A Doppler Coherence Imaging Diagnostic for the Mega-Amp Spherical Tokamak

SCOTT ALAN SILBURN

SILBURN, SCOTT ALAN (2014) A Doppler Coherence Imaging Diagnostic for the Mega-Amp Spherical Tokamak. Doctoral thesis, Durham University.

A Doppler Coherence Imaging Diagnostic for the Mega-Amp Spherical Tokamak

Scott Alan Silburn

A thesis presented for the degree of
Doctor of Philosophy

Centre for Advanced Instrumentation
Department of Physics
University of Durham
UK

September 2014

A Doppler Coherence Imaging Diagnostic for the Mega-Amp Spherical Tokamak

Scott Alan Silburn

Submitted for the degree of Doctor of Philosophy
September 2014

Abstract

Developing a plasma exhaust solution suitable for future high power tokamaks is one of the major challenges facing the development of magnetic confinement fusion as a terrestrial energy source. In order to improve our understanding of the relevant physics, high quality experimental measurements of plasma dynamics in the scrape-off-layer (SOL) and divertor plasma regions are required. This thesis is concerned with the development of diagnostic instrumentation for measuring exhaust plasma flow: an important phenomenon with implications for the control of exhaust particles and heat as well as unwanted impurities.

Coherence imaging spectroscopy (CIS) is a relatively new diagnostic technique which can be used to obtain time resolved 2D imaging of flows using the Doppler shifts of visible ion emission lines. The technique makes use of an imaging polarization interferometer and is based on the concept of Fourier transform spectroscopy. The principle advantages of this over other flow measurement techniques are the very large amount of spatial information collected, and the simple relationship between the measured quantities and spatially varying flows in the plasma.

This thesis presents the development of, and first results from, a CIS ion flow diagnostic for the UK's Mega Amp Spherical Tokamak (MAST). The diagnostic can image flows of intrinsic C^{+}, C^{2+} and He^{+} impurity ions over fields of view between 10° - 45°, at frame rates between 50Hz - 1kHz and with flow resolution typically around 1km/s (compared with measured flows of typically 5 - 30km/s). Spatial resolution is better than 4.5 cm over a 1.4 x 1.4m area of the plasma cross-section.

After reviewing the principles and theory of the coherence imaging technique, the design of a coherence imaging flow diagnostic for MAST is presented in detail. Results of careful laboratory characterization and calibration of the instrument are presented, and the instrument performance is compared to the design calculations.

The diagnostic was used successfully for flow measurements on MAST during an experimental campaign in May - September 2013. On-plasma validation of the instrument performance is presented, as well as examples of novel flow observations made with the diagnostic. These include field-aligned flow structures associated with high field side gas fuelling of the plasma, and the first measurements of spatial flow structure in the divertor associated with the application of resonant magnetic perturbations (RMPs). Possible future improvements to the instrument design and extensions of the present work are suggested.

Contents

List of Figures

List of Tables

Declaration

The work in this thesis is based on research carried out at the Centre for Advanced Instrumentation in the Department of Physics, University of Durham, UK, and experiments performed at Culham Centre for Fusion Energy, UK. No part of this thesis has been submitted elsewhere for any other degree or qualification. The work within this thesis is all that of the author, except where explicitly stated otherwise in the text or captions.

Publications

Elements of this work have been published in:

Silburn, S. A., Harrison, J. R., Howard, J., Gibson, K. J., Meyer, H., Michael, C. A., and Sharples, R. M.
2014. Coherence imaging of scrape-off-layer and divertor impurity flows in the Mega Amp Spherical Tokamak (invited). *Review of Scientific Instruments*, 85(11), 11D703.

Acknowledgements

I would like to thank my supervisor, Prof. Ray Sharples, for giving me the opportunity to carry out this work; for his insightful perspective on many difficult problems, and for his didctaed guidance and support throughout this project. I would also like to thank my supervisors at CCFE: Dr Clive Michael for his pivotal role in defining the direction of this project and getting it off the ground, his technical input and sparking many valuable ideas, and Dr Hendrik Meyer for his support including experiment planning, leading MAST experiments, and support with the ECELESTE diagnostic.

It has been a great pleasure to collaborate with Prof. John Howard at Australian National university during the course of this work, and I would like to thank him for lending his vast technical knowledge and experience, for his part in construction of the instrument and for his general enthusiasm and encouragement. I am also particularly grateful to Dr. James Harrison at CCFE, whose untiring enthusiasm for this project and consistent scientific, technical, logistical and moral support was truly invaluable. I would also like to thank Dr. Martin Dunstan and Ian Fitzgerald at CCFE for their help in making the instrument come together, and the entire MAST team, with whom it was a great pleasure to work.

I would like to thank Paul Clark for his expert technical assistance in the lab at NetPark, and also the Durham Atomic and Molecular physics research group's 'team strontium' for the use of their tuneable laser source. I am grateful to Dr John Bradford at RAL Space for the loan of the Photron APX-RS camera used in this work, and to Prof. Kieran Gibson of York Plasma Institute for his support with hardware, use of the YPI experimental facilities and general support of this project.

I am very grateful to my colleagues and friends in CfAI for many enlightening discussions, both on and off the topic of physics; for many fryups, and for being a fantastic group of friends who've made working in CfAI an absolute pleasure. I would also like to thank my fellow DTN students and the members of York Plasma Institute for their encouragement, friendship and support.

I am very lucky to have met so many good friends during my time at Durham, York and CCFE, and wish to thank them all for their help, support and many good times. While there is not space to name everyone here, special thanks to my housemates during the course of this PhD: Christophe Vaillant, Ophélie Lebrasseur and Penny Lawton, for helping to keep me sane(ish).

Finally I'd like to thank my family, for all their love and support over the years which has got me to where I am today.

List of Acronyms

α-BBO	Alpha barium borate (BaB_2O_4)
ADU	Analog-to-Digital Unit
CIS	Coherence Imaging Spectroscopy
CWL	Centre wavelength
CXRS	Charge exchange recombination spectroscopy
DFT	Discrete Fourier transform
DND	Double null divertor
DSLR	Digital single lens reflex
ELM	Edge-localised mode
FLC	Ferroelectric liquid crystal
FTS	Fourier transform spectroscopy
FWHM	Full width at half maximum
LCFS	Last closed flux surface
LSND	Lower single null divertor
MHD	Magnetohydrodynamic
NBI	Neutral beam injection
PFC	Plasma-facing component
PFR	Private flux region
PMT	Photomultiplier tube
PSF	Point-spread function
RMP	Resonant magnetic perturbation
SART	Simultaneous algebraic reconstruction technique
SNR	Signal-to-noise ratio
SOL	Scrape-off layer

Chapter 1

Introduction

In modern developed countries, communications, infrastructure, business, industry, and people's everyday lives are increasingly dependent on electrical power. As more of the world's population strive for increased living standards and economic prosperity, global electricity demands are increasing and are projected to more than triple by the year 2050 [IAEA, 2013]. Electricity is also projected to account for an increasing fraction of overall global energy demand.

Currently, the majority (67% as of 2011 [U.S. Energy Administration, 2014]) of global electricity generation is based on burning fossil fuels. This is unsustainable, especially in the face of rising demands, due to dwindling and unevenly distributed fuel resources, and the environmental impact of both extracting these fuels and releasing their combustion products into the atmosphere. So-called renewable energy sources such as hydroelectric, wind and solar power show promise as more sustainable methods of power generation, however these are unlikely to meet total demands by themselves due to their intermittency, relatively low power density and dependence on local geography and climate. Nuclear fission offers another alternative, and is attractive in terms of producing no atmospheric pollution, but has significant drawbacks in its current form. These include production of long-lived highly radioactive waste, and limited fuel resources using the current Uranium-based fuel cycles.

If we wish to continue our modern reliance on, and development of, electronic technology, we must find a new reliable, secure, sustainable and affordable method of electrical power generation. It has long been suggested that controlled nuclear fusion could provide such a power source, and research into this area has been ongoing for more than 50 years. A fundamental challenge in this research is obtaining detailed and accurate measurements of the behaviour of experimental fusion devices, in order to advance our understanding towards the ultimate goal of commercially available fusion power. This thesis is concerned with the development of diagnostic instrumentation and techniques for this purpose.

1.1 Fusion Energy

Nuclear fusion is the process in which two atomic nuclei collide and become bound together by the strong nuclear force, forming a single heavier nucleus. For nuclei lighter than iron-56 this process is exothermic, liberating the difference in binding energy between the initial and final nuclei. As an energy source fusion ‘fuel’ has an extremely high energy density: as much energy could be released by fusing 1kg of hydrogen isotopes to form helium as by burning around 45,000 tons of coal. The fusion reaction of most interest as an energy source is the fusion of the heavy hydrogen isotopes deuterium and tritium, resulting in an alpha particle (helium) and a neutron while releasing 17.6MeV of energy:

$$ {}^{2}_{1}\mathrm{D} \; + \; {}^{3}_{1}\mathrm{T} \; \rightarrow \; {}^{4}_{2}\mathrm{He}\,(3.5\mathrm{Mev}) \; + \; {}^{1}_{0}\mathrm{n}\,(14.1\mathrm{Mev}). \tag{1.1.1} $$

Here the energies in brackets are the (kinetic) energies carried by the fusion products. The main reason for this choice of reaction is its particularly large cross-section (probability of occurring) at the range of collision energies expected to be realisable in man-made fusion devices. The necessary deuterium can be extracted from seawater, and so represents a readily available, practically limitless fuel source. Tritium, however, undergoes beta decay with a half-life of 12.3 years and therefore does not occur naturally in meaningful amounts. Instead it can be produced using reactions of fast neutrons with lithium, and it is envisaged that a fusion power plant would be self-sufficient in tritium by using lithium and the fusion-produced neutrons to breed its own fuel. Unlike fission, the fusion reaction itself does not produce long-lived highly radioactive waste, although some short-lived (50 - 100 years) waste will be produced by the fast neutrons causing activation of reactor components in a D-T fusion plant.

In order to fuse, the positively charged fuel nuclei must collide with enough energy to have a significant chance of overcoming their electrostatic repulsion, in order for the strong nuclear force to become dominant. The problem of generating a controlled release of net energy from fusion is therefore to create a device which uses less power to produce this situation than is released by the resulting fusion reactions. The required kinetic energy of the fuel is several orders of magnitude higher than its ionisation energy, hence the electrons are completely stripped from the fuel atoms to form a plasma: a quasi-neutral gas of free electrons and ions.

Many fusion devices based on both thermal and non-thermal plasmas have been proposed and investigated both theoretically and experimentally [Bussard, 1991; Laberge, 2007; Meade, 2010]. The two most promising and well developed schemes, which have been the focus of most research since the 1950s, are based on confining a thermal plasma at sufficient temperatures and densities for large numbers of fusion reactions to occur. In inertial confinement fusion, a small pellet of solid deuterium and tritium is compressed and heated using an intense X-Ray or optical radiation field, produced by a high power laser system. The aim is to achieve sufficient density and temperature that large energy production from fusion occurs before the fuel disassembles due to its outward pressure (i.e. the fuel is only held in place by its own inertia). In magnetic confinement fusion, a much lower density plasma is sustained for longer periods of time, with the outwards pressure balanced by magnetic forces to confine the plasma in an equilibrium configuration. The work in this thesis is concerned with the latter class of device, and specifically with tokamaks.

1.2 Introduction to Tokamaks

The tokamak is currently the most well developed magnetic plasma confinement device, using magnetic fields to confine a toroidal (doughnut shaped) plasma. Originally a Russian design, the name is a Russian acronym which translates roughly as 'Toroidal chamber with magnetic coils'. Tokamak (toroidal) geometry is illustrated in figure 1.1. The major radius of the torus is denoted R_0 and the minor radius is a. The toroidal direction is the long direction around the torus, and the poloidal direction is the direction around the cross-section. The coordinate system which will be used throughout this thesis is a cylindrical system with the origin at the centre of the torus and the Z axis vertically up. A point in this system is described by its major radius R, height Z and toroidal angle φ. Where it is convenient to refer to the poloidal angle: the angle around the poloidal cross-section, this is denoted θ.

At least 210 tokamaks have been constructed since their invention in 1956[1], ranging in size from a few centimetres major radius to around 3m. Almost all major tokamak experiments use pure Deuterium plasmas rather than Deuterium-Tritium, to research how to create the necessary plasma conditions for fusion without the difficulty and expense of dealing with radioactive tritium, and limiting damage and activation of machine components by fusion neutrons. The largest currently operating tokamak is the Joint European Torus (JET) at $R_0 = 2.96$ m, $a = 1.25 - 2.1$

---

[1] www.tokamak.info

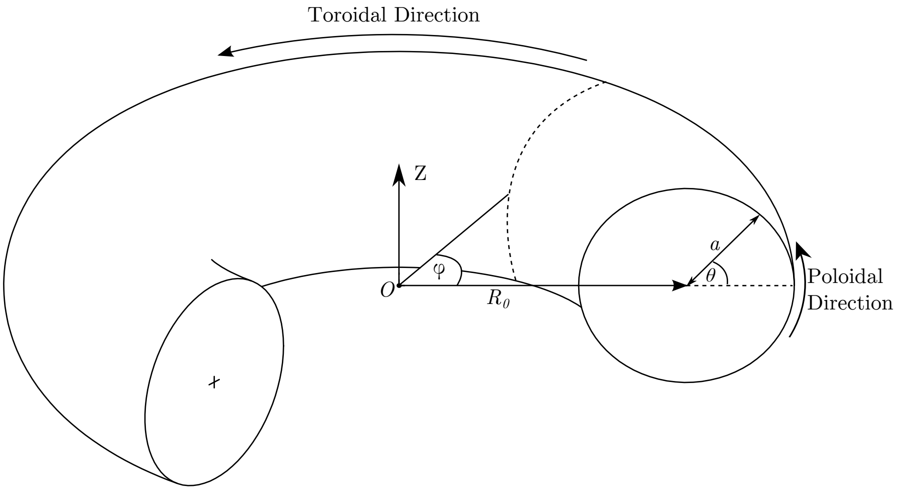

Figure 1.1: Illustration of toroidal geometry and coordinates.

m, located at Culham Centre for Fusion Energy in the UK. JET is the only operational tokamak equipped to run Deuterium-Tritium plasmas. In 1997, it set the (still current) world record for fusion power output using a Deuterium-Tritium plasma, at 16MW of fusion power output while using 25MW of input power to heat the plasma to the necessary temperature [Keilhacker et al., 1999]. This corresponds to a fusion energy gain: the ratio of fusion power produced to the heating power injected into the plasma, of $Q = 0.64$. The next major step in tokamak research is to demonstrate $Q \geq 1$, or 'breakeven'. This is one of the main goals of ITER: a tokamak approximately twice the size of JET currently under construction in southern France, in a collaboration spanning 34 nations. ITER's target is to demonstrate $Q \geq 10$ (generating around 500MW of fusion power), sustained for periods of 300 - 500 seconds, and $Q \geq 5$ in steady state operation [Shimada et al., 2007]. Its other important goal is to test, develop and demonstrate the technology necessary for a power-producing reactor (ITER itself will not generate electricity). If ITER is successful, the next step is then construction of a first demonstration power-producing tokamak, currently known as DEMO.

The magnetic field configuration used to confine a tokamak plasma is illustrated in figure 1.2, and consists of both toroidal and poloidal magnetic field components. The toroidal field B_φ is created by a set of planar current-carrying coils arranged around the torus. It is strongest at the inboard side of the torus and scales as $1/R$, hence the inboard and outboard sides are often called the high field side and low field side, respectively. The poloidal field B_θ is created by driving a toroidal current in the plasma itself. This plasma current is driven by transformer action where the plasma acts as the transformer secondary, the primary being a solenoid inserted through the centre of the torus. The poloidal field is much smaller than the toroidal field, and effectively adds a 'twist' to the toroidal field lines such that they take a helical path around the device (an example field line is shown as a black line in figure 1.2). As they spiral around the device the helical field lines trace out a set of nested, closed toroidal surfaces known as flux surfaces, such that the magnetic field vector lies within the surface at all points. The general principle of the plasma confinement is that because charged particles are constrained to gyrate around magnetic field lines due to the Lorentz force, but are free to move along them, plasma particles are confined to stay on a given flux surface and have difficulty moving perpendicular to the flux surfaces (i.e. perpendicular to the field lines) to escape the device. Additional vertical and radial field components are used to control the position and cross-section shape of the plasma.

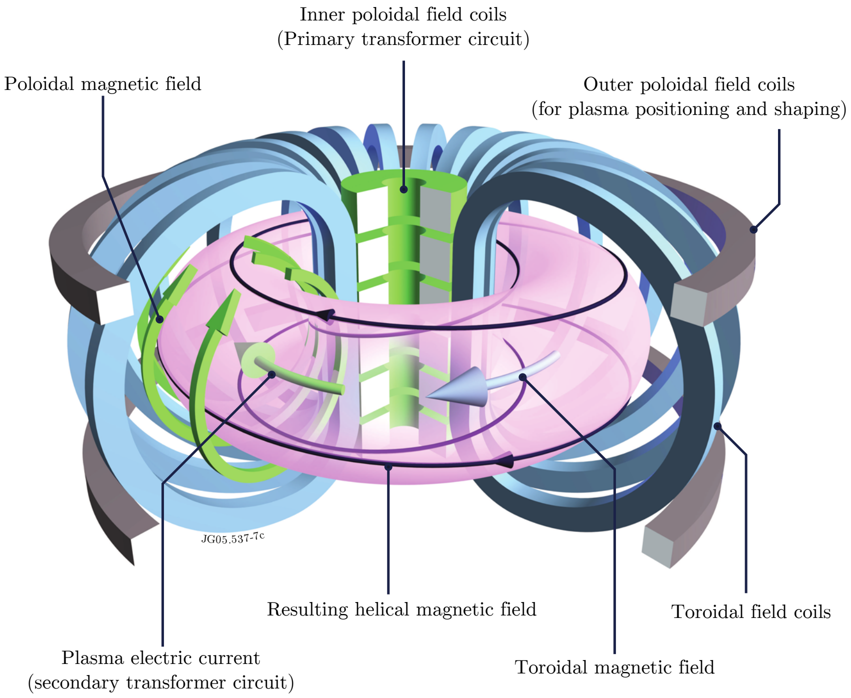

Figure 1.2: Illustration of magnetic field configuration and major components in a tokamak. Courtesy of the EFDA-JET figure database.

Heating of the plasma to achieve fusion relevant temperatures is derived from a variety of sources. The plasma current provides ohmic heating due to the finite resistivity of the plasma, although as the plasma heats up its resistivity decreases, making ohmic heating less effective. An almost ubiquitous non-ohmic heating method is to inject beams of high energy (accelerated) neutral atoms into the plasma, which are then ionised and give up their energy via Coulomb collisions. This is known as

Neutral beam injection (NBI). High power radio frequency heating is also common, in which microwaves resonant with the gyration of the ions or electrons around the magnetic field lines are launched into the torus, directly accelerating the relevant species which then thermalise through collisions. In devices producing significant amounts of fusion power, heating will also occur from the fusion itself: the fusion born energetic alpha particles are trapped by the magnetic field and give up their energy to the plasma via collisions. The extreme of this situation is known as 'ignition': this is the situation where all of the necessary heating comes from fusion alpha particles, and no external heating energy is required to sustain the reaction. In current medium to large scale tokamaks, typical temperatures at the plasma centre are around 10 - 300 million K. Since such numbers are rather unwieldy, it is conventional in plasma physics to quote the charactieristic thermal energy $k_\mathrm{B}T$ in place of temperature (for example, the above temperature range corresponds to around 1 - 26 keV). This convention will be used throughout the remainder of this thesis.

1.2.1 Plasma Exhaust & Divertors

In a real device some particles will always escape across the magnetic field lines and leave the plasma, due to cross-field transport mechanisms such as particle collisions and turbulence. While the associated loss of energy is undesirable since it cools the plasma, some particle loss is always necessary to remove the 'helium ash' from fusion and other unwanted impurities which could otherwise build up in the plasma. This escaping 'exhaust' will impact the walls of the device and cause damage through physical and chemical sputtering, and the resulting sputtered atoms will re-enter the plasma as undesired impurities. The first attempts to mitigate this problem used so-called limiters. A limiter is an object protruding into the edge of the plasma to limit its radial extent. The plasma is then divided into two regions: the confined plasma with closed magnetic field lines and flux surfaces, as already discussed, and the Scrape-off layer (SOL), where the magnetic field lines do not circulate the device indefinitely but intersect the surface of the limiter. The boundary between the two is the Last closed flux surface (LCFS), or separatrix. This situation is shown schematically in figure 1.3(a). Particles escaping across the LCFS into the scrape-off-layer travel rapidly along the helical open field lines until they reach the limiter surface, i.e. the plasma exhaust is 'scraped off' to the limiter before reaching the chamber wall. Limiters have the advantages of limiting the plasma flux to the chamber walls and being very simple to implement, but are not very effective at screening impurities from the main plasma. The high particle flux to the limiter still causes large amounts of sputtering, and the resulting impurities do not have to travel far to penetrate back across the LCFS, since the limiter is directly in contact with the confined plasma region.

The situation can be dramatically improved by using a so-called divertor configuration, which is employed by most modern tokamak experiments. A divertor is created by using additional magnetic coils, above and/or below the plasma, to create a ‘null’ in the poloidal magnetic field. This directs the SOL plasma towards specially designed target plates, situated further away from the confined plasma region. This configuration and some of the terms used to describe the divertor geometry are illustrated in figure 1.3(b).

1.2.2 Low & High Confinement modes

When the input plasma heating power rises above a certain threshold, tokamaks (and indeed all toroidal confinement devices) exhibit a spontaneous transition to an improved mode of plasma confinement known as ‘high confinement’ or H-Mode. This was first observed on the ASDEX (Axially Symmetric Divertor EXperiment) tokamak at Max-Planck-Institut für Plasmaphysik, Garching, Germany in 1982 [Wagner et al., 1982], and the pre-transition state was retroactively named ‘Low confinement’ or ‘L-Mode’. In H-Mode, cross-field particle and energy transport due to turbulence near the edge of the plasma is greatly suppressed, resulting in a factor of around 2 improvement in confinement over L-Mode. Although the physics of the L-H transition itself are not yet well understood, the mechanism of the transport suppression in H-Mode is known to be sheared poloidal flows breaking up the turbulent eddies near the plasma edge, and thereby reducing the distance over which they can transport heat and particles [Terry, 2000]. The transition from L-Mode to H-Mode appears almost instantaneous, and has several experimental signatures including the appearance of steep temperature and density gradients at the plasma edge, increase in core plasma temperature and density, and reduction of visible D_α light emission from the plasma edge and divertor. Figure 1.4 shows wide angle visible light images of a tokamak plasma in (a) L-Mode and (b) H-mode. In L-Mode, the plasma edge appears ‘fuzzy’ and filaments of escaping material are visible, while the H-Mode image shows a very ‘sharp’ plasma edge with no such filaments.

An important feature of H-Mode operation is a quasi-periodic instability known as the Edge-localised mode (ELM). An ELM is triggered when the pressure gradient at the edge of the plasma, which increases due to the improved confinement, rises above a threshold value for stability and triggers a rapid ejection of heat and particles from the plasma edge. These travel rapidly along the SOL field lines to the divertor, and cause very large transient heat and particle loads on the divertor

(a) Limiter configuration

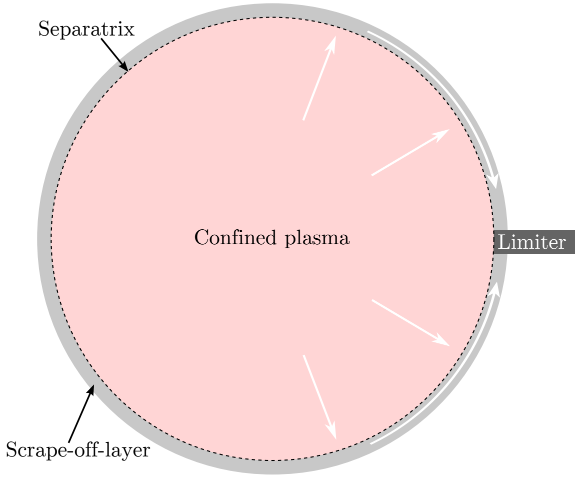

(b) Divertor configuration

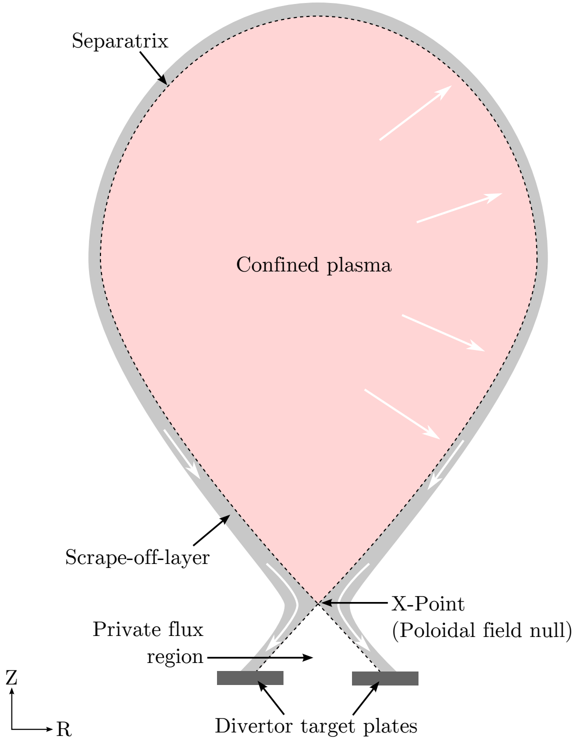

Figure 1.3: Schematics showing plasma exhaust handling and related terminology in limiter (top) and divertor (bottom) plasma configurations. Each schematic shows a poloidal cut-through of the toroidal plasma. White arrows show the plasma exhaust flow projected on to the plane of the figure.

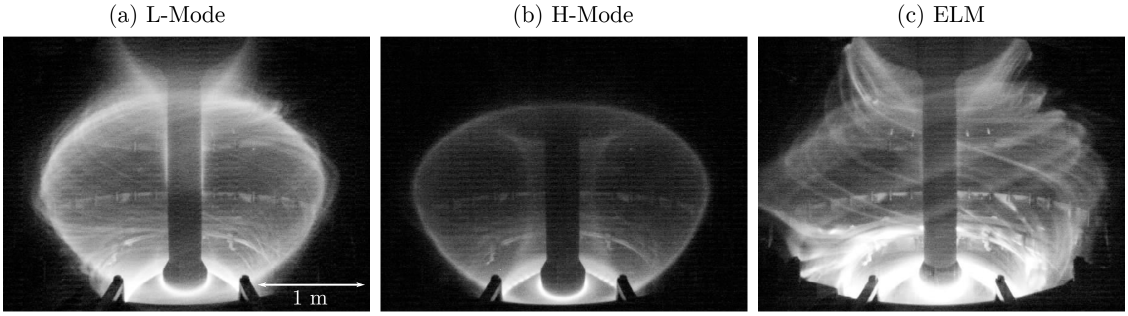

Figure 1.4: Visible light images of a MAST plasma in L-Mode (left), H-Mode (middle) and during an ELM (right). The reduction in edge turbulence can clearly be seen between L- and H-modes. The plasma ejected by the ELM in (c) can be seen interacting with a poloidal field coil and the divertor.

plates. The ejection of material relaxes the pressure gradient which then starts to rise again, repeating the process. Figure 1.4(c) shows a visible light image during an ELM, which shows the violence of this process and intense interaction between the plasma and the divertor and other in-vessel components, towards the bottom of the image. Filaments of material can be seen illuminating the helical magnetic field lines. Although ELMs do not have a large effect on the performance of the core plasma, their damaging effects on plasma facing components such as the divertor plates are expected to cause a significant practical challenge for future devices, and controlling these instabilities is an area of active research [Lang et al., 2013].

1.3 SOL & divertor flow diagnostics

1.3.1 Motivation for measuring flows

In the recent European roadmap for the development of fusion energy, finding a reliable plasma exhaust solution is described as *"probably the main challenge towards the realisation of magnetic confinement fusion"* [Romanelli et al., 2012]. While a conventional divertor design is expected to be sufficient in ITER, for a power generating device the heat and particle loads on the divertor components may well be too high to be compatible with current designs, materials and operation regimes. Substantial effort is therefore being spent to develop alternative divertor concepts and designs, including experiments on, and modifications to, a number of existing tokamaks [Katramados et al., 2011; Mirnov, 2009; Piras et al., 2009; Soukhanovskii et al., 2012]. An important aspect of this work is gaining a better understanding of the physics of SOL and divertor plasmas, and developing predictive models which can be used to inform the design of future devices. Gaining this understanding, and developing and benchmarking numerical models, must start with obtaining high quality experimental data from current devices.

Plasma flows are a key aspect of SOL and divertor physics, and are currently not completely understood or well reproduced by existing models. They play an important role in heat, particle and impurity transport around the SOL, and there is evidence that they have important implications for the core plasma such as affecting toroidal rotation and the power threshold for H-Mode access [LaBombard et al., 2004]. Interaction of the main ion flows with impurities, and the resulting impurity flows, are critically important, since they are required to exhaust helium and other impurities from the main plasma and to prevent these impurities travelling back to the plasma once they have reached the divertor plates. They are also important for understanding and controlling the erosion, transport and subsequent deposition of material from the divertor target plates [Elder et al., 2005], which has important implications for lifetimes of such components. Existing measurements have shown that SOL and divertor flows are highly complex phenomena, with multiple sources and sinks of particles at different plasma locations and flow components both dependent and independent of the magnetic field configuration [Asakura, 2007; LaBombard et al., 2004]. Highly detailed and wide-coverage measurements are therefore required to aid understanding of these phenomena.

1.3.2 Flow diagnostics and Coherence Imaging Spectroscopy

The most commonly used diagnostics for measuring SOL and divertor flows are Mach probes, (dispersive) Doppler spectroscopy, and high speed imaging of injected 'tracer' impurities ('plume imaging'). The capabilities and limitations of these diagnostics are now briefly considered, and the coherence imaging technique introduced in this context.

Mach Probes

Mach probes are a type of electrical probe inserted into the plasma to make in-situ flow measurements of the bulk ions. The simplest examples consist of two probe pins separated by an insulator, with both pins operating at ion saturation, i.e. negatively biased (usually with respect to the vacuum vessel) such that all plasma electrons are repelled but all incident ions are collected. Due to the geometry of the probe head, the two pins collect ions incident from opposite directions, and if there is a net plasma flow in the direction separating the pins, the 'upstream' pin will collect a higher ion flux than the downstream one. The Mach number of the flow (i.e.

the flow speed relative to the plasma sound speed) is determined from the ratio of the currents from the two pins. In practise the theory describing the operation of such probes, i.e. the interaction of the probe with the plasma, is very difficult, and care must be taken when analysing the data since the same measured data can give substantially different flow values depending on the probe theory applied [Chung, 2012].

Mach probe measurements have the advantages that the diagnostic hardware is cheap and relatively straightforward, and they provide a direct measurement of the bulk ion flow which is not possible with spectroscopy or imaging based techniques (for hydrogenic plasmas). The probe head is mounted on a reciprocation, or spatial scanning mechanism, which allows the probe to be quickly inserted and retracted from the plasma (over times of $\sim$ 10s of milliseconds) along a 1D path (typically in the direction across the magnetic field). This is necessary because the high temperature plasma (at least several thousand Kelvin even in the 'cold' divertor plasma) would quickly damage a probe in prolonged contact with the plasma and contaminate the plasma with material removed from the probe. This provides an inherent upper limit to the timescales over which continuous measurements can be made. On the other hand, due to the high sample rates (MHz) at which the probe signal can be digitised, the evolution of fast transient events such as ELMs can be studied [Asakura et al., 2008], and detailed radial profiles obtained during the reciprocation. The measurement is also well localised in 3D space due to the well-known probe location. However, most tokamaks are only equipped with one or at most a few such probes, at fixed (strategically chosen) poloidal locations around the plasma. This means the spatial coverage of these measurements, with respect to the plasma as a whole, is very sparse. In order to build up a picture of the flow behaviour over large scales in the SOL, results often need to be combined from multiple tokamaks, for example see Asakura [2007]. Nevertheless, with well chosen probe locations these measurements can still provide insight into course spatial patterns of flows, and time evolution on fast timescales.

Plume imaging

To measure impurity ion flows and transport, impurities can be deliberately injected into the edge of the plasma, where they are (partially) ionised and emit characteristic spectral lines. The motion of the resulting plume of impurity emission is then observed with high speed, narrow-band filtered cameras. The impurities are most commonly injected by gas puffing, e.g. in Gangadhara and LaBombard [2004] and McLean et al. [2005], but laser ablation [Timberlake, 1983] and electrical spark gaps [Leggate et al., 2013] can also be used. By virtue of their imaging nature, the results can provide intuitive visual indication of, for example, the dominant impurity flow direction, however complete quantitative interpretation of the images requires detailed knowledge of the local plasma conditions and supporting modelling. This is because the dynamics of the emission plume are determined by a range of factors in addition to the flow, such as the exact details of the injection source design and its interaction with the plasma, and the effect of the local plasma conditions on the ionisation and emission behaviour of the injected impurities [Gangadhara and LaBombard, 2004]. The spatial coverage of these measurements is inherently limited to locations where impurity injection can be performed. Like probe measurements, since this technique involves deliberately injecting impurities into the plasma it necessarily perturbs the plasma being measured.

Doppler spectroscopy

Doppler spectroscopy can be used to measure impurity ion flows via the Doppler shift of impurity emission lines. This is possible in the divertor and scrape-off-layer because the plasma is cool enough for the impurities of interest to be only partially ionised. These may be intrinsic impurities (those occurring ‘naturally’ in the device due to the choice of plasma-facing materials or unavoidable contamination), or deliberately injected, e.g. by gas puffing. A typical Doppler spectroscopy system consists of one or more collection lenses, usually outside the vacuum vessel viewing the plasma through a diagnostic port, behind which are arranged a number of optical fibres on to which the impurity emission is imaged. The optical fibres carry the collected impurity light to one or more spectrometers, where the light is dispersed using an element such as a diffraction grating on to a CCD detector to measure the spectrum. Emission line(s) appearing in each fibre’s spectrum can then be fit with suitable line shape functions to find their central wavelengths, and by comparing these with their known rest wavelengths the Doppler shift and therefore line-of-sight ion flows can be determined.

Doppler spectroscopy of intrinsic impurity emission has the advantage over the other techniques described here that it does not perturb the plasma being measured in any way. However, the measurement is poorly localised, since the plasma represents a 3D extended emission source through which the collection optics must view. Each spatial channel therefore obtains the emission integrated along its line-of-sight through the plasma, with the plasma conditions varying along that line of sight. This can make the measured spectra ambiguous and difficult to interpret, especially in the presence of spatially changing spectral features such as Zeeman splitting. This in turn can lead to large uncertainties in derived flow measurements.

The number of sight-lines used for Doppler spectroscopy diagnostics varies greatly from system to system. For example, the systems used in Isler et al. [1999] and Marr et al. [2005] provide 12 and 6 lines of sight, respectively, while those in Gafert et al. [1997] and Morgan [2011] are more sophisticated with totals of 150 and 128 sight lines. However, common to all these configurations is that they only provide 1 dimension of spatial resolution, with the sight-lines arranged in one or more fans through the plasma or viewing a small number of strategically chosen locations. This is clearly not well suited for studying overall SOL and divertor flow behaviour which is known to be 2 - 3 dimensional.

A new technique: Coherence Imaging Spectroscopy

Coherence imaging spectroscopy (CIS) is a relatively new technique which uses the same principle as traditional Doppler spectroscopy: measurement of impurity ion flows using the Doppler shift of their emission lines. However, instead of being a frequency domain measurement of the spectrum, coherence imaging uses an imaging interferometer to perform narrow-bandwidth Fourier transform spectroscopy. This is explained in detail in Chapter 2 of this thesis, but the general idea is that low order spectral information, including the Doppler shift, is encoded on to a 2D filtered image of the plasma in the form of an interference fringe pattern.

Coherence Imaging Spectroscopy (CIS) has a number of key advantages over traditional Doppler spectroscopy. One of the most striking is that it captures at least an order of magnitude more spatial information than even the highest (spatial) resolution traditional systems (for example, the MAST coherence imaging system has the equivalent of over 6,000 sight-lines), and provides 2D spatial resolution. Combined with its wide angle imaging capability, this means coherence imaging is uniquely able to study the 2D flow behaviour in the SOL and divertor at scales ranging from centimetres up to the size of the device, over a continuous 2D field of view. This can also be achieved with simple and relatively cheap hardware compared with complex multi-spectrometer systems with large numbers of optical fibres. A major advantage of the large amount of spatial information is that it becomes possible to tomographically invert the line-integrated measurements (subject to certain symmetry assumptions) to obtain localised flow measurements and untangle the effect of changing plasma conditions along the lines of sight. Additionally, coherence imaging flow measurements are not affected by Zeeman splitting[2], unlike dispersive Doppler

---

[2] See section 2.4.1

spectroscopy systems.

One of the main disadvantages of CIS is that it is very sensitive to spectral contamination from continuum or other line emission close to the line being measured. Where these are present coherence imaging does not capture enough spectral information to separate the effect of the contamination from the measurement, and traditional Doppler spectroscopy is more suitable in these situations. Calibration of coherence imaging instruments is also particularly challenging. A full review of the technique is given in Chapter 2.

1.4 The Mega Amp Spherical Tokamak (MAST)

MAST is a spherical tokamak operated by the UK Atomic Energy Authority, located at the Culham Centre for Fusion Energy. Spherical tokamaks are tokamaks with low aspect ratio $A = R_0/a < 2$, and are of interest because they can stably confine higher pressure plasmas for a given magnetic field strength than traditional 'high aspect ratio' designs [Peng and Strickler, 1986; Sykes et al., 1997]. It is thought this increased confinement efficiency may lead to more compact and cost effective designs for fusion neutron sources and/or fusion power plants, and the goal of MAST is to explore spherical tokamak physics towards these aims. MAST has what is considered a 'medium sized' plasma cross-section in the context of conventional tokamaks, and is one of the two largest spherical tokamaks in the world, the other being the National Spherical Toris EXperiment (NTSX) at Princeston Plasma Physics Laboratory in the US. The key engineering parameters of MAST are listed in table 1.1.

Major Radius	0.85 m
Minor Radius	0.6 m
Plasma Volume	8 m^3
Plasma Current	400 - 900 kA
Toroidal Field	0.5 - 0.6 T
NBI Heating Power	$\leq$ 5 MW
Pulse Duration	0.3 - 0.65 s

Table 1.1: Key engineering parameters of MAST.

A vertical cross-section through MAST is shown in figure 1.5. A cylindrical stainless steel vacuum vessel 4.0 m in diameter and 4.4 m tall forms the basis of its construction. Unlike almost all other tokamaks, there is no close-fitting wall around the plasma, which occupies only around 15% of the internal volume of the vacuum vessel. Combined with a large number of diagnostic access ports, this makes MAST uniquely well suited for optical measurements of the edge and exhaust plasma, since wide angle, relatively unimpeded optical access to the entire plasma boundary can be obtained. An example wide angle colour photograph of a MAST plasma and the vacuum vessel interior, taken from a midplane diagnostic port, is shown in figure 1.6. This is a composite image in which the dynamic range has been greatly compressed to make both the plasma and vessel interior visible. Another unique feature of MAST which can be seen in the photograph and figure 1.5 is that the poloidal field coils used to shape and stabilise the plasma are inside the vacuum vessel.

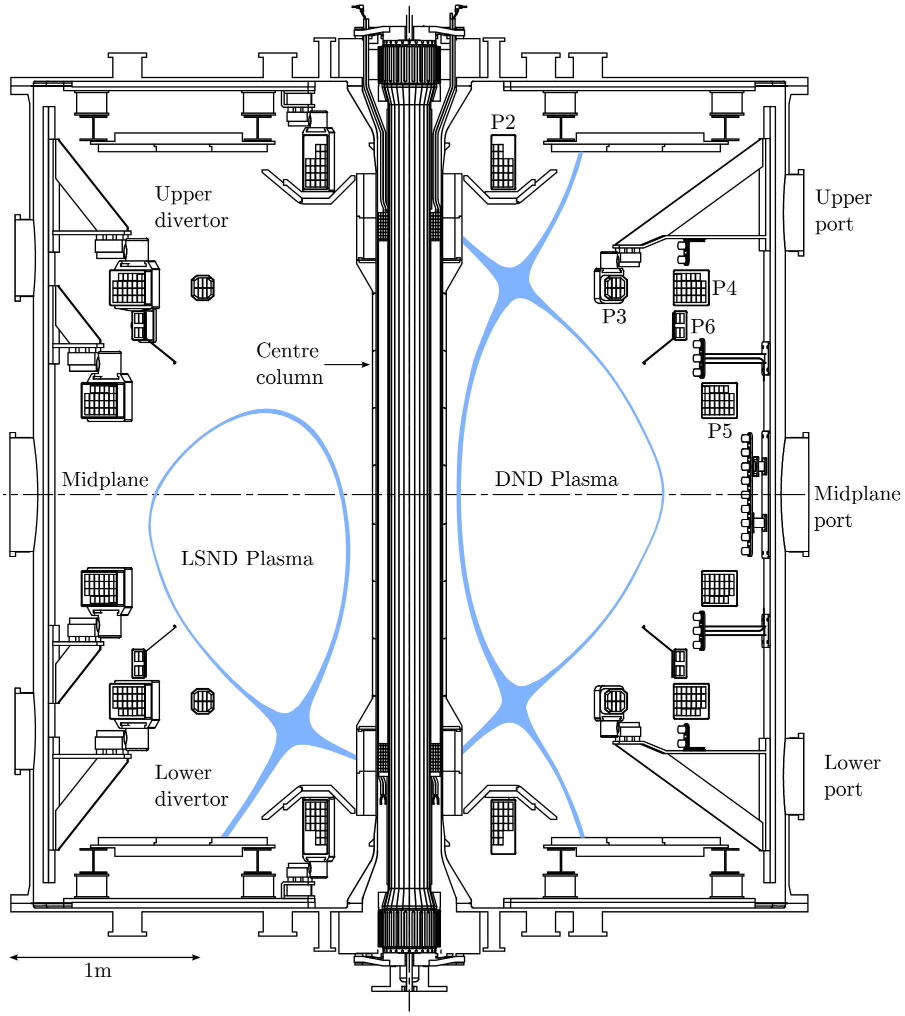

Figure 1.5: Cross-section of MAST showing lower single null divertor (left) and double null divertor (right) plasma shapes, diagnostic port locations and poloidal field coils (labelled P2 - P6 in the upper right of the diagram). MAST vessel cross-section courtesy of CCFE.

