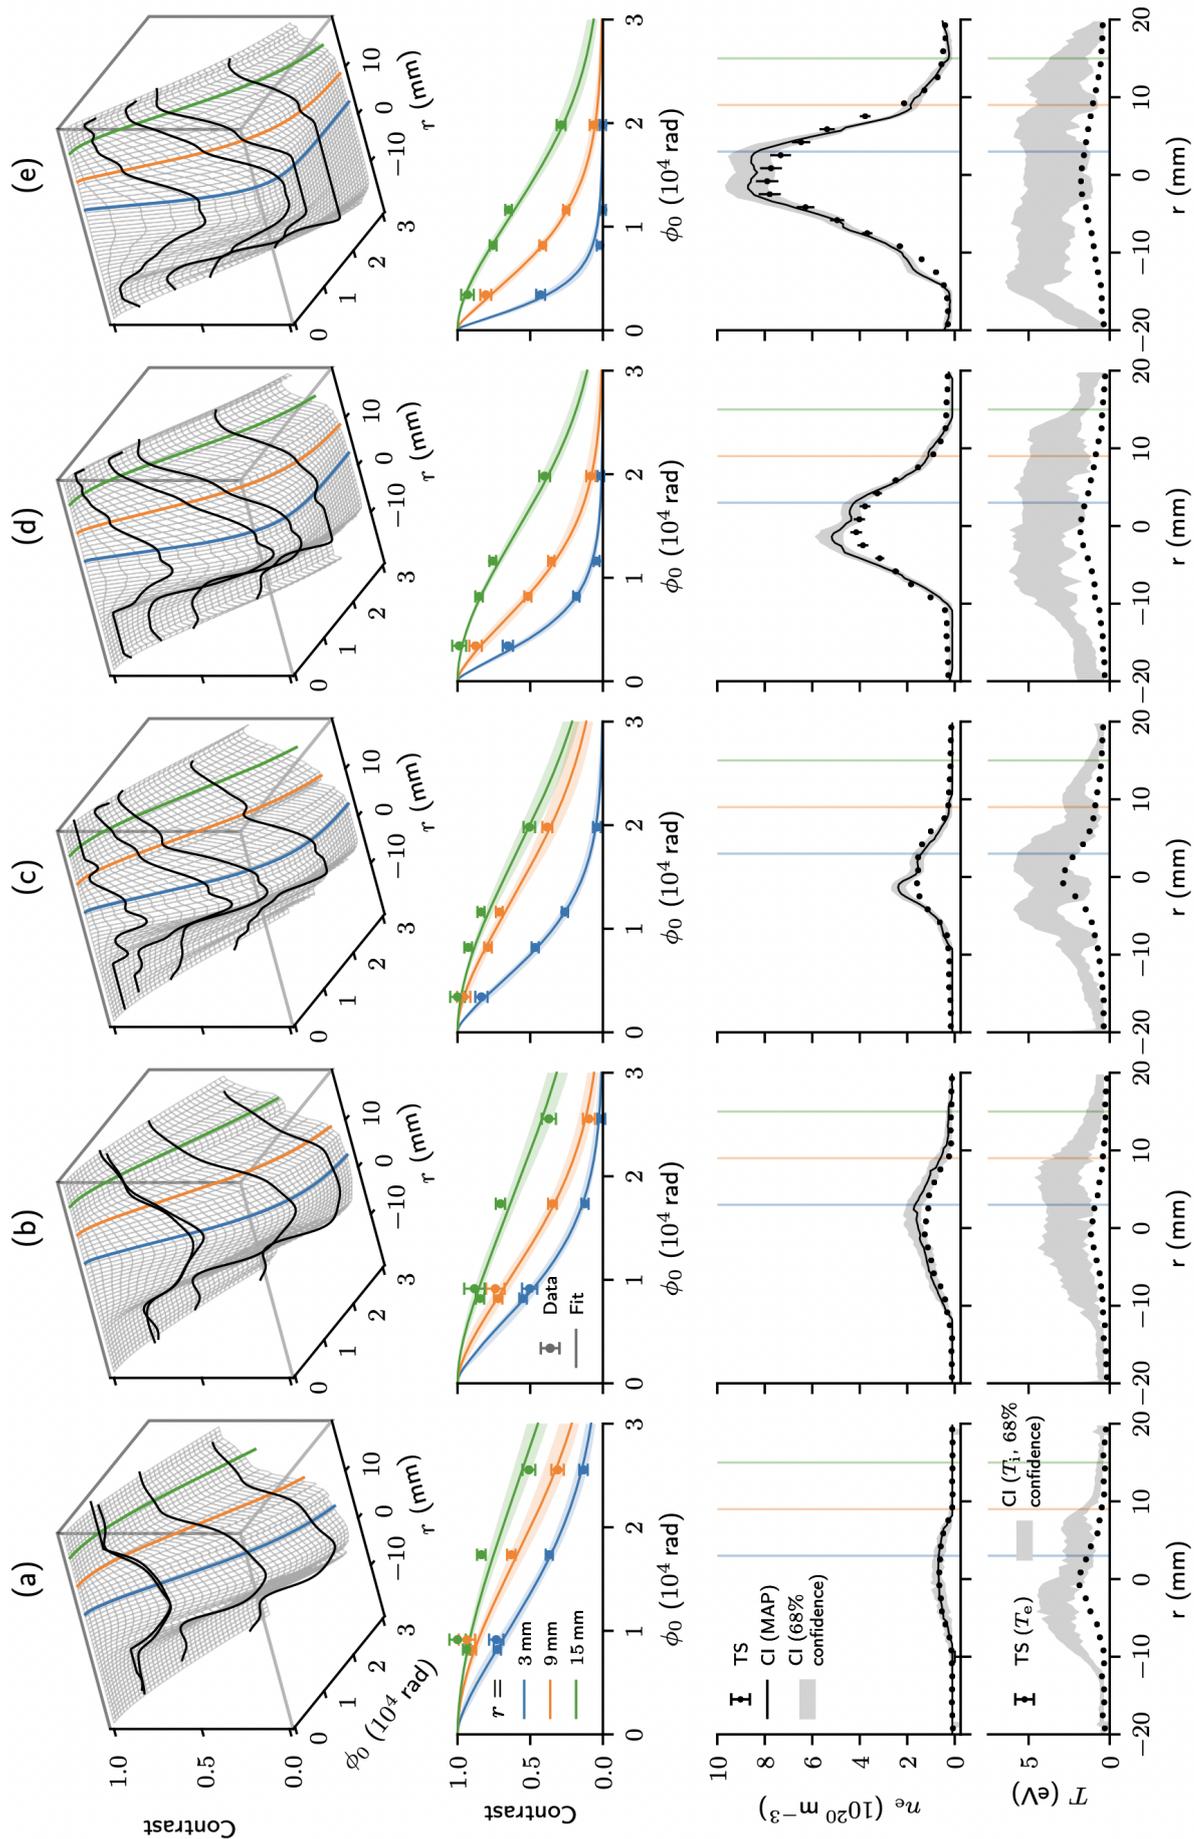


Figure 5.23: Caption on next page.

Figure 5.23: (Image on previous page) Multi-delay CI results observing H_γ emission (434 nm) from Magnum-PSI over a range of plasma conditions. Source gas flow rate is $Q_s = 5$ SLM for all discharges here, while source current I_s and magnetic field strength $|\mathbf{B}|$ are: (a) $I_s = 120$ A, $|\mathbf{B}| = 0.8$ T, (b) $I_s = 130$ A, $|\mathbf{B}| = 0.8$ T, (c) $I_s = 125$ A, $|\mathbf{B}| = 1.2$ T, (d) $I_s = 160$ A, $|\mathbf{B}| = 1.2$ T, (e) $I_s = 200$ A, $|\mathbf{B}| = 1.2$ T. See text for detailed explanation.

overestimation is larger than we would expect by taking the TS n_e and T_e and calculating the expected overestimation (as in Figure 5.11(a)). This could be explained by the $T_i > T_e$ result from the multi-delay CI.

Figure 5.25 plots CI measurements of H_δ emission using an instrument in the Multi-delay 2 configuration (see Table 3.2). These results correspond to $|\mathbf{B}| = 0.8$ T, with gas flow and source current values listed in the figure caption. While the n_e values recorded by the TS diagnostic fall within the range tested for H_γ emission in Figure 5.23, the agreement between CI and TS profiles is consistently worse here. The peak values of the CI n_e profiles are 50–100% higher than those reported by the TS. The reason for such a large discrepancy is unclear. It seems unlikely to be a problem with the lineshape model used, as the agreement between separate lineshape models 2 and 3 in Section 5.1 was found to be equally good for H_γ and H_δ lines. It also seems unlikely to be a problem with the data analysis pipeline, since the same pipeline resulted in a good agreement between the multi-delay H_γ CI and TS results. One possible cause is a systematic error introduced in the contrast calibration step. The H_δ emission (410.2 nm) is further away in wavelength from the Cd calibration line (467.8 nm) than H_γ emission. Measurement of the wavelength dependence of ζ_I would show whether this effect could feasibly be the cause. An alternative explanation is contamination emission reducing the observed plasma contrast and resulting in an n_e overestimate. Analysis of the Magnum-PSI spectrometer data should show evidence of this if it is the cause. The non-ideal performance of the polymer quarter-wave plate used in this work (described in Section 3.4) could also be contributing to this n_e overestimate.

5.4 Summary

This chapter has considered CI measurement of electron density via Stark-broadening of hydrogen Balmer line emission from divertor-relevant plasma conditions. This work is motivated by wanting to make high quality 2-D measurements of n_e in the MAST-U tokamak divertor, with plasma conditions expected to be in the range $5 \times 10^{19} \lesssim n_e \lesssim 5 \times 10^{20} \text{ m}^{-3}$ and $0.5 \lesssim T_e \lesssim 5 \text{ eV}$. First we considered three different

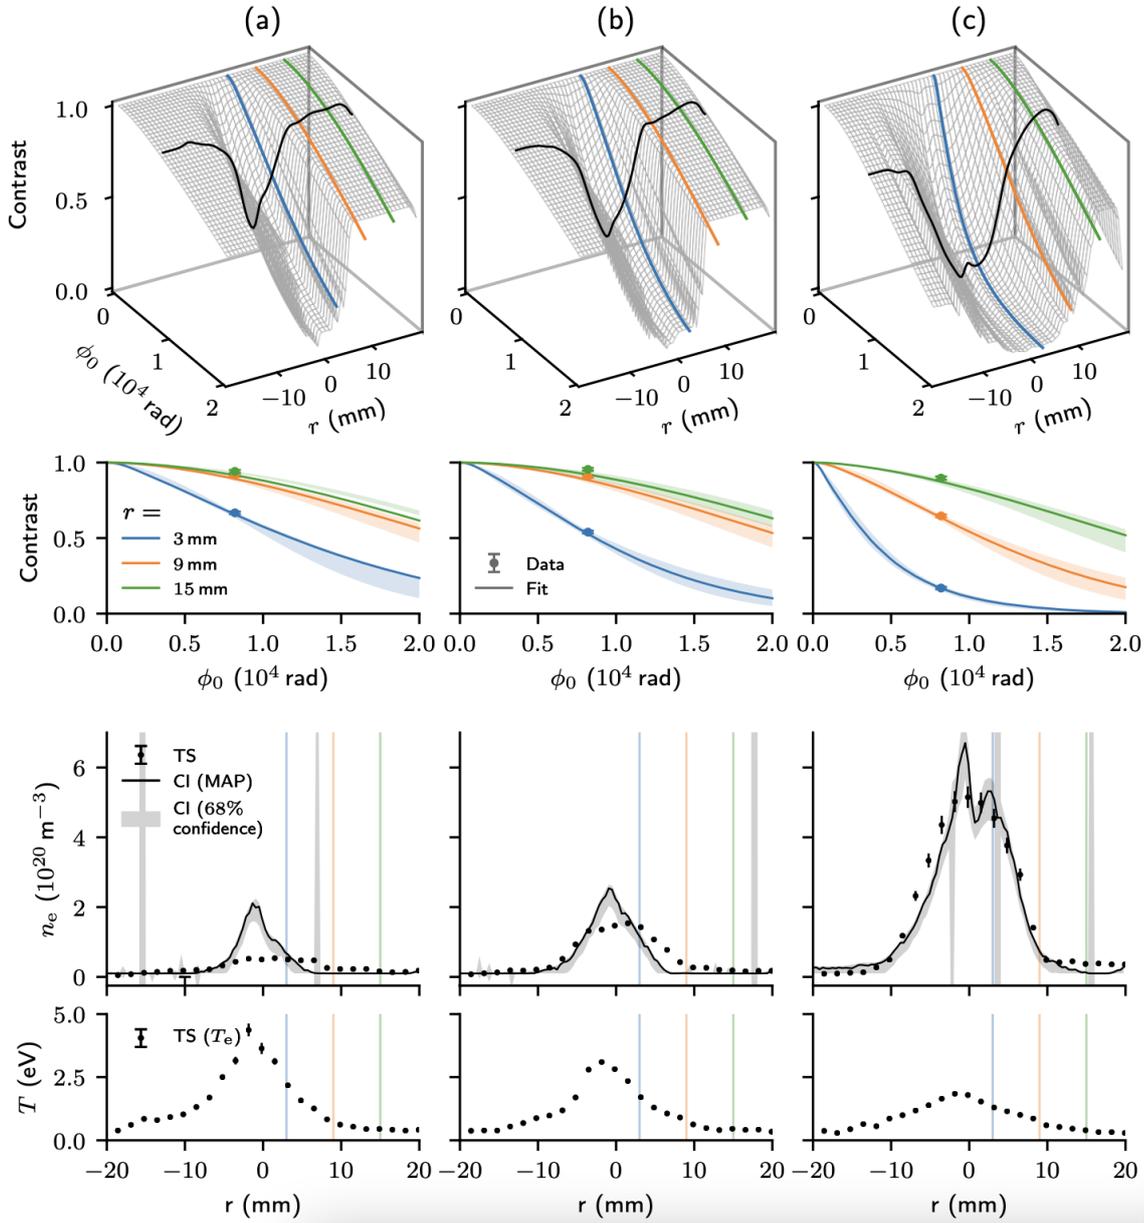


Figure 5.24: Single-delay CI results observing H_γ emission (434.0 nm) from Magnum-PSI over a range of plasma conditions. Source gas flow rate is $Q_s = 4.5$ SLM and magnetic field strength is $|\mathbf{B}| = 1.5$ T for all discharges shown here, while source current I_s is: (a) $I_s = 125$ A, (b) $I_s = 150$ A, (c) $I_s = 190$ A. See text for detailed explanation.