As can be seen in figure 1.5 MAST is up-down symmetric about the midplane, and most commonly runs Double null divertor (DND) plasmas, which are up-down

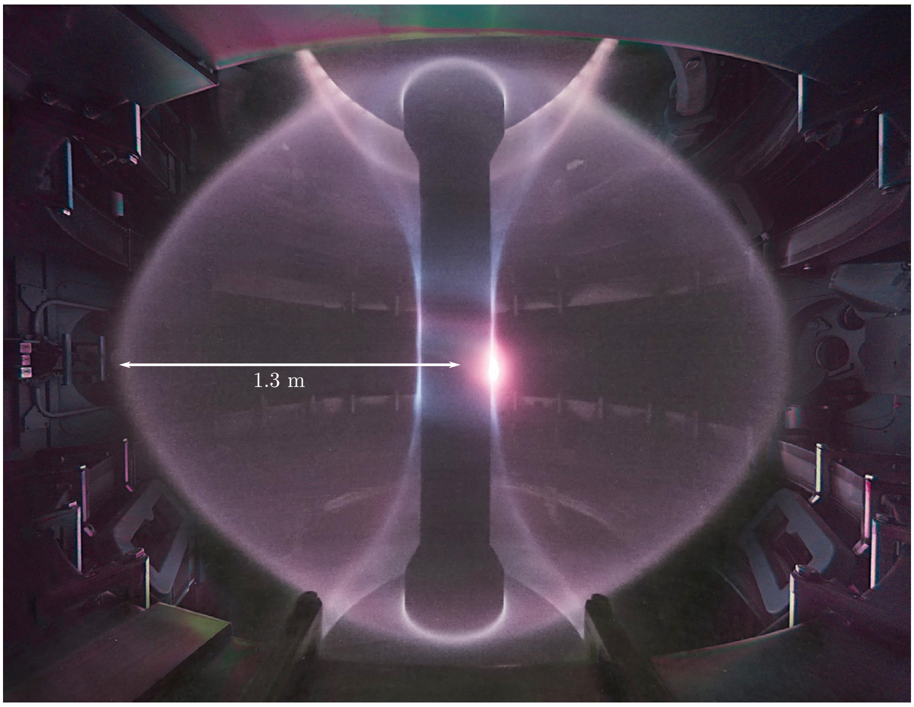

Figure 1.6: Composite colour photograph of a double null divertor MAST plasma and the vacuum vessel interior. The large bright spot on the right of the centre column is due to neutral deuterium gas injection used to fuel the plasma, from a gas nozzle located on the centre column.

symmetric and have divertors at both the top and bottom of the plasma. The cross-sectional profile of such a plasma is shown on the right of figure 1.5. By shifting the plasma vertical position MAST can also routinely run Lower single null divertor (LSND) plasmas, which use only the lower divertor. The profile of such a plasma is shown to the left in figure 1.5. The divertors themselves are a so-called ‘open’ design, meaning they feature no physical structures designed to help baffle impurities and neutral particles, originating at the divertor targets, from diffusing back towards the main plasma. While this is not optimal from a divertor performance standpoint, the open design gives excellent optical access to the divertor plasma. The outer (low field side) divertor targets consist of a set of horizontal graphite ribs attached to the top and bottom plates of the vacuum vessel, while the inner (high field side) targets consist of specially armoured sections of the centre column. For more information about the MAST machine see Cox and MAST Team [1999] and Darke et al. [2005].

1.5 Thesis Outline

This thesis presents the development of, and first results from, a coherence imaging spectroscopy diagnostic for measuring impurity ion flows in the SOL and divertor of MAST. In chapter 2 the theory of the coherence imaging technique is reviewed, along with the various existing implementations of this technique on other plasma devices. Chapter 3 presents the design of the MAST coherence imaging diagnostic in detail, including the development of simple numerical models used to inform the design choices. Chapter 4 describes the data analysis techniques used to obtain both line-of-sight average and tomographically inverted flow measurements from the raw data, and includes numerical estimation of the noise levels and possible systematic errors on the measurements. Chapter 5 presents results from offline characterisation and calibration of both individual diagnostic components and the integrated system, including benchmarking against the design calculations from chapter 3 where appropriate. The procedures for calibrating the instrument for plasma measurements are also described in this chapter. In chapter 6 the results of operating the diagnostic on MAST are presented, including analysis of the diagnostic’s performance and qualitative discussion of selected plasma observations. Chapter 7 summarises the key findings of this thesis and makes suggestions for future extensions to the work.

Chapter 2

The Coherence Imaging Spectroscopy Technique

This chapter reviews the principles and theory of the Doppler CIS technique, as well as previous implementations of of the technique on various plasma devices. The CIS theory presented in this chapter was primarily developed at Australian National University [Howard, 2002; Howard et al., 2001], except for the introduction of the 'multiplet phase' and 'multiplet contrast' in section 2.4.1, which to the best of the author's knowledge is original to this work.

2.1 Impurity line radiation

The Doppler CIS technique uses Doppler shifts of impurity ion emission lines to measure the impurity flow. These emission lines exist in the SOL and divertor because the plasma temperature in these regions: typically a few 10s of eV or less, is not hot enough to completely ionise the impurities, and thus their remaining electrons can undergo radiative transitions and emit characteristic spectral lines. This is illustrated in figure 2.1, which shows how the fractional abundance of different ionisation states of carbon varies with plasma electron temperature (assuming coronal equilibrium: the low plasma density equilibrium situation where collisional process dominate excitation while radiative processes dominate de-excitation; see Hutchinson [2005, pp. 225-227]). The fractional abundances of the C^{1+} and C^{2+} charge states (C II and C III in spectroscopic notation) are largest at temperatures found in the divertor and SOL, thus these ions are of particular interest in this work. In reality the peak emission brightness from each species can occur at a slightly different plasma temperature than suggested by figure 2.1, due to the influence of the local electron density on the emission brightness and also impurities being transported through different temperature plasma regions faster than their populations can reach equilibrium.

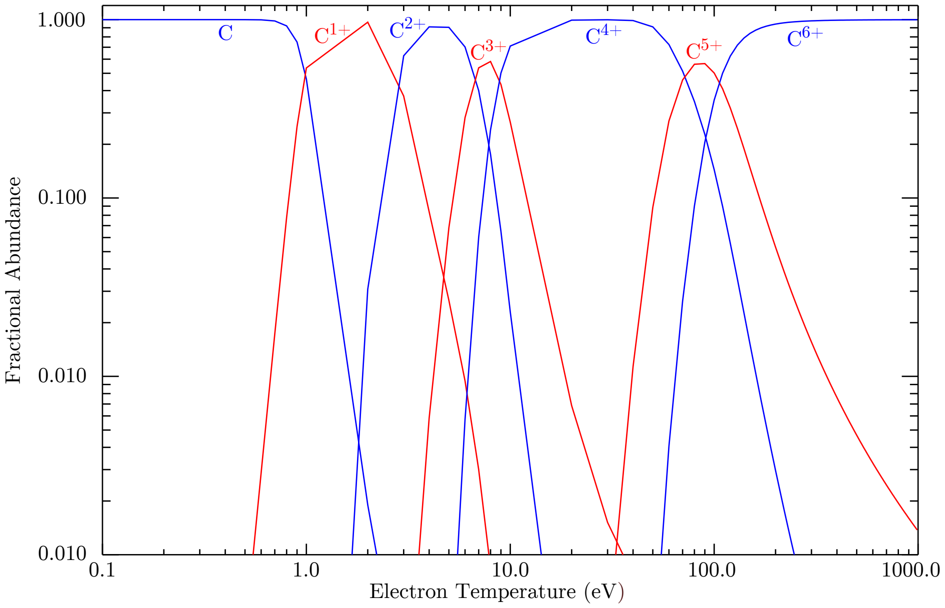

Figure 2.1: Relative abundances of different ionisation states of carbon, as a function of electron temperature, for coronal equilibrium. Figure courtesy of Dr. James Harrison, CCFE; calculated using data from the ADAS system [Summers, 2004].

When considering any spectroscopic diagnostic technique, we must consider what effects can determine the appearance of the spectra under examination. The relevant phenomena for the impurity ion lines in this study are now briefly discussed.

2.1.1 Spectral Line Shapes

Important phenomena which can affect the appearance of emitted impurity lines from plasmas include the Doppler effect, Zeeman splitting and Stark broadening. Each of these will now be briefly discussed. Natural line broadening, due to the finite lifetimes of ion excited states, is typically insignificant compared to these other effects and will not be discussed here. Since the coherence imaging technique operates by measuring an interferogram in the time domain, it will be convenient to discuss spectral line shapes as functions of optical frequency, $\nu = c/\lambda$ where c is the speed of light and λ the light wavelength. Line spectra will be represented in the form $I(\nu) = I_0 g(\nu)$, where $I(\nu)$ is the intensity spectrum, $I_0 = \int I(\nu) d\nu$ is the total intensity, and $g(\nu)$ is a normalised function representing the shape of the spectrum. Furthermore, we will make use of a normalised frequency shift coordinate $\xi = (\nu - \nu_0)/\nu_0$, where ν_0 is the rest-frame centre frequency of the spectral line.

Doppler effect

For any emitting ion moving relative to the observer, the observed frequency ν is related to the frequency in the emitter's rest frame, ν', by the Doppler shift $\nu = (1 - \boldsymbol{v} \cdot \hat{\boldsymbol{l}}/c)\nu'$, where $\boldsymbol{v}$ is the emitter's velocity in the observer's rest frame, c is the speed of light and $\hat{\boldsymbol{l}}$ is the line of sight direction from the observer to the emitter. The Doppler shift can be expressed in terms of the normalised frequency shift as $\xi = -\boldsymbol{v} \cdot \hat{\boldsymbol{l}}/c$. For an emitting plasma species with some velocity distribution function $f(\boldsymbol{v})$, an observer sees a broadened spectral line due to contributions from emitters with different velocities. To obtain the observed line shape $g(\xi)$, we sum over all emitters whose velocity Doppler shifts their emission by an amount ξ:

$$g(\xi) = \int f(\boldsymbol{v})\delta(\xi + \boldsymbol{v} \cdot \hat{\boldsymbol{l}}/c)\boldsymbol{dv}. \tag{2.1.1}$$

Consider an emitting plasma species in thermodynamic equilibrium, and experiencing a bulk plasma flow with flow velocity $\boldsymbol{v}_{\mathbf{D}}$ in the lab frame. In a reference frame drifting at the flow velocity, i.e. a frame where there is zero flow, denoted here by primed quantities, we assume the ion velocity distribution is isotropic and the speed distribution in any direction is given by a Maxwellian:

$$f'(v') = (\pi v_{\mathrm{th}}^2)^{-1/2} \exp\left(-\frac{v'^2}{v_{\mathrm{th}}^2}\right), \tag{2.1.2}$$

where v_{th} is a characteristic thermal speed given by $v_{\mathrm{th}} = \sqrt{2k_{\mathrm{B}}T_{\mathrm{i}}/m_{\mathrm{i}}}$, and T_{i} and m_{i} are the emitting species temperature and mass, respectively. The spectral line shape which would be observed in the drifting frame is then given by:

$$g'(\xi') = \left(\pi\frac{v_{\mathrm{th}}^2}{c^2}\right)^{-1/2} \exp\left(-\xi'^2\frac{c^2}{v_{\mathrm{th}}^2}\right). \tag{2.1.3}$$

This is the well known Gaussian profile of a Doppler broadened spectral line. To obtain the line shape observed in the lab frame, we simply shift the frequency coordinate according to the Doppler shift between the drifting and lab frames, $\xi' = \xi + \boldsymbol{v}_{\mathbf{D}} \cdot \hat{\boldsymbol{l}}/c$:

$$g(\xi) = \left(\pi\frac{v_{\mathrm{th}}^2}{c^2}\right)^{-1/2} \exp\left(-\left[\xi + \frac{\boldsymbol{v}_{\mathbf{D}} \cdot \hat{\boldsymbol{l}}}{c}\right]^2 \frac{c^2}{v_{\mathrm{th}}^2}\right). \tag{2.1.4}$$

This representation of the Doppler line shape will be convenient when considering Doppler coherence imaging measurements.

Stark Broadening

Stark broadening [Hutchinson, 2005, p. 261] arises from perturbations of the emitting ions' energy levels due to the electrostatic fields of neighbouring charged particles, via the Stark effect. The total electrostatic 'micro-field' experienced by the emitting ions varies both with position, and rapidly with time due to the fast motion of neighbouring electrons. The gross effect of the emitting ions experiencing a range of different, rapidly changing micro-fields and therefore a range of energy level perturbations is a broadening of emitted spectral lines, with a line shape resembling a Lorentzian. Since Stark broadening is caused by the interaction of neighbouring particles with the emitter, it is considered to be a type of pressure broadening.

The Full width at half maximum (FWHM) of Stark broadened spectral lines scales as $N_{\mathrm{e}}^{2/3}$, where N_{e} is the plasma electron density. Measurements of this effect can therefore be used as a spectroscopic plasma density diagnostic, with the proper choice of spectral line (for example see Koubiti et al. [2011]; Lischtschenko et al. [2010]; Soukhanovskii et al. [2006]). Using this density scaling, data from Konjević et al. [2002] and typical measurement conditions for the current work, the FWHM of the Stark broadened line shape is expected to be around 40 - 100 times smaller than the Doppler broadening FWHM for the lines used in this work. Stark broadening is therefore not considered in the remainder of this thesis.

Zeeman Effect

In the presence of an external magnetic field, emission lines are observed to split due to the splitting of atomic energy levels in the magnetic field. For atoms which can be described by L-S coupling, the energy level splitting is given by [Sobelman, 1979, p. 189]:

$$\Delta E_{m_J} = \mu_B g m_J B, \tag{2.1.5}$$

where ΔE_{m_J} is the energy shift for a level with total angular momentum projection m_J, μ_B is the Bohr magneton, B is the magnetic field and the Landé g factor is given by:

$$g = 1 + \frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)} \tag{2.1.6}$$

Each level is split into $2j+1$ levels corresponding to $m_J = -J...J$. For electric dipole transitions, allowed transitions between the split states are those with $\Delta J = 0, \pm 1$ and $\Delta m_J = 0, \pm 1$. Transitions with $\Delta m_J = 0$ are called "π" components, and have linear polarisation parallel to $\boldsymbol{B}$. Those with $\Delta m_j = \pm 1$ are "σ^+" and "σ^-" components, and have left and right-hand circular polarisation about the direction of $\boldsymbol{B}$. The relative intensities of the different components are given in Sobelman

[1979, (p. 193)] and reproduced in table 2.1. The observed relative intensities of the π and σ components also depend on the angle θ between the direction of view and $\boldsymbol{B}$ according to $I_{\pi,\text{obs}} = I_\pi \sin^2(\theta)$ and $I_{\sigma,\text{obs}} = I_\sigma[1 + \cos^2(\theta)]$. This model for Zeeman splitting is used in forward modelling calculations in chapters 3 and 4. The expected magnitude of the splitting for some of the impurity lines in this work, in the region of the X-Point on MAST, is similar to the Doppler broadening FWHM. While this would complicate measurements of the Doppler broadening of these lines, it will be shown that this does not significantly effect CIS flow measurements.

ΔJ	I_π	I_σ^+	I_σ^-
0	m_J^2	$\frac{1}{4}(J + m_J)(J + 1 - m_J)$	$\frac{1}{4}(J - m_J)(J + 1 + m_J)$
-1	$J^2 - m_J^2$	$\frac{1}{4}(J + m_J)(J - 1 +_J)$	$\frac{1}{4}(J - m_J)(J - 1 - m_J)$
$+1$	$(J+1)^2 - m^2$	$\frac{1}{4}(J + 1 - m_J)(J - m + 2)$	$\frac{1}{4}(J + 1 + m_J)(J + m_J + 2)$

Table 2.1: Relative intensities of π and σ Zeeman split line components, I_π, I_{σ^+} and I_{σ^-}, for transitions with different $\Delta J = J_{\text{upper}} - J_{\text{lower}}$ vales, from Sobelman [1979].

We now move on to present the theoretical basis for the coherence imaging spectroscopy technique.

2.2 Fourier Transform Spectroscopy

The CIS technique is based on the principles of Fourier transform spectroscopy (FTS). Unlike dispersive spectroscopy which directly records the spectrum in the frequency domain, FTS uses a 2-beam interferometer to record the spectral information in the time domain, which is then related to the desired frequency domain spectrum via the Fourier transform. To illustrate the principle of FTS, consider a polychromatic light source emitting with a complex electric field amplitude $E(t)$, whose spectrum we wish to measure. The light source is considered to be an ensemble of uncorrelated monochromatic emitters of different frequencies, and the resulting E field is a weak-sense stationary random process (i.e. its mean and variance are time independent, and its autocorrelation only depends on the time lag and not absolute time). The light from this source traverses a 2-beam interferometer, where the path difference between the two arms causes a delay τ between the two beams arriving at the output. The complex amplitude at the output will then be $u(t) = \frac{1}{2}(E(t) + E(t + \tau))$, and the corresponding detected intensity is given by:

$$\begin{aligned} S(\tau) = \langle u(t)u^*(t)\rangle = \frac{1}{4}\,\langle\, & E(t)E^*(t) + E(t+\tau)E^*(t+\tau) \\ & +E(t)E^*(t+\tau) + E(t+\tau)E^*(t)\rangle\,, \end{aligned} \tag{2.2.7}$$

where pointy brackets $\langle\rangle$ indicate a time average and * denotes the complex conjugate. The first two terms both equal the total intensity I_0. The second two terms are a quantity plus its own complex conjugate, thus equal to twice the real component. The signal can thus be written:

$$S(\tau) = \frac{1}{2}\left[I_0 + \mathrm{Re}(\Gamma(\tau))\right], \tag{2.2.8}$$

where $\Gamma(\tau) = \langle E(t)E^*(t+\tau)\rangle$ is the time-averaged autocorrelation of the electric field for time lag τ, and is known as the self coherence of the light. It will be convenient to deal with the normalised quantity known as the degree of coherence: $\gamma = \Gamma/I_0$, which is complex quantity characterised by a magnitude ζ where $0 \leq |\zeta| \leq 1$ and a phase $\phi = \arg(\gamma)$. We will see that these define an interferogram contrast (fringe visibility) and fringe phase, respectively. The Wiener-Khinchin theorem [Goodman, 2000, pp. 73-76] states that the Fourier transform of the autocorrelation of a stationary process is given by the power spectral density of that process. In this case, this means the degree of coherence is given by the inverse Fourier transform of the light source's spectral shape:

$$\gamma(\tau) = \int g(\nu)\exp(2\pi i\nu\tau)\mathrm{d}\nu. \tag{2.2.9}$$

From this result and (2.2.8) we see that the interferogram $S(\tau)$ is given by the inverse Fourier cosine transform of the spectrum we wish to measure, and therefore the spectrum can be obtained by applying the Fourier cosine transform to the measured interferogram. This can be understood intuitively by considering the source as an ensemble of monochromatic emitters. For a monochromatic source, 2-beam interference simply gives a set of cosine fringes, with the fringe frequency given by the light wavelength. If we then have an ensemble of such sources at different wavelengths, the observed interferogram will simply be the sum of these individual fringe patterns. The frequency content of the interferogram, i.e. its Fourier transform, is therefore determined by the wavelengths present and their relative intensities, i.e. the light spectrum. This is the basis of conventional FTS, in which $S(\tau)$ is measured over a range of τ values $-\tau_{\max} \leq \tau \leq \tau_{\max}$, determined by the desired spectral resolution by $\Delta\nu = 1/\tau_{\max}$. The interferogram is sampled at intervals set by the minimum wavelength in the spectrum, which determines the maximum frequency in the interferogram which must be at least Nyquist sampled.

Experimental techniques for measuring $S(\tau)$ can generally be classified by the way τ is varied, as either temporal or spatial multiplexing techniques. In temporal multiplexing techniques, the value of τ is changed in time, encoding the spectral information in the time history of the detected signal. The simplest way of achieving this is with a Michaelson interferometer, by equipping one arm of the interferometer with a movable mirror to vary that arm's path length. The mirror position is then simply scanned in time to record $S(\tau)$ over the desired range of τ. Scanning Michaelson systems can easily be used for 2D spectral imaging, by simply imaging a scene on to a 2D detector via the interferometer [Lewis et al., 1995]. However, the speed at which τ can be scanned is limited due to vibrations in the moving mirror reducing the fringe contrast and hence Signal-to-noise ratio (SNR). These systems also require very careful alignment and are highly sensitive to vibrations, ambient temperature changes etc. Such instruments have been widely applied in space-based astronomical and remote sensing applications [Persky, 1995], where variations of the source spectrum are slow and long scanning times (1 - 10s) are therefore acceptable. They have not been widely applied to fusion plasma diagnostics, since they typically cannot meet the high time resolution requirements.

In spatial multiplexing techniques, τ is varied in space along a multi-element detector, allowing the entire interferogram to be recorded simultaneously. Using a 2D detector array these instruments can provide both spectral and 1D spatial information with no scanning or moving parts, by using one detector dimension for spatial information and the other for spectral information (the interferogram). They are therefore much better suited for studying rapidly changing spectral sources where long scanning times are unacceptable, at the cost of spatial information in 1 dimension. Where 2D spectral imaging is required and the spectral scene is not changing too rapidly, e.g. in many remote sensing applications, push-broom scanning of the field of view can be used to build up 2D images from a series of 1D profiles [Smith et al., 1999]. A wide variety of spatial multiplexing FTS designs have been conceived [Harlander et al., 2010; Okamoto et al., 1984; Padgett and Harvey, 1995; Zhang et al., 2002], often featuring monolithic construction of the interferometer to make these systems much more robust to vibration than time multiplexed systems with moving interferometer parts, and in very compact space envelopes.

In all of these conventional FTS instruments, scanning in time of either the interferometer delay or field of view is required to obtain spectral images with 2D spatial resolution, over timescales much longer than are acceptable for fusion diagnostics applications. However, these scanning requirements can be reduced when we are only interested in recovering small amounts of spectral information from individual spectral lines, as is the case for impurity Doppler spectroscopy. This is the basis of the coherence imaging concept, and will now be discussed in more detail.

2.3 Narrow band spectra and fixed delay FTS

We now consider the application of FTS to quasi-monochromatic spectra, i.e. with width $\Delta\nu/\nu_0 \ll 1$, and whose spectral shape $g(\nu)$ can be characterised by a small number of parameters. This is the case when measuring isolated plasma emission lines where the line shape is dominated by a small number of physical phenomena. We start by expressing the spectral shape about its centre-of-mass frequency $\nu_0 = \int \nu g(\nu)\mathrm{d}\nu / \int g(\nu)\mathrm{d}\nu$, such that $g'(\nu - \nu_0) = g(\nu)$. Substituting this into (2.2.9), by the translation property of the Fourier transform the degree of coherence is given by:

$$\gamma(\tau) = \exp(2\pi i \nu_0 \tau)\gamma', \tag{2.3.10}$$

where

$$\gamma'(\tau) = \int g'(\nu)\exp(2\pi i \nu \tau)\mathrm{d}\nu. \tag{2.3.11}$$

Assuming for simplicity of argument that g' is an even function, and therefore γ' is purely real, the interferometer signal is then:

$$S(\tau) = \frac{I_0}{2}\left[1 + \zeta(\tau)\cos(2\pi\nu_0\tau)\right]. \tag{2.3.12}$$

This represents a set of sinusoidal interference fringes oscillating within a slowly varying contrast envelope. The fringe contrast, $\zeta = (S_{\max} - S_{\min})/(S_{\max} + S_{\min})$ is simply given by $\zeta = |\gamma'|$ and is affected only by changes in the line profile g', while the fringe phase is affected only by changes in the centre wavelength ν_0. By measuring the interferogram intensity, contrast and phase at a known value of τ, which requires measurements of $S(\tau)$ over only a single interference fringe, we can recover the line intensity, centre wavelength, and the value of $\gamma'(\tau)$ which carries information about the line shape. This is illustrated in figure 2.2. In order to fully characterise a line profile described by n parameters, we can repeat this process at a small number $\geq n$ of suitable delay values, in order to sufficiently sample $\gamma'(\tau)$.

The need to vary the delay over only a single fringe (at one or more fixed values of τ) instead of the wide range required for broadband FTS enables ‘fixed delay’ interferometer designs, and is the basis for coherence imaging. In temporal multiplexing schemes the very small scanning range can be covered much more quickly than in traditional FTS, and static instruments without moving parts become fea-

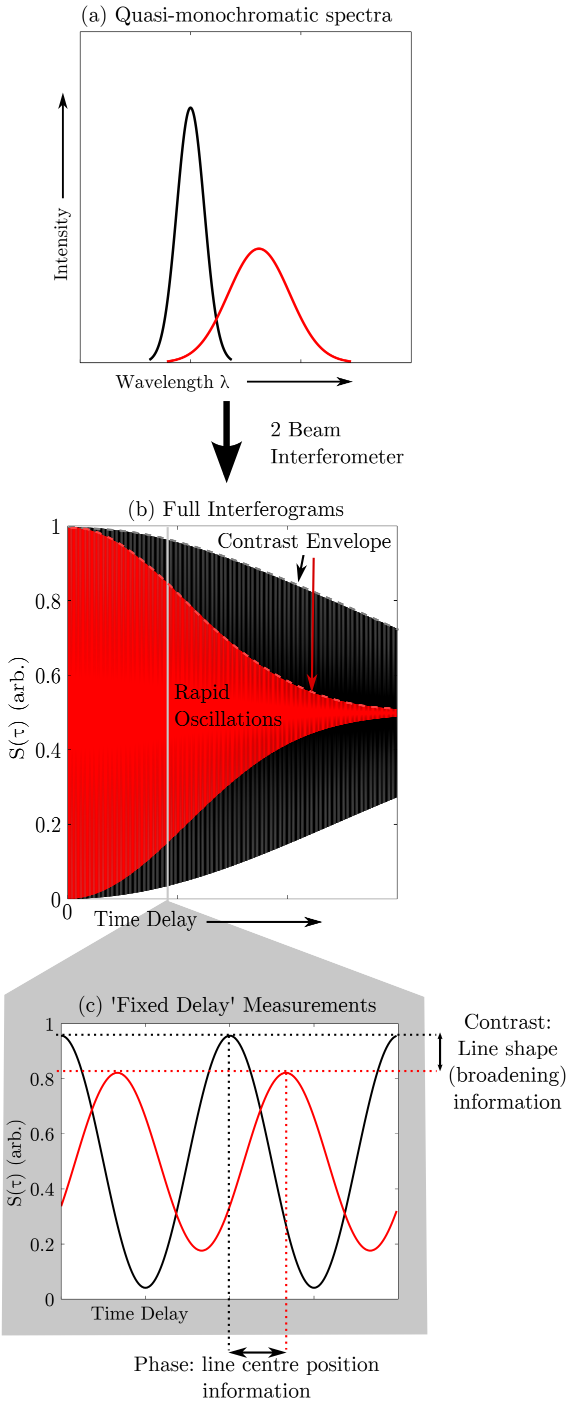

Figure 2.2: Illustration of fixed delay FTS: (a) Spectra of two Gaussian spectral lines of equal brightness but different centre wavelength and broadening. (b) Interferograms corresponding to equation (2.3.12) for these lines. The contrast for the red line decays faster with τ since the broader spectrum has a shorter coherence time. The fringe frequency of the two is slightly different due to the wavelength shift. (c) Magnified region of (b) illustrating measurements of a single fringe at known τ. The broadening and wavelength shift are obtained by the fixed delay measurement without the need to measure the entire interferogram in (b).

sible. For spatial multiplexing instruments, the number of detector pixels required in the spectral direction is greatly reduced, which allows 2D spatial resolution in a single exposure since the spectral information no longer occupies the whole of one detector dimension. Several designs of fixed delay FTS instruments for coherence imaging will be reviewed in section 2.5.

2.3.1 Delay Dispersion

In an ideal instrument the interferometer delay τ would be independent of optical frequency, making equation (2.2.9) a straightforward Fourier transform. However in the instruments used in coherence imaging, which include birefringent crystal components in their designs, dispersion in the crystals means that in fact τ varies with ν. This is included in our analysis by approximating $\tau(\nu)$ as linear about the centre wavelength ν_0:

$$\tau(\nu) = \tau_0 + (\nu - \nu_0) \left.\frac{\partial \tau}{\partial \nu}\right|_{\nu_0}, \tag{2.3.13}$$

where $\tau_0 = \tau(\nu_0)$. Changing the variable of integration in (2.2.9) to the normalised frequency shift ξ, and using $\nu = \nu_0(1+\xi)$ and $\tau = \tau(\nu)$, the degree of coherence is given by:

$$\gamma = \int g(\xi) \exp\{2\pi i \nu_0 [1+\xi]\, \tau(\nu_0[1+\xi])\} \mathrm{d}\xi. \tag{2.3.14}$$

Making the same substitution for ν in (2.3.13) we obtain $\tau(\nu_0[1+\xi]) = \tau_0\,(1 + (\kappa - 1)\xi)$, where κ is given by:

$$\kappa = 1 + \frac{\nu_0}{\tau_0} \left.\frac{\partial \tau}{\partial \nu}\right|_{\nu_0}. \tag{2.3.15}$$

Substituting this into (2.3.14) gives the degree of coherence, including the effect of delay dispersion, as:

$$\gamma(N) = \exp(2\pi i N) \int g(\xi) \exp(2\pi i \kappa N \xi) + \mathcal{O}(\xi^2) \mathrm{d}\xi, \tag{2.3.16}$$

where $N = \nu_0 \tau_0$ is the number of waves phase delay at frequency ν_0, and terms $\mathcal{O}(\xi^2)$ will be neglected since we are considering quasi-monochromatic spectra. It will be convenient to introduce the interferometer group delay $\hat{N} = \kappa N$, and note that the factor κ gives the ratio between the phase and group delays.

2.4 Doppler Coherence Imaging Spectroscopy

We now consider how the measured interferogram phase and contrast relate to the desired plasma parameters when measuring spectral lines whose shape is dominated by the Doppler effect. In this case the line shape g' can be characterised by a single parameter: its width, and both the ion flow and temperature can be obtained from measurements at a single fixed delay. Applying the result in equation (2.3.16) to the Doppler line shape in equation (2.1.4), we obtain the degree of coherence:

$$\gamma = \exp(2\pi i N) \exp\!\left(-2\pi i \hat{N} \frac{\boldsymbol{v} \cdot \boldsymbol{l}}{c}\right) \exp\!\left(-(\pi \hat{N})^2 \left(\frac{v_{\mathrm{th}}}{c}\right)^2\right). \tag{2.4.17}$$

The first exponential represents the phase due to the line centre wavelength in the plasma rest frame, while the second represents an additional phase shift due to the line-of-sight plasma flow (Doppler shift). The third exponential describes the fringe contrast envelope and is a function of the emitting species temperature and mass through the thermal velocity v_{th}, introduced in section 2.1.1. To find the form of the interferometer signal for this degree of coherence, we use equation (2.2.8) and $\gamma = \Gamma / I_0$ to give:

$$S = \frac{I_0}{2}\left[1 + \zeta_{\mathrm{D}} \cos(\phi_0 + \phi_{\mathrm{D}})\right]. \tag{2.4.18}$$

The Doppler fringe contrast ζ_{D} is given by:

$$\zeta_{\mathrm{D}} = \exp(-T_{\mathrm{i}}/T_{\mathrm{C}}), \tag{2.4.19}$$

where T_{i} is the emitting ion species temperature and we have defined a characteristic 'instrument temperature' which is a calibration parameter depending on the ion mass and group delay:

$$T_{\mathrm{C}} = \frac{m_{\mathrm{i}} c^2}{2 k_{\mathrm{B}} (\pi \hat{N})^2}. \tag{2.4.20}$$

This relationship can be used to obtain measurements of the emitting species temperature from the measured fringe contrast. The interferogram phase is composed of the rest frame centre phase $\phi_0 = 2\pi N$ plus a 'Doppler phase' given by:

$$\phi_{\mathrm{D}} = -2\pi \hat{N} \frac{\boldsymbol{v} \cdot \hat{\boldsymbol{l}}}{c}. \tag{2.4.21}$$

We can therefore measure the component of the plasma flow along the instrument line of sight by measuring the interferogram phase, subtracting the value of ϕ_0 to isolate ϕ_{D} and using equation (2.4.21). Like T_{C}, ϕ_0 is characteristic of a given combination of spectral line and fixed delay, and is thus treated as a calibration parameter and called the 'instrument phase' ϕ_{I}.

These results for the Doppler phase and contrast are the basis for interpreting Doppler coherence imaging measurements. We will now consider how other factors can complicate this interpretation.

2.4.1 Multiplet Spectral Lines

So far we have considered only single isolated spectral lines when relating the measured interferogram properties to the desired plasma parameters. However, a number of plasma lines of interest consist of multiple closely spaced line components, due to fine structure and Zeeman splitting. In this case, the measured interferogram is given by the sum of the interferograms of all the line components. We now consider how this affects interpretation of the measured fringe phase and contrast.

For multiplet spectral lines we can consider the spectral shape as the convolution of the Doppler spectral shape (2.1.3) with each of a set of Dirac delta functions representing the locations of the line components. For a multiplet consisting of n individual lines with rest frame centre positions $\xi_{1..n}$ and relative intensities $I_{1..n}$, normalised such that $\sum_{i=1}^{n} I_i = 1$, the overall line shape can be written:

$$g(\xi) = g_{\mathrm{D}}(\xi) \otimes \sum_{i=1}^{n} I_i \delta(\xi - \xi_i), \tag{2.4.22}$$

where g_{D} is given by equation (2.1.3). It follows from the convolution theorem that the degree of coherence for this spectrum is given by $\gamma = \gamma_{\mathrm{D}}\gamma_{\mathrm{M}}$, where γ_{D} is our previous result for Doppler coherence given by equation (2.4.17) and γ_{M} is given by:

$$\gamma_{\mathrm{M}} = \sum_{i=1}^{n} I_i \exp(2\pi i \hat{N} \xi_i). \tag{2.4.23}$$

Since the multiplet structure is not, in general, symmetric about ν_0, γ_{M} contributes both a contrast and phase to the interferogram. These modify our previous result for the interferometer signal (equation (2.4.18)) as multiplicative and additive terms respectively, giving:

$$S = \frac{I_0}{2}\left[1 + \zeta_{\mathrm{M}}\zeta_{\mathrm{D}} \cos(\phi_0 + \phi_{\mathrm{M}} + \phi_{\mathrm{D}})\right], \tag{2.4.24}$$

where the 'multiplet contrast' is given by $\zeta_{\mathrm{M}} = |\gamma_{\mathrm{M}}|$ and the 'multiplet phase' is $\phi_{\mathrm{M}} = \arg(\gamma_{\mathrm{M}})$. The multiplet contrast generally goes through a series of maxima and minima as a function of $\hat{N}$, since the sinusoidal interferograms of the individual line components sum to create a beat pattern. This affects the optimal choice of $\hat{N}$

for an instrument, since it is desirable to operate in a maximum of this beat pattern for the best SNR measurement of the interferogram.

From this analysis we see that if ζ_{M} and ϕ_{M} are well known, it is straightforward to account for the effect of multiplet spectral lines in our interpretation of the measured signal. It is therefore important to ensure that the structure of the line to be measured is well known such that γ_{M} can be calculated, or that there is a method by which ζ_{M} and ϕ_{M} can be measured directly. It is also important that the line structure does not change in time or space, since this will change the local values of ζ_{M}, ϕ_{M} and ϕ_0 (the latter through simply changing the line centre-of-mass ν_0) and result in inaccurate determination of the plasma parameters from the measured signal. A notable example of a lineshape effect which can vary in space is Zeeman splitting, particularly in a spherical tokamak where the magnetic field strength and orientation are strong functions of position, and the resulting line structure can vary over the field of view or even along individual sight lines. However, since Zeeman splitting produces a symmetrical splitting about the line centre, its multiplet phase is always zero and it does not change the line centre-of-mass, so time or spatially varying Zeeman splitting does not affect flow measurements. This is an advantage of the coherence imaging technique, since interpretation of line integrated spectra in the presence of Zeeman splitting is a complication faced by dispersive measurements. Note, however, that Zeeman splitting can still contribute a spatially varying multiplet contrast, which should be considered in interpretation of the measured fringe contrast.

2.4.2 Inhomogeneous Sources

The analysis so far has considered the case for an emitting ion population described by a single value of ion flow and temperature, i.e. a homogeneous plasma. However, in a tokamak plasma the species of interest will have spatially varying emissivity $e_0(\boldsymbol{r})$, flow velocity $\boldsymbol{v}(\boldsymbol{r})$ and ion temperature $T_{\mathrm{i}}(\boldsymbol{r})$. Any optical system viewing the 3D plasma will collect light emitted at all points along its line of sight, resulting in an observed spectral shape $\breve{g}(\xi)$ given by an emissivity weighted average along the line of sight:

$$\breve{g}(\xi) = \frac{1}{\breve{e}_0} \int_L e_0(\boldsymbol{r}) g(\xi, \boldsymbol{r}) dl, \tag{2.4.25}$$

where

$$\breve{e}_0 = \int_L e_0(\boldsymbol{r}) dl. \tag{2.4.26}$$

This will no longer be a simple Doppler lineshape, but will be broadened and likely asymmetrical due to the contributions of different spectral shapes along the line of sight. To properly interpret measurements of inhomogeneous plasmas, we must understand how the interferogram phase and contrast associated with this lineshape relate to the spatially varying ion flow and temperature. Substituting (2.4.25) into equation (2.2.9), we find the degree of coherence of the line integrated spectrum is simply given by the weighted line-average of the local degree of coherence $\gamma(\boldsymbol{r}, N)$, i.e.

$$\breve{\gamma}(N) = \frac{1}{\breve{e}_0} \int_L e_0(\boldsymbol{r}) \gamma(\boldsymbol{r}, N) dl, \tag{2.4.27}$$

where $\gamma(\boldsymbol{r}, N)$ is the degree of coherence of the spectrum emitted at position $\boldsymbol{r}$, and N is the number of waves phase delay.

The line average and local degrees of coherence can be written in polar form, described by their contrast and phase, as $\breve{\gamma} = \breve{\zeta} \exp[i\breve{\phi}]$ and $\gamma(\boldsymbol{r}) = \zeta(\boldsymbol{r}) \exp[i\phi(\boldsymbol{r})]$, respectively. We also consider the local phase value as being the line-average phase plus a local perturbation: $\phi(\boldsymbol{r}) = \breve{\phi} + \tilde{\phi}(\boldsymbol{r})$. Making these substitutions in (2.4.27), we obtain:

$$\breve{\zeta} = \frac{1}{\breve{e}_0} \int_L e_0(\boldsymbol{r}) \zeta(\boldsymbol{r}) \exp[i\tilde{\phi}(\boldsymbol{r})] dl. \tag{2.4.28}$$

Taking the Maclaurin series of the complex exponential to third order, the real part of this equation is:

$$\breve{\zeta} \approx \frac{1}{\breve{e}_0} \int_L e_0(\boldsymbol{r}) \zeta(\boldsymbol{r}) [1 - \tilde{\phi}(\boldsymbol{r})^2/2] dl. \tag{2.4.29}$$

For small $\tilde{\phi}$ such that the $\tilde{\phi}^2$ term can be neglected, this shows the measured contrast is simply the emissivity weighted line average of the local contrast $\zeta(\boldsymbol{r})$. Proceeding similarly with the imaginary part of (2.4.28), and using the result for $\breve{\zeta}$ while neglecting the $\tilde{\phi}^2$ term, we obtain a similar relationship for the phase:

$$\breve{\phi} \approx \frac{1}{\breve{e}_0 \breve{\zeta}} \int_L e_0(\boldsymbol{r}) \zeta(\boldsymbol{r}) [\phi(\boldsymbol{r}) - \tilde{\phi}^3(\boldsymbol{r})/6] dl. \tag{2.4.30}$$

We now have simple line integral relationships between the measured phase and contrast and their local values, under the approximation of $\tilde{\phi} \ll 1$. We can now substitute in the results from section 2.4 for the local Doppler phase and contrast, to obtain the relationships between the measured Doppler contrast $\breve{\zeta}_{\mathrm{D}}$ and phase $\breve{\phi}_D$ and the local ion flow and temperature:

$$\breve{\zeta}_\mathrm{D} \approx \frac{1}{\breve{e}_0}\int_L e_0(\boldsymbol{r})\exp(-T_\mathrm{i}(\boldsymbol{r})/T_\mathrm{C})\mathrm{d}l \tag{2.4.31}$$

$$\breve{\phi}_\mathrm{D} \approx -\frac{2\pi\hat{N}}{\breve{e}_0\breve{\zeta}c}\int_L e_0(\boldsymbol{r})\exp(-T_\mathrm{i}(\boldsymbol{r})/T_\mathrm{C})\boldsymbol{v}(\boldsymbol{r})\cdot\mathbf{d}\boldsymbol{l}, \tag{2.4.32}$$

where T_C is the instrument temperature defined in (2.4.20). Having these simple line integral relationships enables tomographic inversion of the measured interferogram properties to obtain local values of flow and temperature. However, note that the contrast line integral (2.4.29) is weighted by the local emissivity, and the phase line integral (2.4.30) is weighted by both the local emissivity and temperature (contrast). Reconstruction of the temperature therefore requires weightings calculated by reconstructing the emissivity, and reconstruction of the flow requires reconstructions of the both the emissivity and temperature. This amplifies reconstruction errors and can lead to very noisy results, since errors in the reconstruction of each parameter are fed back into the next. If we are only interested in the plasma flow, this can be alleviated by choosing the instrument delay such that $T_\mathrm{i} \ll T_\mathrm{C}$, i.e. the expected range of ion temperatures is small compared to the instrument temperature. The local contrast is then very insensitive to ion temperature and we have $\zeta(\boldsymbol{r}) \approx \breve{\zeta}\ \forall\, \boldsymbol{r}$, reducing (2.4.32) to:

$$\breve{\phi}_D = -\frac{2\pi\hat{N}}{\breve{e}_0\, c}\int_L e_0(\boldsymbol{r})\boldsymbol{v}\cdot\boldsymbol{dl}. \tag{2.4.33}$$

It is then only necessary to reconstruct the emissivity before the flow. Tomographic inversion of flow data based on this result is discussed in section 4.6.