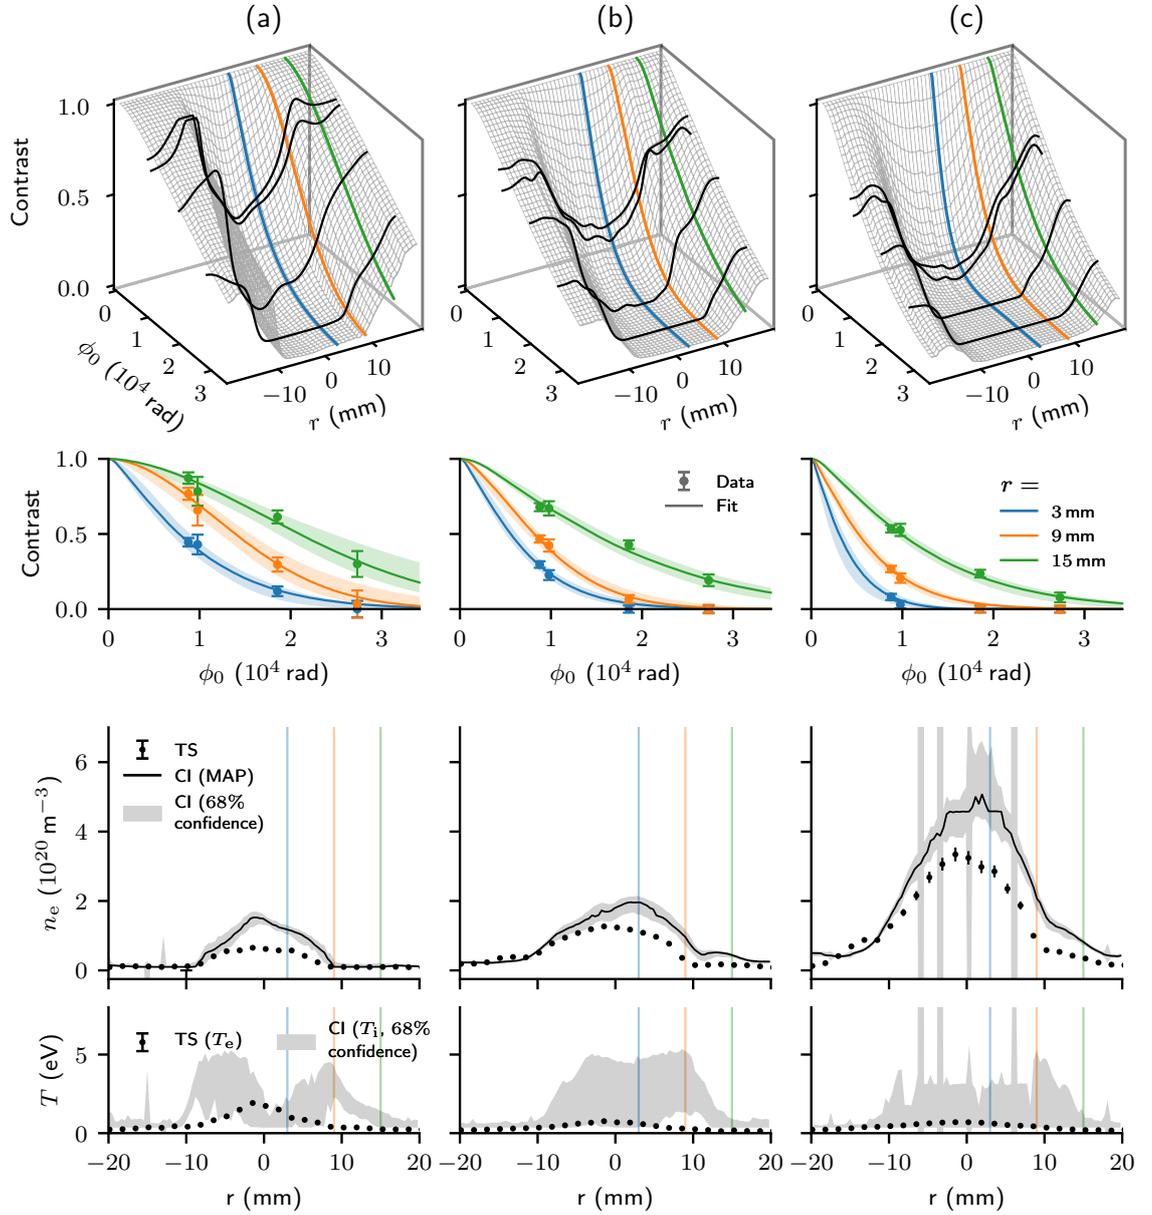


Figure 5.25: Multi-delay CI results observing H_δ emission (410.2 nm) from Magnum-PSI over a range of plasma conditions. Magnetic field strength is $|\mathbf{B}| = 0.8$ T for all discharges shown here, while source current I_s and gas flow rate Q_s are: (a) $I_s = 120$ A, $Q_s = 5$ SLM, (b) $I_s = 140$ A, $Q_s = 6$ SLM, (c) $I_s = 180$ A, $Q_s = 6$ SLM. See text for detailed explanation.

lineshape models for Balmer series lines with varying levels of complexity, each of which has been used in previous spectroscopic investigations. Next, we introduced a fitting tool that uses Bayesian parameter estimation to infer n_e from single-delay or multi-delay CI contrast measurements. Tests of this fitting tool using synthetic data showed that multi-delay instrument configurations ought to significantly improve the dynamic range of the n_e measurement by improving robustness against Doppler broadening effects. This allows n_e at the lower end of the anticipated MAST-U range to be measured without significant systematic error. These tests also showed that the simplified analytical lineshape model for Stark-broadening used in previous work can result in large systematic errors when compared to the more modern numerical lineshape models, which should be used instead.

Experimental CI measurements of divertor-relevant plasma conditions were then presented from the Magnum-PSI linear device at DIFFER in Eindhoven, NL. For measurements of H_γ , good agreement was found between the measured CI n_e profiles and those measured using the Magnum-PSI Thomson scattering diagnostic. Also, the predicted benefit of a multi-delay CI configuration appears to be born out experimentally, as the multi-delay instrument could accurately measure down to at least $n_e \sim 5 \times 10^{19} \text{ m}^{-3}$ while accurate recovery of the beam profile below $n_e \sim 10^{20} \text{ m}^{-3}$ was not achieved. However, there were some unanswered questions too. The inferred T_i results from the multi-delay instrument are a factor ~ 2 higher temperature than the T_e recorded by the TS diagnostic. Additionally, a consistent and significant over-estimate of n_e was found in the results of CI measurements of the H_δ emission line.

The measurements in this Chapter were carried out using CI instruments with the pixelated phase-mask designs introduced in Chapter 3. While the simple geometry of the Magnum-PSI plasma beam is perhaps not the best example for showcasing the improved spatial resolution of the PPM-CI designs, this work represents the first experimental demonstration of this technique.

Chapter 6

Summary and Future Work

In this thesis, we have presented several new developments to the coherence imaging (CI) technique for diagnosing the scrape-off layer (SOL) and divertor plasma of fusion experiments. CI is a narrow-band spectral imaging technique that isolates a single emission line in the visible range and measures its Doppler shift and line width. The Doppler shift encodes the flow velocity component of the emitting species along the instrument line of sight. The line width encodes either the emitting species temperature T_i or the electron density n_e , depending on whether Doppler broadening or Stark broadening respectively is the dominant effect. Broadly, the three new developments presented in this thesis are

1. Incorporation of a pixelated polariser sensor into the CI instrument, improving spatial resolution and robustness.
2. A new technique for absolute calibration of CI flow velocity measurements using standard gas-discharge lamps, with significantly reduced costs compared to existing techniques.
3. Improved methods for 2-D n_e CI measurement in divertor-relevant plasma conditions, demonstrated experimentally on the Magnum-PSI linear plasma experiment with a direct comparison to 1-D Thomson scattering measurements.

In Chapter 1, the tokamak was introduced and we saw that one of the key challenges in designing a fusion reactor is handling the large amounts of power exhausted via the SOL into the divertor target plates. High-quality measurements of the SOL and divertor plasma are required to benchmark numerical simulations of divertor performance, which will lead to a better understanding of current experiments and to more reliable predictions for future experiments. Using the example of MAST-U,

a spherical tokamak due to begin operations later this year, we introduced the different diagnostic tools typically used to make these measurements and discussed their relative merits. CI was introduced as an imaging technique that can measure 2-D maps of important plasma properties over a wide field of view and at a high spatial resolution. CI is desirable as it can observe plasma properties that are not directly observable using other imaging techniques.

Chapter 2 covered CI theory. The CI instrument is a two-beam interferometer, splitting observed light into two beams and introducing a time delay between them. When the beams recombine an interference fringe pattern is formed which carries information about the temporal coherence of the light, and therefore also about its frequency spectrum. If a single emission line from the plasma is isolated, then the observed fringe pattern is determined by a small number of well-understood processes and assumes a particularly simple form (Figure 2.2): fringe phase is determined by the line's central wavelength and fringe contrast is determined by the line's width. The line's central wavelength encodes emitter flow velocity via Doppler shift while line width is determined by different parameters depending on emitting species and plasma conditions. The simple form of the fringe pattern means that the interferometer delay need only be scanned about a single fixed value to recover the parameters of interest, so robust designs without moving parts become feasible. In practice the required fixed delay is introduced using birefringent crystal plates (Figure 2.5), an example of a wider technique in optics called polarisation interferometry. CI has been most widely implemented to measure flow velocity of C and He impurity ions via passive emission in the plasma edge[51, 52, 78, 90]. Applications to Thomson scattering emission[94–96] and charge-exchange recombination emission[93] have been demonstrated but have not yet been widely adopted. CI measurement of n_e via Stark broadening of Balmer emission lines has also been demonstrated, on a linear plasma experiment[97]. Though not the subject of this thesis, CI can also be used for spectro-polarimetry, most prominently used to measure magnetic pitch angle via the motional Stark effect[70].

6.1 Pixelated Phase-mask CI

Pixelated phase mask (PPM) interferometry is a relatively new technique in optics for maximising the spatial resolution of interferometric imaging measurements[102]. It uses a camera with a pixelated polariser array bonded to its sensor at manufacture to produce a pixelated interference fringe pattern. In Chapter 3, we saw that this PPM technique can be incorporated into the CI framework easily, and only requires minor hardware changes. Both single-delay and multi-delay PPM-CI designs were presented in this chapter and experimental images for each design were shown. Demodulation

algorithms for recovering the interferogram phase and contrast were described and tested experimentally and numerically. All of the results presented in this thesis — with the exception of the work in Section 4.2.1 — use the PPM-CI instrument designs presented in this chapter.

6.2 Calibration of Flow Measurements using Gas-discharge Lamps

MAST-U will operate a CI diagnostic for 2-D measurements of impurity flow velocity in the plasma edge, which is expected to be as large as 30 km/s. Chapter 4 of this work investigated phase calibration, which sets the zero point of the CI flow measurements. This calibration requires the phase of a set of unshifted fringes, measured at the rest-frame wavelength of the targeted emission line, to be subtracted from the Doppler-shifted phase of the plasma observations. The problem is that this calibration phase is very sensitive to changes in ambient temperature and alignment, for example changing by 16 km/s equivalent over the period of a week when monitored on MAST in 2013[63]. Phase calibration cannot then be a one-off. It must be carried out regularly — ideally inter-shot — to monitor these changes, and so requires its own hardware. In the last few years, several fusion experiments with permanently installed flow CI diagnostics have installed tuneable laser calibration systems that can provide in-situ monitoring of any calibration changes. DIII-D have installed a custom diode laser, tuneable over the 464–468 nm wavelength range[91, 92] and W7-X have installed a an Optical Parametric Oscillator (OPO) laser, tuneable over the 450–650 nm wavelength range[77]. In both cases a wavemeter is required to measure and control the laser wavelength. These systems represent the most direct method for calibrating flow measurements of emission lines within their tuning ranges. However, they increase the total hardware cost of the diagnostic significantly, being at least comparable to the cost of the CI instrument itself. Chapter 4 of this work sought to answer the question: *is there a cheaper way to calibrate the flow measurement zero point to a high accuracy?*. Where we take ‘high accuracy’ here to mean better than ± 1 km/s error in the measured flow.

Necessity is the mother of invention and in Chapter 4 we showed that the answer is yes. Our method for phase calibration uses only gas-discharge lamps and bandpass interference filters, both cheap and common optical components. While the available gas-discharge lamps do not, generally, emit intense coherent light at the precise rest-frame wavelengths of the targeted plasma lines, they do emit intense coherent light at wavelengths nearby. We show that by measuring the CI phase at

these nearby wavelengths, and fitting an instrument model to the data, the required calibration phase can be predicted to a high accuracy at the plasma wavelength of interest, without the need for a light source at that wavelength. This is analogous to how gas-discharge lamps are used to calibrate the wavelength axis of a slit-coupled spectrograph. To demonstrate the method, we measured phase images at four emission lines produced by Cd and Zn gas-discharge lamps spread over a ~ 40 nm range. On fitting an instrument model to the measurements, we were able to correctly predict the calibration phase image measured at a fifth lamp line, separated ~ 5 nm from the nearest measurement included in the fit, to our goal accuracy of ± 1 km/s flow-equivalent. This was made possible by dealing with two complications that have not been previously considered in detail in the CI literature:

1. Measured phase data is inherently wrapped within the $(-\pi, \pi]$ rad interval. Lamp lines are typically sparse enough in wavelength that the measured phase data cannot be trivially unwrapped. This ambiguity is handled in this work by using wrapped likelihood functions when fitting the model to the wrapped phase data. When this approach is combined with extremely broad prior information, obviously incorrect fit solutions are automatically rejected and the ambiguity is overcome.
2. Over the wavelength shifts between the calibration lamp lines, the available Sellmeier coefficients describing instrument dispersion do not model the CI phase shift with sufficient accuracy. This is predictable based on the discrepancies between available sets of Sellmeier coefficients (Figure 4.4) and is also observed experimentally. We handle this by leaving two of the Sellmeier coefficients as free parameters when fitting the instrument model to the measured phase data.

At least one of the six Cd and Zn lamp lines measured in Chapter 4 is within 5 nm of each of the three impurity lines targeted for measurement on MAST-U (C III at 464.8810 nm, He II at 468.5701 nm and C II at 514.1842 nm). It follows that each can be calibrated using the methods and lamp lines discussed in that chapter.

6.3 Coherence Imaging Measurements of Electron Density

In divertor plasma conditions, Stark broadening is a significant component of the observed lineshapes for the hydrogen Balmer emission series. The Stark width is determined mostly by electron density, and chordal spectroscopic measurement of this

effect is an established technique for diagnosing n_e in the tokamak divertor plasma. The pre-existing work applying this principle to CI was at the proof-of-principle level[97]. One goal of Chapter 5 was to build on this existing work, optimising the n_e CI measurement technique in preparation for implementation on the MAST-U divertor. Two important findings that have come out of the modelling carried out in this work are:

1. That although the simplified Stark-broadened lineshape model used in the previous CI work provides an intuitive analytical treatment for interpreting results, it can introduce significant systematic error into the inferred CI n_e when compared to more sophisticated numerical lineshape models[124, 129, 130] (see Figure 5.12). These numerical models are already the gold standard for grating spectroscopy[21, 23, 24] and should be used where appropriate for all future CI work.
2. Multi-delay CI measurements can significantly extend the dynamic range of n_e measurements to lower densities $< 10^{20} \text{ m}^{-3}$ at moderate temperatures $T_i \sim 3 \text{ eV}$ by fitting to the Doppler broadening component of the lineshape which, in single-delay CI measurements introduces large systematic errors in these plasma conditions.

Experimental CI measurements of n_e were made on the Magnum-PSI linear plasma experiment, which can produce a wide range of plasma conditions ($10^{19} \leq n_e \leq 10^{21} \text{ m}^{-3}$, $0.5 \leq T_e \leq 5 \text{ eV}$, $0 \leq |\mathbf{B}| \leq 2.5 \text{ T}$), easily covering those expected in the MAST-U divertor (see Figure 1.6). To avoid stray magnetic field problems, the CI instrument was mounted 5.5 m from the plasma beam and a telescopic lens was used as the objective lens such that the 20 mm diameter plasma beam filled the field of view. A pixel-based technique was used to tomographically invert the line-integrated, emissivity-weighted projections of CI contrast, assuming the plasma beam to be made up of two cylindrical segments, each with cylindrically symmetrical properties. To infer n_e from the local CI contrast, a fitting tool was developed based on a pre-calculated look-up table of contrast profiles and the framework of Bayesian parameter estimation. The inferred n_e profiles showed good agreement with the Thomson scattering measurements for multi-delay CI measurements of H_γ emission (434.0 nm) across a $5 \times 10^{19} - 8 \times 10^{20} \text{ m}^{-3}$ density range. For a single-delay CI instrument observing the same line, the influence of Doppler broadening makes accurately inferring n_e below $n_e \sim 10^{20} \text{ m}^{-3}$ difficult at the relevant temperatures $\sim 3 \text{ eV}$. For CI measurements of H_δ emission (410.2 nm), an unexplained overestimate is found relative to the Thomson scattering measurements, possibly due to contaminant emission or a systematic calibration error.

6.4 Future Work

The results of this thesis will help optimise CI measurements of the MAST-U SOL and divertor. For calibration of CI measurements of carbon and helium impurity flows, the method presented in Chapter 4 ought to be adopted, in-situ and fully automated. A filter wheel can be used to remotely cycle through the different calibration filters in turn while remote control of the Cd and Zn lamp power can ensure that the appropriate lamp illuminates each filter. On the software side, the instrument model fitting algorithm and necessary phase extrapolation is already automated. While the clear advantage of this method is cost, its disadvantage is the complexity of the data analysis required relative to direct measurement of tuneable laser light. This can be mitigated by publishing the instrument model and fitting algorithm (all written in Python) as an open source software package. While Chapter 4 demonstrated that absolute phase calibration using gas-discharge lamps is possible, no investigation was carried out into how robust the results of the instrument model fit are to changes in mechanical alignment or ambient temperature — instrument phase certainly needs to be regularly tracked for environmental changes at a wavelength near to the targeted emission line but how regularly does the fitting procedure need to be carried out? As mentioned in 4.2.4, the method described could also feasibly be made simpler in a number of ways.

Applying PPM-CI to MAST-U has the advantage that it would maximise spatial resolution and eliminate ringing artefacts from Fourier demodulation techniques, which were a limitation in previous work that used the more typical linear spatial carrier method[51, 63]. In this work, it was found experimentally that the instrumental contrast and phase calibration parameters for pixelated carriers are more robust against environmental effects than the equivalent quantities for linear carriers. The only disadvantage in using the PPM-CI instrument as presented in this work is the limited frame rate of the FLIR Blackfly camera (75 Hz), however high-speed cameras with pixelated polariser sensors are available (e.g. the Photron Chrysta with a frame rate of 7000 Hz) if measurements of transient phenomena are required. Although not the topic of this thesis, the PPM-CI technique could also be used for imaging motional stark effect spectropolarimetry[70], conferring the same advantages as for the purely spectroscopic applications discussed in this work.

For measurements of n_e in the MAST-U divertor, the results presented in this thesis suggest that, in the absence of strong prior information on neutral hydrogen temperature, the multi-delay CI configuration extends the range over which accurate n_e measurements can be made to lower density, higher temperature conditions, relative to the single-delay configuration.

We have already mentioned that multi-delay CI would be required to detect the presence of a non-Maxwellian velocity distribution function (VDF) in the Doppler-broadened lineshape. However, even for a Maxwellian VDF, multi-delay CI could be useful for improving the quality of tomographic inversion of CI data. Since the multiple delays confer more information on the observed, projected lineshape, the inversion would be better constrained. This would likely reduce artefacts most significantly for lines of sight along which the flow and/or broadening shows the most variation.

Although sadly outside the scope of this project, in the last few years coherence imaging measurements have begun to contribute to high-quality physics studies of fusion plasmas[52, 89, 90]. With its relative affordability, coverage and flexibility, this diagnostic technique looks set to make significant contributions to fusion plasma exhaust studies in years to come.

Appendix A

Maths

A.1 Fourier transforms

Convention: continuous transform

In this thesis we use, as a convention, the unitary form of the Fourier transform. The forward transform of the time-domain function $f(t)$ is

$$F(\nu) = \int_{-\infty}^{\infty} f(t) \exp(-2\pi i \nu t) dt, \quad (\text{A.1.1})$$

And the inverse transform of the frequency domain function $F(\nu)$ is:

$$f(t) = \int_{-\infty}^{\infty} F(\nu) \exp(2\pi i \nu t) d\nu. \quad (\text{A.1.2})$$

Shifting property

$$F(\nu - \nu_0) = \mathcal{F}\{f(t) \exp(-2\pi i \nu_0 t)\} \quad (\text{A.1.3})$$

and

$$f(t - t_0) = \mathcal{F}^{-1}\{F(\nu) \exp(2\pi i \nu t_0)\} \quad (\text{A.1.4})$$

Convolution theorem

Convolution in one domain (e.g., time) is equal to point-wise multiplication in the other domain (e.g., frequency).

$$\mathcal{F}\{f_1(t) * f_2(t)\} = \mathcal{F}\{f_1(t)\} \cdot \mathcal{F}\{f_2(t)\} = F_1(\nu) \cdot F_2(\nu) \quad (\text{A.1.5})$$

Gaussian function

The area-normalised Gaussian function (Normal distribution) centred at k_0 in frequency space can be written

$$F(k) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left[\frac{k-k_0}{\sigma}\right]^2\right), \quad (\text{A.1.6})$$

where σ is a scale parameter specifying the standard deviation. Its inverse Fourier transform, in time/space, is given by

$$f(x) = \exp(2\pi i k_0 x) \exp(-2\pi^2 \sigma^2 x^2). \quad (\text{A.1.7})$$

Lorentzian function

The area-normalised Lorentzian function (Cauchy distribution) centred at x_0 can be written

$$f(x) = \frac{1}{\pi\Gamma} \left(\frac{\Gamma^2}{(x-x_0)^2 + \Gamma^2} \right), \quad (\text{A.1.8})$$

where Γ is a scale parameter specifying the half-width at half-maximum. Its Fourier transform is given by

$$F(k) = \exp(-2\pi i x_0 k) \exp(-2\pi\Gamma|k|). \quad (\text{A.1.9})$$

Sampling property

$$\int_{-\infty}^{\infty} f(t)\delta(t-t_0)dt = f(t_0) \quad (\text{A.1.10})$$

$$\int_{-\infty}^{\infty} \delta(\nu) \exp(2\pi i \nu t) d\nu = 1 \quad (\text{A.1.11})$$

A.2 Derivation of the Wiener-Khinchin theorem

Autocorrelation is the correlation of a function with a delayed copy of itself, as a function of that delay. Consider a function $u(t)$ that is either deterministic or is a wide-sense stationary random process. The autocorrelation of $u(t)$ as a function of delay τ is

$$C(\tau) \equiv \int_{-\infty}^{\infty} u^*(t)u(t+\tau)d\tau. \quad (\text{A.2.1})$$

If $\mathcal{F}\{u(t)\} = U(\nu)$, then $u^*(t)$ can be written

$$u^*(t) = \int_{-\infty}^{\infty} U^*(\nu) \exp(-2\pi i \nu t) d\nu, \quad (\text{A.2.2})$$

and $u(t + \tau)$ can be written

$$u(t + \tau) = \int_{-\infty}^{\infty} U(\nu) \exp(2\pi i \nu [t + \tau]) d\nu. \quad (\text{A.2.3})$$

We then substitute these two expressions into Equation A.2.1, taking care to independently label the frequency domains of the shifted and unshifted functions, to get

$$C(\tau) = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} U^*(\nu) \exp(-2\pi i \nu t) d\nu \right] \left[\int_{-\infty}^{\infty} U(\nu') \exp(2\pi i \nu' [t + \tau]) d\nu' \right] d\tau. \quad (\text{A.2.4})$$

Rearranging, and making use of the shifting property and the Fourier transform of the Dirac delta function, we can write:

$$C(\tau) = \int_{-\infty}^{\infty} U^*(\nu) U(\nu') \delta(\nu - \nu') \exp(2\pi i \nu' \tau) d\nu d\nu'. \quad (\text{A.2.5})$$

Next, we can use the sampling property of the Dirac delta function to show that

$$C(\tau) = \int_{-\infty}^{\infty} |\mathcal{E}(\nu)|^2 \exp(2\pi i \nu \tau) d\nu. \quad (\text{A.2.6})$$

Here, $|\mathcal{E}(\nu)|^2$ is called the power spectral density. So, with the stated assumptions, the autocorrelation and the power spectral density of a function are a Fourier transform pair.

Appendix B

Fourier Demodulation of Interference Fringes

In this appendix, we outline how Fourier techniques are used to demodulate an interference fringe pattern, extracting the brightness, contrast and phase signals. We will also consider how noise in the raw image manifests in the demodulated signals. The Fourier method was first proposed for 1-D signals in 1982 by Takeda et al.[150]. It was soon extended to 2-D images in 1986 by Bone et al.[151]. It is now a standard and well-understood technique[112].

The convention for the discrete Fourier transform (DFT) used in this report is chosen to match that of the Numpy Python library. When x and y are the pixel coordinates on a camera sensor array with respective dimensions N and M , the forward transform of an image f_{xy} is:

$$F_{kl} = \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} f_{xy} \exp \left(-2\pi i \left[\frac{xk}{N} + \frac{yl}{M} \right] \right). \quad (\text{B.0.1})$$

And the inverse transform is:

$$f_{xy} = \frac{1}{NM} \sum_{k=0}^{N-1} \sum_{l=0}^{M-1} F_{kl} \exp \left(2\pi i \left[\frac{xk}{N} + \frac{yl}{M} \right] \right). \quad (\text{B.0.2})$$

Henceforth, when Fourier transform is abbreviated to FT it refers to this 2D discrete formulation of the transform. Parseval's theorem is a statement of the conservation of "energy" through the transform. For the 2D DFT as defined, Parseval's theorem can be written:

$$\sum_{x=0}^{N-1} \sum_{y=0}^{M-1} |f_{xy}|^2 = \frac{1}{NM} \sum_{k=0}^{N-1} \sum_{l=0}^{M-1} |F_{kl}|^2. \quad (\text{B.0.3})$$

Consider a measured (real-valued) interferogram of the form

$$s_{xy} = I_{xy} \left[1 + \zeta_{xy} \cos(\phi_{xy} + \phi_c) \right], \quad (\text{B.0.4})$$

where x and y are the pixel coordinates. There are three unknowns in the image: the brightness I_{xy} , the contrast ζ_{xy} and the phase ϕ_{xy} . The carrier term in the interferogram phase is ϕ_c and it is assumed to be known. The purpose of the carrier is to separate the desired contrast and phase information in (spatial) frequency-space from the background brightness of the scene. Assume that the carrier is linear in y , giving us perfectly horizontal fringes:

$$\phi_c = \frac{2\pi k_c}{N} y. \quad (\text{B.0.5})$$

Figure B.1(a) shows an interferogram generated using Equations B.0.4 and B.0.5 (plus noise) with dimensions $N = M = 1000$ and whose linear carrier (Figure B.1(b)) is modulated by inputs I_{xy} , ζ_{xy} and ϕ_{xy} . The Fourier transform of the interferogram $S_{kl} \equiv \mathcal{F}\{s_{xy}\}$ can be written

$$\begin{aligned} S_{kl} &= \mathcal{F}\{I_{xy}\} + \mathcal{F}\{I_{xy}\zeta_{xy} \cos(\phi_{xy} + \phi_c)\} \\ &= \mathcal{F}\{I_{xy}\} + \mathcal{F}\{I_{xy}\zeta_{xy} \cos \phi_{xy} \cos \phi_c\} - \mathcal{F}\{I_{xy}\zeta_{xy} \sin \phi_{xy} \sin \phi_c\} \\ &= \mathcal{F}\{I_{xy}\} \\ &\quad + \mathcal{F}\{I_{xy}\} * \mathcal{F}\{\zeta_{xy}\} * \mathcal{F}\{\cos \phi_{xy}\} * \frac{1}{2} \left(\delta_{l0} \delta_{k(-k_c)} + \delta_{l0} \delta_{kk_c} \right) \\ &\quad + \mathcal{F}\{I_{xy}\} * \mathcal{F}\{\zeta_{xy}\} * \mathcal{F}\{\sin \phi_{xy}\} * \frac{i}{2} \left(\delta_{l0} \delta_{k(-k_c)} + \delta_{l0} \delta_{kk_c} \right). \end{aligned} \quad (\text{B.0.6})$$

Here, we have made use of the convolution theorem, and the standard result for the Fourier transform of $\cos(ay)$. The Kronecker delta δ_{ij} is defined

$$\delta_{ij} = \begin{cases} 0 & (i \neq j), \\ 1 & (i = j). \end{cases} \quad (\text{B.0.7})$$

So the Fourier transformed signal consists of three parts: the brightness image $\mathcal{F}\{I_{xy}\}$ and then two terms centred at the positive and negative carrier frequency $\pm k_c$. If the spectral content of I_{xy} , ζ_{xy} and $\cos \phi_{xy}$ is sufficiently low-frequency, then these three components are clearly separable in the frequency domain. Figure B.1(c) shows the periodogram estimate of the image power spectral density (PSD), found by taking $\frac{|S_{kl}|^2}{NM}$. Using the filter windows $W_{xy}^{(\text{DC})}$ and $W_{xy}^{(c)}$, brightness can be extracted as

$$I_{xy} = \mathcal{F}^{-1} \left\{ W_{kl}^{(\text{DC})} S_{kl} \right\} \quad (\text{B.0.8})$$

and subtracted from the image.