2.5 Coherence Imaging Spectrometer Designs

We now review several designs of fixed delay interferometer which have been developed for coherence imaging, principally at the plasma research laboratory at Australian National University, and deployed on various plasma devices. All of these instruments are polarisation interferometers: rather than using beamsplitters and physically separate optical paths for each arm, a single optical path is used and orthogonal polarisation states act as the two interferometer beams. This design is much less sensitive to misalignment, vibrations, and thermal expansion effects than traditional interferometer designs due to the common optical path, and allows compact, robust and easy to align instrument designs. In order to use orthogonally polarised states as the interferometer beams, birefringent components are used to

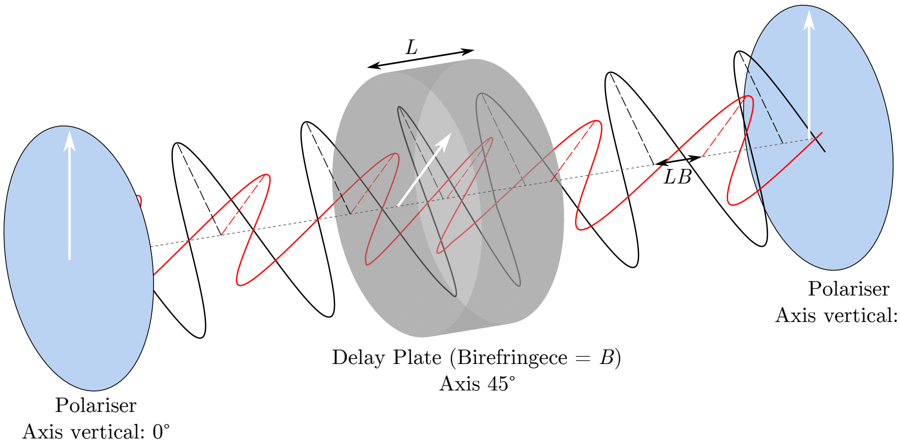

Figure 2.3: Illustration of a simple polarisation interferometer showing the ordinary (black) and extraordinary (red) wave components. The orientation of the polarisers and delay plate fast axis are shown by white arrows.

produce the interferometer delay.

2.5.1 Polarisation Interferometers

The simplest possible polarisation interferometer consists of two linear polarisers with their polarisation axes either parallel or perpendicular to each other, between which is placed a birefringent crystal plate (‘delay plate’) with its optic axis in the plane perpendicular to the optical axis and rotated 45° relative to the polariser axes. This configuration is shown in fig.2.3. The light to be measured is then collimated through this setup. The linearly polarised light transmitted by the first polariser is resolved into two equal amplitude, in-phase components in the delay plate, polarised parallel and perpendicular to the crystal optic axis and called the extraordinary and ordinary components, respectively. These two components experience different refractive indices in the crystal, n_e and n_o, causing a relative phase delay between the two components of:

$$\phi_0 = \frac{2\pi LB}{\lambda_0}, \tag{2.5.34}$$

where $B = n_\mathrm{e} - n_\mathrm{o}$ is the birefringence of the crystal and L is its thickness. The amount of light transmitted to the output by the second polariser depends on the overall polarisation state of the light after emerging from the crystal, which depends on the relative phase between the o and e components.

The action of this system can be analysed using Jones calculus [Hecht, 2002, pp. 377-379], in which the light propagating through the system is described by a 2-element Jones vector, composed of the electric field complex amplitudes for x and y oriented polarisation states:

$$\boldsymbol{E} = \begin{pmatrix} E_x(t) \\ E_y(t) \end{pmatrix}. \tag{2.5.35}$$

Each optical component affecting the polarisation is then represented by a 2x2 Jones matrix, which multiplies the Jones vector. The Jones matrix of a linear polariser is represented by diagonal coefficients h and v which denote the transmission for horizontal (x) and vertical (y) polarisations. A delay plate is represented by diagonal coefficients $e^{i\phi_0}$ and 1. Elements rotated about the optical axis are treated by multiplication of their Jones matrices by the usual 2D rotation matrices, which will be denoted by R_θ where θ is the angle of rotation in degrees. We assume randomly polarised input with $\langle E_x^\dagger E_x \rangle = \langle E_y^\dagger E_y \rangle = I_0/2$. The Jones vector at the output of our system is then given by:

$$\begin{aligned} \boldsymbol{E}_{out} &= \begin{pmatrix} h_2 & 0 \\ 0 & v_2 \end{pmatrix} \cdot \boldsymbol{R_{45}} \cdot \begin{pmatrix} e^{i\phi_0} & 0 \\ 0 & 1 \end{pmatrix} \cdot \boldsymbol{R_{-45}} \cdot \begin{pmatrix} h_1 & 0 \\ 0 & v_1 \end{pmatrix} \cdot \begin{pmatrix} E_x \\ E_y \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} E_x h_1 h_2 (e^{i\phi_0} + 1) + E_y v_1 h_2 (e^{i\phi_0} - 1) \\ E_x h_1 v_2 (e^{i\phi_0} - 1) + E_y v_1 v_2 (e^{i\phi_0} + 1) \end{pmatrix}, \end{aligned} \tag{2.5.36}$$

and the intensity at the output, given by $\left\langle \boldsymbol{E}_{out}^\dagger \boldsymbol{E}_{out} \right\rangle$, is:

$$I_{out} = \frac{I_0}{4}(h_1^2 + v_1^2)(h_2^2 + v_2^2)[1 + \zeta_P \,\mathrm{Re}(\gamma(\phi_0))]. \tag{2.5.37}$$

This is clearly a 2-beam interference pattern comparable to equation (2.2.8), with the contrast factor $|\zeta_P| \leq 1$ given by:

$$\zeta_P = \frac{(h_1^2 - v_1^2)(h_2^2 - v_2^2)}{(h_1^2 + v_1^2)(h_2^2 + v_2^2)}. \tag{2.5.38}$$

This depends on the extinction ratio of the polarisers, which is the ratio of the transmission of the selected and unwanted polarisation states. With the axes of the input and output polarisers parallel, if each has infinite extinction ratio (i.e. transmits only the desired h or v polarisation), $\zeta_P = 1$. With the input and output polarisers crossed, i.e. one vertical and one horizontal, $\zeta_P = -1$, which corresponds to a 180° phase shift of the interferogram. Note also that the overall intensity of the interferogram is (at best) half of that in (2.2.8) - this is a disadvantage of the polarisation interferometer approach since the first polariser rejects half the intensity of the original randomly polarised source. In some cases this rejected light has been put to use by using a polarising beamsplitter as the input polariser, and directing the rejected light to other instruments [Michael et al., 2001].

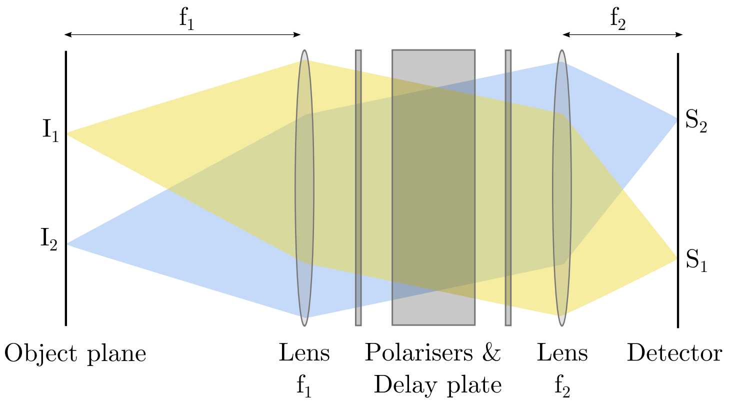

Figure 2.4: Schematic of an imaging polarisation interferometer. The interferograms recorded at positions S_1 and S_2 on the detector give information about the spectra emitted from points I_1 and I_2 in the object plane, respectively. Shaded areas indicate the ray bundles from each point, showing how light from each point is collimated and angularly multiplexed through the interferometer.

For the realistic case of non-ideal polarisers with finite extinction ratios, $|\zeta_P| < 1$ and the measured fringe contrast is reduced. In a real instrument, other imperfections e.g. in the lenses and birefringent components also cause further reduction of the fringe contrast. This is analogous to the instrument broadening in a dispersive spectrometer, since it causes apparent broadening of the spectrum due to imperfections in the optical system. Notice that because the CIS instrument operates in the time rather than the frequency domain, the instrument broadening appears as a multiplicative effect on the measured quantity rather than a convolution. The total reduction in measured contrast due to instrumental effects is described by a multiplicative factor $\zeta_{\mathrm{I}} < 1$, called the instrument contrast.

Imaging

In order to achieve 2D imaging with the type of polarisation interferometer discussed so far, light from different object positions is angularly multiplexed through the interferometer. The simplest configuration to achieve this consists of a lens one focal length away from the source to be imaged, with the interferometer placed behind it. The lens maps light from different points on the object to collimated beams traversing the interferometer at different angles. A second lens is then placed after the interferometer, and a 2D detector array one focal length behind that. The second lens re-images the object on to the detector, resulting in an image where each pixel gives the interferogram measurement for its corresponding point on the object. This setup is illustrated in figure 2.4.

A complication in this type of system is that the refractive index experienced by

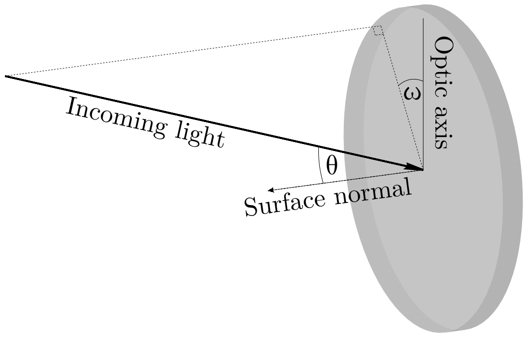

Figure 2.5: Illustration of angles used to describe the incidence direction of light on birefringent components.

the extraordinary ray in the delay plate depends on the angle of incidence. Therefore, light traversing the delay plate at different angles, i.e. light from different spatial points on the object, experiences slightly different phase delays. For light incident on the delay plate at a small oblique angle, the phase delay can be approximated as [Steel, 1983, p. 134],[Veiras et al., 2010]:

$$\phi_0(\theta,\omega) \approx \frac{2\pi LB}{\lambda_0}\left[1 - \frac{\theta^2}{2n_\mathrm{o}}\left(\frac{\cos^2(\omega)}{n_\mathrm{o}} - \frac{\sin^2(\omega)}{n_\mathrm{e}}\right)\right]. \tag{2.5.39}$$

In this expression the angle of incidence is described by θ, the angle between the plate's surface normal and the incident ray, and ω, the angle between the plane of incidence and the plate's optic axis. These coordinates are illustrated in figure 2.5. This angular dependence of the delay can be seen by observing the image at the detector in a setup such as figure 2.4, under narrow band illumination of the entire field of view. Instead of the flat, featureless image expected if the phase delay was the same at all image points, the angular dependence of the delay creates an unwanted parabolic fringe pattern, as shown in figure 2.6. Note that the spacing of the fringes decreases towards the edges of the field of view, i.e. with increasing incidence angle. For wide field of view instruments, this effect can lead to pixels near the image edges spanning a significant fraction of a fringe, and any interferogram measurement from such pixels will suffer from reduced fringe contrast and therefore reduced measurement performance. This ultimately limits the usable field of view of this simple instrument design. 'Field widening' techniques [Michael et al., 2001] can be used to greatly reduce the delay variation with incidence angle and overcome this limit, however these were not employed for the MAST instrument since instead the incidence angle was kept small by the design of the imaging optics (described in Chapter 3).

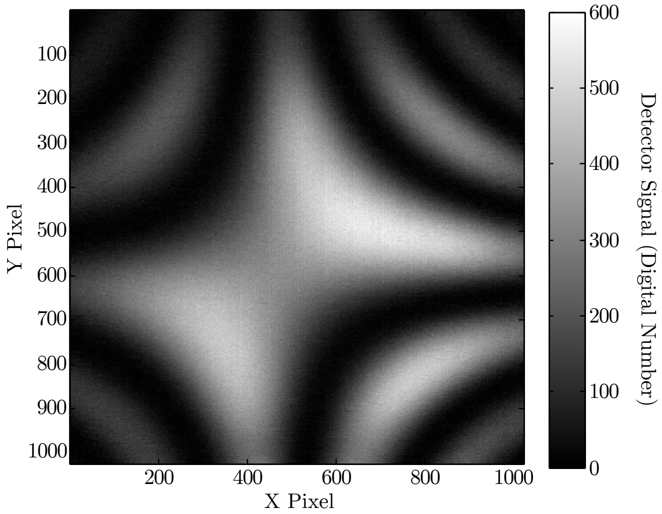

Figure 2.6: Example photograph of the parabolic fringe pattern from a simple polarisation interferometer configuration, taken using a delay plate from the MAST CIS system illuminated by HeNe laser light. The fringe pattern is caused by the angular dependence of the phase delay in the birefringent plate.

The simple imaging polarisation interferometer described so far gives a measurement of $S(\tau)$ over a 2D object, but only at a single value of τ for each image point: to make a spectral measurement for coherence imaging the delay must be scanned over at least 1 fringe. Several different instrument designs for achieving this are now reviewed, and as with broadband FTS, are divided into spatial and temporal multiplexing techniques. For more complete information on polarisation interferometer designs in general, see Françon and Mallick [1971].

2.5.2 Temporal Multiplexing Designs

The first coherence imaging instruments demonstrated on plasma devices used temporal multiplexing techniques. The most straightforward temporal multiplexing designs are based on using a delay plate made from a material which exhibits the linear electro-optic (Pockels) effect. The birefringence of such a material can be altered by application of an electric field, therefore by applying a time-varying voltage to the delay plate (typically of order kV) via attached electrodes the delay can be scanned in time over a small range without any need for mechanical scanning. Examples of instruments using this principle are the MOSS camera on the H1-Heliac [Michael et al., 2001], and a Doppler coherence imaging system used on the the WEGA stellarator [Chung et al., 2005]. The MOSS camera used a sinusoidal sweep of the delay at frequencies of 10s of kHz, and a multi-anode Photomultiplier tube (PMT) as the detector. Although this combination enabled high time resolution, PMT arrays only provide a limited number of detector elements, and the MOSS camera had only 16 spatial channels (the measurement was also only 1D in space due to restrictions of the viewing configuration). The WEGA instrument instead used a 2D CCD camera as the detector, allowing many more spatial channels at the cost of time resolution of the recorded signal: the camera gave 164x164 pixels at approximately 75Hz. In this case rather than being sinusoidally swept, the delay was stepped synchronously with the CCD exposure, providing one complete spectral measurement for every 3 camera frames (A minimum of 3 samples of the interferogram are required to recover the 3 unknowns: brightness, contrast and phase).

An alternative time multiplexing technique is based on using Ferroelectric liquid crystal (FLC) waveplates rather than electro-optic modulation, which removes the need for high voltage amplification and drive circuitry. The FLC plates act as delay plates whose optical axis orientation can be controlled by applying an external voltage. Half- and quarter-wave delay FLCs can be used together in addition to the main fixed delay plate [Howard, 2010] to quadrature sample the interferogram. With no applied voltage, the fast axes of the FLCs are oriented parallel to the main delay plate, and together contribute an extra 3/4 wave delay to the interferometer. With a small applied voltage (order 10V), their optical axes can be rotated through 45° to be parallel with the input and/or output polarisers, in which case they act equally on both polarisation components and do not contribute to the interferometer delay. By switching the two FLCs between these two states in sequence and synchronised with the camera, again a spectral measurement is obtained for every 3 camera frames.

An advantage of time multiplexed techniques is that each detector element records its own independent spectral information for a corresponding point in space, and therefore spatial information is recorded at the full resolution of the detector. However, when using CCD or CMOS imaging sensors to provide high spatial resolution, the time required to obtain a single spectral measurement (typically tens of milliseconds), although much shorter than in conventional FTS, can be long compared to the timescales of plasma phenomena which we may wish to study. This is a problem because changes in the spectrum over the duration of the measurement will produce erroneous results, and it limits the range of plasma phenomena which can be studied. Furthermore, design of the control systems for these instruments is complicated by the need to step the delay synchronously with the camera framing.

2.5.3 Spatial Multiplexing Designs

As was mentioned in section 2.3, since only a few samples of the interferogram are required for CIS measurements it is possible to design instruments where both the spectral and 2D spatial information are recorded simultaneously on a 2D detector. The time resolution is then only limited by the detector technology and the available photon flux. The earliest spatial multiplexing designs, an example of which was demonstrated on the H1 Heliac [Howard, 2006], produced 4 images at different delays side-by-side on the detector with a design based on crossed Wollaston prisms in addition to delay plates. The delay step between each image was 1/4 wave, and the interferogram properties were then obtained from sums and differences of the 4 images. In order to accurately interpret the data, the 4 images had to be carefully registered and any distortion, vignetting etc which could affect the 4 images differently had to be accounted for. The optical system was also relatively complex and required careful alignment, and the spatial resolution was reduced in both image directions by having to fit four images on a single detector.

More recent spatial multiplexing instruments, such as those used for imaging Charge exchange recombination spectroscopy (CXRS) on the TEXTOR tokamak [Howard et al., 2010b] and for flow imaging of the divertor plasma on the DIII-D tokamak [Howard et al., 2011; Weber et al., 2012], have used a ‘spatial heterodyne’ technique, which uses a simpler and more robust optical design. This is the scheme selected for the present work on MAST. In spatial heterodyne instruments the plasma image occupies the whole of the detector, while the delay is scanned by many waves along one dimension of the image (typically by about 1 wave per 10 detector pixels). With the delay scanned in the vertical (y) image direction, and for now ignoring angle of incidence effects of simplicity, the total instrument phase at a point x, y in the image is given by:

$$\phi_{\mathrm{I}} = \phi_0 + \phi_{\mathrm{M}} + \phi_{\mathrm{S}}(y), \tag{2.5.40}$$

where ϕ_0 is the fixed delay as before and $\phi_{\mathrm{S}}(y)$ is the additional phase delay due to the spatial scanning. With ϕ_{S} a linear function of y, this results in a set of horizontal fringes being superimposed on the plasma image, as shown in figure 2.7(a), which shows a raw data image from the MAST system. Since each column of pixels now contains a sinusoidal signal, Fourier or wavelet based processing techniques can be used to find the phase of the signal, i.e. the fringe phase, at each pixel. By subtracting the fringe phase from a suitable known wavelength calibration image, the Doppler phase shifts $\phi_{\mathrm{D}}(x, y)$ carrying the plasma flow information can be isolated.

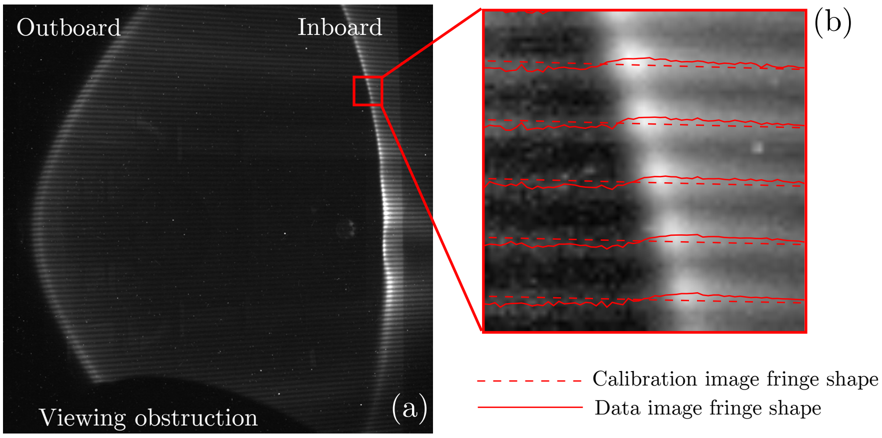

Figure 2.7: Illustration of the spatial heterodyne coherence imaging concept. (a) Raw image from the MAST CIS diagnostic, viewing the main chamber SOL in C III light. The image of the plasma can be seen with a set of horizontal fringes, which carry the spectral information, superimposed. (b) Magnified section showing the distortion of the horizontal fringe pattern due to Doppler shifts. The superimposed red lines are lines of constant fringe phase and trace the shape of the fringes in the plasma image shown (solid lines) compared to a calibration with no Doppler shift (dashed lines). See text for further explanation.

These phase shifts can be seen on the image as distortions of the horizontal fringe pattern, as shown in figure 2.7(b). This shows a close-up section of the raw image across the high field side SOL, and has lines of constant fringe phase (specifically $\phi \mod 2\pi = 0$) superimposed, which trace the shape of the fringes. The solid lines show the phase from the plasma data image while the dotted lines show the calibration; the fringes in the plasma image are seen to be distorted downwards inside the SOL (left of the image) and upwards outside the SOL (right of the image) with respect to the calibration, indicating flow in opposite directions in these two regions. The horizontal fringe pattern can be thought of as a carrier signal which enables the measurement, and is phase modulated by the Doppler phase and amplitude modulated by the emissivity and contrast. The spatial resolution of this type of system is anisotropic: in this example with horizontal fringes, each image column provides independent spectral information, however within an image column, since multiple pixels are required to obtain the 3 unknown local fringe properties (phase, contrast and brightness), the spatial resolution is reduced. It is therefore desirable to orient the fringes to be parallel with the steepest gradients expected in the plasma. Analysis of spatial heterodyne images, including the anisotropy of the spatial response, is discussed in more detail in Chapter 4.

The optical design of the spatial heterodyne system adds only a single compo-

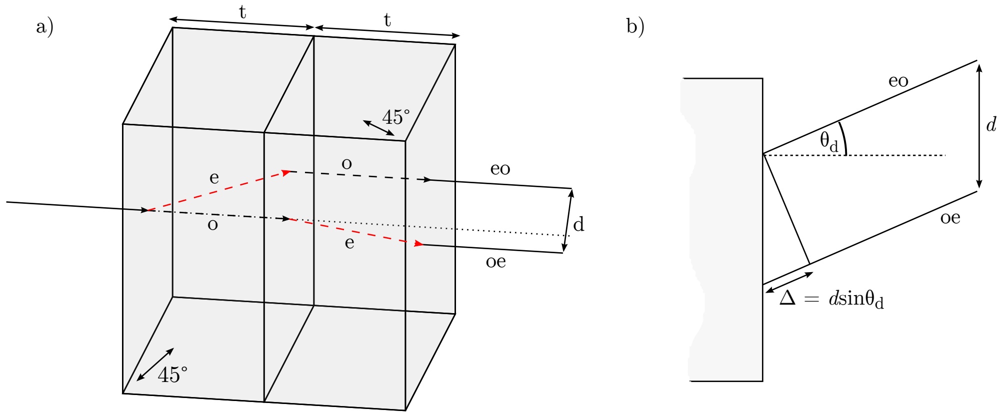

Figure 2.8: Construction and operation of a Savart polariscope. a) Construction and ray paths, showing the paths of the ordinary (black dashed line) and extraordinary (dashed red line) rays in each birefringent plate. The direction of the optical axis in each plate is shown by double-ended arrows. b) Mechanism of angle-dependent path difference generation due to beam separation. In b) the direction of the displacement d is in the plane of the page.

nent to the basic fixed delay polarisation interferometer from section 2.5.1: a Savart polariscope added after the main delay plate. This is a type of polarising beamsplitter [Françon and Mallick, 1971], consisting of two birefringent plates each cut with their optical axis at 45 degrees to the plate faces. For light incident on one such plate, the ordinary and extraordinary rays exit the crystal parallel to each other but laterally separated. As well as the separation, a phase delay is added between the two components due to the birefringence of the plate. In a Savart polariscope two identical plates of this type are cemented together, with the projections of their optic axes on the plate surfaces perpendicular to each other. This arrangement is shown in fig. 2.8(a). The ordinary ray in the first plate becomes the extraordinary ray in the second plate and vice versa, cancelling the phase delay due to the material birefringence. The second plate adds a lateral separation between the two rays at 90 degrees to that from the first, thus the total displacement d is $\sqrt{2}$ times that of a single plate, and is at 45 degrees to the vertical.

The Savart polariscope is inserted in the interferometer after the main delay plate, with the principal section of the first plate parallel to the optical axis of the main delay plate. This displaces the o and e components from the main delay plate in the vertical direction. This displacement results in a physical path difference $\Delta = d\sin(\theta_\mathrm{d})$ between the two beams, where θ_d is the incidence angle of the light on the Savart polariscope in the plane of the beam displacement. This is illustrated in fig. 2.8(b). Since different angles of incidence through the interferometer map to different image plane positions, this provides the required position-dependent delay sweep in the image, along the direction of the Savart plate's beam separation. The scale of the fringes on the detector is set by the Savart polariscope thickness and the focal length of the final imaging lens. Similarly to the operation of the delay plate for off-axis incidence, the displacement of the two beams has some dependence on the direction of the incident light, meaning the fringe pattern deviates from the ideal case of straight horizontal fringes. The phase delay between the *eo* and *oe* beams at the detector is in fact given by [Wu et al., 2007]:

$$\begin{aligned}\phi(\theta,\omega) = \frac{2\pi t}{\lambda}\Bigg[&\frac{n_\mathrm{o}^2 - n_\mathrm{e}^2}{n_\mathrm{o}^2 + n_\mathrm{e}^2}[\cos(\omega)+\sin(\omega)]\sin(\theta) \\ &+ \frac{1}{\sqrt{2}}\frac{n_\mathrm{o}}{n_\mathrm{e}}\frac{n_\mathrm{o}^2 - n_\mathrm{e}^2}{(n_\mathrm{o}^2 + n_\mathrm{e}^2)^{3/2}}[\cos^2(\omega)-\sin^2(\omega)]\sin^2(\theta)\Bigg],\end{aligned} \tag{2.5.41}$$

where t is defined as illustrated in figure 2.8. The coordinates used to describe the incidence direction are the same as for the delay plate, except that since the optic axis is no longer in the plane of the crystal face, ω is the angle between the plane of incidence and the *projection* of the optic axis of the front half of the polariscope on to its front face. The combined angle of incidence effects for the Savart polariscope and delay plate lead to a slightly varying fringe spacing across the image; this can be seen in figure 2.7(a) where the fringes are more closely spaced on the right-hand side of the image compared to the left.

2.6 Summary

In this chapter the principles, theory and existing implementations of the CIS technique have been reviewed. Impurity flows are measured using the Doppler shift of emission lines from partially ionised impurities, which exist in the SOL and divertor due to the relatively low plasma temperature. Various phenomena influencing the spectral line shapes have been briefly reviewed.

The CIS technique uses a 2-beam imaging polarisation interferometer to measure the spectral coherence of the impurity emission, at a particular value of time (or equivalently, phase) delay. It was shown in section 2.3 that the measured signal takes the form of a sinusoidal fringe pattern modulated by a slowly varying envelope. Measurements of a single fringe at a known value of τ can provide three pieces of information, given by the fringe brightness, phase and contrast. The relationship between these measured parameters and the spatially varying plasma flows and temperatures have been derived.

Previous implementations of CIS instruments have been reviewed, including temporal and spatial multiplexing systems. Temporal multiplexing systems potentially offer better spatial resolution at the cost of time resolution and instrument complexity, whereas spatial multiplexing instruments can offer better time resolution and potentially simpler instrument design at the cost of spatial resolution in at least one dimension. The spatial heterodyne technique was chosen for the MAST instrument, and the principles and construction of this type of instrument presented.

In the next chapter, we go on to present the design of a CIS flow diagnostic for MAST in detail.

Chapter 3

The MAST Coherence Imaging Diagnostic

In this chapter, the design & specifications of a CIS flow diagnostic for MAST are presented. We discuss the choice of spectral lines to be measured, as well as the choices and component specifications of the detector, imaging lenses, filters and interferometer components. The final design is a trade-off between optimising flow measurement performance and optimising the instrument cost within the finite available budget.

3.1 Candidate Spectral Lines

The most abundant impurity species naturally present in MAST plasmas are carbon and helium. This is because of the use of all-graphite Plasma-facing components (PFCs), and of helium glow discharge wall conditioning between each plasma shot. The brightest visible emission lines from these impurities, which were targeted for Doppler coherence imaging, are the C III triplet at 464.9nm, C II multiplet at 514.2nm, and He II multiplet at 468.6nm. The C II multiplet has 7 components, while the He II line has 13 components very closely spaced within a range of 0.055nm. The wavelengths, transitions, transition probabilities A_{ki} and statistical weights g_i and g_k for each multiplet component, obtained from the NIST atomic spectra database [Karmida et al., 2013], are given in tables 3.1, 3.2 and 3.3. To assess the approximate range of Doppler shifts exhibited by these lines, and therefore required to be measured by the CIS diagnostic, preliminary C II and C III divertor flow measurements were performed using a high resolution multi-chord spectrometer [Ham, 2012]. These measurements showed a range of line-average flows from almost zero up to around 30km/s for C III near the X-Point, and which were highly

C III ($1s^2 2s3p \rightarrow 2s^2 2s3s$)

λ_{rest} (nm)	Transition	A_{ki} (10^8 s^{-1})	g_k	g_i	$I_{\text{rel, calc}}$	$I_{\text{rel, obs}}$	$\sigma_{\text{I,fitting}}$	$\sigma_{\text{I,observed}}$
464.742	$^3\text{P}_2 \rightarrow ^3\text{S}_1$	0.726	5	3	0.556	0.543	±0.009	±0.010
465.025	$^3\text{P}_1 \rightarrow ^3\text{S}_1$	0.725	3	3	0.333	0.337	±0.009	±0.007
465.147	$^3\text{P}_0 \rightarrow ^3\text{S}_1$	0.724	1	3	0.111	0.120	±0.007	±0.006

Table 3.1: Wavelengths, transitions, transition probabilities, statistical weights (upper level k, lower i), calculated and measured relative intensities and associated errors for components of the C III 464.9nm spectral line. Fitting errors $\sigma_{\text{I,fitting}}$ are typical values for a single measurement, while observed variations $\sigma_{\text{I,observed}}$ are statistical variations of the observations over a range of plasma shots and spectrometer views. Transition data from Karmida et al. [2013]

dependent on position along the 1D fan of sight-lines used.

Since the Doppler CIS technique provides only 3 scalar parameters describing the spectrum for each spatial point (the brightness, fringe contrast and fringe phase), not enough information is obtained to distinguish any spectral contamination in the instrument's passband from the desired signal. Such contamination can include both line emission from other species and broadband continuum light, e.g. from Bremsstrahlung, and can lead to misinterpretation of the 3 measured quantities. We must therefore ensure that the lines of interest can be isolated from the plasma spectrum (using suitable bandpass filters) without significant contamination. Also as we have seen in section 2.4.1, the multiplet structure of the lines must be well characterised and constant in time and space for accurate interpretation of the measurements. To assess whether these conditions are met for the candidate lines on MAST, the visible spectrum near 465 - 468nm and 514nm was investigated using high resolution dispersive spectrographs. Measurements were made with 4 different plasma views: midplane, divertor floor, and wide and narrow angle lower X-point views. The spectrometer views and representative spectra are shown in Figure 3.1. Each spectrum has been normalised such that the peak intensity in the line of interest has a value of 1, since we are concerned with the relative level of contaminating features compared to the line of interest. The most notable contamination is of the He II line at the divertor floor, which shows a significant background including very crowded line structures. This is attributed to emission from atomic and molecular processes in the cold divertor plasma at or close to the divertor strike point, which is significant relative to the He II line because of the weakness of the line itself (this background is also visible on the C III measurements only 3nm away, but that line is sufficiently bright that the background is insignificant). This is expected to greatly reduce the fringe contrast and distort the fringe phase, making accurate

C II ($2s2p(^3\mathrm{P}_0)3p \rightarrow 2s2p(^3\mathrm{P}_0)3s$)

λ_{rest} (nm)	Transition	A_{ki} (10^8 s^{-1})	g_k	g_i	$I_{\text{rel, calc}}$	$I_{\text{rel, obs}}$	$\sigma_{\mathrm{I,fitting}}$	$\sigma_{\mathrm{I,observed}}$
513.295†	$^4\mathrm{P}_{\frac{3}{2}} \rightarrow {}^4\mathrm{P}_{\frac{1}{2}}$	0.389	4	2	0.140	0.289	±0.003	±0.007
513.328†	$^4\mathrm{P}_{\frac{5}{2}} \rightarrow {}^4\mathrm{P}_{\frac{3}{2}}$	0.280	6	4	0.151			
513.726	$^4\mathrm{P}_{\frac{1}{2}} \rightarrow {}^4\mathrm{P}_{\frac{1}{2}}$	0.155	2	2	0.029	0.026	±0.003	±0.003
513.917	$^4\mathrm{P}_{\frac{3}{2}} \rightarrow {}^4\mathrm{P}_{\frac{3}{2}}$	0.124	4	4	0.045	0.042	±0.003	±0.004
514.350	$^4\mathrm{P}_{\frac{1}{2}} \rightarrow {}^4\mathrm{P}_{\frac{3}{2}}$	0.773	2	4	0.139	0.142	±0.003	±0.006
514.517	$^4\mathrm{P}_{\frac{5}{2}} \rightarrow {}^4\mathrm{P}_{\frac{5}{2}}$	0.649	6	6	0.349	0.344	±0.003	±0.015
515.109	$^4\mathrm{P}_{\frac{3}{2}} \rightarrow {}^4\mathrm{P}_{\frac{5}{2}}$	0.416	4	6	0.149	0.157	±0.003	±0.004

Table 3.2: Wavelengths, transitions, transition probabilities, statistical weights (upper level k, lower i), calculated and measured relative intensities and associated errors for components of the C II 514.2nm spectral line. Fitting errors $\sigma_{\mathrm{I,fitting}}$ are typical values for a single measurement, while observed variations $\sigma_{\mathrm{I,observed}}$ are statistical variations of the observations over a range of plasma shots and spectrometer views. †The two shortest wavelength components are combined in measured quantities since they are too closely spaced to be resolved experimentally. Transition data from Karmida et al. [2013]

Doppler imaging of the He II line near the divertor strike points unfeasible. The wide-angle X-point view also shows the same background structure, typically at a lower level than at the divertor floor but highly variable in time. This time variation is due to changes in the He II line brightness while the background stays relatively constant. From the limited amount of data collected with this view it is difficult to assess how large a problem this is generally, or precisely what part of the plasma is affected. Both the divertor floor and wide angle X-point views also show a bright neutral helium line at 471.3nm, however this is well enough separated from He II not to be of concern. The midplane and narrow angle X-point views do not suffer significant contamination of the He II line, and show a very clean spectrum suitable for coherence imaging.

Due to limited availability of the spectrometers used, the C II line was only measured at the narrow angle divertor view. This consistently shows a very clean spectrum, with only a small molecular carbon emission feature at 516.5nm.

The C III measurements show a number of oxygen (O II) emission line components very close to the CIII line, including one within the C III multiplet itself at 464.9nm. These are strongest at the divertor floor and wide-angle X-point views, however even in the worst case the intensity in the 464.9nm O II line is $< 5\%$ of that in the C III multiplet. The magnitude of the error on C III flow measurements due to this contamination is expected to be comparable to the flow measurement noise floor, and is discussed in more detail in section 4.5.

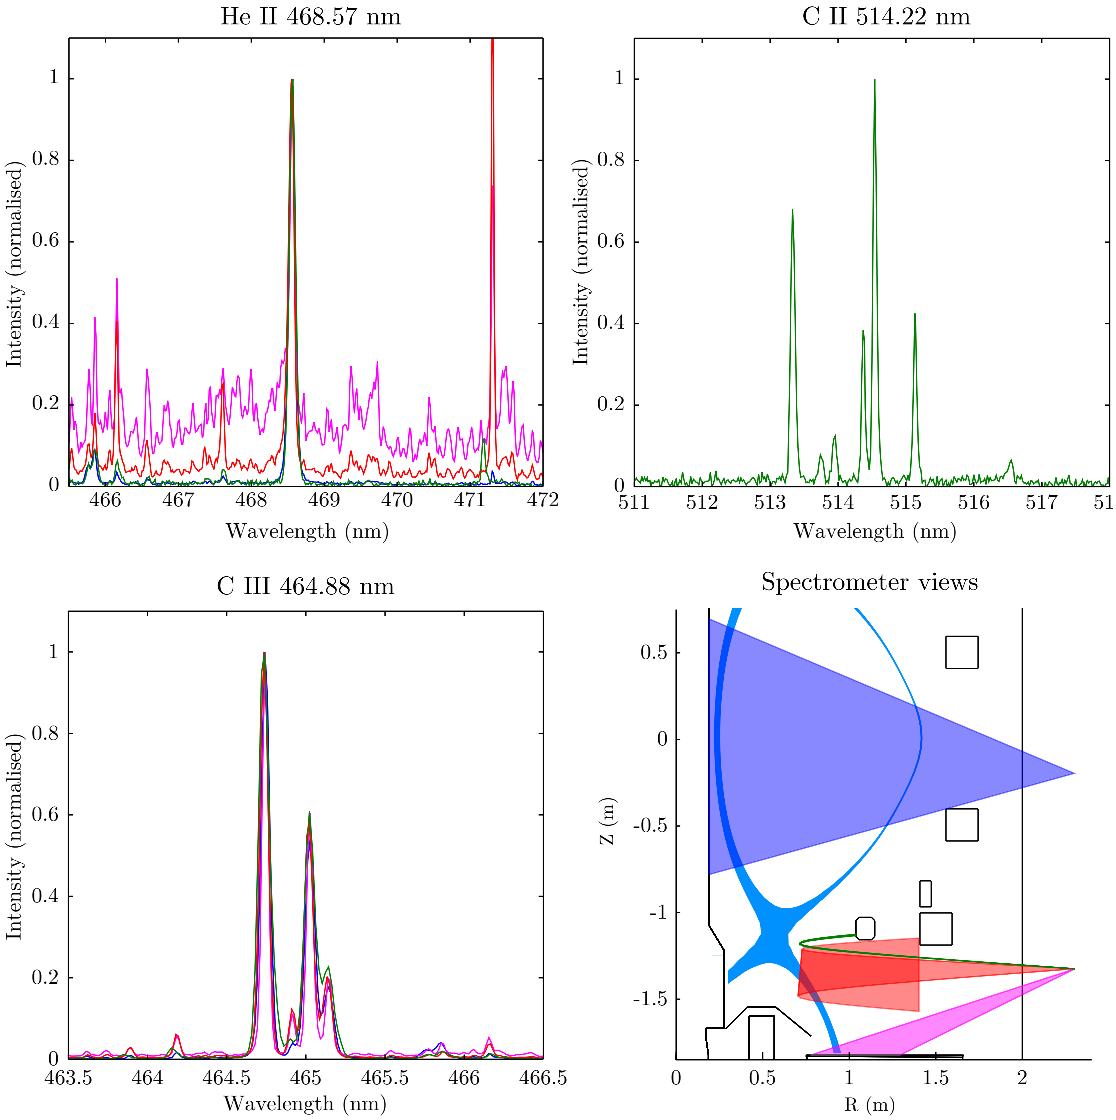

Figure 3.1: Measured spectra of the He II 468.6nm, C II 514.2nm and C III 464.9nm emission lines on MAST, at the midplane (blue lines), outer divertor target (pink lines), and near the X-point with wide and narrow angle views (red and green lines respectively). Spectra are normalised so that the peak of the line of interest has an intensity value of 1. The bottom-right panel shows the spectrometer views projected on to the poloidal cross-section of MAST, colour coded to correspond to the plotted spectra. Narrow angle X-point data were collected by Chiel Ham and James Harrison [Ham, 2012].

He II ($n = 4 \rightarrow n = 3$)

λ_{rest} (nm)	Transition	A_{ki} (10^8 s^{-1})	g_k	g_i	$I_{\text{rel, calc}}$
468.537	$4d\,^2\mathrm{D}_{\frac{3}{2}} \rightarrow 3p\,^2\mathrm{P}_{\frac{1}{2}}$	0.939	4	2	0.082
468.541	$4p\,^2\mathrm{P}_{\frac{3}{2}} \rightarrow 3s\,^2\mathrm{S}_{\frac{1}{2}}$	0.491	4	2	0.043
468.552	$4s\,^2\mathrm{S}_{\frac{1}{2}} \rightarrow 3p\,^2\mathrm{P}_{\frac{1}{2}}$	0.098	2	2	0.004
468.557	$4p\,^2\mathrm{P}_{\frac{1}{2}} \rightarrow 3s\,^2\mathrm{S}_{\frac{1}{2}}$	0.491	2	2	0.021
468.570	$4f\,^2\mathrm{F}_{\frac{5}{2}} \rightarrow 3d\,^2\mathrm{D}_{\frac{3}{2}}$	2.060	6	4	0.269
468.570	$4d\,^2\mathrm{D}_{\frac{5}{2}} \rightarrow 3p\,^2\mathrm{P}_{\frac{3}{2}}$	1.127	6	4	0.147
468.576	$4p\,^2\mathrm{P}_{\frac{3}{2}} \rightarrow 3d\,^2\mathrm{D}_{\frac{3}{2}}$	0.006	4	4	0.001
468.576	$4d\,^2\mathrm{D}_{\frac{3}{2}} \rightarrow 3p\,^2\mathrm{P}_{\frac{3}{2}}$	0.188	4	4	0.016
468.580	$4f\,^2\mathrm{F}_{\frac{7}{2}} \rightarrow 3d\,^2\mathrm{D}_{\frac{5}{2}}$	2.207	7	6	0.384
468.583	$4f\,^2\mathrm{F}_{\frac{5}{2}} \rightarrow 3d\,^2\mathrm{D}_{\frac{5}{2}}$	0.147	6	6	0.019
468.588	$4p\,^2\mathrm{P}_{\frac{3}{2}} \rightarrow 3d\,^2\mathrm{D}_{\frac{5}{2}}$	0.050	4	6	0.004
468.591	$4s\,^2\mathrm{S}_{\frac{1}{2}} \rightarrow 3p\,^2\mathrm{P}_{\frac{3}{2}}$	0.196	2	4	0.009
468.592	$4p\,^2\mathrm{P}_{\frac{1}{2}} \rightarrow 3d\,^2\mathrm{D}_{\frac{3}{2}}$	0.056	2	4	0.002

Table 3.3: Wavelengths, transitions, transition probabilities, statistical weights (upper level k, lower i), and calculated relative intensities for the He II 468.6nm line. Transition data from Karmida et al. [2013]

For both carbon lines, Gaussian peak fitting to each multiplet component was used to determine the relative component intensities. Here these are defined as the fraction of the total light in the multiplet in each component. Data from all of the spectrometer views and from multiple plasma shots were included, to assess any variation of the component ratios. The resulting relative intensities are shown in Tables 3.1 and 3.2. These are compared with the expected relative intensities calculated by assuming the upper states are populated according to their statistical weights, giving the ratio of two line components as [Chung, 2004, p. 23]:

$$\frac{I}{I'} = \frac{A_{ki} g_k \lambda'}{A'_{ki} g'_k \lambda}. \tag{3.1.1}$$

The measured and calculated relative intensities are found to be in very good agreement, and variations observed over different spectrometer views and shots were very small - in most cases very close to the typical fitting error. This is a good indication that the multiplet structure of these lines is as expected and can be considered constant. This analysis was not carried out for the He II line since the very fine multiplet structure was not resolved.

Overall the C III and C II lines on MAST are good candidates for Doppler coherence imaging. The He II line at the midplane is also well suited, however in the divertor care must be taken with this line, particularly at the strike points where accurate Doppler measurements are not expected to be possible due to background contamination.

3.2 Design Overview

The MAST coherence imaging system is based on the spatial heterodyne scheme described in Section 2.5.3. This allows snapshot imaging at high time resolution whilst maintaining a very simple and robust optical design. A schematic of the optical system is shown in Figure 3.2. The first optical element l_1 is a commercial camera zoom lens with focal length f_1 adjustable between 17-70mm at f/2.8-4.5, which produces a real demagnified image of the plasma. This intermediate image is placed at the film/sensor plane of a second camera lens l_2, with $f_2 = 105$mm at f/2.8, and with its focus set to infinity. This arrangement collimates the light through the interferometer components, with light from different points in the field of view traversing the interferometer at different angles. Using this two-lens configuration rather than a single lens acts to reduce the beam angles through the interferometer for a given field of view, allowing good control of field effects in the birefringent crystals and interference filters even for wide fields. It also allows trivial adjustment of focus and field of view using the normal focus and zoom controls of l_1. The choice of specifications for the lenses used is discussed in detail in section 3.4.