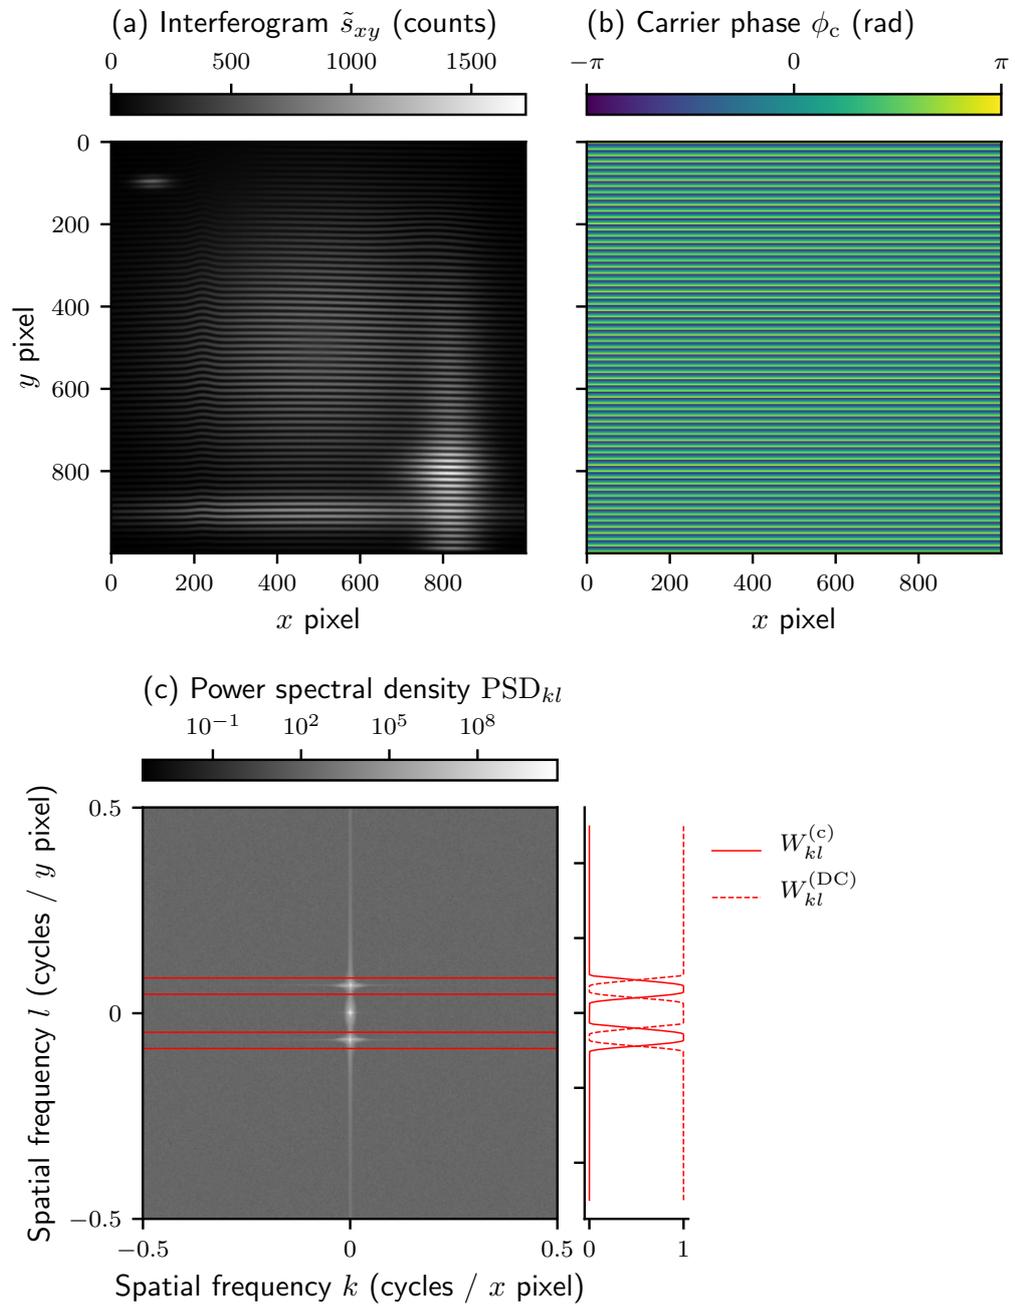


Figure B.1: (a) Synthetic interferogram \tilde{S}_{xy} (with noise). (b) Modelled ideal linear carrier phase ϕ_c . (c) Periodogram estimate of image power spectral density and example window functions W_{kl} for isolating the DC (brightness) and spatial carrier components of the image.

This leaves us with the filtered carrier frequency in the Fourier domain

$$\begin{aligned} W_{kl}^{(c)} S_{kl} &= \mathcal{F}\{I_{xy}\} * \mathcal{F}\{\zeta_{xy}\} * \mathcal{F}\{\cos \phi_{xy}\} * \frac{1}{2} \left(\delta_{l0} \delta_{k(-k_c)} + \delta_{l0} \delta_{kk_c} \right) \\ &+ \mathcal{F}\{I_{xy}\} * \mathcal{F}\{\zeta_{xy}\} * \mathcal{F}\{\sin \phi_{xy}\} * \frac{i}{2} \left(\delta_{l0} \delta_{k(-k_c)} + \delta_{l0} \delta_{kk_c} \right) \end{aligned} \quad (\text{B.0.9})$$

Next, convert this into the analytic signal by doubling the positive frequencies, preserving the zero frequency component and suppressing the negative frequencies. Alternatively this can be done using the Hilbert transform.

$$\begin{aligned} W_{kl}^{(c)} S_{kl}^{(a)} &= \mathcal{F}\{I_{xy}\} * \mathcal{F}\{\zeta_{xy}\} * \mathcal{F}\{\cos \phi_{xy}\} * \delta_{l0} \delta_{kk_c} \\ &+ \mathcal{F}\{I_{xy}\} * \mathcal{F}\{\zeta_{xy}\} * \mathcal{F}\{\sin \phi_{xy}\} * i \delta_{l0} \delta_{kk_c} \end{aligned} \quad (\text{B.0.10})$$

Next, take the inverse Fourier transform:

$$\begin{aligned} \mathcal{F}^{-1}\{W_{kl}^{(c)} S_{kl}^{(a)}\} &= I_{xy} \zeta_{xy} \cos \phi_{xy} \exp i \phi_c + I_{xy} \zeta_{xy} i \sin \phi_{xy} \exp i \phi_c \\ &= I_{xy} \zeta_{xy} \exp(i[\phi_{xy} + \phi_c]) \end{aligned} \quad (\text{B.0.11})$$

From this, it follows that ζ_{xy} can be extracted by taking

$$\zeta_{xy} = \left| \frac{\mathcal{F}^{-1}\{W_{kl}^{(c)} S_{kl}^{(a)}\}}{I_{xy}} \right| \quad (\text{B.0.12})$$

and ϕ_{xy} can be extracted as

$$\phi_{xy} = \arg \left(\mathcal{F}^{-1}\{W_{kl}^{(c)} S_{kl}^{(a)}\} \right) - \phi_c \quad (\text{B.0.13})$$

Figure B.2 shows the input and demodulated brightness, phase and contrast for the synthetic interferogram in Figure B.1.

B.1 Noise Propagation

Next, consider how noise in the raw interferogram propagates into the demodulated brightness, contrast and phase images. The noise variance at pixel x, y in the raw interferogram and in the demodulated output are related by a constant of proportionality that depends on the Fourier domain window function used. This simple result is well-known in fringe pattern analysis[151] but has not yet found use in CI. Some previous CI investigations have used numerical simulation instead to quantify noise in the extracted phase and/or contrast data[63, 70]. In this section, standard signal processing techniques are used to derive the result. Additionally, we will find that the PDFs describing the demodulated phase and contrast will approximate normal distributions only at high SNR. As such, care must be taken when drawing inferences

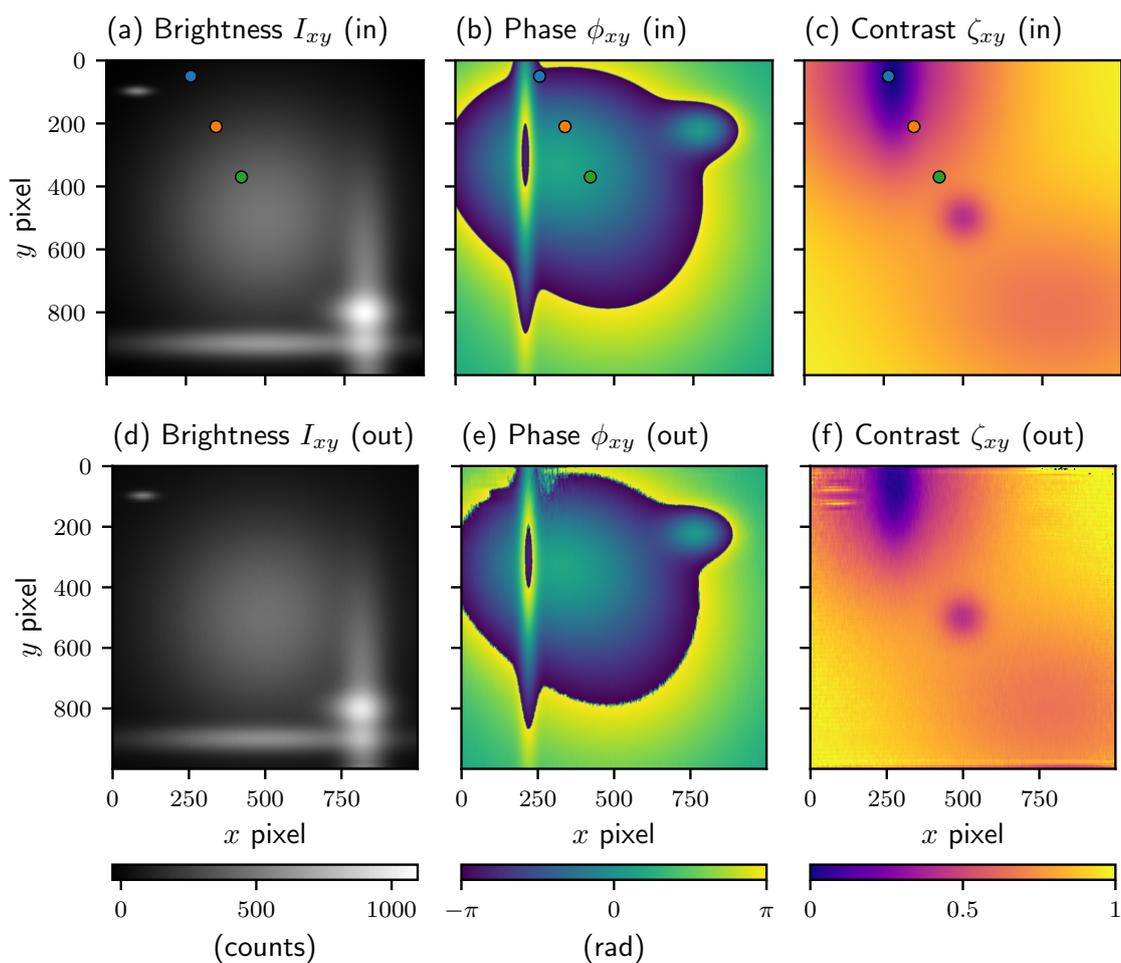


Figure B.2: Input and demodulated output profiles of brightness, phase and contrast from the synthetic interferogram image shown in Figure B.1(a).

from low SNR data.

First, express the measured interferogram \tilde{s}_{xy} as the sum of a deterministic signal and a zero-mean noise term:

$$\tilde{s}_{xy} = s_{xy} + e_{xy}. \quad (\text{B.1.1})$$

Using the linearity of the Fourier transform, we can say that

$$\begin{aligned} \tilde{S}_{kl} &= \mathcal{F}\{\tilde{s}_{xy}\} \\ &= \mathcal{F}\{s_{xy}\} + \mathcal{F}\{e_{xy}\} \\ &= S_{kl} + E_{kl}. \end{aligned} \quad (\text{B.1.2})$$

To understand more about the Fourier-domain noise term E_{kl} , we can write out its full definition as the Fourier transform of e_{xy} :

$$\begin{aligned} E_{kl} &= \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} e_{xy} \cos\left(2\pi\left[\frac{xk}{N} + \frac{yl}{M}\right]\right) \\ &\quad - i \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} e_{xy} \sin\left(2\pi\left[\frac{xk}{N} + \frac{yl}{M}\right]\right). \end{aligned} \quad (\text{B.1.3})$$

Here, the terms inside the summations are themselves random variables. We can invoke the central limit theorem to say that, for sufficiently large image dimensions N_x and N_y , the real and imaginary parts of E_{kl} are each normally distributed, regardless of the noise distribution in the spatial domain[152]. It can be shown that E_{kl} is zero-mean, provided that e_{xy} is too, and that E_{kl} is roughly statistically independent between different pixels and between real and imaginary components[152]. Equation B.1.3 also implies an equal variance of the real and imaginary parts of E_{kl} .

To say more about the nature of the Fourier domain noise, and its relation to the spatial domain, we will consider the autocorrelation of the noise measured at pixel x, y with that of x', y' , defined as

$$C_{xyx'y'} \equiv \mathbb{E}[e_{xy}e_{x'y'}^*] \quad (\text{B.1.4})$$

Where $\mathbb{E}[\dots]$ is the expected value operator and $*$ is the complex conjugate. We will assume that e_{xy} is normally distributed with pixel-dependent variance σ_{xy}^2 . Formally this is written $e_{xy} \sim \mathcal{N}(0, \sigma_{xy}^2)$. Importantly, we will also consider the noise at each pixel to be statistically independent of all others (uncorrelated). This allows us to write $C_{xyx'y'}$ as

$$C_{xyx'y'} = \sigma_{xy}^2 \delta_{xx'} \delta_{yy'}. \quad (\text{B.1.5})$$

We can recast this in terms of the pixel ‘lag’ in each dimension, defining $a \equiv x - x'$

and $b \equiv y - y'$:

$$C_{xyab} = \sigma_{xy}^2 \delta_a \delta_b. \quad (\text{B.1.6})$$

Since σ_{xy}^2 depends on pixel coordinate x, y , the image noise is a ‘non-stationary’ process (in space). For a non-stationary signal evolving in space/time, the power spectral density (PSD) is the Fourier transform of the *space/time-averaged* autocorrelation function[153]. This is a generalised form of the Wiener-Khinchine theorem that underpins Fourier transform spectroscopy (see Section 2.1 and Appendix A.2) that equates the PSD of a stationary signal with the Fourier transform of that signal’s autocorrelation function. So, we can average the autocorrelation in Equation B.1.6 and take its Fourier transform to obtain the PSD:

$$\begin{aligned} \text{PSD}_{kl} &= \mathcal{F} \left\{ \frac{1}{NM} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sigma_{xy}^2 \delta_a \delta_b \right\} \\ &= \frac{1}{NM} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sigma_{xy}^2 \mathcal{F} \{ \delta_a \delta_b \} \\ &= \frac{1}{NM} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sigma_{xy}^2 \end{aligned} \quad (\text{B.1.7})$$

That the PSD of e_{xy} is independent of the spatial frequency indices k, l means that it is constant across all frequencies—we have white noise.

The power spectral density of the noise e_{xy} is defined as:

$$\begin{aligned} \text{PSD}_{kl} &\equiv \frac{1}{NM} \text{E} [|E_{kl}|^2] \\ &= \frac{1}{NM} \text{E} [\Re \{ E_{kl} \}^2 + \Im \{ E_{kl} \}^2] \\ &= \frac{2\sigma^2}{NM}, \end{aligned} \quad (\text{B.1.8})$$

Setting equal Equations B.1.7 and B.1.8, we arrive at a relation between the noise variance in the spatial and frequency domains

$$\sigma^2 = \frac{1}{2} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sigma_{xy}^2 \quad (\text{B.1.9})$$

The same result can be derived more directly by considering the discrete version of Parseval’s theorem (Equation B.0.3).

B.1.1 Fourier-domain Filtering

Next we will consider the effect of Fourier-domain filtering. To separate components in frequency space, we used a (real-valued) filter window function W_{kl} . The power

spectral density of the filtered noise becomes

$$\begin{aligned}
\text{PSD}_{kl}^{(f)} &= \frac{1}{NM} \mathbb{E}[|W_{kl} E_{kl}|^2] \\
&= W_{kl}^2 \frac{2\sigma^2}{NM} \\
&= W_{kl}^2 \frac{1}{NM} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sigma_{xy}^2
\end{aligned} \tag{B.1.10}$$

Writing the generalised Wiener-Khinchine theorem relating the space-averaged auto-correlation and the power spectral density we find that

$$\begin{aligned}
\langle C_{xyab}^{(f)} \rangle &= \mathcal{F}^{-1}\{\text{PSD}_{kl}^{(f)}\} \\
&= \left(\frac{1}{N^2 M^2} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sigma_{xy}^2 \right) \sum_{k=0}^{N-1} \sum_{l=0}^{M-1} W_{kl}^2 \exp\left(2\pi i \left[\frac{ak}{N} + \frac{bl}{M} \right]\right) \\
C_{xyab}^{(f)} &= \frac{\sigma_{xy}^2}{NM} \sum_{k=0}^{N-1} \sum_{l=0}^{M-1} W_{kl}^2 \exp\left(2\pi i \left[\frac{ak}{N} + \frac{bl}{M} \right]\right) \\
C_{xy00}^{(f)} &= \left(\frac{1}{NM} \sum_{k=0}^{N-1} \sum_{l=0}^{M-1} W_{kl}^2 \right) \sigma_{xy}^2.
\end{aligned} \tag{B.1.11}$$

So we have found that the effect of Fourier-domain filtering is to modify the variance of the noise across the whole image by a constant factor (in the brackets). If $0 \leq W_{kl} \leq 1$, then this is a reduction in noise variance. It is also clear from this treatment that the filter profile used determines the correlation of the image noise.

Next, consider noise manifesting in the analytic signal. The phase $\tilde{\phi}_{xy}$ and contrast $\tilde{\zeta}_{xy}$, are inferred via the modulus and argument of complex random variable \tilde{a}_{xy} . Since $\Re(\mathcal{E}_{a,xy})$ and $\Im(\mathcal{E}_{a,xy})$ are uncorrelated, \tilde{a}_{xy} has a circular normal distribution in the complex plane. Therefore, to propagate uncertainty from \tilde{a}_{xy} into $\tilde{\phi}_{xy}$ and $\tilde{\zeta}_{xy}$, we must derive the PDF in polar coordinates of a circular bivariate normal distribution in Cartesian coordinates. Analytic expressions have been derived for these distributions[154], which have found use in radio communications and magnetic resonance imaging[155]. The Rice distribution describes the distribution of the radial polar coordinate r (the modulus):

$$P(\tilde{r} | r, \sigma) = \frac{\tilde{r}}{\sigma^2} \exp\left[-\frac{\tilde{r}^2 + r^2}{2\sigma^2}\right] J_0\left(\frac{r\tilde{r}}{\sigma^2}\right), \tag{B.1.12}$$

where J_0 is the modified Bessel function of the first kind with order zero. When the signal to noise ratio is high ($r \gg \sigma$), the Rice distribution approximates a normal distribution, with standard deviation equal to the Cartesian value: $P(\tilde{r} | r, \sigma) \approx \mathcal{N}(r, \sigma^2)$ When SNR is low ($r \sim \sigma$), the Rice distribution is positively skewed. The

corresponding distribution for the angular polar coordinate θ is:

$$P(\tilde{\theta} | \theta, r, \sigma) = \frac{1}{2\pi} \exp\left(\frac{-r^2}{2\sigma^2}\right) \left[1 + \frac{\sqrt{\pi}}{2} Q \exp(Q^2) [1 + \operatorname{erf}(Q)]\right], \quad (\text{B.1.13})$$

where

$$Q = \frac{r \cos(\tilde{\theta} - \theta)}{\sqrt{2}\sigma}, \quad (\text{B.1.14})$$

and $\operatorname{erf}(\dots)$ is the error function. When $\bar{A} \gg \sigma$, then this distribution can also be approximated with a normal distribution:

$$P(\tilde{\theta} | \theta, r, \sigma) \approx \mathcal{N}(\theta, \sigma^2/r^2). \quad (\text{B.1.15})$$

Finally, Figure B.3 plots histograms for brightness, phase and contrast for 3000 independent realisations of the noisy test image in Figure B.1 for three separate image points (marked). These histograms are compared to the ground truth values and also to the predicted PDFs based on the theory outlined in this appendix and with assumed knowledge of the camera's photon transfer curve (which can be characterised experimentally — see Section 5.1 of [63]). There is good agreement between the numerical test and the analytic predictions.

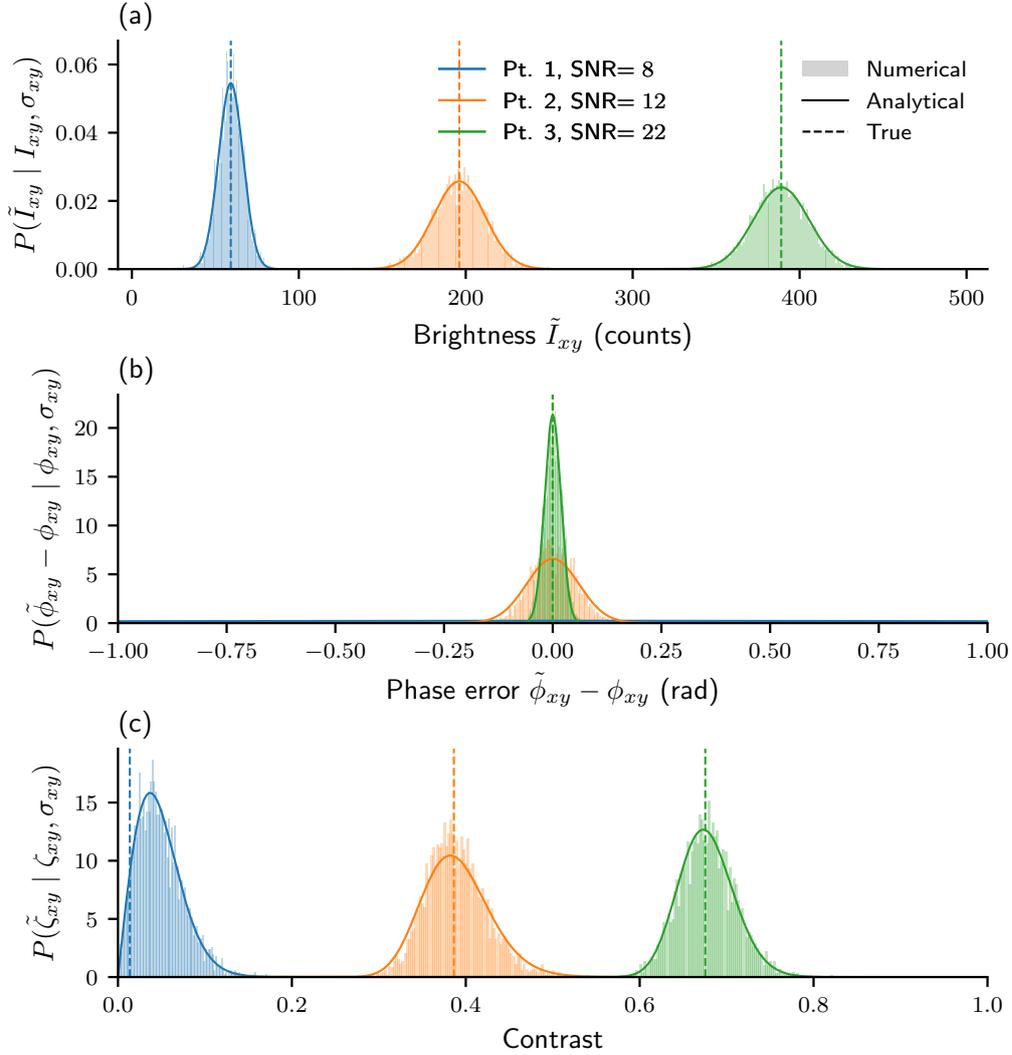


Figure B.3: PDFs for brightness, phase and contrast from analytic prediction (solid lines) and numerical simulation (histogram) for the three points (different colours) marked in Figure B.2. The numerical results are from 3000 independent realisations of the image noise e_{xy} .

s

Appendix C

Bayesian Parameter Estimation

Thanks to increased computational power, Bayesian methods are now widely used in science, and are also finding use in fusion[48, 156–158]. In this work, we use Bayesian methods for model fitting problems in Chapters 4 and 5, and so we very briefly introduce the basics here. See e.g. [137] for more detail.

C.1 Bayes’ Theorem

In the Bayesian picture, the probability of an experimental outcome is interpreted as the ‘degree of belief’ one can reasonably have in that outcome happening. This differs from the frequentist picture, where probability describes the frequency of the outcome, were the experiment to be repeated an infinite number of times from an identical starting point. Consider a model $f(\boldsymbol{\theta})$ that predicts the data that will be measured in an upcoming experiment. This ‘forward model’ is a function of N model parameters $\boldsymbol{\theta} = \{\theta_1, \theta_2, \dots, \theta_N\}$. We denote the joint PDF across all parameters, conditional on any relevant background information I , as $P(\boldsymbol{\theta} | I)$. This is the *prior distribution*. It represents our degree of belief in each possible set of parameter values before any measurements are made. If we then perform the experiment and measure M data points $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_M\}$, we can denote the updated joint PDF as $P(\boldsymbol{\theta} | \mathcal{D}, I)$. This is the *posterior distribution*, representing our degree of belief about the parameter values, updated so as to account for any new information from the experiment. Updating the prior to find the posterior is done using Bayes’ theorem, which can be written¹:

$$P(\boldsymbol{\theta} | \mathcal{D}, I) \propto \mathcal{L}(\mathcal{D} | \boldsymbol{\theta}, I)P(\boldsymbol{\theta} | I). \quad (\text{C.1.1})$$

¹This weak form of the theorem omits a normalising factor that is important for model selection problems. But that’s not important for us here.

Here, $\mathcal{L}(\mathcal{D} \mid \boldsymbol{\theta}, I)$ is called the *likelihood function*. If the experiment has been designed well, the likelihood function will significantly modify the prior distribution.

In general, the likelihood function depends on the measured data, the model $f(\boldsymbol{\theta})$, and on all sources of uncertainty involved in the measurement process. For example, consider a single data point \mathcal{D}_i — how much information does this give us? If we assume that the measurement uncertainty on this data point is additive Gaussian noise with known variance σ_i^2 , then we may write

$$\mathcal{L}(\mathcal{D}_i \mid \boldsymbol{\theta}, I) \propto \exp\left(-\frac{1}{2}\left[\frac{\mathcal{D}_i - f(\boldsymbol{\theta})}{\sigma_i}\right]^2\right). \quad (\text{C.1.2})$$

If all M data points are statistically independent of one another, then the product rule of probability allows us to write the likelihood function for the total dataset as the product of the individual likelihood functions of all data points:

$$\mathcal{L}(\mathcal{D} \mid \boldsymbol{\theta}, I) \propto \prod_{i=1}^M \mathcal{L}(\mathcal{D}_i \mid \boldsymbol{\theta}, I). \quad (\text{C.1.3})$$

By combining Equations C.1.2 and C.1.3 with an appropriate model and prior, the posterior $P(\boldsymbol{\theta} \mid \mathcal{D}, I)$ can be calculated for any possible $\boldsymbol{\theta}$.

C.2 Parameter Estimation

In parameter estimation, the model $f(\boldsymbol{\theta})$ is well understood and the purpose of the experiment is make inferences about the parameter values. The $\boldsymbol{\theta}$ for which $P(\boldsymbol{\theta} \mid \mathcal{D}, I)$ reaches its global maximum (its mode) is the most plausible $\boldsymbol{\theta}$, given all of the available information. This model fitting technique is called maximum a posteriori (MAP) estimation. The definition of the MAP estimate can be written

$$\hat{\boldsymbol{\theta}}_{\text{MAP}} = \arg \max_{\boldsymbol{\theta}} \left[\mathcal{L}(\mathcal{D} \mid \boldsymbol{\theta}, I) P(\boldsymbol{\theta} \mid I) \right]. \quad (\text{C.2.1})$$

A special case occurs if we decide that there is no relevant prior information. In this case all of our information comes from the data and we define the maximum likelihood (ML) estimate as

$$\hat{\boldsymbol{\theta}}_{\text{ML}} = \arg \max_{\boldsymbol{\theta}} \left[\mathcal{L}(\mathcal{D} \mid \boldsymbol{\theta}, I) \right]. \quad (\text{C.2.2})$$

Often, we are more interested in the posterior PDF of a single model parameter θ_1 , than in any of the other $N - 1$ parameters. We therefore want to find the *marginal distribution* $P(\theta_1 \mid \mathcal{D}, I)$. We can find this by integrating the joint posterior PDF

over all of the other model parameters:

$$P(\theta_1 | \mathcal{D}, I) = \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} P(\boldsymbol{\theta} | \mathcal{D}, I) d\theta_2 \dots d\theta_N. \quad (\text{C.2.3})$$

C.3 Markov chain Monte Carlo

Only in very simple cases can Equation C.1.1 be evaluated analytically and so numerical techniques are required. If the model has a small number of parameters ($\lesssim 3$) then a grid-based ‘brute force and ignorance’ approach is often the most practical solution. Here, the posterior is evaluated numerically on a uniform grid in parameter space. Higher dimensional problems require a more efficient sampling of parameter space. Markov chain Monte Carlo (MCMC) is a commonly used sampling technique, in which a random walk is taken through model parameter space. Each proposed step in the walk is randomly generated and is either accepted or rejected using a criterion that favours travel from regions of low probability to high probability[137]. The result is a collection of samples whose density converges to the true posterior PDF. In this work, we use a type of MCMC called Gibbs sampling, which takes its steps a single parameter at a time.