In the collimated region after l_2 is a temperature stabilised enclosure housing the interferometer components: a linear polariser followed by an Alpha barium borate (BaB_2O_4) (α-BBO) delay plate to provide the interferometer fixed delay, and a Savart polariscope to create the delay scan across the detector, followed a second linear polariser with its axis parallel to the first. The delay plate and Savart polariscope optic axes are parallel with each other and at 45° to the polariser axes. A third camera lens l_3 with $f_3 = 150$mm at f/2.8 and with its focus set to infinity forms the final image on the detector: a 1024x1024 pixel CMOS camera capable of frame rates up to 3kHz. An interchangeable 50mm diameter band pass interference filter is mounted on the front of the final lens, to isolate the emission line of interest from the plasma spectrum. The optical system is mounted directly outside the MAST vacuum vessel, and views the plasma through diagnostic ports at either the lower divertor or midplane. The field of view of the instrument for the widest zoom setting of $f_1 = 17$mm is 39°, and can be adjusted down to 9° at $f_1 = 70$mm. At a viewing distance of approximately 2.2m where most light is collected on MAST, this corresponds to pixel scales in the plasma between 1.5mm (wide angle) - 0.4mm (narrow angle), and spatial resolution perpendicular to the fringes of approximately

20mm (wide angle) - 5mm (narrow angle).

In the following sections various aspects of the instrument specifications and design are discussed in more detail.

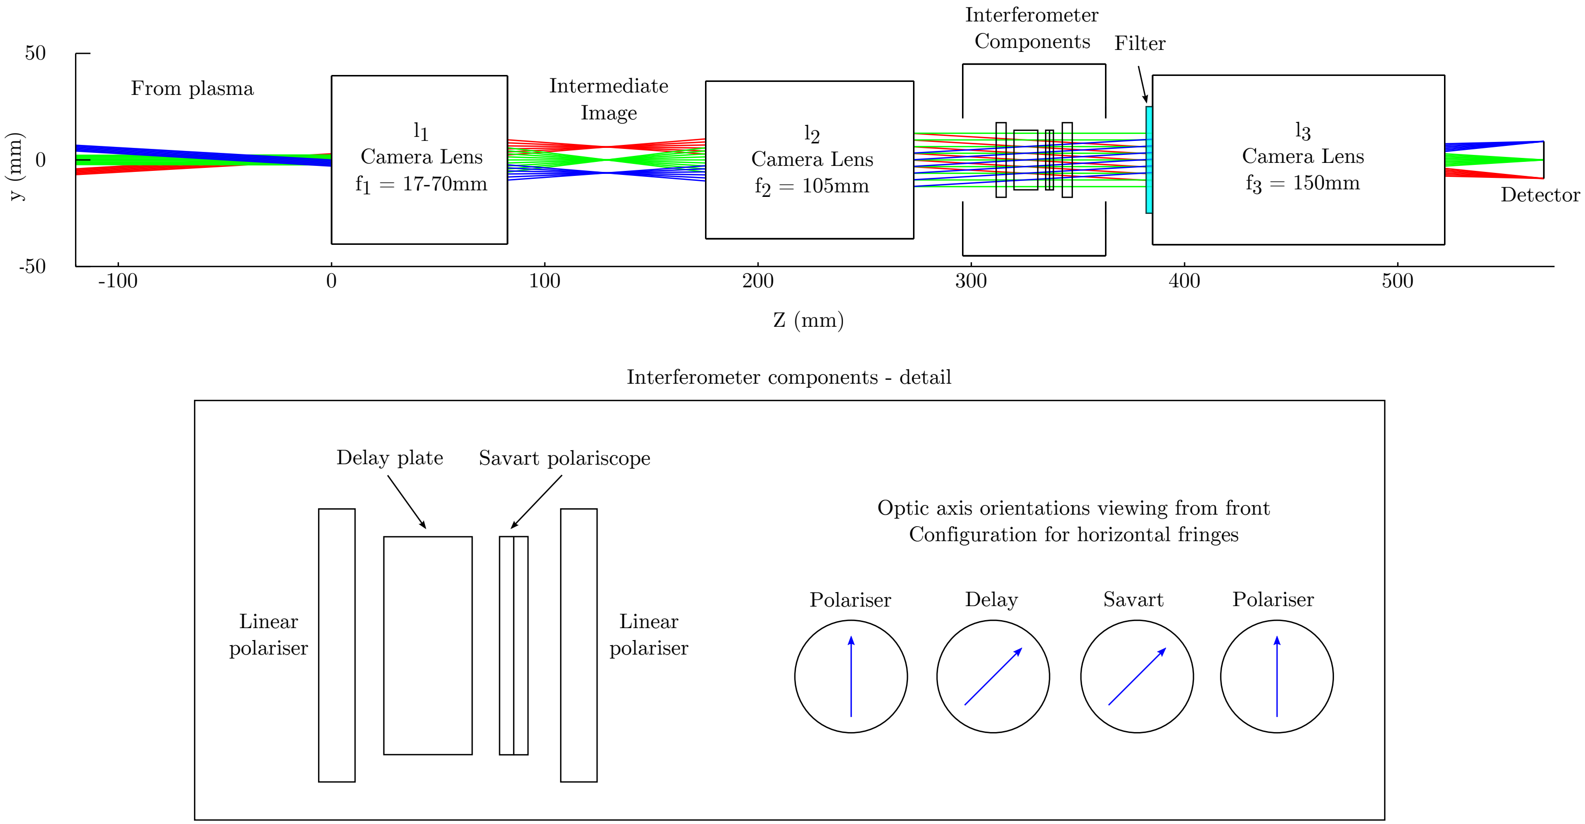

Figure 3.2: Schematic showing the layout of the MAST coherence imaging system.

3.3 Detector

The detector used was a Photron APX-RS high speed video camera, on loan from Rutherford Appleton Laboratories. This uses a proprietary CMOS image sensor with 1024×1024 pixels and $17\mu m$ pixel size to record 10-bit monochrome images, and is capable of a maximum frame rate of 3kHz when reading out the full detector. Images are captured to a 2.6GB buffer onboard the camera, then transferred to the controlling computer via firewire after recording has ended. This camera was chosen primarily for its high frame rate capability, to take advantage of the high light flux available from viewing the plasma directly. In the spectral region of interest the camera has a quantum efficiency of approximately 45% according to the manufacturer's specifications. Since this camera had also been operated on MAST previously for fast visible imaging, it was proven to operate under the magnetic field and radiation conditions adjacent to the MAST machine and was already integrated into the diagnostic triggering and data acquisition systems.

3.4 Imaging Optics

The specifications of the lenses & bandpass filters are discussed here together since they are inter-dependent. This is because the filters must cleanly isolate the spectral lines of interest over the entire image, meaning they must operate satisfactorily at a range of incidence angles, defined by the specifications of the lenses. This must be achieved while also utilising the full available area of the detector, maximising throughput and minimising vignetting, and obtaining the required field of view of $> 35°$ in the plasma. The final choice of lens models and filter specifications were found to provide a good compromise between these concerns, with the zoom lens providing additional flexibility in the field of view. The lenses used were all commercial still photography lenses designed for F-Mount Digital single lens reflex (DSLR) cameras. These were chosen for their ability to cast a large image circle sufficient to fill the image sensor of the Photron camera, which has a diagonal sensor size similar to APS-C format DSLR cameras, and because of their optimisation for high quality wide angle imaging. The Photron APX-RS camera also features an F-Mount interface as standard, allowing trivial interfacing of the final lens with the camera. The line selection filters were 3-cavity interference filters, the specifications of which are discussed in section 3.4.2.

In the following sections specific performance considerations for the imaging system are described in detail. First, the model forming the basis of the design and performance calculations is described.

3.4.1 Imaging optics model

Since the lenses used in the system are commercial photography products, their detailed optical prescriptions were not readily available to produce a full physical model of the imaging system. However the manufacturer of the lenses, Sigma Imaging, was able to provide the locations of the entrance and exit pupils and nodal points for various lenses under consideration, and from this it was possible to produce the estimates for the vignetting, throughput & filter performance discussed in the following sections. The model essentially treats each camera lens separately with a thick lens model moving front-to-back through the system. The geometry of this model is illustrated in figure 3.3, which shows the plasma-facing lens and intermediate image region of the system with light from the edge of the field of view. For each lens the distances of the entrance and exit pupils from the image plane, l_{ent} and l_{ex} in Fig.3.3, were provided by the manufacturer. The image plane is a fixed distance from the lens body given by the flange-to-focal distance for F-Mount lenses of 46.5mm. The entrance pupil size r_{ent} is given by the focal ratio and focal length of the lens, and the exit pupil size r_{ex} is determined from this and the pupil magnification. The pupil magnification was calculated from the pupil and nodal point locations [Jenkins and White, 1981], except for the 17-70mm lens where the nodal point locations were not available. In this case, the value of pupil magnification was estimated by visual inspection of the lens and noting that the exit pupil is just not vignetted by the rear element of the lens for an on-axis source.

For each lens, illumination of the entrance pupil is determined either by the results from the previous lens, or for l_1 assumed to be uniform across the entire pupil. Based on this 2D illumination pattern, a bundle of rays is launched from the exit pupil through a perfect geometrical focus at the image plane (or perfectly collimated in the case of the collimated region). The image height y_{i} is determined from the thin lens equation. The rays are checked for vignetting at each aperture in the system and discounted accordingly, as shown in Fig.3.3 where vignetted rays are shown in grey. The locations of the remaining rays reaching the entrance pupil of the next lens then define the illumination pattern used to model that lens in the same manner. The intensity lost due to vignetting in the complete system is calculated by comparing the starting number of rays initialised at the entrance pupil of l_1 with the number which successfully reach the detector. The results from this model are used in the filter performance and étendue calculations presented in sections 3.4.2 and 3.4.3.

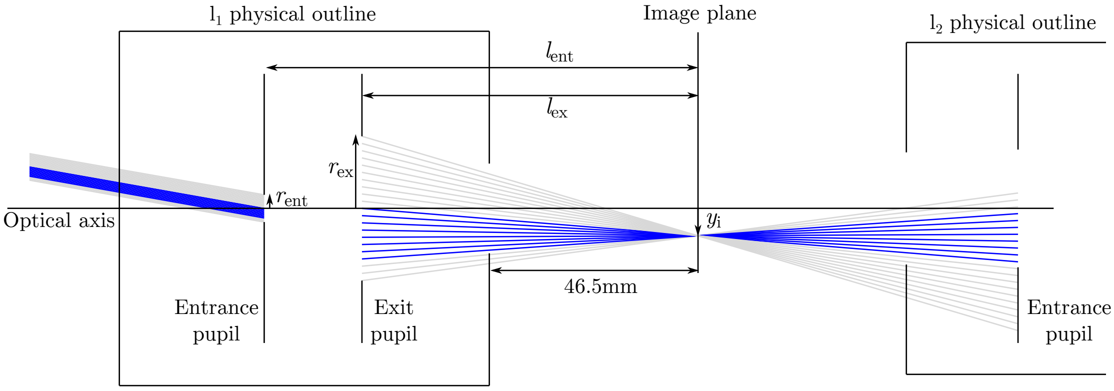

Figure 3.3: Illustration of the model geometry used to determine vignetting and other optical parameters of the diagnostic, illustrating modelling of the light through the intermediate image plane of the system. Rays shown in blue are accepted by the second lens while rays shown in grey are vignetted.

3.4.2 Filter Performance

The bandpass filters used for spectral line selection were 3-cavity interference filters. These are constructed from stacks of dielectric layers of different refractive indices, to form back-to-back optical cavities in which interference between multiple reflections creates a narrow band pass profile. Triple cavity designs such as those used here create a filter profile with a flat top and steep edges. The main performance concern when using such filters in an imaging system is that the peak transmission wavelength is blue-shifted for light incident on the filter at an oblique angle. The shifted bandpass centre λ_θ for an incidence angle θ (measured from the normal to the filter surface) is related to the value at normal incidence λ_0 by [Baillard et al., 2006]:

$$\lambda_\theta = \lambda_0 \sqrt{1 - \frac{\sin^2\theta}{n*^2}}, \tag{3.4.2}$$

where $n*$ is the effective refractive index of the filter. There are three main negative effects the filter shift can have for our application: 1) Vignetting of the image as the band pass is shifted away from the line of interest towards the image edges, 2) Induced changes in the line shape or multiplet component ratios across the image due to the slopes of the filter profile being shifted across the multiplet, and 3) Admission of unwanted spectral features in parts of the image as the band pass shifts. The combination of filter centre wavelength, bandwidth, band pass shape, $n*$ and range of incidence angles determines the size of these effects. Of these, the band pass shape and $n*$ were essentially fixed by the type of filters used, while the rest were free parameters in the design (the range of incidence angles being set by the choice of lenses).

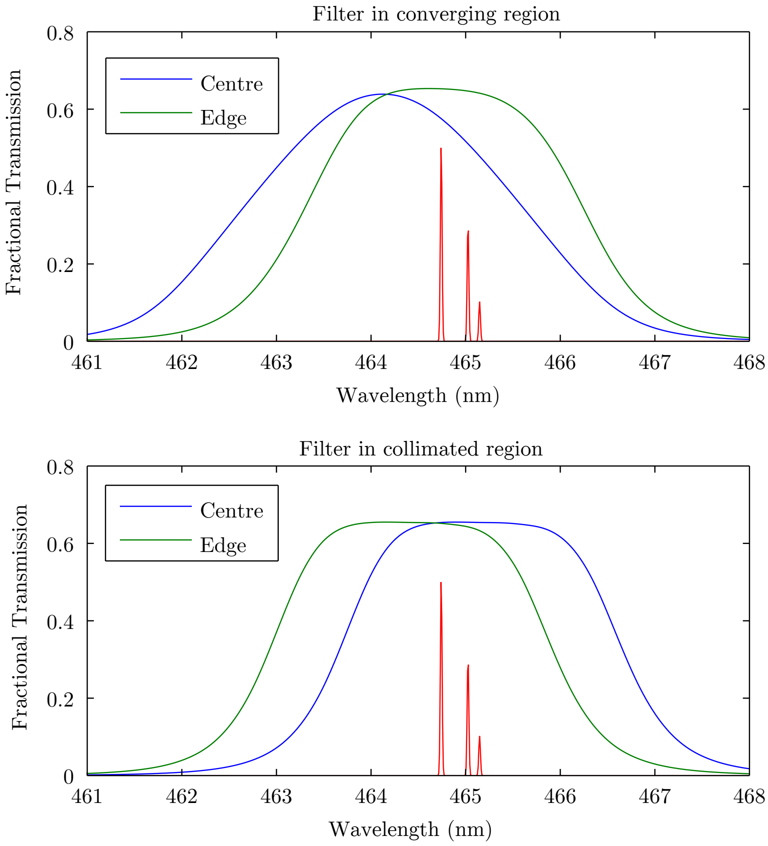

Figure 3.4: Estimated filter transmission profiles for a CIII filter, at the image centre and edges, with the filter placed in the converging (top) and collimated (bottom) regions. The structure of the C III multiplet is shown in red.

Two possible filter locations were considered for the system: near the intermediate image plane in the converging beam region between l_1 and l_2, or in the collimated region between l_2 and l_3 along with the interferometer components. To assess the performance of each option, filter transmission curves were estimated for both cases, starting from a theoretical transmission curve for normal incidence provided by the filter manufacturer, Andover Corp. The calculations were performed for points at the centre and extreme edge of the field of view to assess the uniformity over the image. In the collimated region, profiles for the image edge were estimated as a shift of the provided normal incidence profile by an amount given by equation (3.4.2), using

$$\theta_{\text{edge}} = \tan^{-1}\left(\frac{d}{2f_3}\right) \tag{3.4.3}$$

at the image edge where d is the diagonal sensor size. This neglects broadening of the band pass shape which occurs for oblique incidence, however for the small beam angles used here this is not expected to be a significant effect. For the converging beam region, the model described in section 3.4.1 was used to estimate the range of ray directions passing through the filter at the image centre and edge, and the total bandpass shape was calculated by summing over the shifted profiles for each ray. The resulting filter profiles for C III and $f_3 = 150$mm are shown in figure 3.4. For the converging beam case at the centre of the field of view, the range of ray angles causes the flat top of the pass band to become significantly rounded and the edges to become very shallow compared to the collimated case. Note that the different C III multiplet components experience different transmission, which changes significantly between the centre and edge of the field of view. The resulting apparent line structure change between different image points would be difficult to accurately account for in the data analysis. The rounding effect of the band shape is less pronounced at the edge of the field of view, due to a reduction of the range of ray angles through the filter because of vignetting. Due to the large variation in the profile shape over the image and the rounding of the filter profile, this option was deemed unsuitable and the filter was placed in the collimated region. As can be seen in figure 3.4, for this configuration the entire C III multiplet can be kept in the flat top of the filter over the whole field of view. Having the filter attached to the final lens is also advantageous since it acts to reject ambient light, and makes access for interchanging the filter much easier.

The filters for the MAST system were 'semi-custom' components, meaning the centre wavelength and FWHM could be freely specified. To determine the optimum values, the filter performance was assessed over a 2D parameter space defined by the filter centre and FWHM. The criteria for acceptable performance were chosen to be: (1) Transmission of the line of interest at the field of view edge must be $> 80\%$ of that at the image centre, (2) Distortions of the fringe phase between the image centre and edge, due to changing ratios of multiplet components or spectral contamination, must be equivalent to a flow of < 1km/s and (3) Admission of significant nearby spectral lines must be $< 2\%$ across the image. Fulfilment of these criteria was assessed by taking transmission profiles for the image centre and edge, calculated as above, and multiplying these by a (non Doppler shifted) model spectrum based on the results in section 3.1. The intensity and fringe phase for the resulting spectra were then calculated at the image centre and edge and assessed against the performance criteria. This was performed over the 2D filter parameter space using shifts and width scaling of the theoretical transmission profile. The final filter specifications were fixed by positioning a 2D box representing the filter manufacturing tolerances in the parameter space, such that the area inside the box

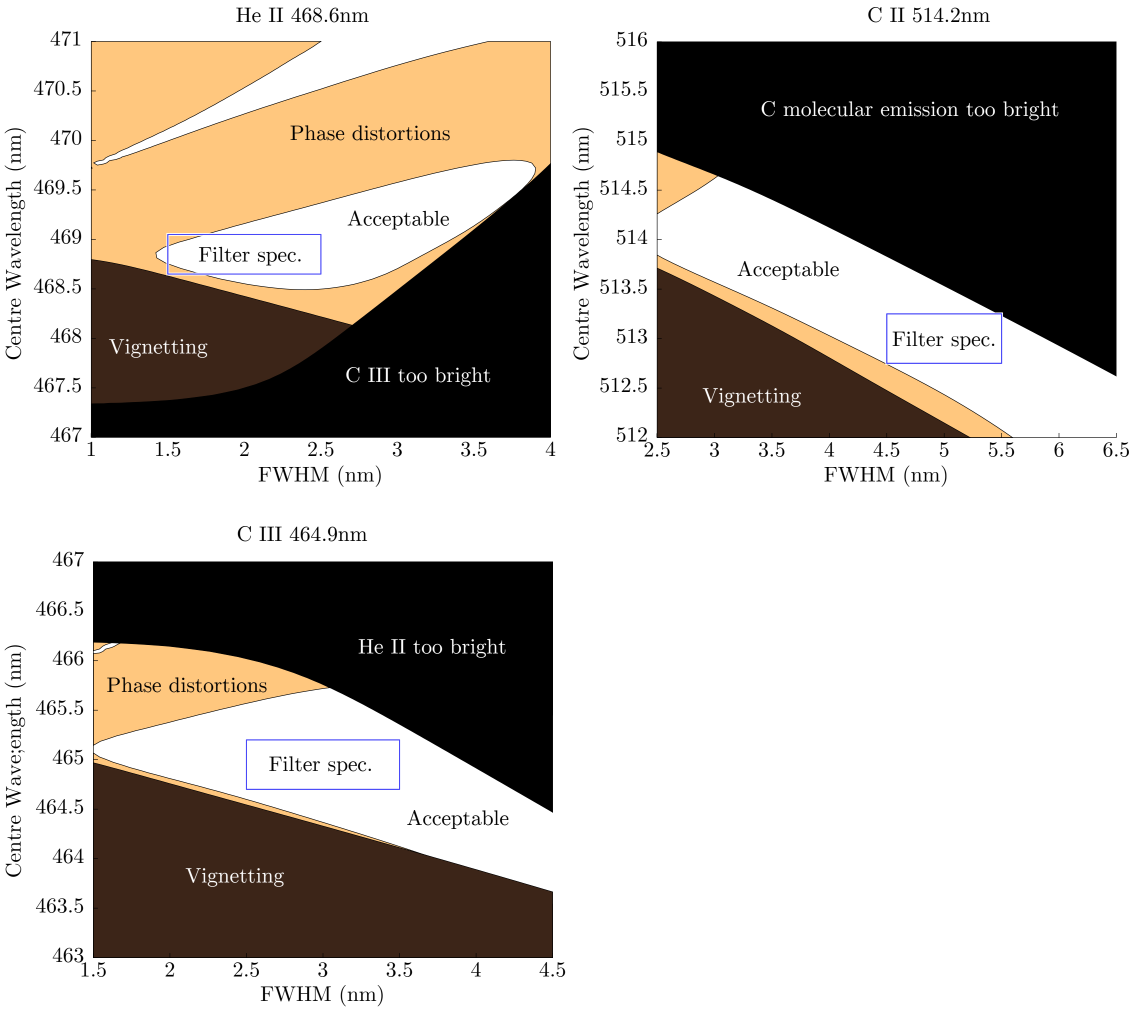

Figure 3.5: Contour plots illustrating the bandpass filter specifications in a parameter space defined by centre wavelength and Full width at half maximum (FWHM). Unshaded regions of the parameter space give acceptable filter performance, and the final filter specifications are shown as blue rectangles.

satisfied the performance criteria. These calculations also contributed to the lens choices, since the chosen lens combination was required to give a usable area of the parameter space large enough to encompass the filter manufacturing tolerances. Plots of the filter parameter spaces for each spectral line with $f_3 = 150$mm are shown in figure 3.5, where the white regions represent areas of acceptable performance and the boxes represent the manufacturing tolerances of the final filter specifications $(+0.5/-0$nm in centre wavelength and $\pm$ 0.5nm in FWHM$)$. It was not possible in all cases to completely avoid regions of poor performance, because the filter FWHM could only be specified in 0.5nm increments. Details of the specifications for each filter are included in Appendix A.

3.4.3 Collecting Power & Vignetting

It is clearly desirable to maximise the light collecting power of the optical system, in order to allow the shortest possible exposure times and therefore maximum frame rates and to maximise the image SNR. It is also desirable to achieve uniform performance across the entire image, i.e. to minimise vignetting.

The light collecting ability of any optical system can be quantified using the étendue, or geometrical extent, of the collected light, which is a measure of the spatial and angular extent of the light which can be collected by the system. It is given numerically by $E = A\Omega$, where A is the area of the light source and Ω is the solid angle subtended by the optical system's entrance pupil, as viewed from the source. The product of the source radiance and étendue give an upper limit for the amount of power transferred from the source to the detector by the system. Étendue is a conserved quantity both for light propagating in free space and in ideal reflection and refraction, and so is conserved throughout a lossless optical system.

In order to estimate the per-pixel throughput of the MAST coherence imaging system over the whole field of view, first the étendue of the light collected by l_1 in isolation was calculated. This is given by $E_{\mathrm{l}_1} = A_{\mathrm{pix}}\Omega_{\mathrm{l}_1}$, where A_{pix} is the area of one detector pixel projected into the plasma, and Ω_{l_1} is the solid angle subtended by a circular aperture the size of l_1's entrance pupil as seen from the plasma. The fraction of this light reaching the detector, F_T, was then estimated using the model described in section 3.4.1. The overall per-pixel effective étendue is then given by $E_{\mathrm{eff}} = E_{\mathrm{l}_1}F_T$. This was calculated for the various combinations of lenses under consideration, for image locations from the detector centre to the far corners, and the final selection of lenses was made based on a trade-off between effective étendue at the image centre, minimising vignetting, and meeting the filter performance criteria discussed in the precious section. The results of the effective étendue calculation for the final system design are shown in figure 3.6, for the extreme zoom settings of l_1. Excluding the image corners, the illumination at the edge of the detector is between 55 - 60% of that at the centre. The difference between the widest and narrowest zoom settings is also relatively small, with the effective étendue at $\mathrm{f}_1 = 70$mm approximately 80% or more of that for $\mathrm{f}_1 = 17$mm over the whole image. Note that this estimated collecting power does not take into account the losses due to the polarisers used in the interferometer or losses due to reflections etc. These results are tested against experiment in chapter 5.

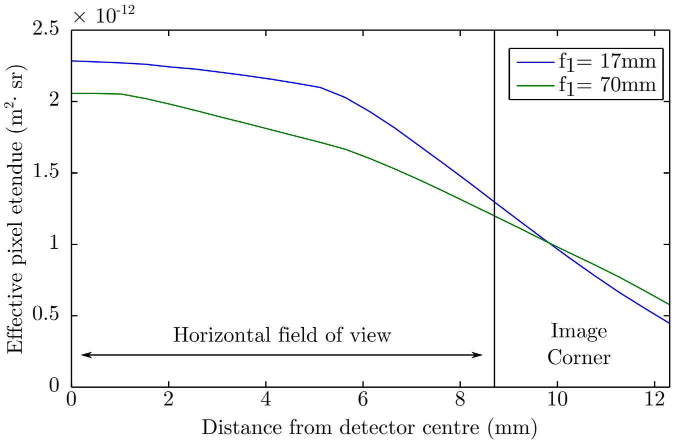

Figure 3.6: Estimated per-pixel effective étendue across the detector for the widest and narrowest zoom settings.

3.5 Interferometer Components

The interferometer components consist of two linear polarisers, a delay plate to produce the fixed interferometer delay and a Savart polariscope to scan the delay across the image, laid out as in figure 3.2. The delay and Savart plates used are made from α-BBO, chosen for its combination of large birefringence ($n_e - n_o \approx -0.12$) and the low sensitivity of its refractive indices to temperature (small thermo-optic coefficients). The latter property is important to maximise the calibration stability of the instrument with respect to changes in ambient temperature: thermal expansion and changes in refractive index with temperature result in a change of the instrument phase ϕ_0, and hence of the absolute flow calibration. In order to minimise this effect, the α-BBO components were mounted (along with the polarisers) in a temperature controlled enclosure. The design of the enclosure and mounting system is one developed at Australian National University specifically for this purpose, based on a modified filter temperature controller from Andover Crop. The enclosure and mounts for the MAST system were purchased from Australian National University. The polarisers and α-BBO plates are mounted in custom mounts with a series of notches at 22.5° around their outer edge; these then slide on to two guide rails inside the enclosure to maintain the correct orientation of the components. This system allows easy swapping of interferometer components and/or reorientation of the components in 22.5° increments. A photograph showing the interior of the enclosure and optic mounts is shown in figure 3.7. The temperature regulation accuracy quoted by the manufacturer of the filter oven is ±0.25°C.

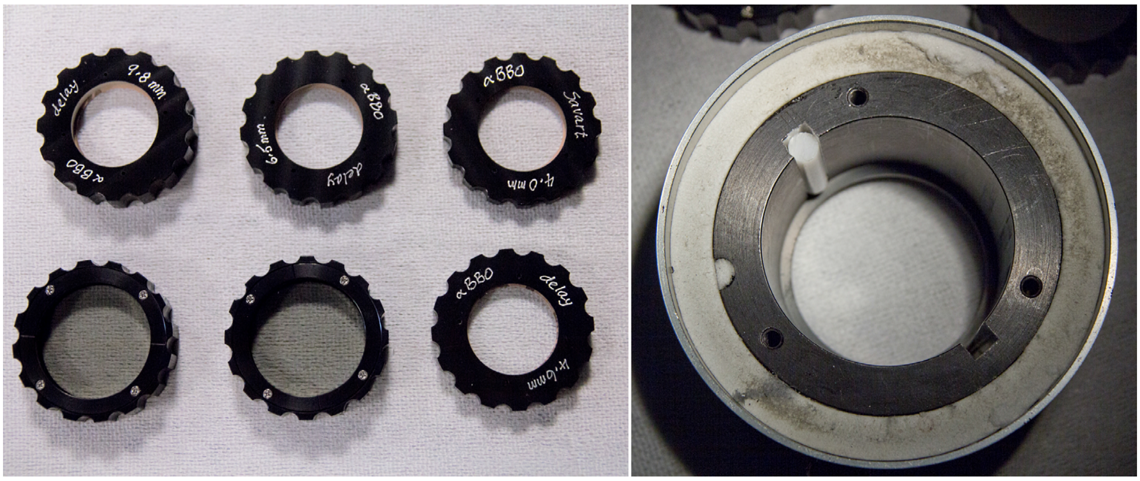

Figure 3.7: Photographs showing (left) the interferometer components in their custom mounts and (right) an end-on view of the temperature stabilised cell in which the components are mounted, with the end cap removed. One of the two white plastic guide rails which hold the orientation of the optics can be seen.

3.5.1 Choice of Fixed Delay (Delay Plate Thickness)

The interferometer group delay $\hat{N}$ is chosen to achieve good sensitivity of the instrument to flows, while keeping the temperature sensitivity low to ensure validity of the line integral (2.4.33) used for tomography. The most important concern for achieving good flow sensitivity when measuring multiplet lines is maximising the fringe contrast, which goes through numerous maxima and minima with increasing $\hat{N}$. From the results in Chapter 2, the expected contrast for multiplet spectral lines dominated by Doppler broadening is given by $\zeta(\hat{N}, T_\mathrm{i}) = \zeta_\mathrm{M}(\hat{N}) \exp[T_\mathrm{i}/T_\mathrm{C}(\hat{N})]$, where $\zeta_\mathrm{M}(\hat{N}) = |\gamma_\mathrm{M}|$ is the multiplet contrast defined by equation (2.4.23), T_i is the ion temperature and T_C is the characteristic instrument temperature.

The expected contrast for each of the candidate lines on MAST was calculated as a function of group delay and ion temperature using this expression, with the multiplet contrast calculated from the results in section 3.1. Zeeman splitting was also included in the calculations of ζ_M using the model in section 2.1.1, for a magnetic field of 0.5T and viewing angle tangential to the magnetic field, however this was found not to have a significant effect on the choice of delays for optimum contrast. The thickness of delay plate corresponding to a given group delay is given by $L = \hat{N}\lambda_0/\kappa_0 B_0$, where subscript 0 signifies quantities evaluated at the centre-of-mass wavelength of the line of interest. The values for birefringence B_0 and the dispersion correction κ_0 in α-BBO were calculated using the Sellmeier equations

[Kato, 1986]:

$$n_o(\lambda) = \left(2.7359 + \frac{0.01878}{\lambda^2 - 0.01822} - 0.01354\lambda^2\right)^{1/2} \tag{3.5.4}$$

$$n_e(\lambda) = \left(2.3753 + \frac{0.01224}{\lambda^2 - 0.01667} - 0.01516\lambda^2\right)^{1/2}, \tag{3.5.5}$$

in which λ is measured in microns. The results are shown in figure 3.8. The beat pattern due to the multiplet structures of the C III and C II lines is seen to dominate for these lines, and is particularly complex for the C II line due to the large number of components and wide spacing.

To maximise flexibility of the system and allow optimisation for the different species, three interchangeable delay plates of different thickness were purchased, indicated in Fig.3.8 by dashed vertical lines. The first of these at 4.6mm thick corresponds to $\hat{N} \approx 1400$ waves at 465nm, and was optimised for the contrast of the C III interferogram. Note that the contrast changes very slowly with temperature for this delay value, which ensures the validity of the simplified tomography problem in equation (2.4.33). The second plate was 6.5mm thick and optimised for a narrow contrast peak in C II at approximately $\hat{N} \approx 1700$ waves, again with very low temperature sensitivity. This is also expected to produce good results for C III: although the fringe contrast for this delay is lower than the 4.6mm plate, the increased phase sensitivity to flows at the larger delay counteracts this effect and the overall flow sensitivity is expected to be at least as good. The third delay plate at 9.8mm is not expected to be optimal for flow measurements, and was chosen to investigate temperature measurements and the effect of the larger temperature sensitivity on the flow results. It also targets the high contrast peak for the C II line at $\hat{N} \approx 2500$ waves. The temperature sensitivity of this delay is particularly large for He II.

3.5.2 Choice of fringe period (Savart polariscope thickness)

As discussed briefly in Section 2.5.3, the nature of the spatial heterodyne measurement scheme leads to reduced spatial resolution in the direction perpendicular to the superimposed fringes. In fact, as will be seen in Chapter 4, the spatial resolution in this direction is related to the spatial period of the fringes (i.e. the number of pixels per fringe), and smaller fringe periods are desirable to obtain high quality images in the presence of sharp image details. However, smaller fringe periods also result in a reduced instrument contrast, lowering the SNR of the fringes, due to 1) Reduced performance of the final imaging lens at higher spatial frequencies, and 2)

(a) C III 464.9 nm

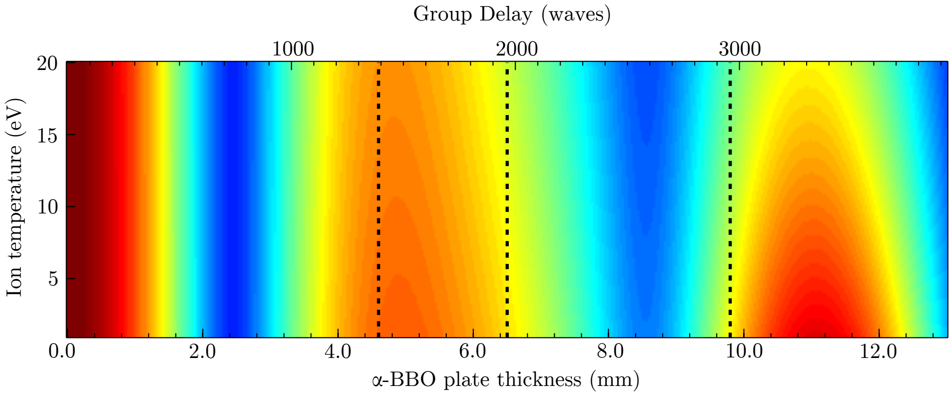

(b) C II 4512.4 nm

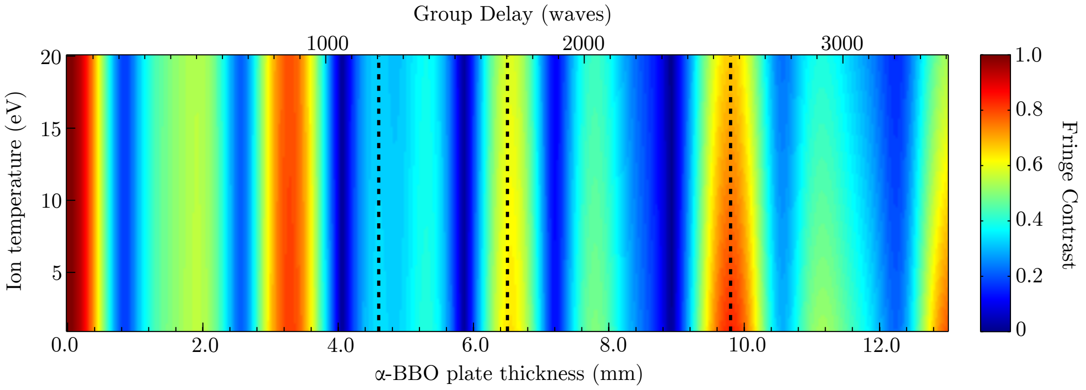

(c) He II 468.6 nm

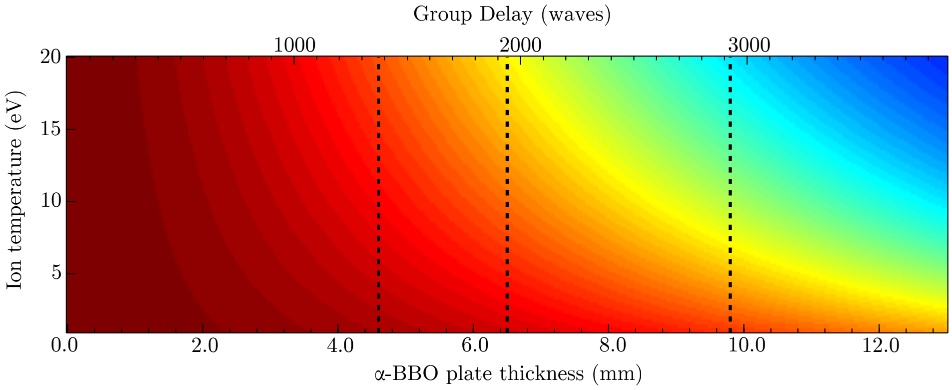

Figure 3.8: Calculated fringe contrast for the candidate spectral lines on MAST, as a function of delay plate thickness (group delay) and ion temperature. Black dotted lines indicate the chosen delay plate thicknesses.

Each detector pixel spanning, and therefore averaging over, an increasing fraction of a fringe. Choice of fringe period is therefore a compromise between spatial resolution and noise handling. Previous Doppler coherence imaging systems have used periods of approximately 10 pixels [Howard et al., 2011], which on the MAST system gives a spatial resolution of a few centimetres for the widest angle views. From equation (2.5.41), taking only first order terms and $\omega = 45°$, the overall thickness of the Savart polariscope required to obtain a given fringe spacing at the centre of the detector is:

$$L_{\mathrm{sp}} = \sqrt{2} f_3 \left(\frac{\lambda_0}{dp}\right)\left(\frac{n_o^2 + n_e^2}{n_o^2 - n_e^2}\right), \tag{3.5.6}$$

where d is the desired fringe period in pixels and p is the pixel size. For the MAST system and a 10 pixel fringe spacing, this gives a thickness of 7.7mm. Such a component was not available, however 4mm and 2.2mm parts were used together to obtain a fringe period of approximately 12 pixels at 465nm. Due to the longer wavelength, the fringe spacing is increased by approximately 10% for the C II line relative to C III and He II.

3.6 Mechanical Design & Plasma Views

The diagnostic was designed to be compatible with pre-existing camera mounting locations on MAST. These use sections of Newport X48 optical rail, commonly cantilevered from the MAST vacuum vessel, from which cameras are supported using standard rail carriers. The coherence imaging system was therefore designed as a single unit which could be supported from an X48 optical rail. The completed diagnostic occupies a compact space envelope of 84cm × 23cm × 22cm (L×W×H), and has a total mass of 11.9kg. The optics and camera are mounted on a 10mm thick custom anodised aluminium base plate. The system is then supported from the X48 rail with two brackets attached to the base plate, one near the rear of the base plate using two rail carriers and a smaller one near the front with a single rail carrier. Photographs of the instrument mounted for testing in the lab and at the lower divertor view on MAST are shown in figure 3.9.

To describe the viewing locations of diagnostic systems on MAST, the vacuum vessel is divided toroidally into 12 segments (‘sectors’), each with upper, lower and equatorial ports for diagnostic and electrical feedthrough access. The primary goal of the diagnostic was to investigate divertor flows, and as such the primary plasma view was a wide angle view of the lower divertor covering both strike points and divertor legs up to the X-Point. This used a window in the lower port on sector 7,

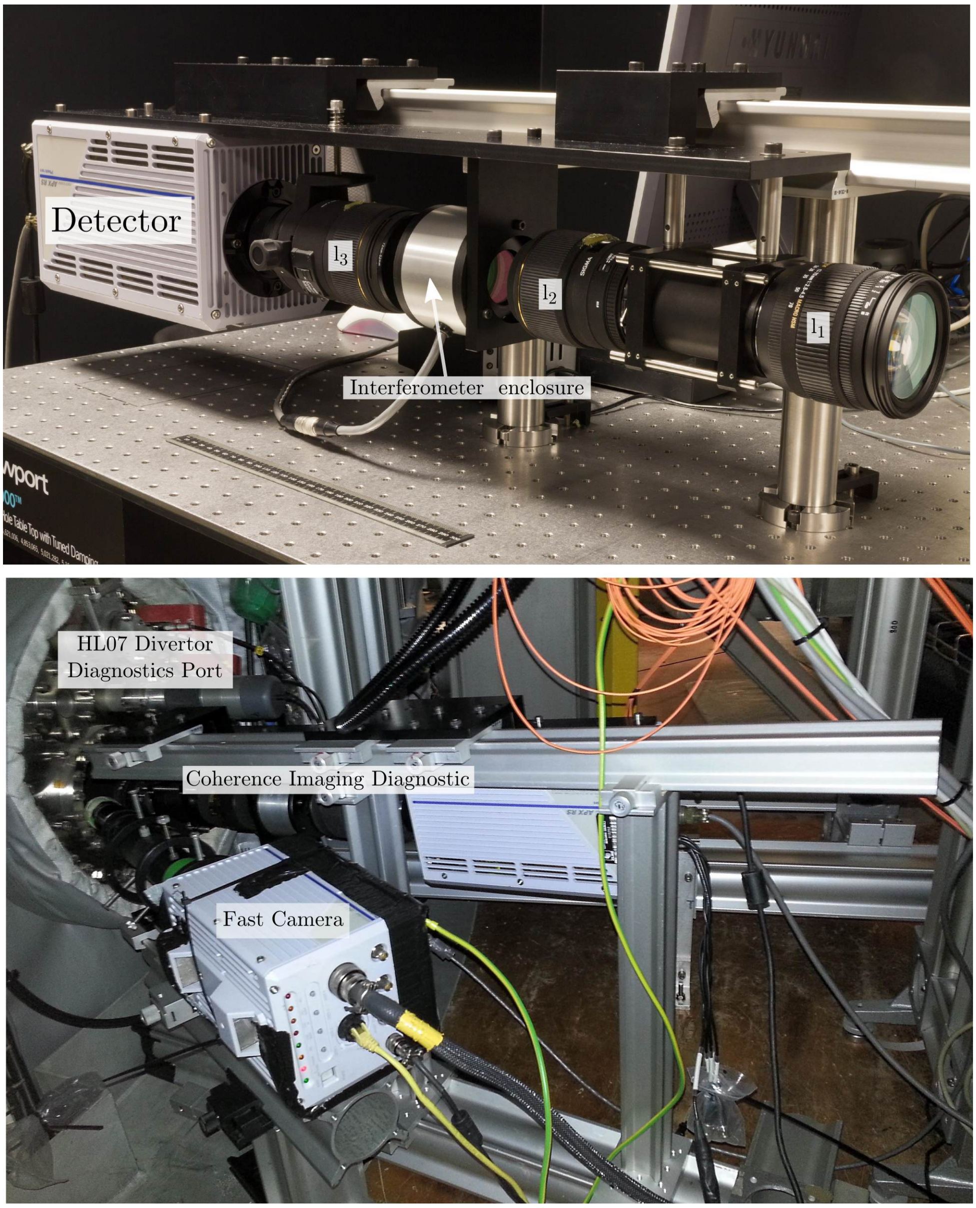

Figure 3.9: Photographs showing the completed MAST coherence imaging diagnostic. Left: set up in the lab, indicating the optical components and with a 30cm ruler for scale. Right: Mounted at the MAST divertor.

denoted HL07, usually occupied by a filtered divertor camera (DivCam). The view of the X-point in DND plasmas was somewhat limited in this view by the P3 and P4 poloidal field coils which obstructed the view at the top of the image. This was less problematic for LSND plasmas since the X-point is typically lower in the vacuum vessel in this configuration.