Bibliography

- [1] D. J. Ward et al. ‘The economic viability of fusion power’. In: *Fusion Engineering and Design* 75-79.SUPPL. (2005), pp. 1221–1227. DOI: 10.1016/j.fusengdes.2005.06.160.
- [2] R. Toschi. ‘Nuclear fusion, an energy source’. In: *Fusion Engineering and Design* 36.1 (1997), pp. 1–8. DOI: 10.1016/S0920-3796(97)00007-0.
- [3] Jeffrey P. Freidberg. *Plasma physics and fusion energy*. Vol. 9780521851077. 2007, pp. 1–671. DOI: 10.1017/CB09780511755705.
- [4] J. Ongena and G. Van Oost. ‘Energy for future centuries Will fusion be an inexhaustible, safe and clean energy source?’ In: *Fusion Science and Technology* 41.2 (2002), pp. 3–14. DOI: 10.13182/fst02-a11963498.
- [5] R. A. CAIRNS. *Tokamaks 3rd Edition by John Wesson, Oxford University Press 2004, 749pp, Hardback 0 19 850922, £ 135*. Vol. 71. 3. 2005, pp. 377–377. DOI: 10.1017/s0022377804003058. arXiv: arXiv:1011.1669v3.
- [6] Kenneth S. Krane and William G. Lynch. ‘Introductory Nuclear Physics’. In: *Physics Today* 42.1 (1989), pp. 78–78. DOI: 10.1063/1.2810884.
- [7] M. E. Sawan and M. A. Abdou. ‘Physics and technology conditions for attaining tritium self-sufficiency for the DT fuel cycle’. In: *Fusion Engineering and Design* 81.8-14 PART B (2006), pp. 1131–1144. DOI: 10.1016/j.fusengdes.2005.07.035.
- [8] R. Betti and O. A. Hurricane. *Inertial-Confinement fusion with lasers*. 2016. DOI: 10.1038/NPHYS3736.
- [9] V. A. Rozhanskii. *Impurity Transport in Tokamaks*. Tech. rep. 4. Max Planck IPP, 1980, pp. 465–470. DOI: 10.1016/b978-0-08-021989-9.50026-9.
- [10] F. Wagner et al. ‘Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak’. In: *Physical Review Letters* 49.19 (1982), pp. 1408–1412. DOI: 10.1103/PhysRevLett.49.1408.

- [11] M. Keilhacker. ‘JET deuterium-tritium results and their implications’. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 357.1752 (1999), pp. 415–442. DOI: 10.1098/rsta.1999.0335.
- [12] ITER Org. ‘ITER Research Plan within the Staged Approach’. In: *ITER Research Plan* ITR-18-003 (2018).
- [13] Ronald D. Stambaugh et al. ‘The spherical tokamak path to fusion power’. In: *Fusion Technology* 33.1 (1998), pp. 1–21. DOI: 10.13182/FST33-1.
- [14] Daniel Clery. *A piece of the sun: the quest for fusion energy*. Vol. 51. 07. Gerald Duckworth & Co. Ltd, 2014, pp. 51–3896–51–3896. DOI: 10.5860/choice.51-3896.
- [15] Jason Parisi and Justin Ball. *The Future of Fusion Energy*. 2019. DOI: 10.1142/q0160.
- [16] A. Loarte et al. ‘Chapter 4: Power and particle control’. In: *Nuclear Fusion* 47.6 (2007). DOI: 10.1088/0029-5515/47/6/S04.
- [17] T Donn . *European Research Roadmap to the Realisation of Fusion Energy*. Tech. rep. EUROfusion, 2018. URL: https://www.euro-fusion.org/fileadmin/user%7B%5C_%7Dupload/EUROfusion/Documents/TopLevelRoadmap.pdf.
- [18] G. F. Matthews. ‘Plasma detachment from divertor targets and limiters’. In: *Journal of Nuclear Materials* 220-222 (1995), pp. 104–116. DOI: 10.1016/0022-3115(94)00450-1.
- [19] R. A. Pitts et al. ‘Physics basis for the first ITER tungsten divertor’. In: *Nuclear Materials and Energy* 20 (2019). DOI: 10.1016/j.nme.2019.100696.
- [20] R. A. Pitts et al. ‘Status and physics basis of the ITER divertor’. In: *Physica Scripta T*. Vol. T138. 2009. DOI: 10.1088/0031-8949/2009/T138/014001.
- [21] S. Potzel et al. ‘A new experimental classification of divertor detachment in ASDEX Upgrade’. In: *Nuclear Fusion* 54.1 (2014). DOI: 10.1088/0029-5515/54/1/013001.
- [22] A. Huber et al. ‘Impact of the ITER-like wall on divertor detachment and on the density limit in the JET tokamak’. In: *Journal of Nuclear Materials* 438.SUPPL (2013), S139–S147. DOI: 10.1016/j.jnucmat.2013.01.022. URL: <http://dx.doi.org/10.1016/j.jnucmat.2013.01.022>.
- [23] K. Verhaegh et al. ‘Spectroscopic investigations of divertor detachment in TCV’. In: *Nuclear Materials and Energy* 12 (2017), pp. 1112–1117. DOI: 10.1016/j.nme.2017.01.004. arXiv: 1607.04539.

- [24] J. R. Harrison et al. ‘Characterisation of detached plasmas on the MAST tokamak’. In: *Journal of Nuclear Materials*. Vol. 415. 1 SUPPL. 2011. DOI: 10.1016/j.jnucmat.2010.12.226.
- [25] H. Zohm et al. ‘On the physics guidelines for a tokamak DEMO’. In: *Nuclear Fusion* 53.7 (2013). DOI: 10.1088/0029-5515/53/7/073019.
- [26] A. Kallenbach et al. ‘Impurity seeding for tokamak power exhaust: From present devices via ITER to DEMO’. In: *Plasma Physics and Controlled Fusion* 55.12 (2013). DOI: 10.1088/0741-3335/55/12/124041.
- [27] R. A. Pitts et al. ‘Physics conclusions in support of ITER W divertor monoblock shaping’. In: *Nuclear Materials and Energy* 12 (2017), pp. 60–74. DOI: 10.1016/j.nme.2017.03.005.
- [28] D. Brunner et al. ‘Surface heat flux feedback controlled impurity seeding experiments with Alcator C-Mod’s high-Z vertical target plate divertor: Performance, limitations and implications for fusion power reactors’. In: *Nuclear Fusion* 57.8 (2017). DOI: 10.1088/1741-4326/aa7923.
- [29] G. Fishpool et al. ‘MAST-upgrade divertor facility and assessing performance of long-legged divertors’. In: *Journal of Nuclear Materials* 438.SUPPL (2013), S356–S359. DOI: 10.1016/j.jnucmat.2013.01.067. URL: <http://dx.doi.org/10.1016/j.jnucmat.2013.01.067>.
- [30] E. Havlíčková et al. ‘SOLPS analysis of the MAST-U divertor with the effect of heating power and pumping on the access to detachment in the Super-x configuration’. In: *Plasma Physics and Controlled Fusion* 57.11 (2015). DOI: 10.1088/0741-3335/57/11/115001.
- [31] William Morris et al. ‘MAST accomplishments and upgrade for fusion next-steps’. In: *IEEE Transactions on Plasma Science* 42.3 (2014), pp. 402–414. DOI: 10.1109/TPS.2014.2299973.
- [32] J. R. Harrison et al. ‘Enhancements to MAST upgrade to address the EUROfusion plasma exhaust strategy’. In: *43rd European Physical Society Conference on Plasma Physics, EPS 2016*. 2016.
- [33] I. H. Hutchinson. *Principles of Plasma Diagnostics*. 2002. DOI: 10.1017/cbo9780511613630. arXiv: arXiv:1011.1669v3. URL: <http://ebooks.cambridge.org/ref/id/CB09780511613630>.
- [34] William Morris et al. ‘MAST Upgrade Divertor Facility: A Test Bed for Novel Divertor Solutions’. In: *IEEE Transactions on Plasma Science* 46.5 (2018), pp. 1217–1226. DOI: 10.1109/TPS.2018.2815283.

- [35] J. Lovell et al. ‘A compact, smart Langmuir Probe control module for MAST-Upgrade’. In: *Journal of Instrumentation*. Vol. 12. 11. 2017. DOI: 10.1088/1748-0221/12/11/C11008.
- [36] D. Eldon et al. ‘Initial results of the high resolution edge Thomson scattering upgrade at DIII-D’. In: *Review of Scientific Instruments*. Vol. 83. 10. 2012. DOI: 10.1063/1.4738656.
- [37] R. Scannell et al. ‘A 130 point Nd:YAG Thomson scattering diagnostic on MAST’. In: *Review of Scientific Instruments*. Vol. 81. 10. 2010. DOI: 10.1063/1.3460628.
- [38] M. Kempenaars et al. ‘Enhancement of the JET edge LIDAR Thomson scattering diagnostic with ultrafast detectors’. In: *Review of Scientific Instruments*. Vol. 79. 10. 2008. DOI: 10.1063/1.2969078.
- [39] B. Kurzan and H. D. Murmann. ‘Edge and core Thomson scattering systems and their calibration on the ASDEX Upgrade tokamak’. In: *Review of Scientific Instruments* 82.10 (2011). DOI: 10.1063/1.3643771.
- [40] F. Glass et al. ‘Upgraded divertor Thomson scattering system on DIII-D’. In: *Review of Scientific Instruments* 87.11 (2016). DOI: 10.1063/1.4955281.
- [41] T. N. Carlstrom et al. ‘Initial operation of the divertor Thomson scattering diagnostic on DIII-D’. In: *Review of Scientific Instruments* 68.2 (1997), pp. 1195–1200. DOI: 10.1063/1.1147893.
- [42] J. Hawke et al. ‘Outline of optical design and viewing geometry for divertor Thomson scattering on MAST upgrade’. In: *Journal of Instrumentation* 8.11 (2013). DOI: 10.1088/1748-0221/8/11/C11010.
- [43] Reuben Epstein. ‘Principles of Plasma Spectroscopy, by Hans R. Griem. Cambridge University Press, 1997, 366 pages. ISBN 0 521 45504 9. \$100.00.’ In: *Journal of Plasma Physics* 60.1 (1998), pp. 203–207. DOI: 10.1017/s0022377898219805.
- [44] Erdal Yiğit. *Introduction to Plasma*. Berlin: Springer, 2018, pp. 1–19. DOI: 10.1007/978-3-319-62006-0_1.
- [45] V. A. Soukhanovskii et al. ‘Electron density measurements in the National Spherical Torus Experiment detached divertor region using Stark broadening of deuterium infrared Paschen emission lines’. In: *Review of Scientific Instruments*. Vol. 77. 10. 2006. DOI: 10.1063/1.2336456.

- [46] S. Potzel et al. ‘Electron density determination in the divertor volume of ASDEX Upgrade via Stark broadening of the Balmer lines’. In: *Plasma Physics and Controlled Fusion* 56.2 (2014). DOI: 10.1088/0741-3335/56/2/025010.
- [47] B. A. Lomanowski et al. ‘Inferring divertor plasma properties from hydrogen Balmer and Paschen series spectroscopy in JET-ILW’. In: *Nuclear Fusion* 55.12 (2015). DOI: 10.1088/0029-5515/55/12/123028.
- [48] C. Bowman et al. ‘Development and simulation of multi-diagnostic Bayesian analysis for 2D inference of divertor plasma characteristics’. In: *Plasma Physics and Controlled Fusion* 62.4 (2020). DOI: 10.1088/1361-6587/ab759b.
- [49] T. Farley et al. ‘Filament identification in wide-angle high speed imaging of the mega amp spherical tokamak’. In: *Review of Scientific Instruments* 90.9 (2019). DOI: 10.1063/1.5109470.
- [50] W. A.J. Vijvers et al. ‘Conceptual design and proof-of-principle testing of the real-time multispectral imaging system MANTIS’. In: *Journal of Instrumentation* 12.12 (2017). DOI: 10.1088/1748-0221/12/12/C12058.
- [51] S. A. Silburn et al. ‘Coherence imaging of scrape-off-layer and divertor impurity flows in the Mega Amp Spherical Tokamak (invited)’. In: *Review of Scientific Instruments* 85.11 (2014). DOI: 10.1063/1.4891165.
- [52] C. M. Samuelli et al. ‘2D imaging of helium ion velocity in the DIII-D divertor’. In: *Physics of Plasmas* 25.5 (2018). DOI: 10.1063/1.5017999.
- [53] Joseph W Goodman. *Statistical Optics*. Wiley, 2015. DOI: 10.1017/s0263574701253491.
- [54] W.M Steen. *Principles of Optics M. Born and E. Wolf, 7th (expanded) edition, Cambridge University Press, Cambridge, 1999, 952pp. 37.50/US \$59.95, ISBN 0-521-64222-1*. Vol. 32. 5. 2000, p. 385. DOI: 10.1016/s0030-3992(00)00061-x. arXiv: arXiv:1011.1669v3.
- [55] W H Steel. *Interferometry*. Second. Cambridge University Press, 1987.
- [56] Sumner P. Davis, Mark C. Abrams and James W. Brault. ‘Nonideal (real-world) interferograms’. In: *Fourier Transform Spectrometry*. 2001, pp. 67–80. DOI: 10.1016/b978-012042510-5/50005-6.
- [57] Pierre Connes. ‘Early history of fourier transform spectroscopy’. In: *Infrared Physics* 24.2-3 (1984), pp. 69–93. DOI: 10.1016/0020-0891(84)90052-6.

- [58] A. A. Michelson. ‘Application of Interference Methods to Spectroscopic Measurement’. In: *Publications of the Astronomical Society of the Pacific* 4 (1892), p. 190. DOI: 10.1086/120489.
- [59] Sumner P. Davis, Mark C. Abrams and James W. Brault. ‘Introduction’. In: *Fourier Transform Spectrometry*. 2001, pp. 1–16. DOI: 10.1016/b978-012042510-5/50001-9.
- [60] R. L. Hilliard and G. G. Shepherd. ‘Wide-Angle Michelson Interferometer for Measuring Doppler Line Widths*’. In: *Journal of the Optical Society of America* 56.3 (1966), p. 362. DOI: 10.1364/josa.56.000362.
- [61] Christoph R. Englert, David D. Babcock and John M. Harlander. ‘Doppler asymmetric spatial heterodyne spectroscopy (DASH): Concept and experimental demonstration’. In: *Applied Optics* 46.29 (2007), pp. 7297–7307. DOI: 10.1364/AO.46.007297.
- [62] G. G. Shepherd et al. *The wind imaging interferometer (WINDII) on the upper atmosphere research satellite: A 20 year perspective*. 2012. DOI: 10.1029/2012RG000390.
- [63] Scott Alan Silburn. *A Doppler Coherence Imaging Diagnostic for the Mega-Amp Spherical Tokamak (PhD thesis)*. Durham University, 2014.
- [64] Clive Alvin Michael. *Doppler spectroscopy of argon plasmas in H-1NF using a coherence imaging camera (PhD thesis)*. Australian National University, 2003. URL: <https://openresearch-repository.anu.edu.au/handle/1885/151238>.
- [65] A. Thetford. *Polarization Interferometers*. Vol. 19. 8. Wiley-Interscience, 1972, pp. 716–716. DOI: 10.1080/716099328.
- [66] Eugene Hecht. *Hecht’s Response*. Vol. 48. 1. 2010, pp. 5–6. DOI: 10.1119/1.3274347.
- [67] William Shurcliff. ‘Polarized Light, Production and Use By W. A. Shurcliff’. In: *Oyobuturi* 33.2 (1964), 140a–141. DOI: 10.11470/oubutsu1932.33.140a.
- [68] Russell A. Chipman, Wai-Sze Tiffany Lam and Garam Young. *Polarized Light and Optical Systems*. 2018. DOI: 10.1201/9781351129121.
- [69] Jinil Chung. ‘Time resolved 2D Doppler imaging of ion dynamics’. PhD thesis. University of Greifswald, 2004.
- [70] Alexander Thorman. *Polarisation Coherence Imaging of Electric and Magnetic Fields in Plasmas (PhD thesis)*. Australian National University, 2018. URL: <https://openresearch-repository.anu.edu.au/handle/1885/148724>.

- [71] Russell A Chipman. ‘Chapter 22: Polarimetry’. In: *Handbook of Optics. Vol1.* 2009, pp. 22.1–22.37. DOI: 10.1117/1.JBO.19.7.
- [72] Maximino Avendaño-Alejo and Orestes N. Stavroudis. ‘Huygens’s principle and rays in uniaxial anisotropic media II Crystal axis orientation arbitrary’. In: *Journal of the Optical Society of America A* 19.8 (2002), p. 1674. DOI: 10.1364/josaa.19.001674.
- [73] Maximino Avendaño-Alejo, Orestes N. Stavroudis and Ana Rosa Boyain y Goitia. ‘Huygens’s principle and rays in uniaxial anisotropic media I Crystal axis normal to refracting surface’. In: *Journal of the Optical Society of America A* 19.8 (2002), p. 1668. DOI: 10.1364/josaa.19.001668.
- [74] María C. Simon and Karin V. Gottschalk. ‘Waves and rays in uniaxial birefringent crystals’. In: *Optik* 118.10 (2007), pp. 457–470. DOI: 10.1016/j.ijleo.2006.03.032.
- [75] Francisco E. Veiras, Liliana I. Perez and María T. Garea. ‘Phase shift formulas in uniaxial media: An application to waveplates’. In: *Applied Optics* 49.15 (2010), pp. 2769–2777. DOI: 10.1364/AO.49.002769.
- [76] I. N. Bronshtein et al. *Handbook of mathematics, sixth edition.* 2015, pp. 1–1211. DOI: 10.1007/978-3-662-46221-8.
- [77] Dorothea Gradic et al. ‘A new calibration implementation for Doppler Coherence Imaging Spectroscopy’. In: *Fusion Engineering and Design* December 2018 (2019), pp. 1–4. DOI: 10.1016/j.fusengdes.2019.01.138. URL: <https://doi.org/10.1016/j.fusengdes.2019.01.138>.
- [78] D. Gradic et al. ‘Doppler coherence imaging of divertor and SOL flows in ASDEX upgrade and Wendelstein 7-X’. In: *Plasma Physics and Controlled Fusion* 60.8 (2018). DOI: 10.1088/1361-6587/aac4d2.
- [79] Lei Wu, Chunmin Zhang and Baochang Zhao. ‘Analysis of the lateral displacement and optical path difference in wide-field-of-view polarization interference imaging spectrometer’. In: *Optics Communications* 273.1 (2007), pp. 67–73. DOI: 10.1016/j.optcom.2006.12.034.
- [80] John Howard. ‘Electro-optically modulated polarizing Fourier-transform spectrometer for plasma spectroscopy applications’. In: *Applied Optics* 41.1 (2002), p. 197. DOI: 10.1364/ao.41.000197. URL: <http://www.opticsinfobase.org/abstract.cfm?%7B%5C%7Ddid=67208>.
- [81] Jinil Chung et al. ‘Time resolved coherence-imaging spectrometer on WEGA stellarator’. In: *Plasma Physics and Controlled Fusion* 47.6 (2005), pp. 919–940. DOI: 10.1088/0741-3335/47/6/012.

- [82] John HOWARD et al. ‘Spatial Heterodyne Spectro-Polarimetry Systems for Imaging Key Plasma Parameters in Fusion Devices’. In: *Plasma and Fusion Research* 5.January (2010), S1010–S1010. DOI: 10.1585/pfr.5.s1010. URL: <http://joi.jlrc.jst.go.jp/JST.JSTAGE/pfr/5.S1010?from=CrossRef>.
- [83] W. H. Steel. ‘A Demonstration Fourierscope’. In: *Applied Optics* 9.7 (1970), p. 1721. DOI: 10.1364/ao.9.001721.
- [84] Chunmin Zhang et al. ‘A static polarization imaging spectrometer based on a Savart polariscope’. In: *Optics Communications* 203.1-2 (2002), pp. 21–26. DOI: 10.1016/S0030-4018(01)01726-6.
- [85] Tobin R. Weber et al. ‘Impurity flow measurements at DIII-D using a coherence imaging spectrometer’. In: *Journal of Nuclear Materials* 438.SUPPL (2013), S1257–S1260. DOI: 10.1016/j.jnucmat.2013.01.279. URL: <http://dx.doi.org/10.1016/j.jnucmat.2013.01.279>.
- [86] J. Howard et al. ‘Coherence Imaging of Flows in the DIII-D Divertor’. In: *Contributions to Plasma Physics* 51.2-3 (Mar. 2011), pp. 194–200. DOI: 10.1002/ctpp.201000062. URL: <http://doi.wiley.com/10.1002/ctpp.201000062>.
- [87] T. R. Weber, S. L. Allen and J. Howard. ‘C-III flow measurements with a coherence imaging spectrometer’. In: *Review of Scientific Instruments* 83.10 (2012), pp. 1–4. DOI: 10.1063/1.4728311.
- [88] C. M. Samuell et al. ‘Verification of doppler coherence imaging for 2D ion velocity measurements on DIII-D’. In: *Review of Scientific Instruments* 89.9 (2018). DOI: 10.1063/1.5039367.
- [89] I. Waters et al. ‘Field aligned flows driven by neutral puffing at MAST’. In: *Nuclear Fusion* 58.6 (2018). DOI: 10.1088/1741-4326/aab7ad.
- [90] V. Perseo et al. ‘Direct measurements of counter-streaming flows in a low-shear stellarator magnetic island topology’. In: *Nuclear Fusion* 59.12 (2019). DOI: 10.1088/1741-4326/ab4320.
- [91] C. M. Samuell et al. ‘Absolute calibration of Doppler coherence imaging velocity images’. In: *Journal of Instrumentation* 12.8 (2017). DOI: 10.1088/1748-0221/12/08/C08016.
- [92] S. L. Allen et al. ‘Laser calibration of the DIII-D coherence imaging system’. In: *Review of Scientific Instruments* 89.10 (2018). DOI: 10.1063/1.5038739.
- [93] J. Howard et al. ‘Imaging charge exchange recombination spectroscopy on the TEXTOR tokamak’. In: *Plasma Physics and Controlled Fusion* 52.12 (2010). DOI: 10.1088/0741-3335/52/12/125002.

- [94] T. HATAE et al. ‘Development of Polarization Interferometer Based on Fourier Transform Spectroscopy for Thomson Scattering Diagnostics’. In: *Plasma and Fusion Research* 2 (2007), S1026–S1026. DOI: 10.1585/pfr.2.s1026.
- [95] J. Howard. ‘Application of polarization interferometers for Thomson scattering’. In: *Plasma Physics and Controlled Fusion* 48.6 (2006), pp. 777–787. DOI: 10.1088/0741-3335/48/6/005.
- [96] J. Howard and T. Hatae. ‘Imaging interferometers for analysis of Thomson scattered spectra’. In: *Review of Scientific Instruments*. Vol. 79. 10. 2008. DOI: 10.1063/1.2969421.
- [97] O. Lischtschenko et al. ‘Density measurements using coherence imaging spectroscopy based on Stark broadening’. In: *Review of Scientific Instruments* 81.10 (2010), pp. 1–5. DOI: 10.1063/1.3490023.
- [98] John Howard. ‘Modulated optical solid-state spectrometer applications in plasma diagnostics’. In: *Review of Scientific Instruments* 70.1 II (1999), pp. 368–371. DOI: 10.1063/1.1149468.
- [99] J. Howard et al. ‘Stokes-Doppler coherence imaging for ITER boundary tomography’. In: *Review of Scientific Instruments* 87.11 (2016). DOI: 10.1063/1.4963712.
- [100] Tomohiro Yamazaki et al. ‘Four-directional pixel-wise polarization CMOS image sensor using air-gap wire grid on 2.5- μm back-illuminated pixels’. In: *Technical Digest - International Electron Devices Meeting, IEDM*. 2017, pp. 8.7.1–8.7.4. DOI: 10.1109/IEDM.2016.7838378.
- [101] Viktor Gruev and Rob Perkins. ‘A 1 MPixel CCD image sensor with aluminum nanowire polarization filter’. In: *ISCAS 2010 - 2010 IEEE International Symposium on Circuits and Systems: Nano-Bio Circuit Fabrics and Systems*. 2010, pp. 629–632. DOI: 10.1109/ISCAS.2010.5537513.
- [102] James E. Millerd et al. ‘Pixelated phase-mask dynamic interferometer’. In: *Interferometry XII: Techniques and Analysis*. Vol. 5531. 2004, p. 304. DOI: 10.1117/12.560807.
- [103] Matt Novak et al. ‘Analysis of a micropolarizer array-based simultaneous phase-shifting interferometer’. In: *Applied Optics* 44.32 (2005), pp. 6861–6868. DOI: 10.1364/AO.44.006861.

- [104] Brad Kimbrough and James Millerd. ‘The spatial frequency response and resolution limitations of pixelated mask spatial carrier based phase shifting interferometry’. In: *Interferometry XV: Techniques and Analysis* 7790. August 2010 (2010), 77900K. DOI: 10.1117/12.860751.
- [105] Katherine Creath and Goldie Goldstein. ‘Dynamic quantitative phase imaging for biological objects using a pixelated phase mask’. In: *Biomedical Optics Express* 3.11 (2012), p. 2866. DOI: 10.1364/boe.3.002866.
- [106] Tatsuki Tahara et al. ‘High-speed three-dimensional microscope for dynamically moving biological objects based on parallel phase-shifting digital holographic microscopy’. In: *IEEE Journal on Selected Topics in Quantum Electronics* 18.4 (2012), pp. 1387–1393. DOI: 10.1109/JSTQE.2011.2178062.
- [107] Bradley T. Kimbrough. ‘Pixelated mask spatial carrier phase shifting interferometry algorithms and associated errors’. In: *Applied Optics* 45.19 (2006), pp. 4554–4562. DOI: 10.1364/AO.45.004554.
- [108] M. Servin and J. C. Estrada. ‘Error-free demodulation of pixelated carrier frequency interferograms’. In: *Optics Express* 18.17 (2010), p. 18492. DOI: 10.1364/oe.18.018492.
- [109] M. Servin, J. C. Estrada and O. Medina. ‘Fourier transform demodulation of pixelated phase-masked interferograms’. In: *Optics Express* 18.15 (2010), p. 16090. DOI: 10.1364/oe.18.016090.
- [110] J. J. MICHEL. ‘Mesure De La Visibilité D’Une Raie De Plasma’. In: *Le Journal de Physique Colloques* 28.C2 (1967), pp. C2–109–C2–112. DOI: 10.1051/jphyscol:1967222.
- [111] Peter Urlings. *Multiple Delay Coherence Imaging Charge Exchange Recombination Spectroscopy (MSc thesis)*. Eindhoven University of Technology, 2015. URL: <https://pure.tue.nl/ws/portalfiles/portal/46924582/841497-1.pdf>.
- [112] Manuel Servin, J. Antonio Quiroga and Moises Padilla. *Fringe Pattern Analysis for Optical Metrology: Theory, Algorithms, and Applications*. Vol. 9783527411528. 2014, pp. 1–327. DOI: 10.1002/9783527681075.
- [113] Moises Padilla. ‘Synchronous Demodulation of Pixelated Phase-Shifted Interferograms for Dynamic Interferometry’. PhD thesis. Centro de Investigaciones en Optica, 2014, p. 63. URL: <https://cio.repositorioinstitucional.mx/jspui/bitstream/1002/342/1/15988.pdf>.