In order to perform measurements of the main chamber SOL including diagnostic performance and calibration assessments, three views of the main plasma from equatorial ports were also available. These included one from sector 7 viewing small major radii R either side of the centre column (‘HM07 radial’), one from the same port viewing large enough R to capture the both the inboard and outboard edges of the plasma (‘HM07 tangential’), and one viewing small R from sector 2 (HM02). The availability of these ports during the campaign and the viewing angles available were determined by interactions with other diagnostic systems.

The different plasma views described above are illustrated in Figure 3.10. For each view the extent of the field of view, and the extent of the ‘tangency plane’ (projected on the plasma poloidal cross-section), are shown. The tangency plane is a curved surface normal to the camera lines of sight, where the lines of sight are tangent to the toroidal direction. Images from the diagnostic are dominated by light emitted close to this surface, hence its extent in the poloidal cross-section defines the parts of the plasma for which useful data is obtained. Note that while the tangency plane extends on both sides of the centre column for some views, its extent in the poloidal cross-section is always shown as single-sided in Fig. 3.10. The irregular shapes of the tangency plane outlines are due to viewing obstructions by in-vessel components not shown in these simplified diagrams. The focus distance of the diagnostic is set to coincide with the tangency plane, which on MAST is typically 2.1m from the diagnostic.

3.7 Summary

In this chapter design and specifications of the MAST CIS flow diagnostic have been presented. The best candidate spectral lines for CIS on MAST are the C III 465nm triplet, C II 514nm multiplet and He II 468nm multiplet, due to their high brightness. High resolution spectral measurements of these lines have been performed and they have found to be suitable in terms of line structure and lack of contamination. A simplified model of the imaging system was developed and used to choose the optimal combination of lenses and filters for the instrument. The main interferometer delay is provided by 3 interchangeable α-BBO delay plates

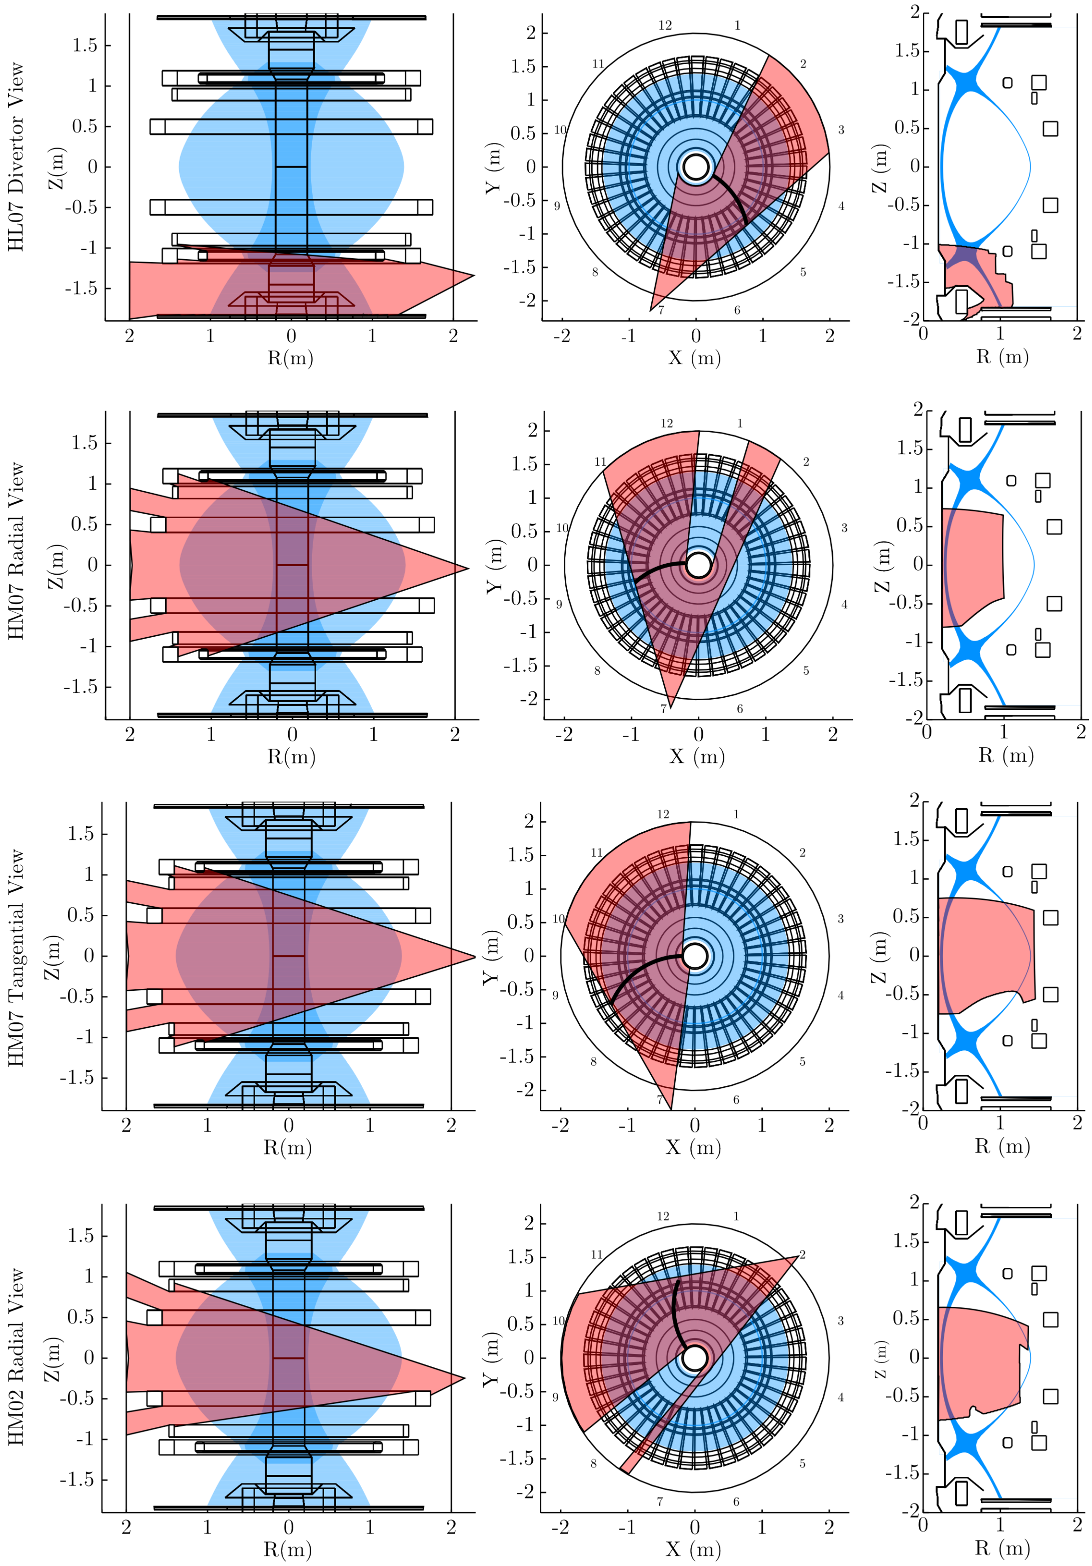

Figure 3.10: Illustration of the CIS plasma views used on MAST, showing (left) side views of the field of view, (centre) top views of the field of view (thick black lines show the tangency plane location), and (right) the extent of the tangency plane projected on the poloidal cross-section.

(used one at a time) of 4.6mm, 6.5mm and 9.8mm thickness, chosen to optimise flow measurements based on calculations of the fringe contrast. The spatial heterodyne fringe pattern is produced by two Savart plates together producing a fringe spacing of 12 pixels. A table summarising all of the component specifications is given in Appendix A. Four different plasma viewing configurations are available for CIS measurements: three for the main SOL and one for the lower divertor, as shown in figure 3.10.

The next chapter presents the data analysis techniques for extracting flow information from the raw diagnostic images, some of which were used to analyse the results in this chapter. Some of the results from this chapter will also be used in forward modelling of the diagnostic to estimate flow measurement errors.

Chapter 4

Data Analysis & Measurement Uncertainties

This chapter describes the techniques used to extract the spectral information of interest from MAST CIS data, and presents numerical investigations of related aspects of the instrument performance. Both of these areas require the use of simulated data containing features similar to those expected in the real measurements. The method for generating such simulated data is therefore presented first.

4.1 Generation of Simulated Data

4.1.1 Line-integrated spectra

For the testing of tomographic inversion techniques it is desirable to have synthetic line-integrated data corresponding to known profiles of local plasma conditions. Using realistic plasma profiles means that such images should have similar features to those seen in the real data, making them suitable for testing all stages of the analysis chain. To this end, simulated line-integrated spectra were generated based on MAST simulation results from the OSM-EIRENE code package [Lisgo et al., 2005]. This is not a predictive code but an interpretive model constrained by experimental data. It uses an onion skin model (OSM) for the deuterium plasma and the Monte-Carlo neutral hydrogen code EIRENE for neutrals. The particular simulation used here is for MAST shot 25028 at t = 0.31s; this was an L-Mode LSND plasma with 1MW of NBI heating, performed to investigate divertor detachment on MAST [Harrison, 2010]. The simulation results were provided by Dr James Harrison at CCFE. While they do not include impurities, which are what we actually measure with the CIS diagnostic, the general structure in the 2D profiles should be sufficiently similar to what appears in real data for the development and testing purposes here. The simulation outputs of interest were ion temperature, parallel flow speed and D_α emissivity, used in place of impurity emissivity. The vector magnetic field was obtained from the EFIT magnetic equilibrium reconstruction code for the same shot and time. Both the plasma parameters and magnetic field profiles were specified in 2D in the poloidal (R, Z) plane and assumed to be toroidally symmetric. Line of sight vectors for the simulated diagnostic were generated using a simple pinhole camera model, with the design values for the field of view and positioning of the camera. Since the main application of the diagnostic was in the divertor, the simulated diagnostic view considered here corresponds to the wide-angle lower divertor view from port HL07.

Calculation of the line integrated spectra proceeded by iterating over each line-of-sight (i.e. each virtual detector pixel), and for each one stepping along the line of sight away from the camera. The lines of sight were considered to be narrow pencil beams, i.e. the finite depth of field and pixel size of the real system were not taken into account. At each step along the sight-line, the current position in the R, Z plane was calculated, and the local plasma parameters were obtained by interpolating the input plasma and B field data to that R, Z location. The positions and intensities of the spectral line components emitted at that plasma position were determined from the line multiplet structure, Zeeman splitting model described in 2.1.1 and line-of-sight Doppler shift, determined by the parallel flow velocity and the angle between the local magnetic field and the sight-line. Note that the flow is assumed to be entirely parallel to the magnetic field for the purposes of generating the simulated spectra. Gaussian line profiles, with widths given by Doppler broadening for the local ion temperature, were then used to represent each line component. These local spectra along the line of sight were all calculated at the same sampling points in wavelength, and the final line integrated spectrum was computed by integrating the intensity at each wavelength along the line of sight. An example line integrated spectrum for C III is shown as the blue line in Fig. 4.1, along with the spectrum which would be emitted from a homogeneous source with the same average flow, temperature and magnetic field for comparison (grey line). The line-integrated profile is highly distorted, and visibly consists of a double peak at each line component: this is due to contributions from distinct brightly emitting plasma regions with different flow speeds. While this example is an extreme case of distortion of the line shape, the same principle applies to all sight-lines and illustrates why line-of-sight spectroscopic measurements can be very difficult to interpret when only limited spatial coverage of the plasma is available.

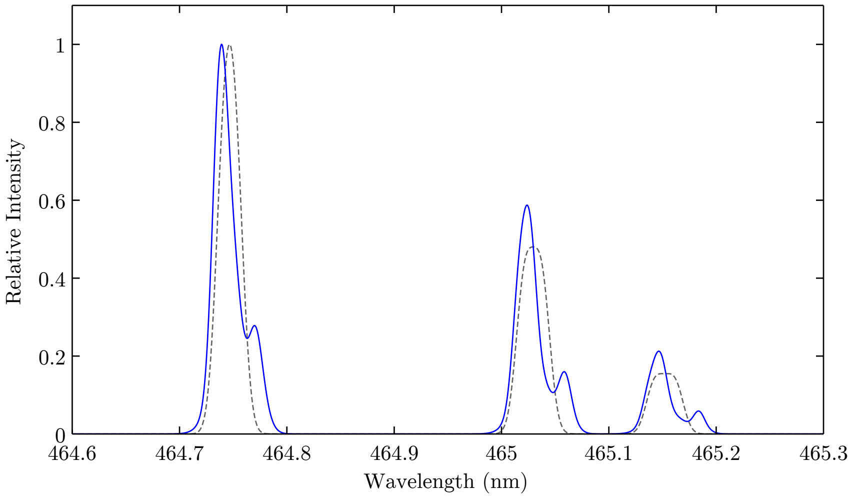

Figure 4.1: Example synthetic line integrated spectrum for C III in the MAST divertor (blue line). The dotted grey line shows the spectrum for a single point with the average conditions along the line-of-sight, for comparison.

After repeating this procedure for all simulated sight-lines, the result is a set of simulated line-integrated spectra for the light arriving at each virtual detector pixel. In order to use these results for data analysis development, they must then be converted into simulated data images.

4.1.2 Image generation from simulated spectra

Simulated data images were created pixel-by-pixel from the calculated line-integrated spectra. For each pixel, the position on the detector was mapped to incidence angles through the interferometer components (θ, ω) according to $\theta = \arctan(\sqrt{x^2+y^2}/f_3)$ and $\omega = \arctan(x/y)$, where x, y are the pixel's position on the detector in the same units as f_3. For each wavelength sample in the spectrum, the detected spectral intensity at that wavelength was calculated according to:

$$S(\lambda) = \frac{I(\lambda)}{2}(1 + \zeta_I \cos(\phi_{\text{Delay}}(\theta, \omega, \lambda) + \phi_{\text{Savart}}(\theta, \omega, \lambda)), \tag{4.1.1}$$

where ϕ_{Delay} and ϕ_{Savart} are given by equations (2.5.39) and (2.5.41) respectively and $I(\lambda)$ is the line integrated spectral intensity from the plasma. Filter effects can be included by multiplying $I(\lambda)$ by the filter transmission profile, based on data from the manufacturer and the θ-dependent band pass shift from equation (3.4.2). The instrument contrast ζ_I is an input parameter accounting for the real instrument contrast, since realistic contrast reducing effects in the optics were not included in these simulations. The total intensity in the pixel is then given by $\int S(\lambda) d\lambda$. Once this has been performed for all pixels the complete image is scaled to the desired brightness, typically similar to that seen in the real plasma data. Read noise and shot noise can be included by adding Gaussian and Poisson distributed random numbers, respectively, at levels consistent with the real detector as measured in Section 5.1. An example simulated C III image is shown in figure 4.2(a). As with the real data, to demodulate these simultaed images a suitable calibration reference is required, i.e. a flat field image with no Doppler shift. This was generated in the same way as the simulated data image, using the un-shifted spectrum of the line of interest for every pixel rather than the simulated line-integrated spectra.

Interferogram Property Images

For testing of interferogram demodulation, the image brightness, fringe contrast and fringe phase which would be recovered by an ideal demodulation algorithm must be known. These are not provided by the above simulated images, since the detected intensity at each pixel is calculated directly without calculating the fringe phase and contrast. Phase and contrast images were instead calculated from the line-integrated spectra by numerical evaluation of equation (2.2.9), to find the degree of coherence γ of the spectra and thereby the phase $\phi = \arg(\gamma)$ and contrast $\zeta = |\gamma|$. The brightness image was obtained by simply integrating the spectra over wavelength. The desired Doppler phase was isolated from the calculated phase ϕ by subtracting the phase calculated in the same way for the un-shifted spectrum. Brightness, phase and contrast images for the OSM plasma simulation with $\hat{N} = 1404$ waves are shown in figure 4.2. Data analysis code was tested by extracting the fringe properties from the image in fig. 4.2(a) and comparing the results with the data in fig. 4.2(b),(c) and (d).

The interferogram demodulation technique is now presented, including results of tests carried out using simulated data generated as described here.

4.2 Interferogram Demodulation Technique

As formulated in Chapter 2, the recorded detector signal down a column of pixels (perpendicular to the horizontal fringes) is of the form:

$$S(y) = \frac{I_0(y)}{2}[1 + \zeta(y)\cos(\phi_{\mathrm{I}}(y) + \phi_{\mathrm{D}}(y))], \tag{4.2.2}$$

where $\phi_{\mathrm{I}}(y)$ is the instrument phase defined in equation (2.5.40), $\zeta(y)$ is the fringe contrast including instrumental, multiplet structure and Doppler broadening effects,

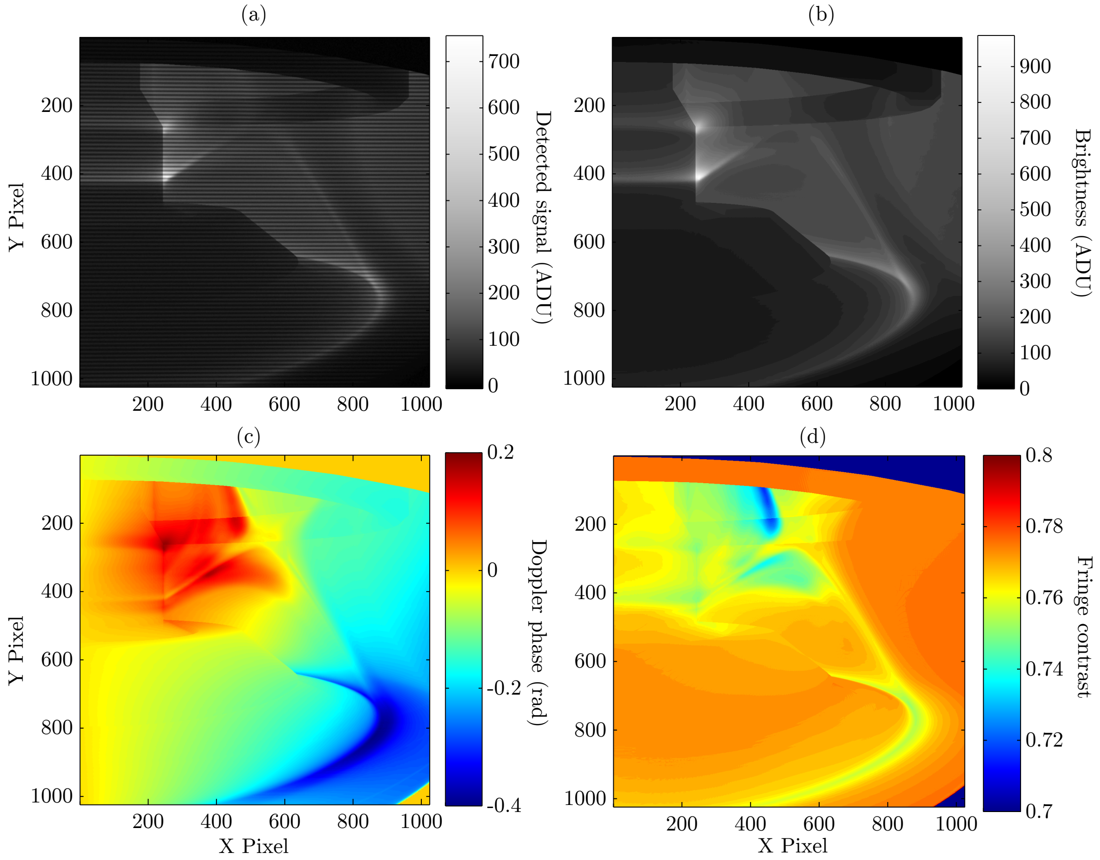

Figure 4.2: Examples of simulated data images based on OSM-EIRENE simulations for shot 25028: (a) Simulated camera image, (b) Brightness image, (c) Doppler phase image, (d) Fringe contrast image.

and the brightness $I_0(y)$ and Doppler phase $\phi_D(y)$ contain the information of most interest for flow measurements. Extraction of the brightness, contrast and phase from the interferogram amounts to demodulation of a phase and amplitude modulated carrier signal with a varying background level.

For convenience and consistency, the following descriptions all assume horizontal fringes. In some MAST experiments, the fringe pattern was rotated on the detector (for example see section 6.2.4). In these cases, both the calibration and data images were rotated such that the fringes were horizontal before applying the same data processing. Tests with simulated data images indicate this rotation does not have a significant effect on the accuracy of demodulation.

4.2.1 Removal of neutron effects from real images

During plasma discharges using NBI auxiliary heating, D-D fusion reactions between energetic beam ions, and between beam ions and the background plasma, produce yields of 2.45MeV neutrons of typically around 10^{13}/s [Stammers and Loughlin, 2006]. The interaction of such neutrons with the un-shielded detector produces bright spots and streaks in the images when NBI heating is operational. An example of a data frame from MAST showing this can be seen in in Fig. 4.3(a). These features must be removed before demodulation of the images to prevent them from producing artefacts in the processed images. The following algorithm was applied row-wise to the images to achieve this:

1. A 3-point median filter is applied to remove single hot and cold pixels, caused both by neutrons and hot and cold pixels on the detector.
2. The intensity gradient along the row is estimated with a 3-point central difference.
3. Peak finding is used on the resulting gradient to detect positive and negative spikes larger than an empirically set threshold value. A pair of positive and negative peaks indicate the edges of a bright feature to be removed. The threshold value was usually set to twice the standard deviation of the gradient array.
4. Where a positive gradient peak is followed by a negative peak within a certain threshold distance (empirically set and representing the maximum size of feature to be removed), the pixel values between the two peak locations are linearly interpolated from the pixel values either side of the feature. The size threshold is usually set to 4 pixels.

This process was applied to all images from MAST before interferogram demodulation. Fig 4.3 shows a section of a data image before and after this process is applied, showing the reduction in bright artefacts in the image. This filtered image is then subject to the interferogram demodulation, which proceeds column-by-column across the image.

4.2.2 Intensity extraction

The first parameter to be extracted from the images is the brightness image I_0, by removing the sinusoidal carrier fringe pattern from the image. Since the fringe pattern oscillates around a mean value of $I_0/2$, in principle this can be achieved using boxcar averaging down each image column with an averaging period of exactly 1 fringe. In practise, however, the fringe period is not an integer number of pixels and therefore this approach leaves residual small amplitude fringes in the images.

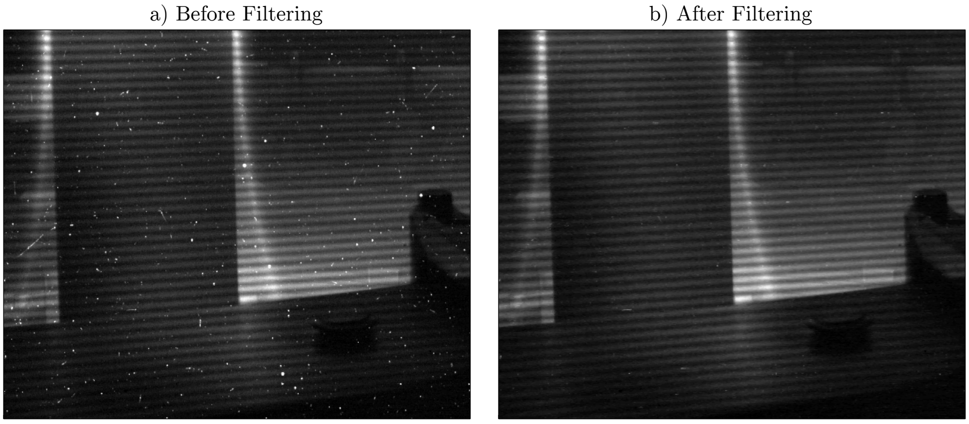
a) Before Filtering
b) After Filtering

Figure 4.3: Section of a data frame from the HM02 midplane view before and after filtering to remove radiation effects, showing an in-vessel poloidal field coil and part of the centre column / high field side scrape-off-layer plasma. The left image (a) is directly from the camera, while (b) has had the filtering applied. Exposure time as 16.7ms and the NBI heating power was 3.3MW.

Another approach is to filter out the carrier fringe frequency in the Fourier domain. While this approach produces much better results for smooth intensity variations, it suffers from large ringing artefacts around sharp changes in brightness in the image. To alleviate this, the two fringe removal techniques were used together: Fourier domain filtering of the carrier frequency was applied first, then boxcar averaging was applied to reduce the remaining ringing artefacts at sharp image features. An example of the results for a single image column is shown in Fig. 4.4, which shows a simulated input data column (based on column 750 in the images in figure 4.2), the corresponding recovered I_0, and the ideal I_0 response. The intensity is well recovered except at sharp edges and peaks, which show smoothing consistent with a top-hat kernel with a width of 1 fringe. For the example in figure 4.4, the RMS fractional difference between the ideal and recovered profiles is 5.3% overall, and just 1.2% when excluding the four sharp features around pixels 60, 140, 240 and 670. Some small residual ringing artefacts are also seen at the sharp edges. As can be seen from figure 4.2, these sharp image features arise at the edges of in-vessel machine components such as poloidal field coils. Ultimately the smoothing of these sharp features is due to limitations in extracting spatial information at smaller scales than the fringe period, from a fringe pattern with unknown brightness, amplitude and phase, and could be improved by using a smaller fringe period. Sudden large jumps in fringe phase (which can also occur at the edges of viewing obstructions) can also cause artefacts in the recovered intensity. The effect of errors in the brightness extraction on flow image demodulation will be discussed in the next section.

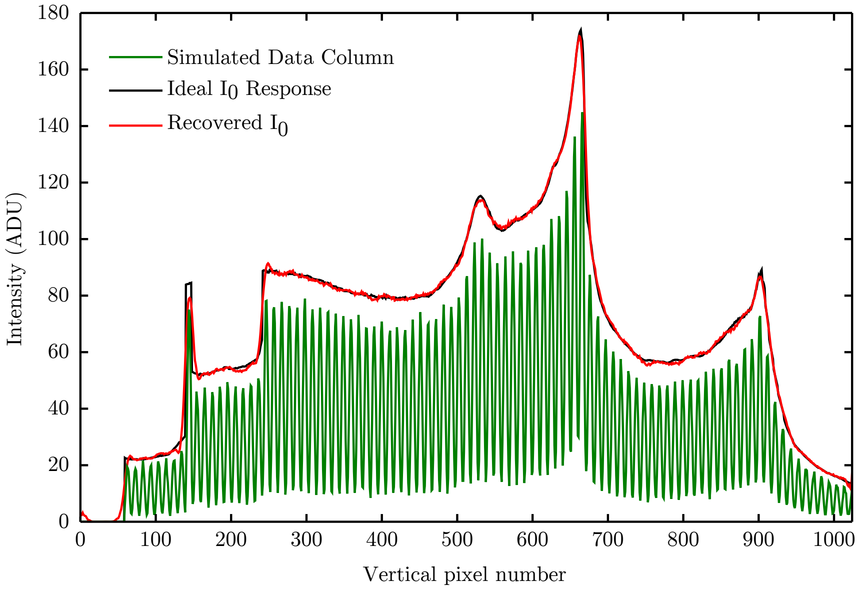

Figure 4.4: Extraction of intensity profile from a simulated image column. Input I_0 for the data generation (black line), the simulated data column (green line) and recovered intensity profile I_0 (red line) are shown. The recovered profile resembles a smoothed version of the input profile with a smoothing period of 1 fringe.

4.2.3 Flow (phase) extraction

To obtain the column signal in a form suitable for phase and contrast demodulation, the extracted intensity is first factored out of the signal according to $S'(y) = [2S(y)/I_0(y)] - 1$. Using equation (4.2.2) we see that $S'(y) = \zeta(y)\cos[\phi_\mathrm{I}(y) + \phi_\mathrm{D}(y)]$, i.e. it is a sinusoidal signal with zero mean, phase modulated by the Doppler phase ϕ_D and amplitude modulated by the fringe contrast. Phase and contrast extraction are then performed via the analytic signal representation of $S'(y)$. The analytic signal was first introduced in Gabor [1946]. The general concept is that due to the Hermitian symmetry (i.e. $f(-x) = -f(x)^*$) of the Fourier transform of a real-valued signal, the negative frequency components of such a transform can be discarded without losing any information about the original signal. The inverse Fourier transform then yields a complex representation of the original signal which makes certain properties, particularly the instantaneous phase and amplitude, more accessible. The analytic representation of the discrete signal S' can be generated easily by taking the Discrete Fourier transform (DFT) of the image column to obtain $\tilde{S}'[f]$ and setting:

$$\tilde{S}'[f] = \begin{cases} \tilde{S}'[f] & : f = 0 \\ 2\tilde{S}'[f] & : f > 0 \\ 0 & : f < 0 \end{cases}$$

before performing the inverse transform [Lawrence Marple Jr., 1999]. The resulting analytic signal will be denoted S_a. The real part of S_a is equal to the original input signal, while the complex part is equal to its Hilbert transform [Gabor, 1946]. Since the effect of the Hilbert transform is a 90° phase shift of all frequency components, the real and imaginary parts of S_a are in phase quadrature and can be used to recover the instantaneous phase and contrast according to $(\phi_\mathrm{I}(y) + \phi_\mathrm{D}(y)) = \arg(S_\mathrm{a}(y))$ and $\zeta(y) = |S_\mathrm{a}(y)|$.

Features in the brightness image with high spatial frequencies, which are not accurately recovered by the intensity extraction as noted in the previous section, cannot be accurately removed from the data when calculating S' and cause ringing artefacts in the demodulated phase. High spatial frequency components of image noise also appear strongly in the resulting phase images. This is illustrated in figure 4.5(b), which shows the ϕ_D image demodulated from noisy test data using the scheme described so far, alongside the ideal result in 4.5(a). The main features of concern in the demodulated image are the high level of noise on the image, and the ringing artefacts at the edges of the poloidal field coil structures. These undesirable effects are reduced (at the cost of fine detail in the image) by a combination of two techniques. The first is windowing in Fourier space when calculating the analytic signal. Using a window function centred at the carrier fringes' spatial frequency, high spatial frequency contents of the recovered Doppler phase, including the ringing artefacts and some of the image noise, are attenuated. This operation is illustrated in Figure 4.6, which shows the magnitude of the DFT of S' and a multiplicative window function which is applied before the inverse transform to yield S_a. Indicated on the figure are f_c, the spatial frequency of the carrier fringes, and

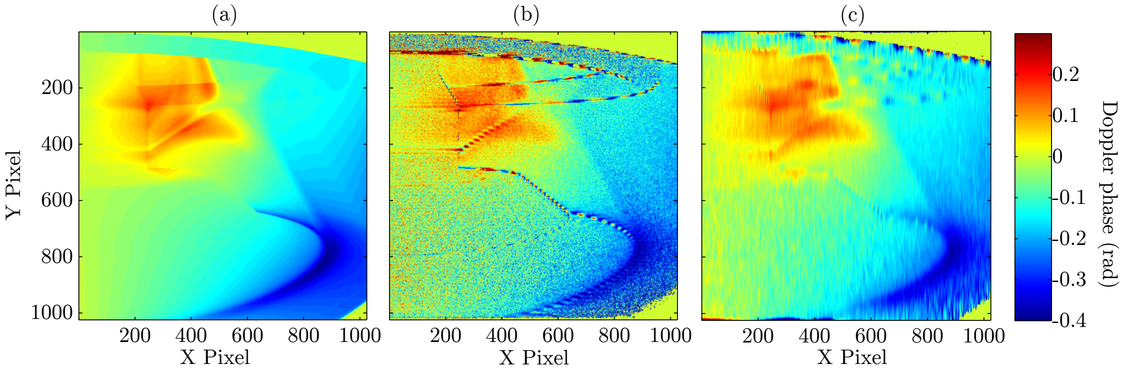

Figure 4.5: Illustration of the effect of windowing when calculating the analytic signal and of apodisation of the column signal at sharp edges. (a) Ideal phase response for simulated noisy divertor data, (b) Demodulated phase without windowing or apodisation, (c) Demodulated phase with windowing and apodisation.

W, the total width of the window function. Choice of the particular window function and width is a compromise between removing high spatial frequency artefacts and noise while preserving desired high spatial frequency information (i.e. fine details) in the signal. By visual inspection of Doppler phase images extratced from simulated test data, and comparison between the demodulated and ideal images, a Blackman-Harris window (as shown in figure 4.6) with width $W = f_c$ was found to produce good robustness against noise and artefacts while preserving features of interest in the simulated data. The second technique, used to reduce the size of the ringing artefacts, is to apodise the column signal S' around locations of sudden changes in I_0, before calculating the analytic signal. This removes the sharp jumps in the signal which give rise to the ringing. To do this, the gradient of I_0 down the image column is estimated with a 3-point central difference, and locations where the gradient has a numerical value greater than an empirically chosen fraction of I_0 are considered sharp edges. This fraction is usually set to 5%. Inverted Hann windows are then applied (in real space) to the column signal centred at these locations. This process is illustrated in figure 4.7, which illustrates the process of demodulating the phase from a single image column. Using both the windowing and apodisation, the demodulated Doppler phase from noisy test data is shown in figure 4.5(c).

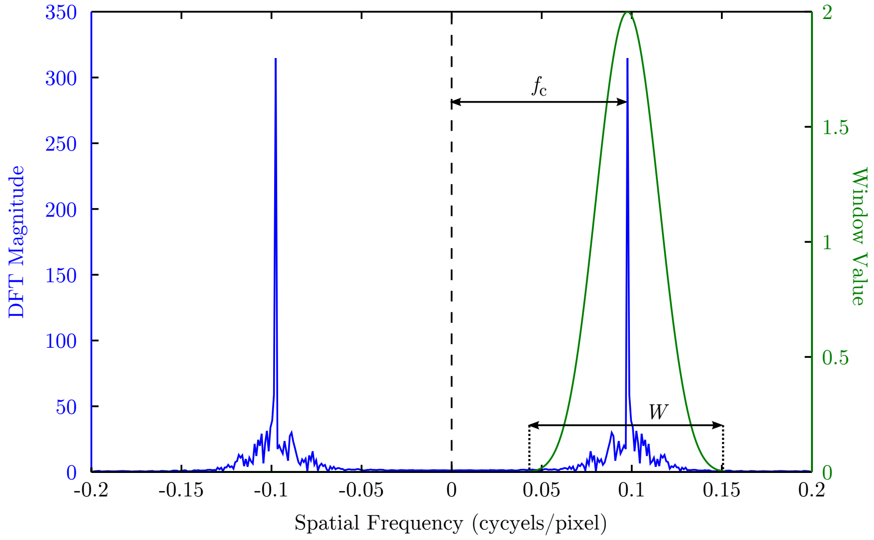

Figure 4.6: Application of a window function in Fourier space when calculating the Analytic signal using the DFT of S'. The carrier fringe frequency f_c and total window width W are indicated.

The output of the phase demodulation for a given input image is the instantaneous phase at each pixel, in the interval $[-\pi, \pi]$. An example section of a simulated (noiseless) image and corresponding demodulated phase are shown in figure 4.8. For

(a) Noisy Input Data S & Extracted Intensity I_0

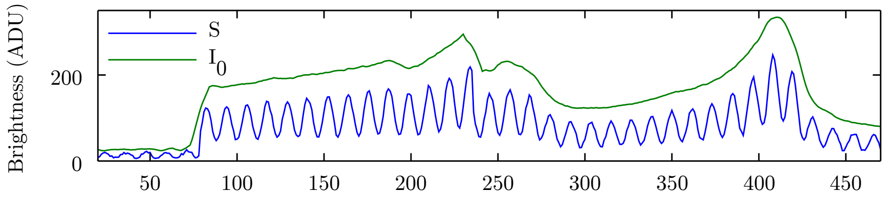

(b) S' (Before apodisation)

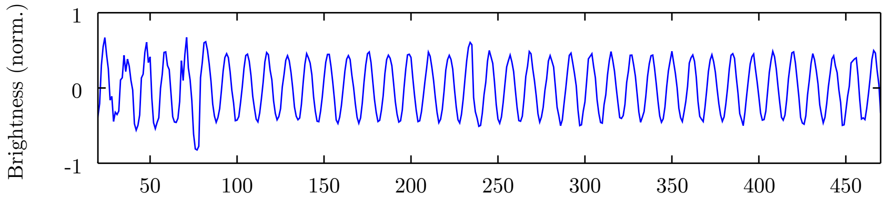

(c) S' (After apodisation)

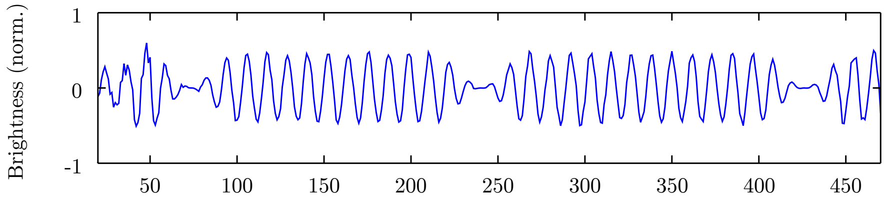

(d) Analytic signal S_a

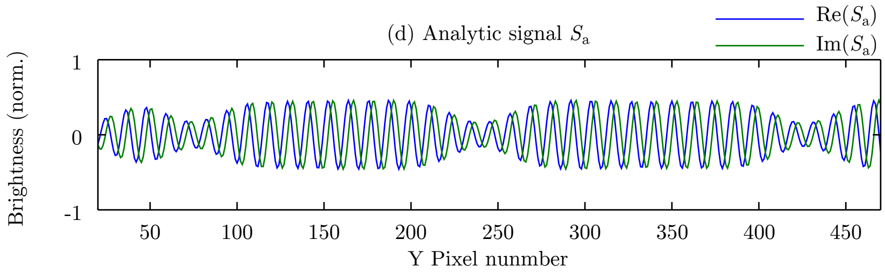

Figure 4.7: Illustration of the steps involved in phase demodulation of a single image column. (a) Input noisy test data and extracted brightness profile. Multiple regions of large brightness gradient are seen. (b) Data with brightness profile factored out, referred to as $S\prime$ in the next. (c) $S\prime$ after apodisation at locations of large brightness gradient, (d) The calculated analytic signal corresponding to the signal in (c), from which the fringe phase is determined.

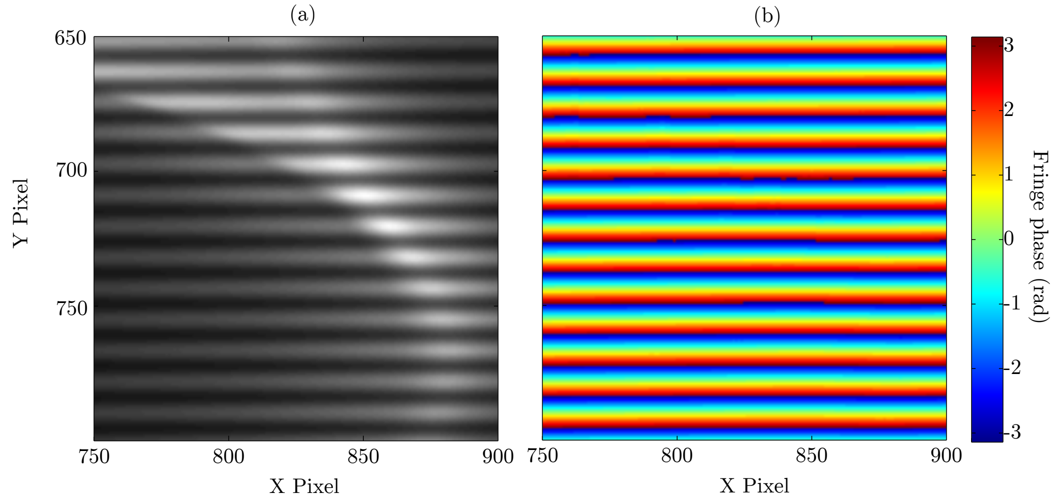

Figure 4.8: Illustration of fringe phase demodulation results for a small section of a noiseless simulated image: (a) Input image and (b) demodulated fringe phase.

calibration images taken in the lab to be used as an instrument phase reference, this phase image is unwrapped to reveal the shape of the instrument phase across the field of view, as will be described in section 5.3.2. For plasma data, the calibrated instrument phase ϕ_I is subtracted directly from the demodulation output and the result is wrapped into the interval $[-\pi, \pi]$, to isolate the Doppler phase ϕ_D. Note that this limits the largest (line-average) flows which can be measured, without further processing to detect phase wrapping in the ϕ_D images, to $\bar{v}_\mathrm{max} = (c/2\hat{N})$ (where c is the speed of light and $\hat{N}$ is the instrument group delay). For the MAST system $\bar{v}_\mathrm{max} \approx 100$km/s, around a factor of 3 larger than any observed flow speeds in the SOL and divertor.

As was derived in section 2.4.2, the measured Doppler phase is related to an emissivity weighted line average of the line-of-sight ion flow. This line-average flow is obtained straightforwardly from the demodulated ϕ_D according to:

$$\bar{v} = \frac{c\phi_\mathrm{D}}{2\pi\hat{N}}, \tag{4.2.3}$$

and from equation 2.4.33 this is related to the local plasma flow by:

$$\bar{v} = \frac{1}{\breve{e}_0}\int_L e_0(\boldsymbol{r})\boldsymbol{v}\cdot\boldsymbol{dl}. \tag{4.2.4}$$

While $\bar{v}$ is not necessarily suitable for quantitative interpretation of the measurements directly, these images can already provide qualitative insight into flow patterns and dynamics. Before going on to discuss tomographic inversion of $\bar{v}$ images to obtain poloidal $v_{||}$ profiles, issues of noise and measurement errors in the line-integrated values are addressed.

4.3 Image Noise Propagation

To properly interpret flow images it is desirable to characterise the amount and appearance of noise on the images. Due to the Fourier domain windowing applied during the phase demodulation process, noise in the demodulated phase images takes on a smoothed appearance perpendicular to the fringe direction. Since each image column is demodulated independently there is no such smoothing in the horizontal direction, and the noise appears elongated in the vertical direction. An example is shown in Fig. 4.9(a), which shows the demodulated Doppler phase from a noisy test image with $\phi_D = 0$ everywhere. The elongated structure of the noise can clearly be seen in the image, and horizontal and vertical cuts through the image are shown in Fig. 4.9(b) which clearly illustrates the anisotropy.