- [114] Kanti V. Mardia and Peter E. Jupp. *Directional Statistics*. 2008, pp. 1–432. DOI: 10.1002/9780470316979.
- [115] A. Kramida et al. *NIST Atomic Spectra Database (ver. 5.6.1)*, [Online]. Available: <http://physics.nist.gov/asd> [11-Aug-2014]. National Institute of Standards and Technology, Gaithersburg, MD. 2019. DOI: 10.18434/T4W30F. URL: <http://physics.nist.gov/asd%7B%5C%%7D0Ahttps://doi.org/10.18434/T4W30F%7B%5C%%7D0Ahttp://physics.nist.gov/asd%7B%5C%%7D0Ahttps://doi.org/10.18434/T4W30F>.
- [116] Kazuyoshi Itoh. ‘Analysis of the phase unwrapping algorithm’. In: *Applied Optics* 21.14 (1982), p. 2470. DOI: 10.1364/ao.21.002470.
- [117] Zundu Luo and Yidong Huang. ‘Laser and Physical Properties of Materials’. In: *Springer Series in Materials Science*. Second. Vol. 289. McGraw-Hill, 2020, pp. 309–343. DOI: 10.1007/978-981-32-9668-8_10.
- [118] D. Eimerl et al. ‘Optical, mechanical, and thermal properties of barium borate’. In: *Journal of Applied Physics* 62.5 (1987), pp. 1968–1983. DOI: 10.1063/1.339536.
- [119] K. Kato. ‘Second-Harmonic Generation to 2048 Å in β -BaB₂O₄’. In: *IEEE Journal of Quantum Electronics* 22.7 (1986), pp. 1013–1014. DOI: 10.1109/JQE.1986.1073097. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1073097>.
- [120] K. Kato, N. Umemura and T. Mikami. ‘Sellmeier and thermo-optic dispersion formulas for β -BaB₂O₄ (revisited)’. In: *Nonlinear Frequency Generation and Conversion: Materials, Devices, and Applications IX* 7582 (2010), p. 75821L. DOI: 10.1117/12.841913. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.841913>.
- [121] Seong-yeol Kim et al. ‘Longitudinal Electron Bunch Shaping Experiments At’. In: *Proceedings of FEL2015, Daejeon, Korea* 5 (2015), pp. 694–697.
- [122] Michael Andrew Dumbson John Robert Taylor, Chris D. Zafirators. *Chapter 03: Atoms*. 2002, pp. 85–124. DOI: 10.1887/0750308397. URL: <http://stacks.iop.org/0750308397>.
- [123] H. R. Griem and Gillian Peach. *Spectral Line Broadening by Plasmas*. Vol. 28. 2. Academic Press, 1975, pp. 61–62. DOI: 10.1063/1.3068826. URL: <http://physicstoday.scitation.org/doi/10.1063/1.3068826>.

- [124] C. Stehlé and R. Hutcheon. ‘Extensive tabulations of Stark broadened hydrogen line profiles’. In: *Astronomy and Astrophysics Supplement Series* 140.1 (1999), pp. 93–97. DOI: 10.1051/aas:1999118.
- [125] Marco Antonio Gigoso. *Stark broadening models for plasma diagnostics*. 2014. DOI: 10.1088/0022-3727/47/34/343001.
- [126] Felix Paulus et al. *Comparative electron density measurements on MAGPIE using a high-resolution spectrometer and a coherence imaging system (MSc thesis)*. December. applied Science and Art Hildesheim/Holzminden/Göttingen, 2017.
- [127] A. Brissaud and U. Frisch. ‘Theory of Stark broadening-II exact line profile with model microfield’. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 11.12 (1971), pp. 1767–1783. DOI: 10.1016/0022-4073(71)90021-5.
- [128] J. L. Terry et al. ‘Experimental determination of the volume recombination rate in tokamak divertors’. In: *Journal of Nuclear Materials* 266 (1999), pp. 30–36. DOI: 10.1016/S0022-3115(98)00812-5.
- [129] J. Rosato et al. ‘Stark broadening of Balmer lines with low and moderate quantum number in dense divertor plasmas’. In: *Contributions to Plasma Physics* 58.6-8 (2018), pp. 578–582. DOI: 10.1002/ctpp.201700100.
- [130] J. Rosato, Y. Marandet and R. Stamm. *A new table of Balmer line shapes for the diagnostic of magnetic fusion plasmas*. 2017. DOI: 10.1016/j.jqsrt.2016.10.005.
- [131] H. Kuramoto et al. ‘Measurement of the poloidal magnetic field with high time resolution Zeeman polarimeter in the JIPP T-IIU tokamak’. In: *Fusion Engineering and Design* 34-35 (1997), pp. 285–288. DOI: 10.1016/S0920-3796(96)00537-6.
- [132] D. M. Thomas. ‘Poloidal magnetic field measurements and analysis with the DIII-D LIBEAM system (invited)’. In: *Review of Scientific Instruments*. Vol. 74. 3 II. 2003, pp. 1541–1546. DOI: 10.1063/1.1526928.
- [133] R. Lester et al. ‘Coherence imaging for ion temperature and flow measurements in a low-temperature helicon plasma source’. In: *Plasma Sources Science and Technology* 25.1 (2016). DOI: 10.1088/0963-0252/25/1/015025.
- [134] Sascha Trippe. ‘Polarization and polarimetry: A review’. In: *Journal of the Korean Astronomical Society* 47.1 (2014), pp. 15–39. DOI: 10.5303/JKAS.2014.47.1.15. arXiv: 1401.1911.

- [135] H. P. Summers. *The ADAS User Manual, version 2.6*. Tech. rep. March. 2004. URL: <http://www.adas.ac.uk>.
- [136] A. Meigs et al. ‘Enhancement of JET’s mirror-link near-ultraviolet to near-infrared divertor spectroscopy system’. In: *Review of Scientific Instruments*. Vol. 81. 10. 2010. DOI: 10.1063/1.3502322.
- [137] G. Winokur and P. Clayton. *Family history studies. I. Two types of affective disorders separated according to genetic and clinical factors*. Vol. 9. 1966, pp. 35–50.
- [138] John Scholten et al. ‘Operational status of the Magnum-PSI linear plasma device’. In: *Fusion Engineering and Design*. Vol. 88. 9-10. 2013, pp. 1785–1788. DOI: 10.1016/j.fusengdes.2013.05.063.
- [139] G. De Temmerman et al. ‘High heat flux capabilities of the Magnum-PSI linear plasma device’. In: *Fusion Engineering and Design*. Vol. 88. 6-8. 2013, pp. 483–487. DOI: 10.1016/j.fusengdes.2013.05.047.
- [140] H. J.N. Van Eck. ‘The linear plasma generator Magnum-PSI’. PhD thesis. Technische Universiteit Eindhoven, 2013, p. 147. DOI: 10.6100/IR749136. URL: <https://pure.tue.nl/ws/files/3405151/749136.pdf>.
- [141] G. G. Van Eden et al. ‘The effect of high-flux H plasma exposure with simultaneous transient heat loads on tungsten surface damage and power handling’. In: *Nuclear Fusion* 54.12 (2014). DOI: 10.1088/0029-5515/54/12/123010.
- [142] T. W. Morgan et al. ‘Liquid metals as a divertor plasma-facing material explored using the Pilot-PSI and Magnum-PSI linear devices’. In: *Plasma Physics and Controlled Fusion* 60.1 (2018). DOI: 10.1088/1361-6587/aa86cd.
- [143] G. F. Matthews et al. ‘Testing of a high temperature radiatively cooled Li/Ta heat pipe in Magnum-PSI’. In: *Fusion Engineering and Design* (2019). DOI: 10.1016/j.fusengdes.2018.12.096.
- [144] W. A.J. Vijvers et al. ‘Multiple discharge channels in a cascaded arc to produce large diameter plasma beams’. In: *Fusion Engineering and Design* 84.7-11 (2009), pp. 1933–1936. DOI: 10.1016/j.fusengdes.2008.12.102.
- [145] Hans J.N. Van Eck et al. ‘A 2.5-T, 1.25-m Free Bore Superconducting Magnet for the Magnum-PSI Linear Plasma Generator’. In: *IEEE Transactions on Applied Superconductivity* 28.3 (2018). DOI: 10.1109/TASC.2017.2779510.

- [146] H. J. Van Der Meiden et al. ‘Advanced Thomson scattering system for high-flux linear plasma generator’. In: *Review of Scientific Instruments* 83.12 (2012). DOI: 10.1063/1.4768527.
- [147] A. E. Shumack et al. ‘Rotation of a strongly magnetized hydrogen plasma column determined from an asymmetric Balmer- β spectral line with two radiating distributions’. In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* (2008). DOI: 10.1103/PhysRevE.78.046405.
- [148] Daniel D. Hickstein et al. ‘A direct comparison of high-speed methods for the numerical Abel transform’. In: *Review of Scientific Instruments* 90.6 (2019). DOI: 10.1063/1.5092635. arXiv: 1902.09007.
- [149] K. Ješko et al. ‘Studying divertor relevant plasmas in the Pilot-PSI linear plasma device: Experiments versus modelling’. In: *Plasma Physics and Controlled Fusion* 60.12 (2018). DOI: 10.1088/1361-6587/aae80d.
- [150] Mitsuo Takeda, Hideki Ina and Seiji Kobayashi. ‘Fourier-Transform Method of Fringe-Pattern Analysis for Computer-Based Topography and Interferometry.’ In: *Journal of the Optical Society of America* 72.1 (1982), pp. 156–160. DOI: 10.1364/JOSA.72.000156. URL: <https://www.osapublishing.org/abstract.cfm?URI=josa-72-1-156>.
- [151] Donald J. Bone, H.-A. Bachor and R. John Sandeman. ‘Fringe-pattern analysis using a 2-D Fourier transform’. In: *Applied Optics* 25.10 (1986), p. 1653. DOI: 10.1364/ao.25.001653. URL: <http://www.ncbi.nlm.nih.gov/pubmed/18231393>.
- [152] J. Schoukens and J. Renneboog. ‘Modeling the Noise Influence on the Fourier Coefficients After a Discrete Fourier Transform.’ In: *IEEE Transactions on Instrumentation and Measurement* IM-35.3 (1986), pp. 278–286. DOI: 10.1109/TIM.1986.6499210.
- [153] George R Cooper and Clare D Mcgillem. ‘Probabilistic Methods of Signal and System Analysis’. In: (1988), p. 496.
- [154] Roger S. Holmes, Elizabeth M. Algar and Peter B. Mather. *Biochemical and genetic studies on mouse aldehyde dehydrogenases*. Vol. 2. 1. 1985, pp. 67–71. DOI: 10.1016/0741-8329(85)90018-7. URL: <http://portal.acm.org/citation.cfm?id=541365>.
- [155] HáKon Gudbjartsson and Samuel Patz. ‘The Rician distribution of noisy MRI data’. In: *Magnetic Resonance in Medicine* 34.6 (1995), pp. 910–914. DOI: 10.1002/mrm.1910340618. arXiv: NIHMS150003.

- [156] R. Fischer, A. Dinklage and E. Pasch. ‘Bayesian modelling of fusion diagnostics’. In: *Plasma Physics and Controlled Fusion* 45.7 (2003), pp. 1095–1111. DOI: 10.1088/0741-3335/45/7/304.
- [157] J. Svensson and A. Werner. ‘Current tomography for axisymmetric plasmas’. In: *Plasma Physics and Controlled Fusion* 50.8 (2008). DOI: 10.1088/0741-3335/50/8/085002.
- [158] G. T. Von Nessi et al. ‘Evidence cross-validation and Bayesian inference of MAST plasma equilibria’. In: *Physics of Plasmas* 19.1 (2012). DOI: 10.1063/1.3677362.
- [159] J. L. Terry and M. L. Reinke. ‘Diagnostic tools for studying divertor detachment: Bolometry, spectroscopy, and thermography for surface heat-flux’. In: *Plasma Physics and Controlled Fusion* 59.4 (2017). DOI: 10.1088/1361-6587/aa5997.
- [160] S. L. Allen et al. ‘Laser calibration of the DIII-D coherence imaging system’. In: *Review of Scientific Instruments* 89.10 (2018). DOI: 10.1063/1.5038739.
- [161] John I.B. Wilson. ‘Tokamaks, 4th edn., by John Wesson’. In: *Contemporary Physics* 53.5 (2012), pp. 450–451. DOI: 10.1080/00107514.2012.720285.
- [162] E A Azizov. ‘Tokamaks: from A D Sakharov to the present (the 60-year history of tokamaks)’. In: *Physics-Uspekhi* 55.2 (2012), pp. 190–203. DOI: 10.3367/ufne.0182.201202j.0202.
- [163] Jinil Chung et al. ‘Time resolved coherence-imaging spectrometer on WEGA stellarator’. In: *Plasma Physics and Controlled Fusion* 47.6 (2005), pp. 919–940. DOI: 10.1088/0741-3335/47/6/012.
- [164] Sergey N. Savenkov. ‘Jones and Mueller matrices: structure, symmetry relations and information content’. In: *Light Scattering Reviews* 4. 2009, pp. 71–119. DOI: 10.1007/978-3-540-74276-0_3.
- [165] Tatsuki Tahara et al. ‘Parallel phase-shifting color digital holographic microscopy’. In: *3D Research* 1.4 (2010), pp. 1–6. DOI: 10.1007/3DRes.04(2010)05.
- [166] R. Prentice et al. ‘A two color mm-wave interferometer for the JET divertor’. In: *Review of Scientific Instruments*. Vol. 66. 2. 1995, pp. 1154–1158. DOI: 10.1063/1.1145996.
- [167] B. Lipschultz et al. ‘Variation of the divertor geometry in Alcator C-Mod’. In: *16th IAEA Fusion Energy Conference*. 1996, p. 18.

- [168] Joe Milnes et al. ‘MAST upgrade - Construction status’. In: *Fusion Engineering and Design*. Vol. 96-97. 2015, pp. 42–47. DOI: 10.1016/j.fusengdes.2015.03.002.
- [169] M. J. van de Pol et al. ‘Operational characteristics of the superconducting high flux plasma generator Magnum-PSI’. In: *Fusion Engineering and Design* 136 (2018), pp. 597–601. DOI: 10.1016/j.fusengdes.2018.03.033.
- [170] Francisco L. Tabarés. ‘Present status of liquid metal research for a fusion reactor’. In: *Plasma Physics and Controlled Fusion* 58.1 (2015). DOI: 10.1088/0741-3335/58/1/014014.
- [171] E. E. Mukhin et al. ‘Physical aspects of divertor Thomson scattering implementation on ITER’. In: *Nuclear Fusion* 54.4 (2014). DOI: 10.1088/0029-5515/54/4/043007.
- [172] D. D. Ryutov. ‘Geometrical properties of a "snowflake" divertor’. In: *Physics of Plasmas* 14.6 (2007). DOI: 10.1063/1.2738399.
- [173] P. M. Valanju et al. ‘Super-X divertors and high power density fusion devices’. In: *Physics of Plasmas* 16.5 (2009), pp. –7. DOI: 10.1063/1.3110984.
- [174] J. M. Padilla, M. Servin and J. C. Estrada. ‘Synchronous phase-demodulation and harmonic rejection of 9-step pixelated dynamic interferograms’. In: *Optics Express* 20.11 (2012), p. 11734. DOI: 10.1364/oe.20.011734.
- [175] Kevin Verhaegh. *Spectroscopic investigations of detachment on TCV (PhD thesis)*. September. University of York, 2018.
- [176] R. Neu et al. ‘Tungsten: An option for divertor and main chamber plasma facing components in future fusion devices’. In: *Nuclear Fusion* 45.3 (2005), pp. 209–218. DOI: 10.1088/0029-5515/45/3/007.