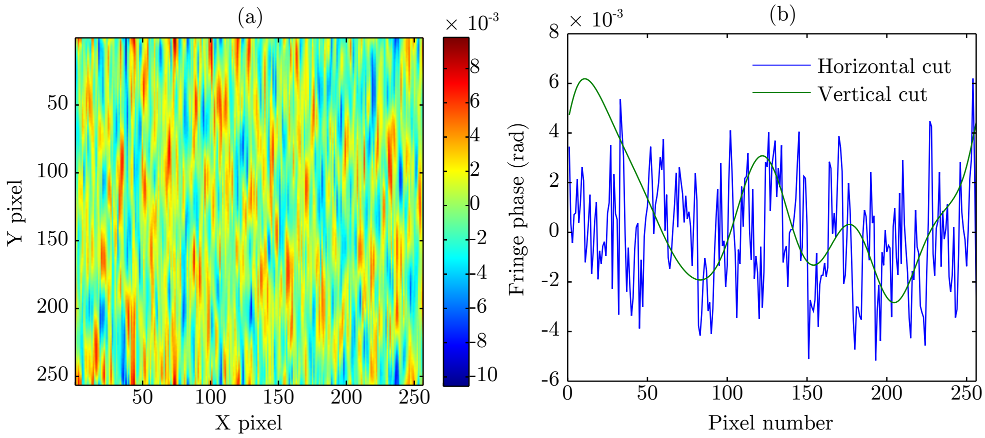

Figure 4.9: Illustration of spatially anisotropic noise on demodulated phase images. Left: demodulated phase image from a zero phase simulated image including noise. Right: Horizontal and vertical profiles through the image.

The amount of ϕ_D noise from a given raw image is determined by the SNR of the fringes, and therefore for given detector parameters is a function of both image brightness and fringe contrast. This dependence was investigated numerically using simulated images. First a fringe pattern with no noise was produced and demodulated to be used as a phase reference image. Then, the same image was generated at different brightness and contrast values, with shot noise and camera noise added according to the detector characterisation results in section 5.1. The difference between the phase extracted from the noiseless image and each noisy image was then purely due to the image noise. The standard deviation of the resulting phase noise, as a function of I_0 and ζ, is shown in Fig. 4.10. The standard deviation is given both in terms of phase and the equivalent line-average flow $\bar{v}$ for $\hat{N} = 1385$ waves, corresponding to C III with the 4.6mm delay plate, to give an indication of the size of the effect on measurements. As expected, the flow noise decreases with increasing signal and increasing fringe contrast. The operating point of the instrument within this I_0, ζ parameter space is determined by the instrument contrast (imperfections in the optical system), spectral line structure, ion temperature, plasma emissivity and the integration time of the camera (related to the desired time resolution). The typical I_0 and ζ operating range for the MAST instrument is shown as the dashed box on the figure, which shows a noise level of ≤ 1.1km/s. The results in figure 4.10 can be used as a look-up-table when analysing data images to estimate the noise in a given portion of the image from the local intensity and contrast.

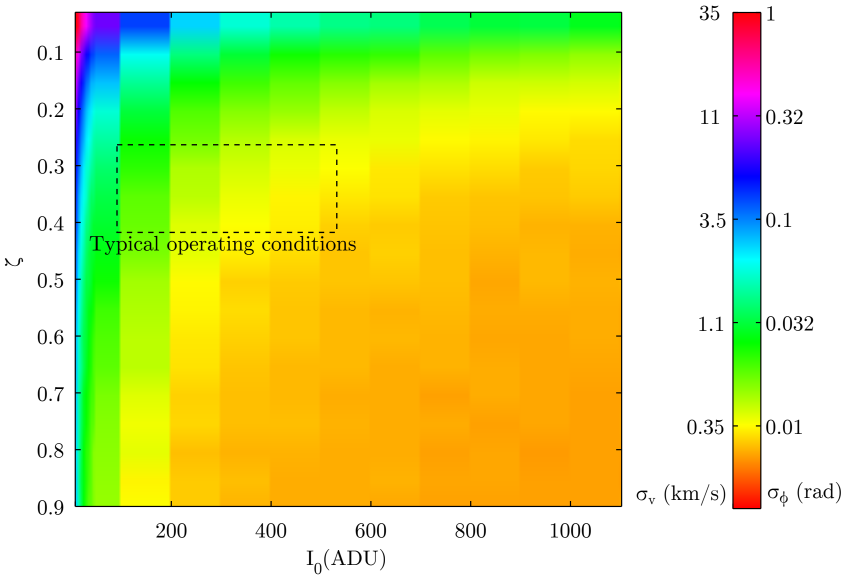

Figure 4.10: Standard deviation of demodulated phase noise as a function of brightness and fringe contrast. Noise level is given in terms of both phase and flow equivalent for $\hat{N} = 1385$ waves.

4.4 Spatial Response

Due to the 1D fringe pattern used to encode the spectral information, the spatial resolution of the diagnostic is highly anisotropic. In the direction perpendicular to the fringes (vertical direction), the spatial response is dominated by the effect of the fringes and demodulation, which essentially apply a smoothing in this direction. Parallel to the fringes (horizontal direction) there is no such effect, and the spatial response is only limited by the imaging performance and finite depth of field of the imaging system.

The effect of the fringes and demodulation scheme design on the spatial response of the measurements can be described by a convolution of the 'true' flow and emissivity images with suitable filter kernels, essentially describing 'effective Point-spread functions (PSFs)' for emissivity and flow imaging. However, unlike the PSF in a conventional imaging system these are not related to optical effects in the imaging system, but purely describe the effects of the spatial heterodyne measurement scheme. Since a 1D fringe pattern is used for the modulation, the corresponding effective PSFs are delta functions in the horizontal direction (parallel to the fringes) and extended in the vertical direction (perpendicular to the fringes). They are also not the same for brightness and flow imaging, since different procedures are used to extract these two quantities from the raw images.

To establish the form of the effective PSFs for I_0 and $\bar{v}$, synthetic data columns using both simulated divertor data and artificial phantoms were used. Taking I_0 as an example, image columns with varying I_0 but constant phase and contrast (to avoid artefacts in the extracted intensity due to rapid changes in these other parameters), were demodulated. The resulting recovered I_0 profiles were fit by the ideal I_0 profile convolved with the effective PSF to be determined, with the width of the PSF being the free parameter in the fitting. Different shapes of PSF were tested to minimise the fitting residuals. A similar procedure was used for the flow imaging, varying ϕ_D down the image column and keeping I_0 constant.

For I_0 imaging, the PSF best describing the response was a top-hat function with a width of 1 fringe period. On the MAST system this is ≈ 12, pixels, or 1.2cm at the focal plane in the plasma for the wide angle views. For flow imaging, the best fit was achieved for a Gaussian PSF with a FWHM of 2.7 fringe periods, giving 32 pixels or 4.4 cm in the plasma for wide-angle MAST measurements. The Nyquist-limited resolution in the orthogonal spatial direction (i.e. along the fringes) is around 3 mm for the same viewing configuration.

4.5 Spectral Contamination

Measurements of the MAST spectral lines in section 3.1 revealed a number of contaminating features within the filter pass bands. Using simulated spectra including these features it is possible to assess their potential impact on flow measurements.

The effect due to each contaminating feature depends mainly on the feature's centre wavelength and brightness. For each ion line of interest, all contaminating features visible in the spectral measurements from Chapter 3 were fit with Gaussian functions in data from multiple shots and times. This gave distributions of observed intensity of these features relative to the line of interest. To explore the effect on the instrument phase spectra were generated consisting of the line of interest at its rest wavelength, plus the known contaminating features each with relative intensity chosen from a normal distribution with a mean and standard deviation given by the fitted relative intensities for that feature. Random Doppler shifts of the contaminating lines corresponding to flows of up to ±30km/s (i.e. the maximum divertor impurity flow speeds measured on MAST) relative to the species of interest were also added, however the results were found to be much more sensitive to the contamination intensity than to these shifts. The spectra were multiplied by the band pass filter profiles so that exclusion of the contaminants by the filters was accounted for.

The fringe phase associated with the simulated contaminated spectrum, relative to the uncontaminated line, was then calculated for each delay plate, and converted to an apparent flow value. This was repeated 5000 times to give a distribution of apparent (false) flow caused by the spectral contamination. The mean of this distribution gives an indication of the mean error caused by the contamination, while the standard deviation indicates the spread of the distribution of errors caused if the contaminating features change brightness (and position) relative to the line we are trying to measure. A small standatd deviation would indicate an almost fixed systematic error, while a large standard deviation signifies that the error could change a lot from frame-to-frame, location-to-location or shot-to-shot. The amount of intensity variation of the contaminant lines used in this calculation is likely slightly over-estimated, since fitting errors were not accounted for (hence the final flow errors are likely to be a slight over-estimate). Despite this, this method is useful to determine whether contamination errors are likely to be significant compared to typical measured values. Continuum background contamination was not included since the observed background level was negligible compared to the line brightness, except for He II near the divertor targets.

Of the 3 targeted spectral lines, C III is the most affected by contamination from nearby lines, specifically at 463.9nm, 464.2nm, 464.9nm, 465.9nm, and 466.2nm. The positions and relative intensities of all except the 465.9nm line are consistent with the O II $3s^22p^23p \rightarrow 2s^22p^23s$ multiplet, meaning these lines are expected to shift and change intensity together. The 465.9nm line may be a C IV charge exchange feature [Zaniol et al., 2001]. Calculated as described above, the mean contamination errors were -1.0km/s, 1.0km/s and -0.6km/s for the 3 delay plates, indicating that for the 4.6mm crystal the contamination tends to cause an apparent blue-shift, while with the 6.5mm crystal it causes an apparent red-shift of a similar magnitude. The standard deviations were 1.2 km/s, 0.9km/s and 0.7km/s for measurements with the 4.6mm, 6.5mm and 9.8mm delay plates respectively. These errors and their variability are an order of magnitude smaller than typical H-Mode C III flows measured with the instrument, and are therefore not considered to pose a serious problem for such measurements, however their potential presence could become significant in some L-Mode cases where measured C III flow speeds can be as small as a few km/s.

He II is also affected by the same O II multiplet as contaminates the C III line, and to a lesser extent the bright He I line at 471.3nm and an unidentified feature at 466.5nm. The narrow bandwidth of the He II filter and distance of these features from the He II line greatly reduces their effect. The distribution of calculated apparent flow for He II with these features had mean and standard deviation $\approx$ 0.1km/s for all delay plates (complete results are given in table 4.1), consistently below the measurement noise limit and not significant. Note that while this result holds for the midplane and some divertor data, it does not take into account the situation seen in divertor floor measurements in section 3.1, where the He II line can be very weak and swamped by very complex, very high relative brightness contamination (in these cases the contamination is sufficient to make the data useless.)

The results for C II give mean and standard deviations of $\leq$ 0.5km/s for the thicker two delay plates, and approaching 1km/s for the 4.6mm plate. Since the 4.6mm plate is poorly optimised for C II contrast and therefore not used, contamination errors are not a concern for the C II line.

Overall, based on the observations of spectral contamination on MAST and these numerical calculations, the errors due to spectral contamination are expected to be small or negligible for most measurements. These results are summarised in table 4.1.

4.6 Tomographic Inversion

While the line-integrated emissivity and line-average flow obtained from the instrument can provide substantial qualitative insight into flow behaviour, they are not suited for direct quantitative comparison with theoretical or modelling results,

Spectral Line	Estimated error from contamination (km/s)		
	$L_{\mathrm{Delay}} = 4.6\mathrm{mm}$	$L_{\mathrm{Delay}} = 6.5\mathrm{mm}$	$L_{\mathrm{Delay}} = 9.8\mathrm{mm}$
C III	-1.0 ± 1.2	1.0 ± 0.9	-0.6 ± 0.7
C II	0.7 ± 0.9	-0.5 ± 0.4	0.3 ± 0.3
He II	-0.1 ± 0.1	0.1 ± 0.1	0.0 ± 0.1

Table 4.1: Summary of estimated flow measurement errors due to spectral contamination, for each spectral line and delay plate combination. The results are presented in the format mean ± standard deviation of the distribution of simulated errors, as described in the text.

which instead give spatially localised values. There are two possible approaches for enabling such comparisons. The first is to start from the output of a plasma model which produces 2-3D profiles of the impurity emissivity and flow which we would like to compare with the measurements. In the manner described in section 4.1, these model plasma profiles can be used to produce simulated line-integrated data, which can then be directly compared with the measurements. The second approach is to start from the line-integrated measurements, and using tomographic inversion techniques infer the local emissivity and flow profiles which can then be directly compared to the model outputs. The latter approach is now discussed in more detail.

Tomographic reconstruction of 2D-3D profiles from line-integrated measurements is a common problem in remote sensing measurements, both in fusion diagnostics [Boswell et al., 2001; Odstrčil et al., 2014; Tal et al., 2013] and other fields such as geophysics [Patella and Patella, 2009] and atmospheric science [Flores and Gradinarsky, 2000; Tucker, 1998]. As such, several well established techniques exist for accomplishing this. Furthermore, inversion of CIS results to obtain poloidal flow and emissivity profiles for comparison with modelling has previously been demonstrated for C III flows on DIII-D [Howard et al., 2010a; Weber et al., 2012]. The techniques used in the present work take a similar approach.

4.6.1 The Tomography Problem

The technique used in this work for tomographic reconstruction belongs to the class of so-called ‘pixel based’ techniques. The emissivity and flow are first assumed to be toroidally symmetric, and therefore can be completely described by a 2D profile in the poloidal (R, Z) cross-section of the tokamak. The poloidal cross-section is then discretised into a regular grid with n spatial pixels (grid cells), each having its own value of local emissivity and flow which we wish to determine, and the whole grid representing the 2D profile we wish to recover. A model of the diagnostic response is then used to calculate the effect of each grid cell's emissivity and flow on each of the n line-integrated measurements provided by the diagnostic (n being the number of pixels in the line-integrated images), resulting in two systems of linear equations: one for emissivity and one for flow, relating the m line-of-sight measurements and the emissivity and flow in the n grid cells. Such systems can be written compactly in matrix form as $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}$, where $\boldsymbol{y}$ is an m-element column vector containing the measurements, $\boldsymbol{x}$ is an n-element vector containing the profile we wish to recover, $\boldsymbol{b}$ is a vector representing the measurement errors and $\boldsymbol{A}$ is an m-by-n matrix, which will be referred to as the response matrix, describing the diagnostic response. Solving for $\boldsymbol{x}$ amounts to a large least squares fitting problem.

The first step in setting up the problem is generating the reconstruction grid in the (R, Z) plane. For the present work, the grid was chosen such that its extent completely contained all the diagnostic lines of sight. If some lines of sight are allowed to extend outside the reconstruction grid, their response cannot be fully described by the response matrix which can lead to artefacts in the reconstructions. A typical reconstruction grid for the MAST divertor is shown in Figure 4.11, along with the projections of several of pixels' typical sight-lines on to the poloidal cross-section (blue lines). Square grid cells of typically 1cm - 1.5cm on each side were used, and were found to produce good quality inversions with reasonable computational requirements for the inversion. We now consider construction of the response matricies for emissivity and flow, given the reconstruction grid and instrument sight lines (calibration of the sight lines is discussed in section 5.2.3). Note that not all grid cells shown in figure 4.11 are well constrained or necessarily included in the inversions, as will be discussed later.

4.6.2 Response Matrix Generation

The response matrix describes the relationship between the n inversion grid cells and m line integrated measurements given by the demodulated brightness and $\bar{v}$ images. We wish to construct response matrices for both emissivity and flow reconstructions, which will be denoted $\boldsymbol{\mathcal{E}}$ and $\boldsymbol{\mathcal{V}}$ respectively. The measurements were typically binned 4x4 (to 256x256 pixels) before inversion, to reduce the size of the required response matrix and therefore the computational requirements of the inversion process. Due to the spatial resolution of the instrument discussed in section 4.4, no spatial resolution is lost in the direction perpendicular to the fringes due to this binning, although spatial resolution is lost in the direction parallel to the fringes.

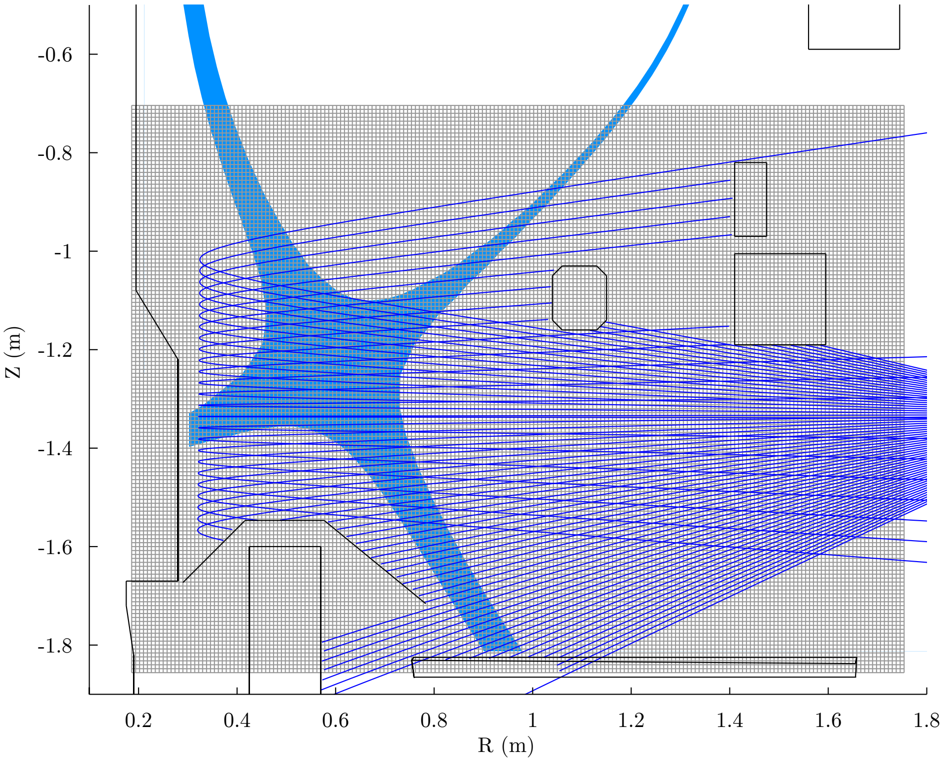

Figure 4.11: Poloidal cross-section of the MAST divertor showing a typical reconstruction grid for emissivity and flow tomography. The grey grid is the reconstruction grid, and the blue curved lines are camera sight-lines from a particular image column projectsed on to the grid. The blue shaded area represents the approximate shape of the divertor plasma from which emission is expected to be observed.

In calculating the system response the finite depth-of-field of the imaging system is ignored, and each line of sight is assumed to collect light from a narrow ‘pencil’ beam through the plasma, with the collection power constant along the line-of-sight. This is a reasonable approximation for inversions of the MAST CIS images, because 1) After the 4x4 binning of the images, the binned pixel size is larger than the expected blur radius for out-of-focus areas of the plasma, and 2) Due to the tangential view of the plasma, the pixels contributing most information to the inversion are dominated by light emitted close to the in-focus object plane. Given this pencil beam model, and without concern for absolute calibration, the i^{th} brightness measurement is given by:

$$S_i \propto \sum_{j=1}^{n} e_{0,j} L_{ij}, \tag{4.6.5}$$

where $e_{0,j}$ is the plasma emissivity in the j^{th} inversion grid cell and L_{ij} is the length of the i^{th} sight-line which falls within that grid cell. This equation constitutes one row of the response matrix, which is given by $\boldsymbol{\mathcal{E}}_{ij} = L_{ij}$ if all detector pixels have the same effective collecting power (or equivalently if the data images are flat field corrected before inversion). The response matrix is calculated by finely discretising each sight-line j into line segments along its length and calculating the R, Z coordinates of each segment’s centre point, to determine which reconstruction cell the segment falls within. The length L_{ij} is then given by $L_{ij} = \sum_k l_k$, where the sum is over all sight-line segments k which fall within the j^{th} grid cell and l_k is the length of the k^{th} line segment. Since a simplified vacuum vessel geometry is used in the sight-line calculation, some sight-lines which are unobstructed in the sight-line model are in fact blocked by in-vessel components in the real measurements. To account for this, image masks are created manually which indicate which sight-lines are affected by this. These sight-lines are then excluded from the fitting process, i.e. they are not included as rows in the response matrix. In addition some grid cells do not encounter any diagnostic sight-lines and are completely unconstrained (e.g. some grid cells fall within machine components or within vessel walls). While these are not explicitly excluded from the matrix generation or fitting, they are represented by a row of zeros in the response matrix and are not involved in the inversion. The brightness response matrix is only a function of the viewing geometry, and therefore needs to be recalculated only when the viewing geometry is changed.

For flow reconstruction the instrument response is more complex, as we see from equation (4.2.4) that the line-integrated flow measurement depends on both the magnitude and direction of the local flow, in addition to the local emissivity. With a single camera, insufficient information is available to recover both the magnitude and 3D direction of the flow. Instead, we assume that the impurity flow is primarily parallel to the magnetic field due to the confining effect of the field, i.e. $\boldsymbol{v} \approx v_{||}\hat{\boldsymbol{B}}$, where $v_{||}$ is the parallel flow speed and $\hat{\boldsymbol{B}}$ is a unit vector in the direction of the local magnetic field. Since the magnetic field on MAST is routinely reconstructed from magnetic probe measurements using the EFIT code, this can provide a constraint for the flow direction and we only have to solve for the scalar quantity $v_{||}$. This is the same approach as used in Howard et al. [2010a]. The response matrix elements can then be calculated as:

$$\boldsymbol{\mathcal{V}}_{ij} = \frac{e_{o,j}}{\sum_{j'=1}^{n} \boldsymbol{\mathcal{E}}_{ij'} e_{0,j'}} \sum_{k} \hat{\boldsymbol{B}}_k \cdot \boldsymbol{l}_k, \tag{4.6.6}$$

where again the sum over k is over each sight-line segment falling within the j^{th} reconstruction grid cell, $\boldsymbol{l}_k$ is the vector line segment along the line of sight, and $\hat{\boldsymbol{B}}_k$ is a unit vector in the direction of the magnetic field at the centre point of $\boldsymbol{l}_k$. Note that the emissivity e_0 and magnetic field direction $\hat{\boldsymbol{B}}_k$ are time-varying properties of the plasma; this means that individual flow response matrices must be generated for each image we wish to invert. The values of e_0 are obtained from inversion of the emissivity, and therefore any reconstruction errors in the emissivity cause inaccurate calculation of the flow response matrix leading to further errors in the flow inversion. Furthermore, not all measurement sight-lines from the line-integrated image contain useful flow information. Specifically, sight-lines where no light is detected cannot carry any spectral information, and must be excluded from the fitting (along with sight-lines with sufficiently low SNR that the flow information is dominated by noise). Therefore only sight-lines which show brightness above a threshold value are included when generating the response matrix.

4.6.3 Inversion Algorithm

The response matrices for the inversions are large, sparse matrices, with typically around 10^9 elements of which $< 1\%$ are non-zero. This sparsity occurs because any given line-of-sight only passes through a small fraction of the inversion grid cells. The system is also overdetermined, i.e. the number of measurements is larger than the number of unknowns in the profile. The inversion problem is therefore to solve a noisy (i.e. inconsistent), overdetermined, sparse system, and several numerical methods exist to achieve this [Censor, 1983; Paige and Saunders, 1982]. The technique used in this work belongs to the group of algebraic reconstruction techniques, which find the solution for $\boldsymbol{x}$ in $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}$ iteratively by minimising the residuals

$\boldsymbol{y}-\boldsymbol{Ax}$. These techniques have the advantage that *a-priori* constraints such as non-negativity or smoothness of the profiles can be easily imposed at each iteration. The particular algorithm used in this work is closely based on the Simultaneous algebraic reconstruction technique (SART), first proposed in Andersen and Kak [1984]. After x is initialised as $x_j^0 = \exp(-1) \forall j$, each iteration of the algorithm updates each grid cell according to:

$$x_j^{k+1} = x_j^k + \lambda \frac{\sum_{i=1}^{m} a_{ij}(y_i - \sum_{j=1}^{n} a_{ij} x_j^k)}{|\sum_{i=1}^{m} a_{ij}|}, \tag{4.6.7}$$

where λ is an empirically chosen relaxation parameter. In this update step x_j is corrected with a weighted sum of the fitting errors from all the measurements y_i which are influenced by x_j. The sum is weighted according to the dependence of each y_i on x_j. In other words, the error in fitting every measurement to which x_j contributes is considered simultaneously and used to correct x_j's' value. By considering the fitting error in all measurements at once in this way, each iteration is robust to noise in individual measurements. The denominator normalises the update by the sum down the j^{th} column of the response matrix, i.e. by the total contribution of x_j to the recovered profile. Such 'column normalised' techniques are known to reduce reconstruction artefacts relative to row action algebraic reconstruction techniques [Willis, 2000]. Additionally for emissivity reconstructions, non-negativity of the emissivity is enforced by setting any grid cells for which $x_j^k < 0$ to zero at each iteration. Although this is the only constraint usually used for the reconstructions in this work, in principle it is possible to apply further constraints such as smoothness of the reconstructed image by using a relaxation *image* $\boldsymbol{\lambda}$ rather than a single parameter for all grid cells. Iterations are stopped when the fractional change in the Euclidean norm of the error $\boldsymbol{y} - \boldsymbol{Ax}$ is less than 10^{-4} between successive iterations.

4.6.4 Reconstruction performance

The reconstruction process was tested using the simulated data from section 4.1. First, the 'ideal' line-integrated brightness and phase images in figure 4.2(b) and (c) were inverted to establish the performance of the inversions with idealised input data. This also tests the validity of the flow line integral in equation (4.2.4), since this relationship was not assumed in the generation of the data. The results are shown in fig. 4.12(b), with the original input profiles in (a). The emissivity profile appears generally well reconstructed, but is somewhat noisy in the Private flux region (PFR) and outside the outer divertor leg. Artefacts following the paths of sight-lines can be seen in the PFR adjacent to the sloping section of divertor coil

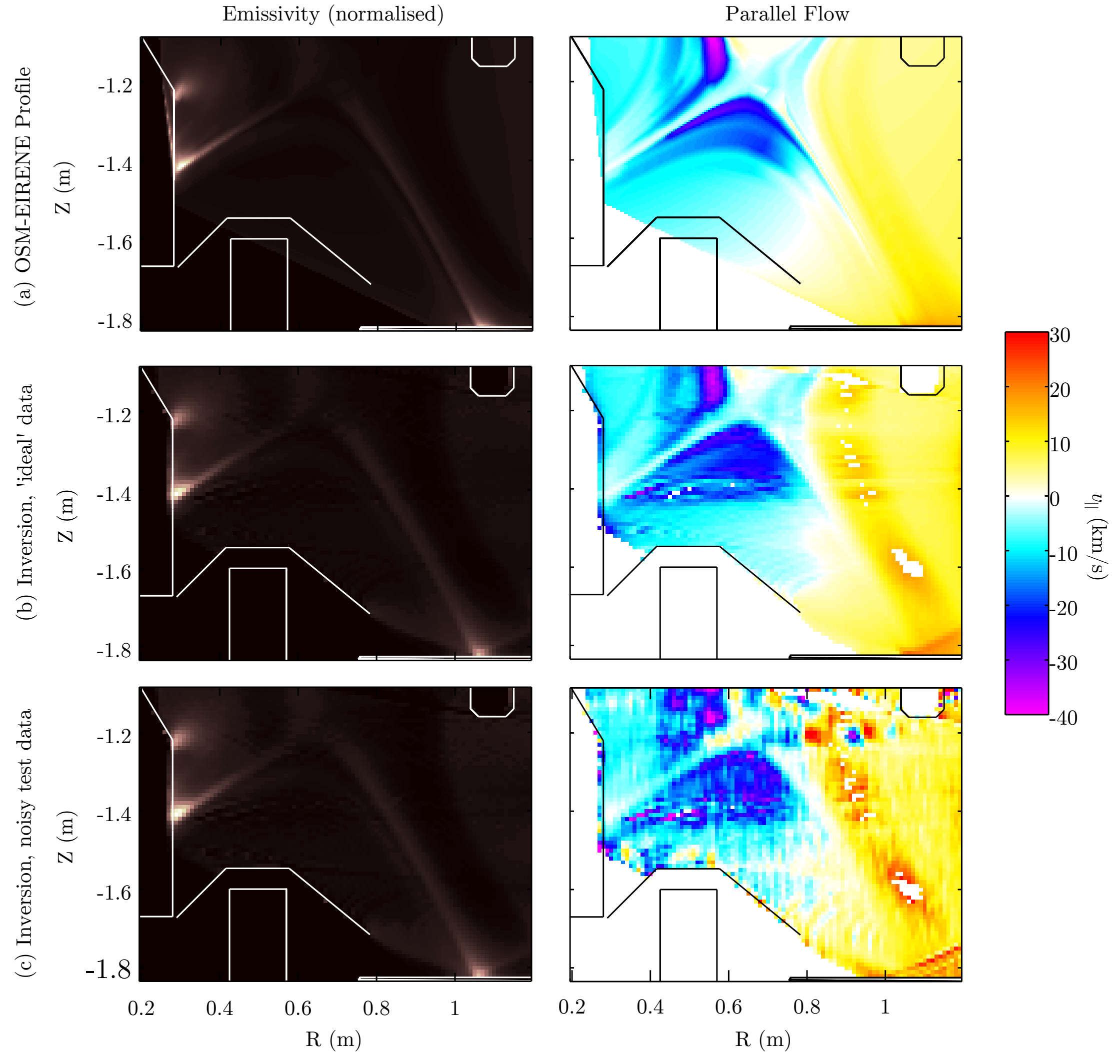

Figure 4.12: Testing of tomographic inversion of emissivity and flow. (a) Top: OSM-EIRENE profiles used to generate test data. (b) middle: Profiles reconstructed from simulated phase & brightness images in figure 4.2(b) and (c). (c) Bottom: Profiles reconstructed from demodulated, noisy test data, as in figure 4.2(a).

armour and leading out from the outboard strike point. These are very similar to artefacts seen in tomographic inversions of camera data from the same plasma view on MAST in independent work [Lisgo et al., 2009], suggesting these features are related to the viewing geometry. The flow profile is also generally well reconstructed, however significant differences from the ideal profiles are seen in the private flux region. These may be partially caused by the errors in the emissivity inversion in this region. Deviations from the ideal profile are mainly in grid cells with very low emissivity, and which therefore contribute little flow information to the line-integral measurements and are poorly constrained in the inversion. This behaviour is expected, since naturally flow information can only be obtained from parts of the plasma where light is being emitted.

The reconstruction was also tested as part of the complete data analysis chain: the brightness and phase were demodulated from the noisy simulated data in figure 4.2(a), and then inverted. The overall appearance of the emissivity reconstruction is not much different from the ideal data inversion, however the artefacts seen in that case appear more pronounced. The flow inversion appears significantly worse than the ideal data case, both in recovery of spatial features and the presence of noise and artefacts. Again as expected this is particularly the case for parts of the profile where the emissivity is low, both due to the poor constraint of these grid cells as previously stated, and the low SNR of the corresponding measurements when noise is included in the simulated data. When viewing inversions of real data, for this reason the flow inversion is only considered where the emissivity profile lies above a threshold value. Some artefacts in the image coincide with those seen in the emissivity reconstruction, indicating the effect of the emissivity reconstruction on the flow. The effects of the anisotropic fringe phase noise can be clearly seen in the inverted profile. Apart from the artefacts mentioned for the ideal data inversions, ‘blobs’ of positive and negative flow appear near the poloidal field coil in the top right of the image ($R \approx 0.8 - 1$m, $Z \approx -1.2$m). These are due to residual ringing artefacts at the coil edges from the phase demodulation.

The agreement between the ideal and reconstructed flow profiles is illustrated in figure 4.13, which shows horizontal slices through the flow profiles above and below the X-Point. Both the idealised data inversions (green lines) and full analysis chain tests (red lines) generally show good quantitative agreement with the input profiles, however the noisy data inversion does show loss of fine spatial detail (e.g. in fig. 4.13(b), $R = 0.4 - 0.55$m) and spurious flow artefacts (these slices also do not illustrate the ringing artefacts previously mentioned). The flow noise level in parts of the profiles away from any artefacts is not significantly amplified over the noise level on the line integrated measurements. Future possibilities for improving the quality of reconstructions from noisy data include applying smoothness constraints to the flow reconstruction, and accounting for the anisotropic spatial resolution when building the response matrix. These tests also do not account for possible errors from uncertainty in the viewing geometry calibration, errors in EFIT, the presence of reflections in real images or the presence of cross-field or non-axisymmetric flows.

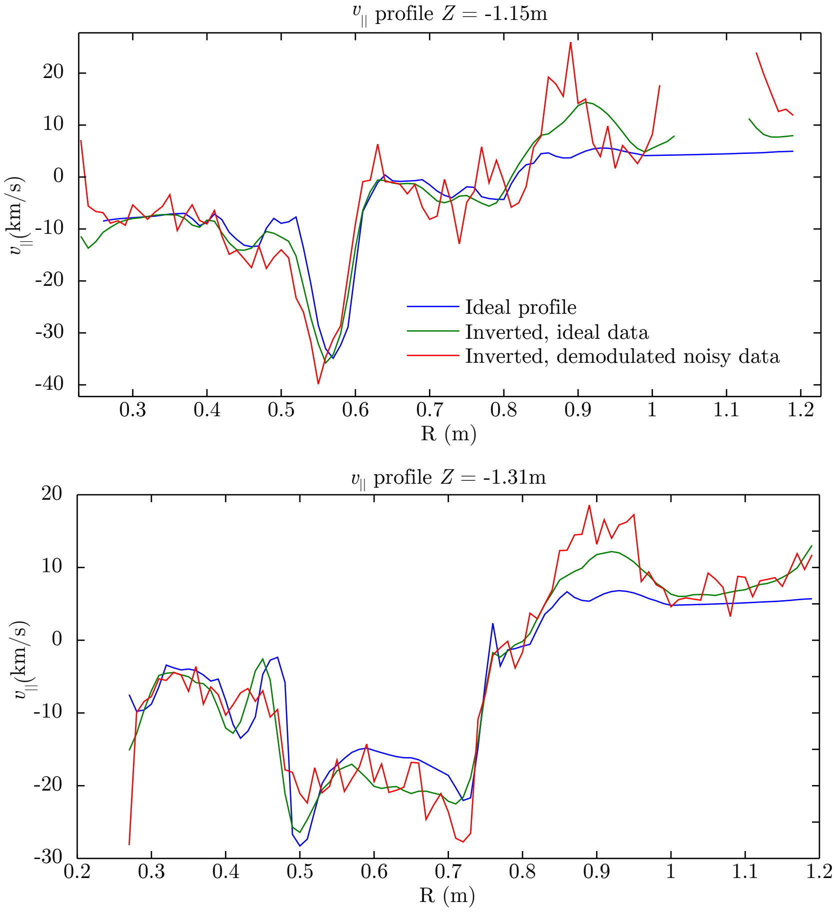

Figure 4.13: Horizontal slices of the $v_{||}$ profiles shown in figure 4.12 at $Z = -1.15$m (top) and $Z = -1.31$m, showing ideal and inverted profiles.

4.7 Summary

This chapter has presented the data analysis methods used with MAST CIS data, including numerical experiments investigating sources of measurement error. This was achieved using simulated line-integrated data based on plasma profiles from OEM-EIRENE simulations of MAST.

The data analysis process begins by removal of bright spots and streaks in the images due to D-D fusion neutrons in NBI heated discharges, before the underlying brightness image is extracted using frequency domain filtering and boxcar smoothing perpendicular to the fringes. The extracted brightness is then factored out, and the flow information (fringe phase) is extracted using the analytic signal representation of the resulting signal. Windowing in Fourier space and apodisation of the input data around sharp jumps in brightness are used to reduce artefacts due to sharp spatial features and image noise, at the cost of smoothing the spatial response to flows.

Pixel-based tomographic inversion using the SART algorithm has been implemented in MATLAB and tested with simulated divertor data, including integrated testing with the interferogram demodulation. The results show good quantitative agreement between the plasma profiles used to generate the test data and the recovered profiles from the simulated measurements. However, for noisy data the flow inversion shows some ringing artefacts at sharp image edges, and larger errors for parts of the plasma profile with low emissivity. These must be kept in mind when applying the techniques to real data, and inverted flow profiles can only be obtained from parts of the plasma which are emitting sufficiently brightly.

In the next chapter, we go on to present experimental characterisation of the MAST CIS diagnostic and its constituent components.

Chapter 5

Instrument Characterisation & Calibration

This chapter presents the characterisation of various individual components of the MAST CIS diagnostic and of the integrated system, and also describes the methods developed for calibration of the instrument for plasma measurements.

5.1 Detector Characteristics

5.1.1 Linearity

Accurate recording of the interference fringe pattern, absolute intensity measurements, and accurate tomographic inversion of both emissivity and flows all require the detector's response to the incoming light flux to be well known, and ideally linear. To test the linearity of the camera, images were recorded with the instrument viewing an absolutely calibrated integrating sphere and white light source without a band pass filter. The camera exposure settings remained constant while the integrating sphere luminance was changed between images, by opening and closing a slit between the sphere and the light source. The mean signal in a 100x100 pixel area at the centre of the image (after dark signal subtraction) was measured as a function of sphere luminance, and is shown in Fig. 5.1. This is not well fit by a straight line through the origin (reduced $\chi^2 = 54$) as would be expected for a linear detector response, but is much better approximated by a cubic polynomial (reduced $\chi^2 = 1.5$). This cubic fit indicates an integral non-linearity of approximately 4%. In order to correct for this when analysing subsequent data from the camera, a linearising look-up-table was created based on the difference between the linear and cubic fits at each Analog-to-Digital Unit (ADU) value. This correction was tested using a second set of images where the exposure time was varied while keeping the integrating sphere luminance constant. When performing a linear fit of camera signal vs exposure time on this data, the linearity correction resulted in a reduction of the reduced χ^2 value from 44 to 2.5. In the remainder of this work, unless otherwise stated this linearisation has been applied to all data from the camera.

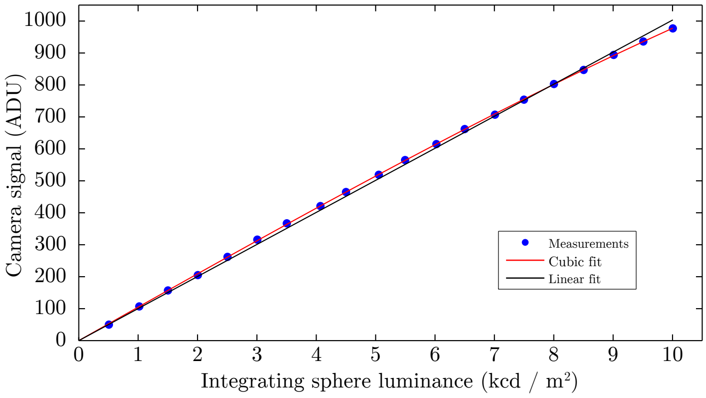

Figure 5.1: Linearity testing of the detector: measured camera signal in integrating sphere flat field images as a function of the integrating sphere luminance, and the best linear and cubic polynomial fits to the data. Error bars are too small to be visible on this scale.

5.1.2 Conversion Gain & Camera Noise

To evaluate the absolute light sensitivity of the instrument and to assess the relative contribution of different noise sources in the raw images (i.e. camera noise vs photon shot noise), it is necessary to know the conversion gain K of the camera. This relates the number of photoelectrons collected by a pixel, S_e, to the camera's digital output S_ADU, according to $S_\mathrm{ADU} = S_\mathrm{e}/K$. This can be measured using the statistical properties of photon shot noise, using the well known mean-variance method. This method is based on the fact that photoelectron counting noise has a poisson distribution, which for large numbers of photoelectrons is well approximated by a Gaussian distribution with variance equal to its mean, i.e. $\sigma^2_{S_\mathrm{e}} = \bar{S}_\mathrm{e}$. Using the definition of K above, this can be written in terms of the digital output as $\sigma^2_{S_\mathrm{ADU}} = \bar{S}_\mathrm{ADU}/K$. The gain can therefore be determined by measuring signal variance as a function of mean signal level, which should give a straight line plot with a gradient of $1/K$. To perform this measurement for the Photron APX-RS camera, a series of images were taken with a variable DC white light source (ThorLabs OSL1-EC fibre illuminator) illuminating a flashed opal diffuser which was viewed by the CIS system with no band pass filter present. A video sequence of 800 frames was captured while slowly increasing the illumination level, and analysis was performed on a 96x96 pixel region of the images where the illumination was most uniform. Images were processed pair-wise through the sequence: for each pair of consecutive frames S_1 and S_2, $\bar{S}_{\mathrm{ADU}}$ was obtained from the mean pixel value in the region of interest including pixels in both images in the pair. A difference image $\Delta S_{\mathrm{ADU}} = S_2 - S_1$ was then created, and the signal variance corresponding to $\bar{S}_{\mathrm{ADU}}$ calculated as $\sigma^2_{ADU} = \mathrm{Var}(\Delta S_{\mathrm{ADU}})/2$, where $\mathrm{Var}(\Delta S_{\mathrm{ADU}})$ was determined from fitting of the histogram of ΔS_{ADU}. This pair-wise processing reduces the effect of any fixed pattern noise and small irregularities in illumination from the measurements. The resulting mean-variance plot is shown in Fig. 5.2, where grey points represent the data without any camera linearity correction and blue points have had the linearising look-up-table applied.

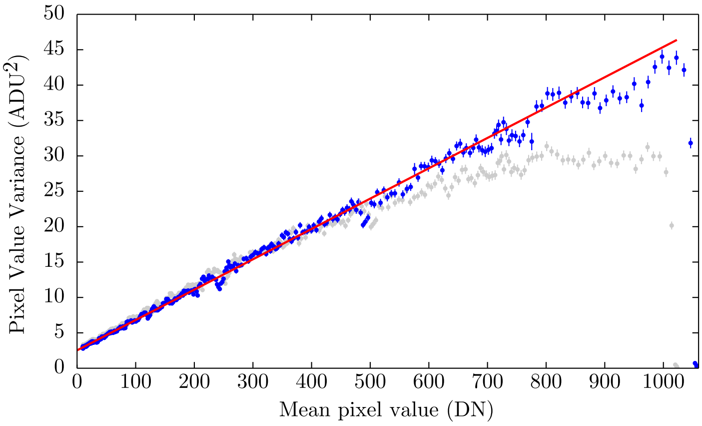

Figure 5.2: Measured mean-variance plot for the detector. Grey points have not been corrected for nonlinearity, blue points have had the linearising look-up-table applied. The solid red line indicates the best linear fit to the blue points for mean values below 700 ADU.

Without the linearisation applied, substantial variation of the gain (i.e. differential nonlinearity) is observed, particularly towards the high end of the signal range. With the linearising look-up-table from the previous section applied this effect is largely removed, and the data below $\bar{S}_{\mathrm{ADU}} = 700$ ADU can be fit with a straight line to give the conversion gain. The best fit to the data in Fig 5.2 gives a value of $K = (23.3 \pm 0.2)$ e$^-$/ADU. The y intercept of the best fit line gives an estimate of the total noise associated with the camera under no illumination, i.e. including read noise, dark current noise & other electronics noise, of $\sigma_{\mathrm{cam}} = (1.59 \pm 0.01)$ ADU. This is in good agreement with camera noise measurements for the same exposure settings using dark frames from a different data set, which gave a value of $\sigma = (1.609 \pm 0.008)$ ADU. Using the measured conversion gain, this is a camera noise level of $\sigma_{\mathrm{cam}} = (37.0 \pm 0.4)$ e$^-$ at 20ms exposure time. While this is large compared to high performance CCD detectors, it is not unusual for this class of high speed CMOS detector. At this level camera noise is the dominant noise source below about 59 ADU, beyond which shot noise starts to dominate. This camera noise level was used in all subsequent calculations involving image noise; whilst the noise is exposure dependent due to the contribution of dark current noise, the 20ms exposure value provides a 'worst case' measurement and hence a conservative estimate of the noise performance of the instrument.

5.1.3 Offset Level Drift

During various experiments the camera was observed to exhibit a drifting signal offset with changes in ambient temperature. This is attributed to sensitivity of both the dark current level and camera electronics in the un-cooled camera to changes in temperature, and the fact that dark level calibration is performed as a one-off procedure prior to taking data. Changes in temperature after the calibration therefore result in a drift of recorded pixel values. To investigate the sensitivity of this effect, dark frames with the camera's dark level subtraction disabled were captured every 5 minutes over a 37 hour period while the ambient temperature was recorded with a data logger. The resulting time histories of temperature and recorded dark signal are shown in Figure 5.3, which clearly shows the offset level tracking ambient temperature variations. From this data, the temperature sensitivity of the offset was estimated to be approximately 1.8 ADU/°C. Furthermore, after the camera was first powered on an increase in the dark level by up to 22ADU was observed over the first 2.5 hours of camera operation as the camera heated up. This effect does not pose a significant problem for plasma measurements, since dark level recalibration is usually performed between each plasma shot. However, it can cause a problem for calibration measurements, as discussed in more detail in section 5.3.3.

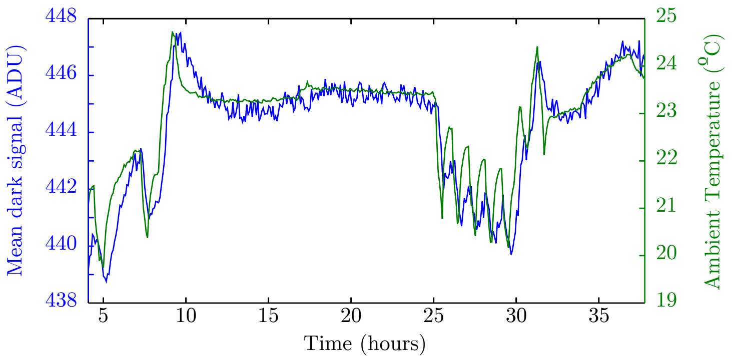

Figure 5.3: Time history of the recorded dark signal level (blue) at 20ms exposure under varying ambient temperature (green).

5.2 Imaging Optics Characterisation

5.2.1 Band Pass Filters

The delivered band pass filters were all within the Centre wavelength (CWL) and FWHM specifications defined in section 3.4.2. The manufacturer provided detailed measurements of the filter transmission profiles for normal incidence for each filter, which are shown along with the target ion line profiles in figure 5.4. By multiplying these transmission curves with synthetic spectral line profiles, the fractional transmission of the plasma lines through each filter (at the image centre) is expected to be 72%, 67% and 55% for the C III, C II and He II lines respectively. By repeating the calculations used in section 3.4.2 to define the filter specifications and using these measured profiles, all 3 delivered filters were found to meet the desired requirements for vignetting, apparent flow changes within the field of view and rejection of spectral contamination.

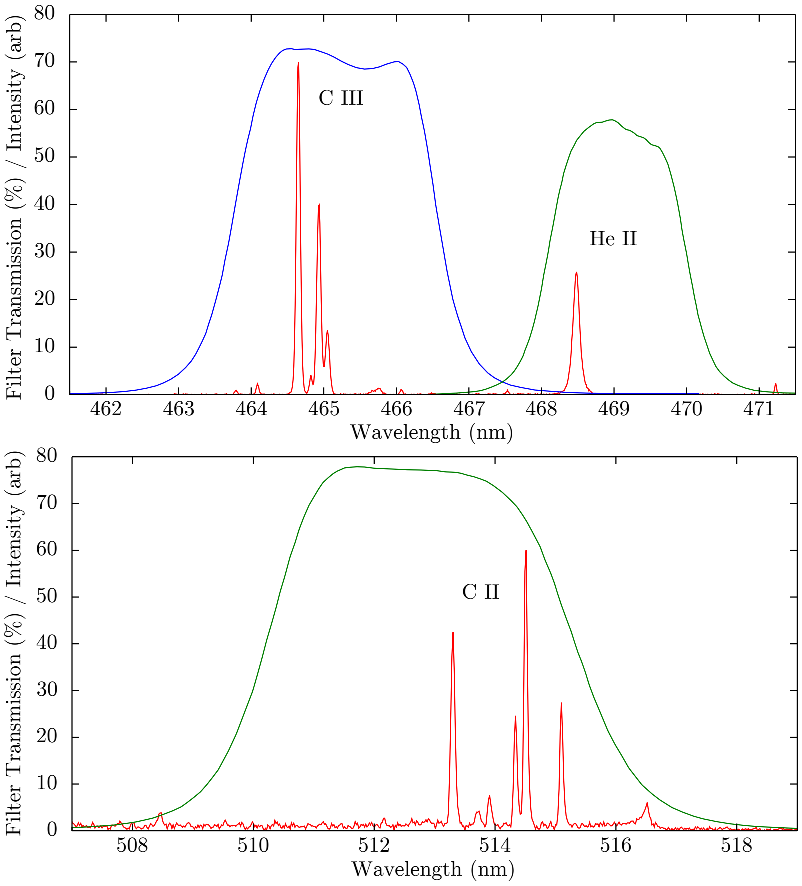

Figure 5.4: Transmission profiles of the delivered filters for each spectral line, with the line structures plotted for reference.

In figure 5.4 it is apparent that the C II line is far off-centre in the filter profile and has some components in the long wavelength wing of the filter. This is due to both the design to reject the parasitic feature at 516.5nm, and the fact that the delivered filter is at at the very short wavelength end of the specified tolerances. The possible effect on flow measurements was estimated by multiplying synthetic C II line shapes, with a range of Doppler shifts, by the filter transmission profile and calculating the fringe phase associated with the resulting spectrum. The effect was found to be negligibly small ($< 0.5\%$ flow error), showing a surprising robustness to the position of the spectral line within the filter passband.

An unexpected phenomenon observed when the filters were first inserted into the diagnostic was the appearance of structure in flat field images not present without the filters. The form of this structure for each filter is shown in figure 5.5, where the effects of vignetting have been removed to isolate only the patterns introduced with the filters. These patterns are not radially symmetric about the image centre as would be expected for an angle of incidence dependent effect in the filters, but show somewhat square shapes (despite the absence of any square apertures or optical components besides the detector). The cause of this effect is not clear from the current data; more detailed measurements of the spatial and angular dependence of the filter transmission would be desirable to attempt to understand it. There is no evidence that this has a significant effect on the fringe phase i.e. Doppler shift measurements, since no similar spatial structure is seen in either calibration or plasma fringe phase measurements. The patterning in the intensity images is accounted for in plasma data by using flat field images with the filters in place, such as those in fig. 5.5, to perform flat field correction of the raw data images.

5.2.2 Light Sensitivity & Signal Calibration

The overall light sensitivity of the instrument was tested against the design values by taking images of an absolutely calibrated integrating sphere and light source combination through each of the 3 band pass filters. The expected output signal S_{DN} for each case, as a function of distance from the detector centre, d, was estimated as:

$$S_{\mathrm{DN}}(d) = \eta E(d) T_p \frac{1}{K} \frac{\lambda_0}{hc} \int L(\lambda) T_f(\lambda) d\lambda. \tag{5.2.1}$$

Here $\eta = 0.45$ is the detector quantum efficiency from the camera specifications, $E(d)$ is the design estimate of the effective étendue from section 3.4.3, $T_p = 0.28$ is a factor accounting for losses due to the polarisers (from the polariser specifications),

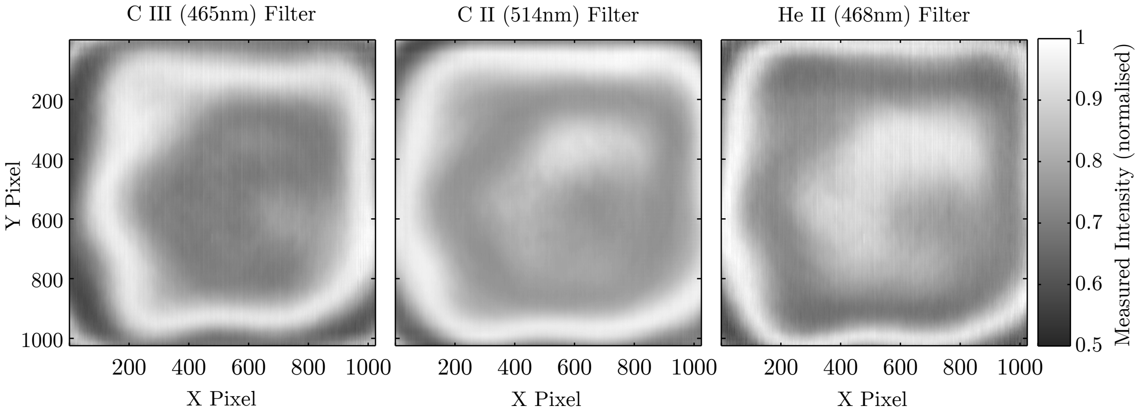

Figure 5.5: Intensity structure introduced into flat field images when the band pass filters are inserted into the system. Each image has been divided by a flat field image with no filter in order to remove the effects of vignetting; only variations associated with the presence of the filters are shown. The horizontal/vertical field of view corresponds to incidence angles through the filter of ±3.3°.

K is the camera conversion gain from section 5.1.2, hc/λ_0 is the photon energy at the centre of the filter pass band, $L(\lambda)$ is the known spectral radiance of the integrating sphere and $T_f(\lambda)$ is the filter transmission profile. The measured profiles from the integrating sphere images with each filter were then compared with these estimates. The results are shown in Fig. 5.6 (left), where for clarity of display the profiles have been normalised such that all the predicted profiles are equal to 1 near the image centre. Figure 5.6 (right) shows the ratios of the predicted and measured signals for each filter. The measured signal levels are in reasonable agreement with the predicted profiles given the simplicity of the calculations, with the C II and He II results within approximately 20% over the entire field. Due to the anomalous structure in the images associated with the filters, discussed in the previous section, there are significant departures from the predicted shape of the profile compared to the calculations. These are highlighted in figure 5.6 (right) were the minima and maxima and of these structures are clearly visible, and particularly affect the throughput of the C III filter close to the image centre.

Known radiance integrating sphere images such as those used in this section can also be used to calibrate the plasma data to determine the absolute impurity emissivity, although this was not performed routinely for the MAST measurements which are primarily concerned with the flows.

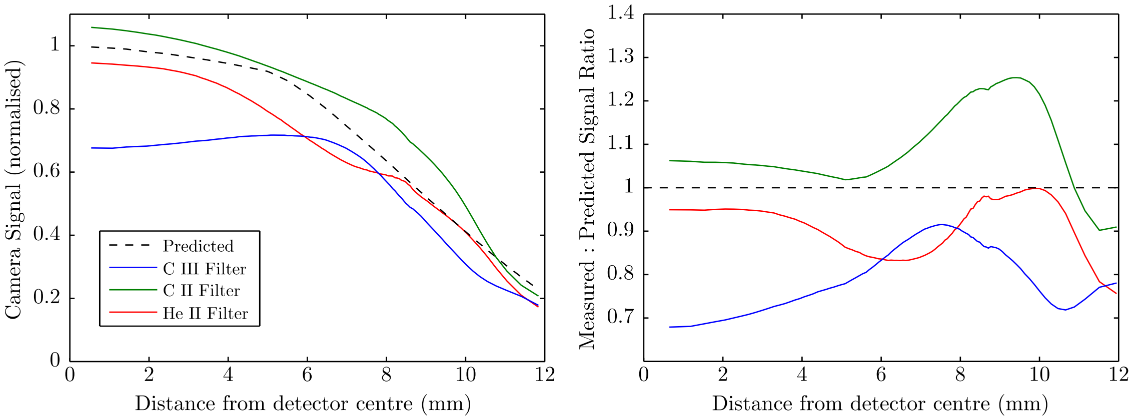

Figure 5.6: Measured intensity profiles across the detector for a calibrated integrating sphere with $f_1 = 17$mm. Left: profiles for each filter compared with design prediction. Right: ratio of measured signal to design prediction.

5.2.3 Pixel Line of Sight Calibration

Proper interpretation of plasma data requires accurate knowledge of the mapping between image pixel coordinates and physical locations in the tokamak. Specifically, for tomographic inversion it is necessary to know, for each pixel in the image, the pixel's line-of-sight vector in 3D space. This is achieved by applying a pinhole camera model to the imaging system, the parameters of which are measured as part of the diagnostic calibration. This type of calibration is common in machine vision and photogrammetry applications, for example see Zhang [1999], Heikkila and Silvén [1997] and references therein. The camera model parameters are divided into two groups: intrinsic and extrinsic camera parameters. The former describe the properties of the optical system e.g. focal length and distortion, while the latter describe the position and orientation of the instrument with respect to the tokamak. The following sections discuss the camera model and how the parameters were calibrated.

Intrinsic Parameters

To describe the imaging properties of the optical system we use a pinhole camera model including distortion. This is formulated in a Cartesian coordinate system with the camera pinhole at the origin and the z axis pointing in the camera's direction of view; this coordinate system will henceforth be referred to as the camera frame or camera coordinate system. Consider the projection of a point in space $\boldsymbol{P}$, with coordinates X, Y, Z in the camera coordinate system, on to the image. The normalised perspective projection for an ideal pinhole camera is simply given by:

$$\begin{pmatrix} x_n \\ y_n \end{pmatrix} = \begin{pmatrix} X/Z \\ Y/Z \end{pmatrix}. \tag{5.2.2}$$

The effects of distortion in the real optical system are included using the Brown-Conrady distortion model [Brown, 1966; Conrady, 1919], which describes radial (i.e. barrel or pincushion) and tangential (wedge-prism like, due to decentring or tilting of optical components) distortion. The normalised projection including distortion then becomes:

$$\begin{pmatrix} x_d \\ y_d \end{pmatrix} = \left[1 + k_1 r^2 + k_2 r^4\right] \begin{pmatrix} x_n \\ y_n \end{pmatrix} + \begin{pmatrix} 2p_1 x_n y_n + p_2(r^2 + 2x_n^2) \\ p_1(r^2 + 2y^2) + 2p_2 xy \end{pmatrix}, \tag{5.2.3}$$

where $r^2 = x_n^2 + y_n^2$, and k_n and p_n are radial and tangential distortion coefficients, respectively. The polynomial in r^2 in the first term describes the radial distortion while the second term represents tangential distortion. Finally, this normalised projection is related to the pixel coordinates x_p, y_p of the projected point by multiplication with the 'camera matrix':

$$\begin{pmatrix} x_p \\ y_p \\ 1 \end{pmatrix} = \begin{pmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_d \\ y_d \\ 1 \end{pmatrix}. \tag{5.2.4}$$

Here f_x and f_y are the effective focal length of the imaging system measured in units of detector pixels in the horizontal and vertical directions, and are expected to be equal for square pixels and non-anamorphic optics. c_x and c_y are the pixel coordinates of the centre of the perspective projection on the sensor, expected to be close to the detector centre.

The parameters f_x, f_y, c_x, c_y, k_n and p_n together constitute the intrinsic camera parameters, and can be measured by capturing multiple images of a planar checkerboard pattern, of known scale, held in a range of positions in front of the camera. The grid corner positions are extracted from each image, and the positions of the grid pattern in 3D space in each image, and the intrinsic parameters, are solved for simultaneously. The solving procedure minimises the error when the resulting grid locations in the camera frame are projected on to the images using the camera model. For the MAST system this procedure was performed using the MATLAB camera calibration toolbox: a collection of freely available MATLAB code designed

Parameter	Value
f_x	(1598 ± 17) px
f_y	(1602 ± 17) px
c_x	(564 ± 22) px
c_y	(532 ± 22) px
k_n	$k_1 = -0.172 \pm 0.03, k_2 = 0.4 \pm 0.2$
p_n	$p_1 = (3.3 \pm 31) \times 10^{-4}, p_2 = (3.6 \pm 2.8) \times 10^{-3}$

Table 5.1: Typical calibrated intrinsic camera parameters for the MAST CIS diagnostic.

for calibration of machine vision systems based on the model described here[1]. Typical calibration parameters at the widest field of view ($f_1 = 17$mm) are shown in table 5.1. The calibrated focal length is equal in x and y as expected and is equivalent to (27.2 ± 0.3) mm. The centre of the image projection is somewhat offset from the detector centre, particularly in the x direction, however the p_n values indicate there is very little discernible tangential distortion present.

Extrinsic Parameters

To obtain the camera lines of sight through the tokamak, the relationship between the camera and laboratory coordinate systems must be known. The conventional Cartesian laboratory coordinate system on MAST has the origin at the centre of the tokamak at the vertical midplane, with positive z vertically up and the positive x axis between machine sectors 12 and 1 (see figure 3.10). The position of the camera pinhole (i.e. the camera coordinate system's origin) in this system is denoted by $x_{\mathrm{cam}}, y_{\mathrm{cam}}, z_{\mathrm{cam}}$. The orientation of the camera coordinate system with respect to the lab coordinate system is described by yaw, pitch and roll angles α, β and γ, respectively. The transformation between coordinates in the camera frame X, Y, Z and the lab frame $X', Y'Z'$ is then given by:

$$\begin{pmatrix} X' \\ Y' \\ Z' \end{pmatrix} = \boldsymbol{R_x}(\gamma)\boldsymbol{R_y}(\beta)\boldsymbol{R_z}(\alpha) \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} + \begin{pmatrix} x_{\mathrm{cam}} \\ y_{\mathrm{cam}} \\ z_{\mathrm{cam}} \end{pmatrix}, \tag{5.2.5}$$

where $\boldsymbol{R_{x,y,z}}(\theta)$ are the 3D rotation matrices for anti-clockwise rotation about their respective axes by an angle θ. This relationship and the pinhole camera model together completely specify the mapping between image pixel locations and physical space in the lab frame.

---

[1] Code and documentation available from: `www.vision.caltech.edu/bouguetj/calib_doc` (accessed 8th May 2013).

The extrinsic parameters $x_{\text{cam}}, y_{\text{cam}}, z_{\text{cam}}, \alpha, \beta, \gamma$ of the MAST system were measured using images of the tokamak interior taken with the instrument mounted in position and the band pass filter removed, and with the vacuum vessel interior illuminated with a high power flash-lamp. This was usually carried out during engineering breaks when the diagnostic was moved or re-mounted. In some cases references images could be obtained from plasma data, using frames where the vessel is brightly illuminated by plasma disruptions or the initial plasma breakdown. The extrinsic parameters were determined using an IDL utility developed by Dr. James Harrison at CCFE. Features visible in the calibration image e.g. in-vessel structural components or wall tile gaps were matched with the corresponding points on a 3D CAD model of MAST. Given the intrinsic camera parameters as an input, the extrinsic parameters are then solved for by minimising the error when the CAD model reference points are projected on to the image. Typical projection errors after this minimisation were between 2 - 5 pixels. While it is possible to solve for both the intrinsic and extrinsic parameters simultaneously using this feature matching technique, calibrating the intrinsic parameters separately reduces the required number of reference points in the vacuum vessel images needed to obtain a good result (since fewer variables have to be determined). This is particularly useful for situations where there are few suitable reference points in the vessel images. An example of a vessel image used for calibration of the lower divertor view is shown in Figure 5.7, with outlines of the major components projected on to the image using the fitted parameters. The application of the view registration to tomographic inversion is discussed further in Chapter 4.

5.3 Interferometer Calibration

For accurate flow measurements, careful calibration of the interferometer components is required. The quantities of interest are the interferometer group delay $\hat{N}$, the instrument phase $\phi_{\text{I}}(x, y)$ and the instrument contrast ζ_{I}. The group delay is the constant of proportionality relating plasma flow and measured Doppler phase ϕ_{D} according to equation (2.4.21), and the instrument phase represents the un-shifted fringe image against which ϕ_{D} is measured. The combination of these two quantities constitutes a complete wavelength calibration of the instrument, necessary for absolute Doppler shift measurements. The instrument contrast ζ_{I} is analogous to the instrumental broadening of a dispersive spectrometer: it increases the apparent spectral width of the line being measured due to imperfections in the optical system. Although line width measurements to obtain ion temperatures are not the focus of

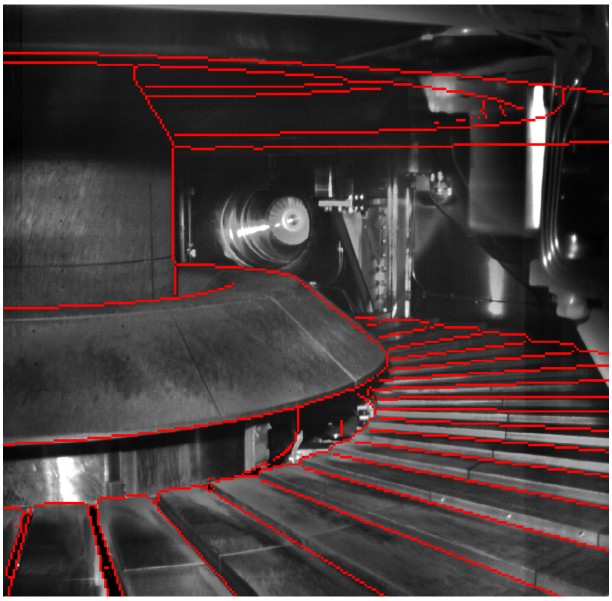

Figure 5.7: Flashlamp illuminated image of the MAST lower divertor (HL07 view) used for calibration of the camera extrinsic parameters. Red lines show outlines of the major components projected on to the image using the fitted camera parameters.

the present work, it is desirable to have a calibration for ζ_I since it can provide additional insight into the operation and performance of the instrument. The following sections consider the calibration of $\hat{N}$, ϕ_{I} and ζ_{I} in more detail.

5.3.1 Group Delay

The interferometer group delay $\hat{N}$ is determined by the delay plate thickness, its birefringence and the dispersion of the birefringence. In order to calibrate the group delay produced by each of the three delay plates, the response of the fringe phase to small, well known changes in wavelength was measured. This also allows calibration of the direction of the fringe shift (sign of the phase shift) relative to the sign of the wavelength change (red-shift or blue-shift direction). A frequency doubled, tunable diode laser (Toptica DL-SHG) was used as the adjustable wavelength light source for the measurements, with the laser frequency monitored with a Fizaeu wavemeter with an absolute accuracy of $\Delta\nu \sim$ 200MHz and resolution ~10MHz. Light from an optical fibre coupled to this laser was incident on a flashed opal diffuser placed in front of the CIS diagnostic, with a spot size sufficient to illuminate a central 512×512 pixel region of the detector. Images were captured at a series of discrete wavelength steps each equivalent to a Doppler shift of a few km/s with respect to the previous. An example of one of the raw images showing the spatial heterodyne fringe pattern and uniform illumination is shown in figure 5.8. Using the image at the shortest wavelength as the 'zero Doppler shift' reference, fringe phase shifts were extracted from the images as described in Chapter 4. The mean of the extracted phase shift values over the 512x512 pixel image was taken as the fringe phase shift $\Delta\phi$ for each wavelength. Measurements were repeated for each of the three delay plates. Written in terms of wavelength changes $\Delta\lambda$ about a wavelength λ_0, the change in fringe phase $\Delta\phi$ is expected to be given by $\Delta\phi = 2\pi\hat{N}(\Delta\lambda/\lambda_0)$. The fixed delay for each plate was therefore determined by fitting straight lines to plots of $\Delta\phi$ as a function of $\Delta\lambda/\lambda_0$, giving a gradient of $2\pi\hat{N}$. The phase measurement results and best fit lines are shown in Figure 5.9.

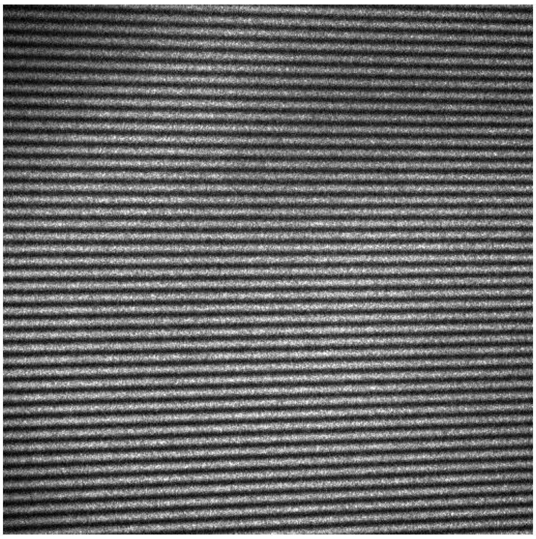

Figure 5.8: Example 512x512 pixel raw image used for group delay calibration, showing the spatial heterodyne fringe pattern. The noisy appearance of the image is mainly due to laser speckle.

In order to use these measured group delay values to calibrate plasma data, they must be extrapolated to the correct wavelengths for the MAST ion lines from the calibration wavelength $\lambda_0 = 460.85$nm. Using the expression for the phase delay of a birefringent plate in equation (2.5.34) and the definition of group delay, the delays at two different wavelengths are related by:

$$\hat{N}(\lambda_2) = \frac{\lambda_1}{\lambda_2}\frac{B(\lambda_2)}{B(\lambda_1)}\frac{\kappa(\lambda_2)}{\kappa(\lambda_1)}\hat{N}(\lambda_1), \tag{5.3.6}$$

where the birefringence $B(\lambda)$ and dispersion correction $\kappa(\lambda)$ are calculated from the Sellmeier equations (3.5.5). The group delays for each emission line calculated using this relationship are shown in table 5.2, along with the design values from the previous chapter. These measurement-based values were used to calibrate the

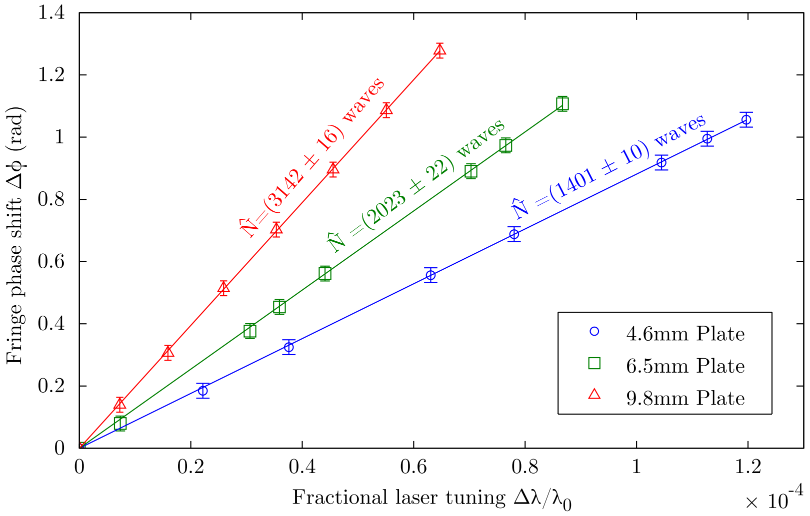

Figure 5.9: Measured fringe phase shift as a function of input wavelength change (laser tuning) for each delay plate. Solid lines are linear fits to the experimental data points, with the corresponding values of $\hat{N}$ given.

Delay plate	$\hat{N}$, C III (464.89nm)		$\hat{N}$, C II (514.24nm)		$\hat{N}$, He II (468.57nm)	
	Design	Calib.	Design	Calib.	Design	Calib.
4.6mm	1389	1391 ± 10	1204	1207 ± 9	1373	1376 ± 10
6.5mm	1962	1997 ± 22	1701	1731 ± 19	1940	1974 ± 21
9.8mm	2958	3102 ± 16	2565	2690 ± 13	2924	3067 ± 15

Table 5.2: Calibrated group delay (in waves) for each delay plate when operating for each MAST ion line, compared with the design values.

MAST plasma data.

The calibrated group delay for the 4.6mm crystal matches extremely well with the design specifications, while the thicker plates deviate by around 1.8% and 4.9% for the 6.5mm and 9.8mm plates respectively. The discrepancy in the 9.8mm plate operating in C II is enough that the delay for this plate misses the narrow high contrast region of 2500 - 2600 waves which it was intended to target. This is expected to result in a fringe contrast less than 65% of the design value. While this reduces the usefulness of this crystal for experimental contrast measurements in C II, it does not effect the baseline flow measurement scenarios or possible contrast measurements in the other species (which use the other two plates). It was found that the varied discrepancies between the design and measured delays for each of the 3 crystals was partly due to the plate thicknesses. The 9.8mm plate was in fact found to have the most accurate delivered thickness at 9.79mm, with the other two crystals measuring 6.35mm and 4.48mm thick. Using these measured thicknesses to calculate the expected group delays, it was found that the Sellmeier equations used in this work under-estimate the group delay by 2.8%, 4.3% and 4.9% for the nominally 9.8mm, 6.5mm and 4.6mm crystals respectively.

5.3.2 Instrument Phase

Accurate calibration of the instrument phase, i.e. the un-shifted fringe pattern, is a challenging aspect of the coherence imaging technique. Ideally laboratory sources of each of the target ion lines, at their rest wavelengths, would be used to obtain instrument phase images which could be directly compared to the plasma image phase. However, in practise such sources of these lines are not available. Previous work studying carbon lines on DIII-D used the flash of light when the plasma is initiated to produce a reference phase image [Howard et al., 2011], assuming that the net motion of the plasma at this time is close to zero. This gives a calibration image with the correct spectral line characteristics and offers per-shot monitoring of any variation of ϕ_I. However, this was not possible on MAST since no bright flash of impurity light is reliably observed at the start of the shot. Instead, a hybrid technique is used where an initial phase image is obtained offline at a different wavelength, and then adjusted for use with the plasma data using information from the plasma light itself. For the purposes of this calibration the instrument phase is considered to consist of two parts: a 'phase shape' image which is a function of image position $\phi_{\text{shape}}(x, y)$, and a 'phase offset' ϕ_{offs}. The phase shape gives the form of the carrier fringe pattern, i.e. the shear of the instrument phase across the detector, while the phase offset is a spatially constant offset specifying the absolute zero flow point of the flow measurement. The total instrument phase is then $\phi_I = \phi_{\text{shape}} + \phi_{\text{offs}}$. This form is adopted in order to split the instrument phase into a part which can be calibrated infrequently and offline: ϕ_{shape}, and a part adjusted frequently using information from the plasma: ϕ_{offs}. The offline part of the phase calibration was performed whenever the crystal plate configuration or zoom / focus settings of the instrument were changed. The methods for calibrating ϕ_{shape} and ϕ_{offs} are now presented in detail.

Phase Shape Calibration

Initial phase shape images were obtained using a cadmium spectral lamp. Cd I has well isolated lines in the wings of both the He II and C II filters, at 467.815nm and 508.58nm respectively. The filter transmission at these wavelengths is 11% for the He II filter and 4% for the C II filter. Calibration images were obtained offline with the diagnostic in the lab, where a Cd spectral lamp was used to illuminate a small ($\sim$15cm) integrating sphere into which the diagnostic was pointed to take images.

Since the camera was limited to a maximum integration time of 20ms, the raw images from this setup had a very poor SNR (< 6). To improve this, stacks of 2048 images taken in succession with 20ms exposure were used to form each calibration fringe image. Once such a fringe image was obtained, it was demodulated to obtain a wrapped fringe phase image as described in section 4.2.3. To obtain the phase shape ϕ_{shape}, this phase image is unwrapped and the phase at the image centre is set to zero. An example fringe pattern image and the corresponding unwrapped phase shape are shown in figure 5.10; the phase shape image is clearly dominated by the up-down phase shear from the Savart polariscope which forms the horizontal fringe pattern. Phase shape images obtained in this way still showed considerable high spatial frequency noise due to noise on the fringe pattern images, although this is not clearly visible in figure 5.10 due to the colour scale. To reduce the effects of this noise, the phase shape images were fit with a 2D, 3$^{\text{rd}}$ order polynomial in pixel coordinates (x, y), and the fitted polynomials then used to model ϕ_{shape}.

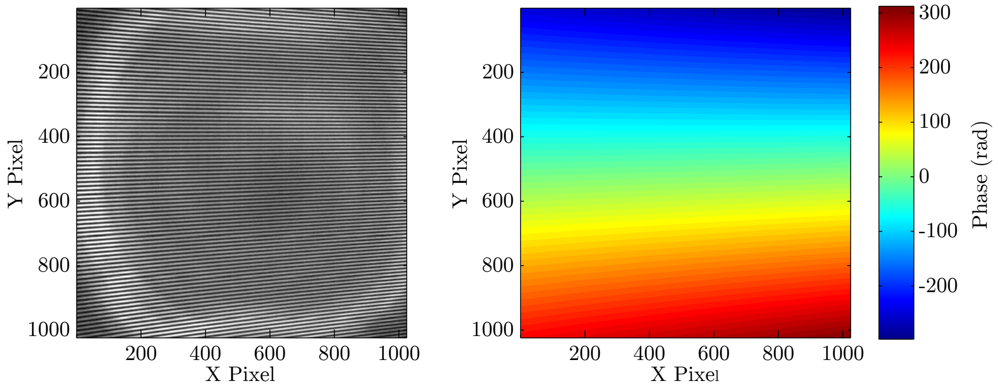

Figure 5.10: Illustration of phase shape calibration using the Cd I 467.8nm line. Left: calibration image produced from a stack of 2048 20ms exposures, showing the calibration fringe pattern. Right: demodulated and unwrapped phase shape for this image.

Phase shape images obtained with the Cd lamp setup must be corrected for the difference in wavelength between the calibration lamp lines and plasma ion lines before they can be used with plasma data. A theoretical form for the phase shape at a given wavelength can be obtained from equations (2.5.39) and (2.5.41). Approximating $\sin(\theta) \approx \theta$ and neglecting terms in θ^2 (these approximations are justified since $\theta < 5°$ for the entire field of view in the MAST system), the phase shape is approximately:

$$\phi_{\text{shape}}(\theta, \omega) \approx \frac{2\pi\theta t}{\lambda} \left(\frac{n_o^2 - n_e^2}{n_o^2 + n_e^2} [\cos(\omega) + \sin(\omega)] \right), \tag{5.3.7}$$

where t is the Savart polariscope thickness and for horizontal fringes θ and ω are related to the image plane position via $\theta \approx \sqrt{x^2+y^2}/f_3$ and $\omega = \arctan(y/x)+3\pi/4$. From this expression, the phase shapes at two different wavelengths are related by:

$$\phi_{\text{shape}}(x,y,\lambda_2) \approx \left[\frac{\lambda_1}{\lambda_2} \left.\frac{n_o^2+n_e^2}{n_o^2-n_e^2}\right|_{\lambda_1} \left.\frac{n_o^2-n_e^2}{n_o^2+n_e^2}\right|_{\lambda_2} \right] \phi_{\text{shape}}(x,y,\lambda_1). \tag{5.3.8}$$

This correction was applied to the Cd lamp phase shape images to obtain the phase images used for plasma data. The accuracy of the correction was tested using ϕ_{shape} images recorded for the Cd I 467.8nm and 508.6nm lines. The phase shape at 467.8nm was extrapolated to 508.6nm using equation (5.3.8), and compared with ϕ_{shape} measurements at 508.6nm. The difference between ϕ_{shape} at the different wavelengths, before and after the wavelength correction are shown in figure 5.11. It is seen that the correction procedure accurately removes the dominant vertical shear in the phase difference over most of the image, although with some residuals at the top and bottom image edges. The phase difference after correction is $\leq$0.38 rad over the entire image (and much less near the image centre), corresponding to an error in the correction of $\leq$1.5%. For the much smaller wavelength differences over which the correction is used for the actual diagnostic calibration, this corresponds to a correction error of $\leq$ 0.03 rad or $\leq$ 1 km/s in the flow calibration

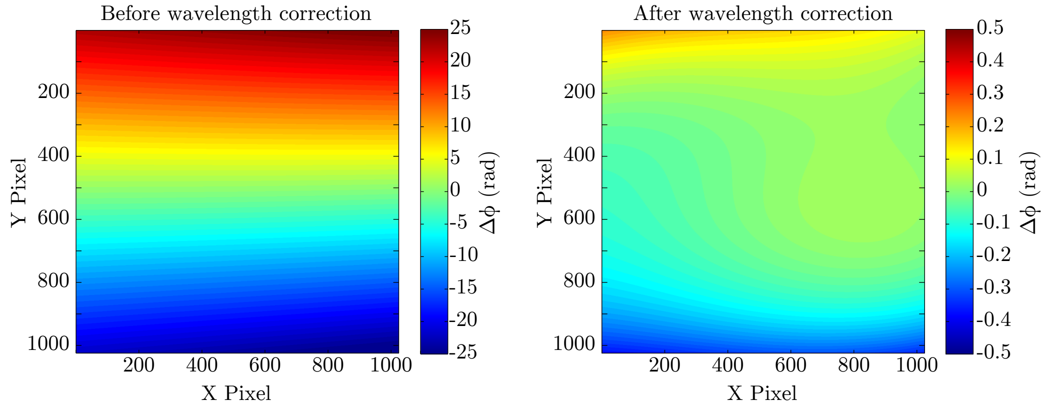

Figure 5.11: Difference between phase shape images for Cd I 467.8nm and 508.6nm spectral lines, without (left) and with (right) correction for the difference in wavelength between the two lines using equation (5.3.8).

Phase Offset

With reference to the notation of chapter 2, the phase offset is expected to be given by $\phi_{\text{offs}} = (\phi_0 + \phi_M + \phi_{\text{drift}}) \mod 2\pi$. In principle $\phi_0 + \phi_{\text{M}}$ could be obtained from the Cd lamp measurements if the delay plate characteristics, calibration wavelength and multiplet structure of the plasma lines were known sufficiently accurately, while ϕ_{drift} represents time-dependent calibration drifts which can occur in the real instrument e.g. due to changes in ambient temperature. For the MAST system the time variation of ϕ_{drift} was sufficiently large that one-off calibration of ϕ_{offs} was insufficient, and required monitoring during plasma operations (see section 5.3.5). This was possible without additional in-situ calibration hardware, because the wide angle views on MAST allow ϕ_{offs} to be obtained from the plasma data itself. The technique is based on the fact that flows in the radial direction in the plasma, perpendicular to the magnetic field, are expected to be very slow compared to the flows parallel to the field which are being measured, due to the confining effect of the magnetic field. Therefore for any line-of-sight viewing the plasma radially, $\phi_{\text{D}} \approx 0$ is expected to be observed. The most commonly used plasma views for CIS on MAST all included such radial sight lines. Letting the image coordinates of such a sight line be denoted $(x_{\text{ref}}, y_{\text{ref}})$, the phase offset can be obtained from the measured fringe phase image $\phi(x, y)$ according to $\phi_{\text{offs}} = \phi(x_{\text{ref}}, y_{\text{ref}}) - \phi_{\text{shape}}(x_{\text{ref}}, y_{\text{ref}})$. In practise for every image this was calculated for all sight lines viewing the plasma close to radially, and the mean value of ϕ_{offs} from these sight-lines was used. This is then sufficient to completely constrain the absolute flow calibration over the entire field of view. This technique is advantageous because it allows the calibration to be monitored on a shot-to-shot and even frame-to-frame basis, however it is only applicable when suitable reference sight lines (typically ones which view the plasma radially) are available in the images.

Temperature Dependence

The interferometer delay varies with the temperature of the birefringent components, due to differential changes of the ordinary and extraordinary refractive indices and thermal expansion of the crystal plates. In general both the phase shape and phase offset are subject to this effect, although the phase offset change is much larger due to the effect of the thick delay plate. For small changes in temperature, such as those encountered during operation with the temperature-stabilised cell, the change in phase shape can be neglected and the temperature effect considered as a phase offset change due to the main delay plate. Under a small change in temperature ΔT, thermal expansion will cause the plate thickness to increase by an amount $\Delta L = L\alpha\Delta T$, where α is the linear thermal expansion coefficient of the crystal material perpendicular to the optical axis. The sensitivity of the refractive indices to temperature is given by the thermo-optic coefficients $\frac{dn_{\text{o}}}{dT}$ and $\frac{dn_{\text{e}}}{dT}$. The corresponding rate of change of birefringence is given by $\frac{dB}{dT} = \frac{dn_\mathrm{e}}{dT} - \frac{dn_\mathrm{o}}{dT}$. Substituting $L = L_0 + \Delta L$ and $B = B_0 + \Delta T \frac{dB}{dT}$ in equation (2.5.34), the expected phase delay change due to a small temperature change is given by:

$$\frac{\Delta\phi_0}{\Delta T} \approx \phi_0 \left(\alpha + \frac{1}{B_0} \frac{dB}{dT} \right). \tag{5.3.9}$$

Since we are ultimately interested in the apparent flow due to a temperature change, note that a shift in ϕ_0 corresponds to an apparent flow of $\Delta v = c\Delta\phi_0/(2\pi\hat{N})$. Using this we can write down an estimate for the flow calibration's sensitivity to temperature, which depends on the material properties and is independent of the interferometer delay:

$$\frac{\Delta v_\mathrm{calib}}{\Delta T} \approx \frac{c}{\kappa} \left(\alpha + \frac{1}{B} \frac{dB}{dT} \right). \tag{5.3.10}$$

This has a weak wavelength dependence, due to the dispersion of B and κ. Using measurements in the literature for α-BBO [Eimerl et al., 1987] of $\alpha = 4 \times 10^{-6}$/K, $\frac{dn_o}{dT} = -1.66 \times 10^{-5}$/K, $\frac{dn_e}{dT} = -9.3 \times 10^{-6}$/K and the Sellmeier equations (3.5.5), the calculated calibration sensitivity is plotted as a solid black line in figure 5.12. This is remarkably large at $>$ 13km/s/°C throughout the visible range, emphasising the need for temperature stabilisation of the interferometer optics.

In order to measure the temperature sensitivity of the calibration experimentally, fringe images were recorded while the temperature of the interferometer optics was adjusted using the temperature controlled cell. Multiple wavelengths were sampled by using a 633nm HeNe laser and two lines from a zinc spectral lamp (468.2nm and 509.4nm), where available. The mean phase change over the entire image between a reference frame and each temperature step was plotted as a function of plate temperature, and a straight line fit used to obtain the temperature sensitivity $\frac{\Delta\phi}{\Delta T}$. This was converted to flow calibration sensitivity according to $\frac{\Delta v_\mathrm{calib}}{\Delta T} = (c\frac{\Delta\phi}{\Delta T})/(2\pi\hat{N})$, where $\hat{N}$ was determined from table 5.2 and equation (5.3.8). The results are shown in figure 5.12.

The observed temperature sensitivity is lower than predicted by equation (5.3.10) for all wavelengths and delay plates by between 14% - 28%. This is most likely due to insufficient accuracy of the thermo-optic coefficients used for the calculation, to which the result is very sensitive. The thermo-optic coefficients used do not have explicitly stated uncertainties, and are an average of measurements at different wavelengths. Errors in the coefficients of between 3 - 4% would be sufficient to explain the discrepancy found here. The experimentally measured sensitivity is also in good agreement with similar measurements carried out at Australian National

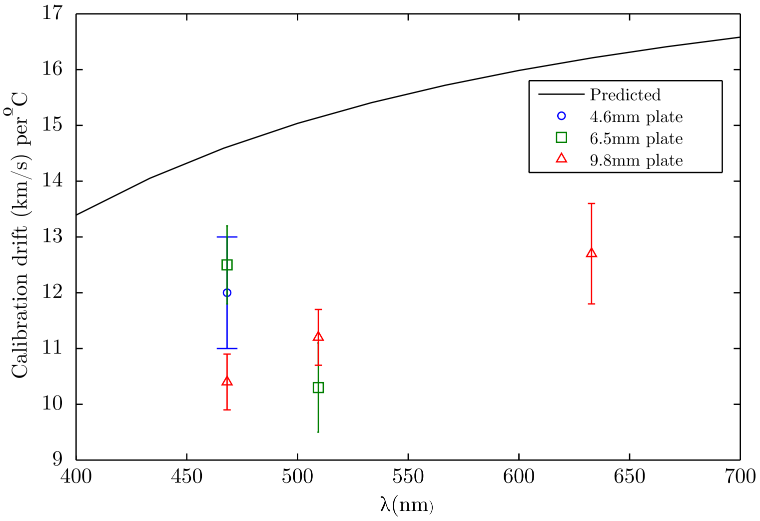

Figure 5.12: Dependence of calibration offset (in terms of flow speed) on interferometer component temperature over the visible spectrum. The black line is calculated from equation 5.3.10 while the points are measured.

University [Howard, J. private communication, 2012].

5.3.3 Instrument Contrast

It was found that the instrument contrast ζ_I could not be reliably calibrated using the Cd lamp setup described previously for the instrument phase calibration. This is shown by the fact that the measured contrast in the plasma data was higher than the calibrated ζ_I in some shots, which should not be possible for an accurate calibration since ζ_I by definition is the highest contrast achievable by the instrument. There are two main possible reasons for this. The first is a difference in illumination of the interferometer pupil between the calibration and plasma light, which can lead to a difference in measured contrast because of inhomogeneity in the birefringent components (see section 5.3.4). However, using the integrating sphere configuration with the instrument in its on-machine configuration is expected to minimise this difference. The second reason is related to the detector. Due to the need to stack many frames to increase the SNR on the calibration images, accumulation of read noise (which can only cause positive fluctuations in pixel values after dark subtraction, since the data is truncated at zero ADU) artificially lowers the contrast in the calibration images. In order to avoid this effect, the on-camera dark signal subtraction was disabled and stacks of 2048 dark frames (with the calibration lamp off) were subtracted from stacks of 2048 frames with the lamp on to obtain calibration images. While this removed the effect of read noise on the contrast, due to the very low signal level any small changes in background or offset signal could cause a significant change in apparent contrast. For example, with a peak signal level in an individual calibration frame of 15 ADU as typically seen in the Cd lamp calibrations, a background level increase of 1 ADU between the dark and signal frames would result in a 10% drop of the apparent contrast. Such changes can occur due to drifts in the camera offset level in the time taken to capture both the dark and signal images (∼20 minutes), particularly if the camera has not been allowed to warm up for at least 3 hours before performing the calibration. Furthermore, since the calibration signal level varies across the image due to vignetting and the structure from the filters, the size of this apparent contrast change varies in the image and can introduce artificial structure into the calibration. Because of these issues, calibration of ζ_I could not be performed sufficiently accurately for quantitative use. Mean values of ζ_{I} over the field of view ranged from 0.28 - 0.6 in different calibrations.

Despite the lack of a quantitatively useful calibration of ζ_{I}, some consistently observed qualitative properties of the contrast do provide insight into the properties of the instrument. An example contrast calibration image is shown in figure 5.13. Note that the contrast in the centre of the field of view is almost a factor of two lower than that at the image edges. This is the opposite of what would be expected if ζ_{I} was dominated by the performance of the imaging lenses: the contrast would be expected to fall off towards the image edges due to increasing abberations typically seen towards the field of view edges in camera lenses. This result therefore indicates that another effect is dominating the instrument contrast. This appears to be inhomogeneity in the birefringent components, and is discussed in more detail in the next section.

5.3.4 Crystal Uniformity

The uniformity of the interferometer components over their aperture, both in terms of surface quality and refractive index uniformity within the crystal, plays an important role in the instrument performance and calibration. Ideally, light passing through any part of the crystal aperture at a given angle should experience an identical phase delay. In reality, thickness variations and refractive index inhomogeneity within the crystal mean that light passing through different parts of the aperture experiences different delays. This has two important effects, the first of which is lowering of the instrument contrast. This occurs because the light arriving at a given point in the image plane has passed through a range of locations in the crystal aperture, and therefore experienced a range of phase delay. The larger the range of phase delay, the lower the fringe contrast at that point in the image. The second

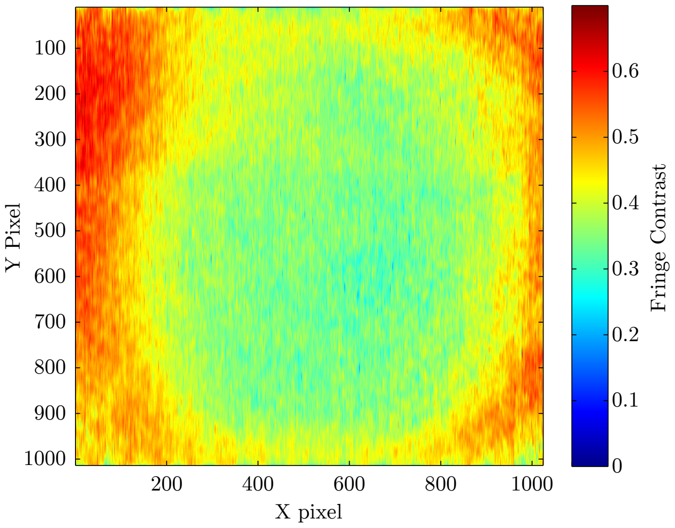

Figure 5.13: Example of a contrast calibration image using the Cd lamp calibration configuration. The contrast is observed to increase significantly towards the image edges.

effect is to make the interferogram calibration highly sensitive to the pupil illumination. If incoming calibration light does not fill the crystal aperture in the same way as the plasma light, it will sample a different range of delays and therefore could give a different instrument phase and contrast.

In order to measure the non-uniformity of the delay produced by the birefringent components, an experiment was carried out to record the fringe patterns created when only illuminating small sub-apertures of the interferometer. A small square aperture was placed in front of the Cd calibration lamp and imaged on to the interferometer components, which were mounted in the temperature controlled cell but removed from the rest of the diagnostic. This illuminated a small square area of the polarisers and crystals of around 2mm x 2mm (the full crystal aperture is circular with 28mm diameter). The temperature controlled cell was mounted on a 2D linear translation stage setup, such that the illuminated area could be scanned over the components' clear aperture. A camera with its lens focused at infinity was placed after the cell to record the fringe pattern, using the same fringe scale at the detector as in the complete diagnostic. The illuminated area was then raster scanned over the aperture of the interferometer components using the translation stages, recording the fringe pattern at each location. These images were demodulated to determine the difference in fringe phase between the different positions in the aperture. Since the illumination used had a constant spectrum for all points, these measured differences in fringe phase correspond to the variations in the interferometer phase delay across the aperture. Measurements were performed for the baseline crystal configuration consisting of the 6.5mm delay plate, 4mm and 2.2mm Savart polariscopes, and for each of these components separately (always including the polarisers before and after the crystals in order to produce the fringe pattern). Since at least one Savart polariscope is required to produce the fringe pattern which enables the measurement, the delay plate only measurement was obtained using the delay plate and 4mm Savart polariscope together, before subtracting the polariscope measurements. The resulting delay profiles across the crystal apertures are shown in figure 5.14.

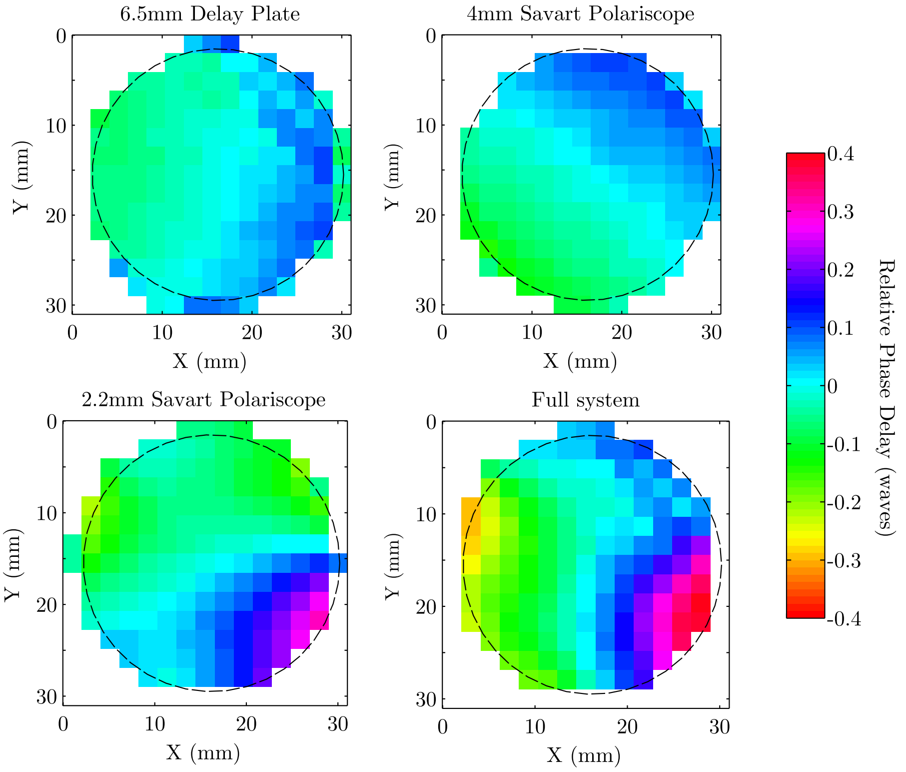

Figure 5.14: Measured variations in phase delay across the interferometer aperture, for each component in the baseline interferometer configuration and for the complete system (consisting of one delay plate and two Savart polariscopes).

For the full baseline crystal configuration, the range of the delay variation is approximately 0.7 waves over the aperture. The largest contribution to this is from the 2.2mm Savart polariscope, which displays the largest non-uniformity and the most complex spatial structure, attributed to particularly large manufacturing imperfections in this component (despite all the crystals being provided by one supplier). In terms of the instrument calibration this amount of phase variation is very large (0.7 waves fringe phase change would be equivalent to a flow measurement of order 100km/s), thus this illustrates the need to ensure the crystal aperture is correctly illuminated (i.e. not under-filled) when carrying out calibrations. These results and the known properties of the instrument's lens configuration also provide a qualitative explanation for the observed instrument contrast structure seen in figure 5.13, i.e. the lower contrast near the image centre than at the edges. This occurs because at the centre of the field of view, the entire crystal aperture is illuminated, hence this light samples the full range of phase delay and the fringe contrast is lowest. Towards the image edges, vignetting due to the imaging lens configuration means a smaller area of the crystals is illuminated, thus the light samples a smaller range of delay and the fringe contrast is higher. This effect appears to be dominant in determining the fringe contrast.

The measurements were repeated with the temperature controlled cell set to different temperature values, and the non-uniformity was found to show a weak temperature dependence, with the RMS phase variation across the aperture increasing from 0.97 rad at $T_{\text{crystals}} = 33.3°\text{C}$ to 1.07 rad at 36.9°C (without significant change to the spatial structure of the variations). The instrument contrast is therefore expected to decrease slightly with increasing temperature of the crystals, due to increased non-uniformity at higher temperatures. Such an effect is indeed observed, as will be seen in the next section.

5.3.5 Calibration Stability Tests

In order to test the overall stability of the interferogram calibration to ambient conditions, images of a diffuser illuminated by a HeNe laser were recorded every 5 minutes for 48 hours in the lab. For the first 24 hours the temperature stabilised cell was turned off to observe the stability without active temperature stabilisation; for the second 24 hours the active temperature stabilisation was enabled. Measurements of the ambient temperature and interferometer optics temperature were also obtained using thermocouples. Time histories of the ambient temperature, crystal plate temperature, fringe phase and fringe contrast are shown in figure 5.15. Phase and temperature measurements are the mean from a central 100x100 pixel region of the image.

The ambient temperature shows similar levels of fluctuation with a range of 4 - 5°C over both the stabilisation off and on periods. The plate temperature shows fluctuations of the same magnitude with the temperature control off, however the response to faster temperature changes is smoothed out due to the (passive) thermal insulation of the cell. With the temperature stabilisation enabled, the plate temperature varies with a range of 0.4°C, consistent with the controller specifications.

With the temperature control off, the fringe phase varies with a range of 2.2 ra-

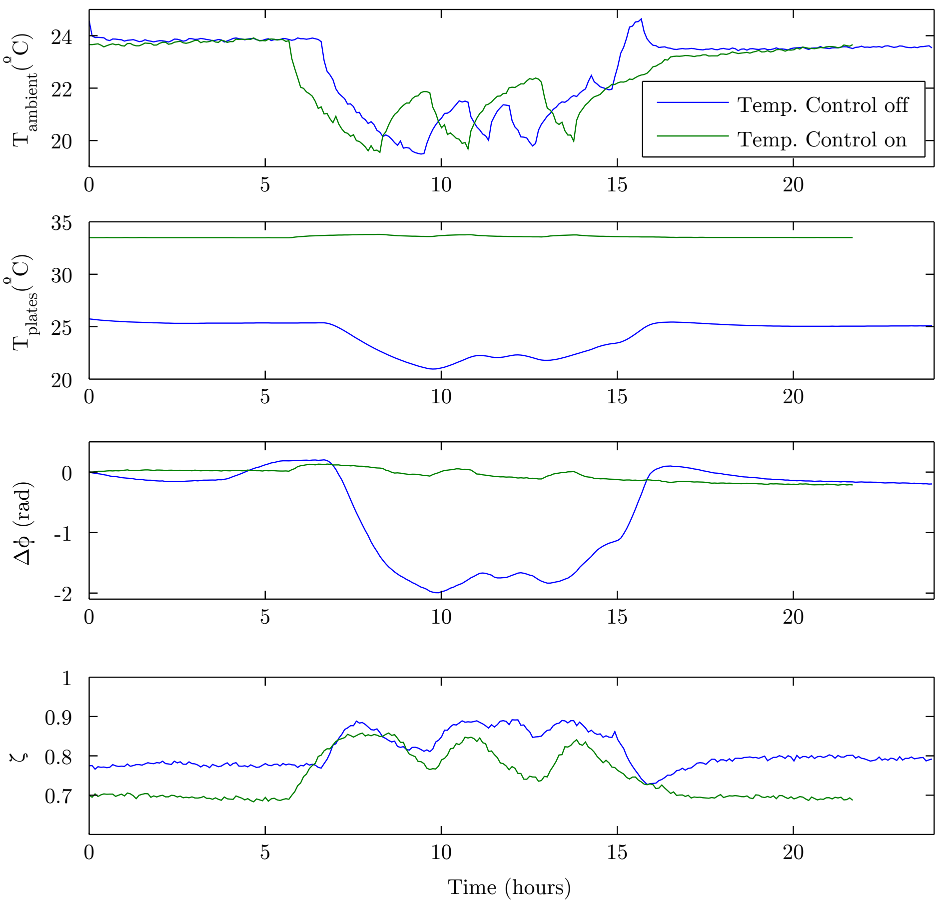

Figure 5.15: Time histories of ambient temperature, polarisation optics temperature, fringe contrast and fringe phase over 24 hour periods with active temperature stabilisation of the optics off (blue lines) and on (green lines).

dians, equivalent to a flow calibration shift of almost 80km/s. With the temperature control on this is reduced to 0.3 rad or 12km/s in flow, somewhat larger than the 5km/s range expected from the stand-alone temperature sensitivity measurements. This could be due to effects other than changes in the crystals themselves, such as thermal expansion in the mechanical structure of the system causing the crystals to shift slightly relative to the detector. The amount of fringe movement on the detector for a phase range of 0.3 rad at 633nm is 0.74 pixels, of which approximately 0.17 rad is expected due to changes in the crystals.

The fringe contrast shows approximately the same amount of variation with the temperature control on as off. This indicates that the contrast calibration drift is dominated by changes in the camera offset level, which will be the same in the stabilisation on and off cases. The range of contrast variations in both cases is approximately 16 percentage points. Since this effect is primarily due to camera offset drifts, it is not expected to be observed during plasma operations since offset level calibration is performed per-shot, whereas it was performed only once for these stability measurements. It is also observed that the contrast is consistently lower with the temperature control on than off. This is due to increased delay inhomogeneity in the crystals when their temperature is raised from ambient to a higher stabilised temperature, and this behaviour is consistent with the measurements in section 5.3.4. It is therefore desirable to operate the temperature controller at the lowest temperature where good stabilisation can be achieved, in order to avoid lowering the fringe contrast and therefore SNR unnecessarily.

Overall the stability measurements show considerable variation in the calibration parameters, at a level at which calibration monitoring during plasma operations is required. In order to improve the stability, it would be desirable to more carefully account for thermal expansion effects in the mechanical design of any future instrument, and to investigate methods of better stabilising the thermal effects in the plates.

5.4 Summary

In this chapter, various aspects of the behaviour and performance of the MAST CIS diagnostic hardware have been investigated experimentally. The detector response and noise level have been measured, and in the next chapter will be shown to result in flow measurement noise of around 1km/s under typical measurement conditions on MAST. Difficulties with accurate contrast calibration due to dark level drift of the camera preclude ion temperature measurements with the MAST instrument.

The line selection band pass filters were found to meet the requirements set out in chapter 3, and are not expected to cause any significant flow measurement errors. The absolute light sensitivity and vignetting of the assembled diagnostic were measured and were within approximately 20% of the design calculations.

The interferometer group delay has been calibrated and found to agree well with the design specifications for two of the three delay plates, although the 9.8mm plate produces a delay around 5% larger than the specifications. This makes the 9.8mm plate less useful for C II measurements but does not affect the baseline flow measurement configuration. Errors on the group delay calibration are $\leq 1.1\%$ and are not expected to be a significant source of measurement error. Measurements of the fringe contrast show the contrast to be highest at the edges of the field of view, which is qualitatively explained by non-uniformities in the interferometer components. Measurements of the calibration stability in the lab confirmed the need for active temperature stabilisation of the birefringent interferometer optics, and the need for monitoring of the calibration during plasma operations. The procedures for calibration of the instrument have been presented, including a hybrid scheme using online and offline measurements to calibrate the instrument phase.

Chapter 6

Coherence Imaging Measurements on MAST

The MAST CIS diagnostic was successfully operated throughout MAST's 2013 experimental campaign, making flow measurements in both dedicated experiments and parasitically during other discharges. Due to the limitations preventing accurate contrast calibration for line width measurements, discussed in chapter 5, only flow measurements were obtained. All four available plasma views (see figure 3.10) were exploited during the campaign. This chapter discusses the on-plasma performance of the instrument, and provides some examples of the complex flow patterns observed.

6.1 Instrument Performance

6.1.1 Signal level, noise & measurement uncertainties

Exposure times used for the collection of plasma data were typically 1 - 4ms for C III, 4 - 16ms for C II and 8 - 20ms for He II. The observed signal level and required exposure time varied strongly depending on the plasma configuration. Significantly longer exposure times were required when using the HM02 plasma view, since the vacuum port window did not have a shutter installed to protect it from deposition during glow discharge cleaning of the vacuum vessel (unlike the other windows used), and exhibited around 5 times lower transmission compared to the sector 7 windows. Due to the shorter exposure times required and therefore high frame rates possible, C III measurements were typically favoured and are of the most interest for studying flow dynamics.

The noise on measured $\bar{v}$ images due to image noise in the raw data was estimated using the results from section 4.3. An example for C III divertor measurements in a LSND plasma, using 2 ms exposure time, is given in figure 6.1, which shows the raw camera image (left) and the estimated noise on the flow measurement (right). The profiles of the divertor legs are clearly visible in both images. In the brightest regions of the image the estimated flow noise is as small as 0.2 - 0.5km/s (the magnitude of the measured C III flows in these areas are around 2.5 - 15 km/s in this example), rising to a few km/s in dimly emitting regions. Noise estimates such as this were validated by taking 1D slices through the measured $\bar{v}$ images, and subtracting a smoothed version of the profile to estimate the actual noise level on the image. The standard deviation of the measured noise level was found to be typically within 20% of the predicted standard deviation, indicating that images of the type shown in fig. 6.1(b) provide a reasonable estimate of the noise level on flow images.

In brightly illuminated image areas with estimated noise levels well below 1km/s, image noise is not likely to dominate measurement error, and spectral contamination (discussed in section 4.5) and errors in instrument phase calibration (discussed here and in section 5.3.2) are around the same magnitude. An important aspect of the phase calibration error to consider is the determination of ϕ_{offs} from radial sight-lines; these sight-lines often show low signal levels since they integrate straight through a thin shell of emission in the SOL. Typically ϕ_{offs} was obtained by averaging over at least several thousand image pixels viewing the centre column at $R < 5$ cm and symmetrical about the radial direction (i.e. close-to-radial sight lines). By looking at fluctuations in the time history of ϕ_{offs} during a shot (having eliminated other influences such as mechanical vibrations; see section 6.1.2), the typical random error on ϕ_{offs} was estimated to have $\sigma < 0.6$ km/s.

6.1.2 Calibration Stability

Phase Offset

As was discussed in section 5.3.2, the flow calibration offset ϕ_{offs} is subject to drifts due to ambient temperature changes and mechanical disturbance of the instrument, and is monitored using radial sight-lines as a zero flow reference. During plasma operations, offset shifts were monitored over two different timescales. The first were intra-shot, i.e. frame-to-frame changes, due to mechanical vibration of the instrument during the discharge. When the system was first installed on MAST it was supported on a rail cantilevered from the tokamak vacuum vessel. During plasma shots, movement and vibrations of the vacuum vessel were transmitted to the instrument and caused movement of the interferometer components relative to the

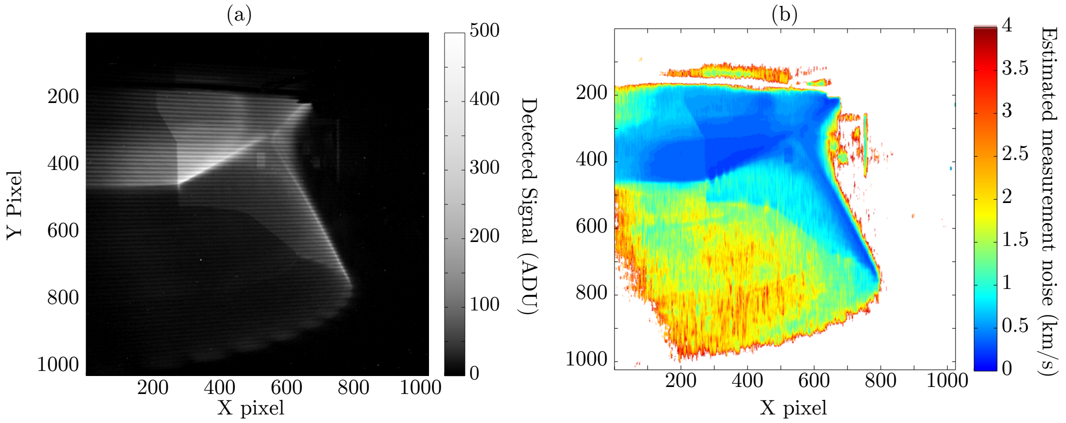

Figure 6.1: Signal level and flow noise images for LSND plasma shot 28841 during H-Mode, using the C III filter and 2 ms exposure time. (a) Raw image from the camera and (b) estimated noise on the flow measurement, estimated from the results in section 4.3.

detector, making the entire fringe pattern move across the detector. This appears in the data as large oscillations in the measured flow (with peak-to-peak amplitudes of up to 40km/s) with a frequency of around 100Hz. The dominant cause was identified as vibration of the temperature control cell and its mount, which was secured at one end from the main base plate. This motion is illustrated in figure 6.2. The problem was mitigated by mounting the instrument from the floor next to

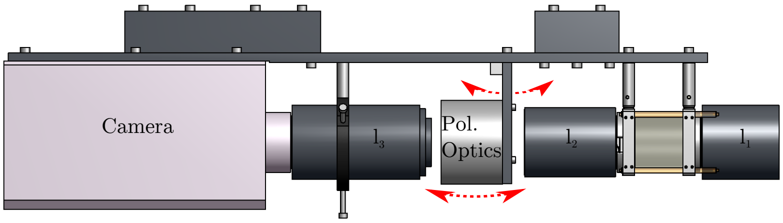

Figure 6.2: Side view of a CAD model of the instrument, illustrating the mechanical vibrations causing frame-to-frame calibration oscillations. Motion is indicated by red arrows.

the tokamak rather than the vacuum vessel, and by adding an additional support for the temperature cell at the end not supported by the primary mount. Floor mounts could only be implemented for the HL07 divertor and HM02 midplane views, due to space restrictions and clashes with other structures & diagnostics. Fig 6.3 shows the flow measured on the ϕ_{offs} reference chords, relative to its mean over the data shown, during two similar plasma shots on the HL07 divertor view: one before and one after the modifications to alleviate the vibration problem. The peak-to-peak amplitude of the oscillations with no mitigation corresponds to a fringe motion of up to 1.4 pixels on the detector, or an angular motion of the crystals of up to 0.01°, and is reduced to 0.15 pixels or 0.001° after the modifications. The sensitivity of the system to such small changes in the angular alignment of the crystals is due to the use of small beam angles through the collimated region, driven by minimising filter blue-shift effects. Larger oscillations were observed on data using the HM07 midplane views, with peak-to-peak amplitudes as high as 40km/s.

In data from the views where floor mounts could not be implemented (and to remove small remaining oscillations seen in some of the data even after the modifications), subtraction of the oscillating ϕ_{offs} usually provided satisfactory correction for the vibrations across the whole field of view, since the motion of the polarisation optics was essentially constrained to 1 axis and relatively small amplitude, making the fringe pattern move systematically up and down on the detector.

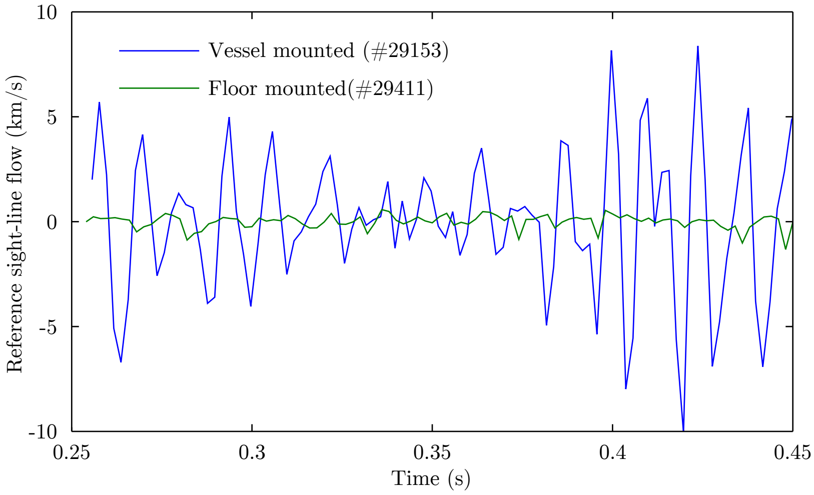

Figure 6.3: Effect of mechanical vibrations on flow measurements before subtraction of ϕ_{offs}, during two comparable LSND plasma discharges viewed from the HL07 divertor port. The mean flow has been subtracted to show only the vibration effect.

Variations of the calibration offset on longer timescales, i.e. between different shots over the course of days and weeks, were monitored using the mean offset within each shot. Variation over a one week period, in km/s, is shown in figure 6.4. On any given day, the offset drifted by between around 2km/s - 15km/s. The range of all values shown in figure 6.4 is around 16km/s.

Both the inter- and intra- shot variations in calibration offset demonstrate the need for the calibration monitoring of the MAST instrument, and potential improvements in future diagnostic design. Namely, it would be desirable to improve both the thermal stabilisation (using an improved temperature controller, and/or passive

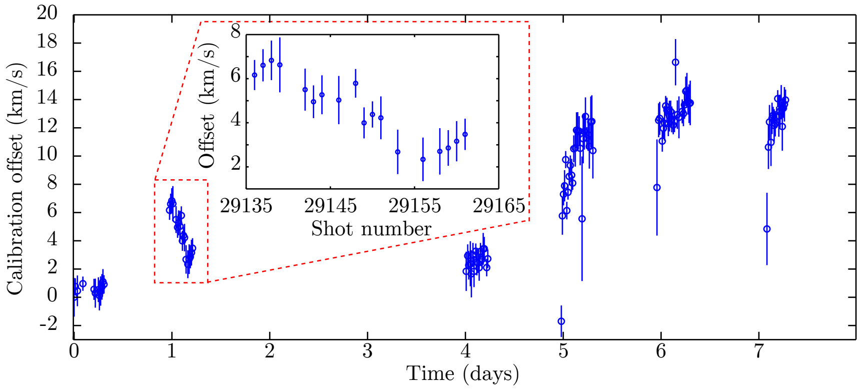

Figure 6.4: Calibration offset variations over a 1 week period. Inset shows the offsets over a single day.

thermal stabilisation using multiple plate materials) and mechanical design in order to try to stabilise the calibration. In addition, an in-situ calibration system not relying on plasma light is highly desirable to maintain accurate calibration of the diagnostic independent of the plasma view.

Phase Shape

As described in section 5.3.2, calibration of the instrument phase shape ϕ_{shape} for a given run of measurements was performed offline in the lab before the diagnostic was installed on the tokamak. This calibration method relies on the phase shape not being significantly affected by thermal drifts or mechanical disturbance during the installation of the instrument or subsequent plasma operations. To test this, phase shape calibrations were performed in the lab before and after a 7 day period of operation on the wide angle divertor view, and compared after subtracting the mean offset over the image. Discrepancies between the two calibrations were found to be up to 3 km/s, and showed the form of a gradient across the image perpendicular to the fringes. As such the largest deviations were at the image corners, and in the central part of the image where most of the spatial flow information is obtained the discrepancies were up to around 1km/s. This shows that the phase shape does in fact suffer changes during installation and operation of the diagnostic, and in-situ, preferably per-shot calibration over the whole image frame would be desirable for accurate calibration of future instruments. Since the spatial structure of the changes is a gradient across the entire image, this does not significantly effect observations of more complex spatial structure in the flows or flow structure at fine spatial scales.

6.1.3 Comparison with dispersive Doppler spectroscopy

In order to validate flow measurements from the CIS system against other diagnostics, measurements of toroidal He II flow at the low field side midplane were performed simultaneously using CIS and dispersive Doppler spectroscopy. The dispersive system used was part of the ECELESTE edge Doppler spectroscopy diagnostic [Morgan, 2011]: a 1m Czerny-Turner spectrograph with a 1680 line/mm grating and 300µm entrance slit, with 60 lines of sight viewing the plasma tangentially between $1.36 \leq R \leq 1.46$m at the midplane. Using the HM07 tangential view, equivalent sight-lines (i.e also viewing the plasma tangentially at the midplane, at the same radius but different toroidal location) were obtained with the CIS instrument. Doppler shift measurements of the He II line were then made in a series of shots based on an L-Mode DND discharge, where the plasma vertical position was scanned vertically between shots. The goal was to attempt to modify the flows between shots by changing the primary X-point from upper to lower, i.e. changing which divertor the near SOL field lines connect with.

Unfortunately these plasmas exhibited sawtooth crashes (periodic reorganisation of the confined plasma where the hot core is ejected, due to an internal kink instability) at a frequency close to the frame rate of the CIS measurements, and which appeared clearly in the ECELESTE measurements which used a much shorter integration time (2ms compared to 20ms for the CIS). The measurements from both diagnostics were averaged over extended lengths of time (40 - 100ms) in order to ensure these variations were averaged over in the same way (or as closely as possible) for both instruments. Another difficulty with these measurements was that the radial zero flow calibration chord usually used for ECELESTE was not available, and the CIS view was not wide enough to capture both the outboard plasma edge and fully radial chords for a phase offset reference. With these restrictions, it is only possible to compare the measured *difference* in flow between consecutive shots with the plasma shifted up and down with any confidence. In order to do this, the mean offset between the flow measurements from the two instruments was found for one of the plasma shots (#28910), and the same offset applied to the other shot. The results are shown in figure 6.5.

Both instruments show little variation in flow as a function of major radius, and a difference of approximately 4km/s between the two shots. This provides only a basic comparison between the two instruments, and using plasma shots with a larger difference between the two shots, more clear spatial structure in the flows and no Magnetohydrodynamic (MHD) activity would be desirable to provide a better test. Data are shown over the radial range where bright He II emission was observed.

The measured shot-to-shot flow difference is similar between ECELESTE and the CIS instrument, and when averaged in space over 1.36m $\leq R \leq$ 1.43 m, is around 1.1km/s larger in the CIS data compared to ECELESTE. Although this discrepancy between the two instruments is fractionally large at around 25%, the absolute value is around the magnitude of expected measurement errors on the CIS measurements, and could also include contributions due to small differences in the sight-line geometry (including differences in contamination such as reflections) and time integration between the two instruments. An improved experiment using more suitable plasma discharges, and preferably with the two instruments sharing common sight-lines and integration time, would be desirable for a more informative comparison.

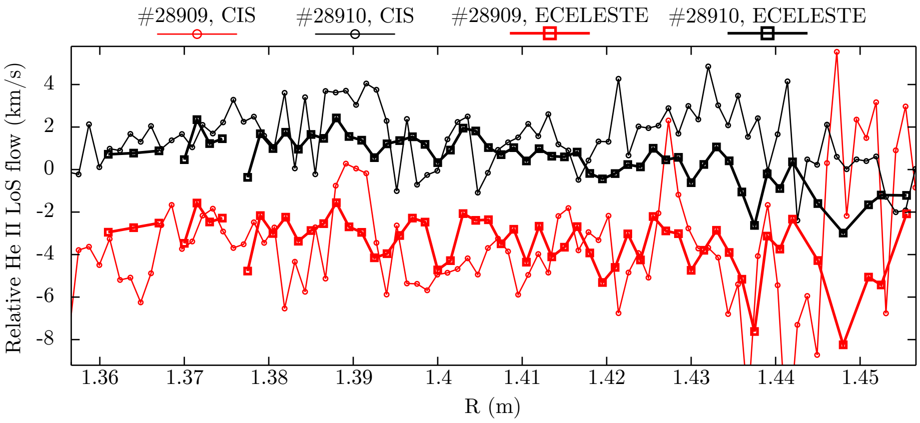

Figure 6.5: Comparison between measurements made in up- and down- shifted plasmas with the ECELESTE and CIS diagnostics. An offset has been added to each instrument for both shots in order to make the comparison. There aren't any error bars on here yet, because I'm not sure what the best way to think about the errors on this is.

6.2 Plasma Observations

6.2.1 Limited Plasmas

At the start of a MAST plasma discharge, the plasma is initially formed in a limiter configuration with the centre column armour acting as the limiter, before connecting to the divertor(s) at later times. During this early limited phase, CIS measurements show a distinctive flow pattern qualitatively consistent with what might be expected for impurities following the main ion flow. Examples of such measurements, in all three species using the HM07 radial view, are shown in figure 6.6. The line integrated brightness and line average flow $\bar{v}$ are shown together in each image:

the image brightness (luminance) represents the line integrated brightness while the colour represents the flow. The positive flow direction is defined as flow away from the camera i.e. redshift of the emission line, and negative flow is towards the camera. Representing the flow and brightness images together in this way helps to give spatial context to the flow information and make the images easier to interpret. The full cross-section of the small limited plasma is clearly visible in figure 6.6, with the centre column visible slightly to the right of the image centre. Due to the viewing configuration, the measured flow is predominantly the toroidal projection of the total flow. The flows appear symmetric about the centre column, as expected for a toroidally symmetric plasma, thereby giving confidence in the instrument phase calibration. The dark, apparently high positive flow feature on the right hand edge of the image towards the top is due to vignetting from a viewing obstruction just in front of the diagnostic. The fringes in the raw data were oriented horizontally, i.e. the best spatial resolution is in the horizontal direction.

First concentrating on the carbon measurements, plasma closest to the centre column shows oppositely directed toroidal flow projections in the top and bottom halves of the plasma: anti-clockwise around the torus (viewed from above) above the vertical midplane and clockwise below the midplane. Assuming this measured flow is predominantly the toroidal component of flow along the magnetic field lines, these flow directions correspond to flow along the field lines towards the high field side midplane in both the upper and lower parts of the plasma. The interpretation of this flow pattern is that the carbon is being dragged along by frictional forces with strong background plasma flow towards the particle sink at the limiter surface (the centre column) on the open field lines in the SOL. In addition to the up-down flow reversal, in the C III image the lower half of the plasma also shows clear radial structure, with a layer rotating anti-clockwise inside (in terms of minor radius) the clockwise layer adjacent to the limiter surface. This appears to indicate a co-current (with respect to the plasma current) toroidal rotation of C III impurities in the confined plasma, with the change in flow sign in the lower half of the plasma essentially providing a visualisation of the separatrix location. This is visible in the C III image since the relatively low temperature of the plasma early in the shot means C III emission is visible both just inside and outside the separatrix. The co-current rotating layer is not seen in the C II image, since C II emits at lower temperatures and is limited to the SOL. In He II only the co-current rotation is seen (no clear up-down direction reversal exists), since He II emits at higher temperatures than C III and is seen almost entirely inside the separatrix at this time.

A feature not consistent with this description is the negative flow region to the

(a) C III (#28751, t = 25ms)

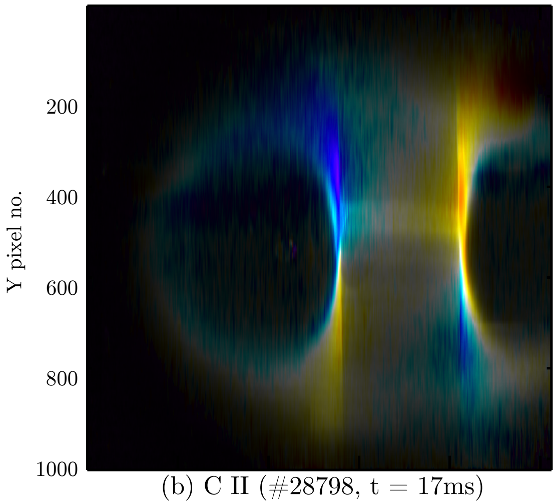

(b) C II (#28798, t = 17ms)

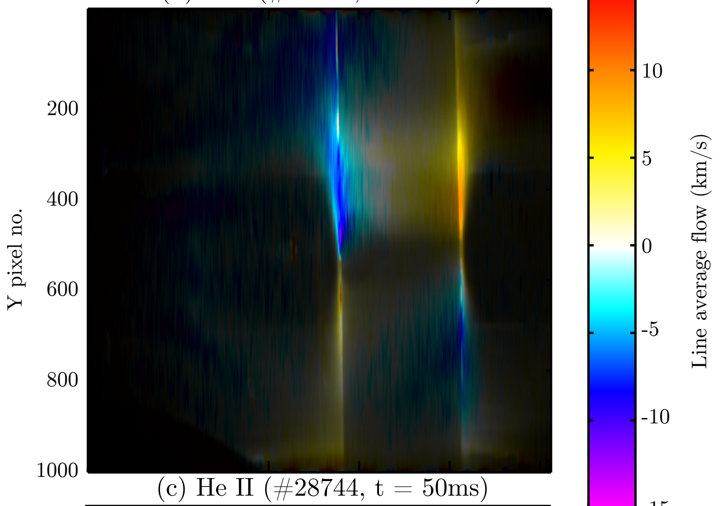

(c) He II (#28744, t = 50ms)

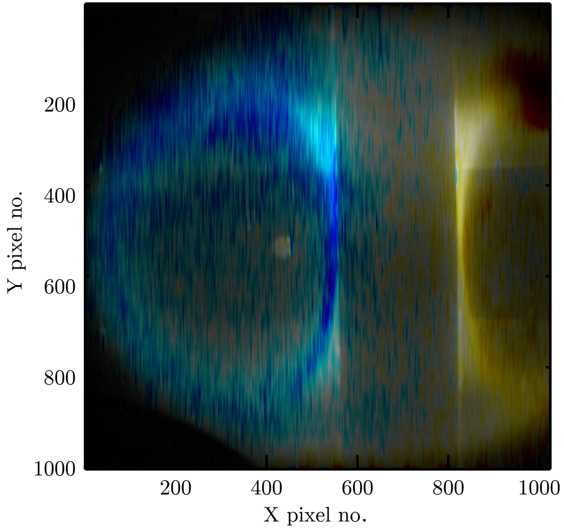

Figure 6.6: Measured flow patterns in limited plasmas at early times during MAST discharges. (a) C III, (b) C II and (c) He II data.

right of the centre column around $y = 300$ pixels in the C III image. Since no similar positive flow feature is seen on the left of the centre column, this cannot be explained by any toroidally symmetric feature of the plasma. This may be an instrumental effect related to reflections from the P5 poloidal field coil which can be seen behind the plasma close to this location, or due to the nearby viewing obstruction mentioned earlier.

6.2.2 SOL flow response to high field side gas fuelling

Novel flow structures have also been observed with CIS during gas puff fuelling of MAST plasmas from the high field side midplane. This is the first time that impurity flows at a gas fuelling location have been observed with an instrument able to provide the level of spatial information given by the CIS technique. MAST has two gas nozzles on the centre column for such fuelling: one on the vertical midplane and one around 100mm below the midplane, approximately opposite each other toroidally (only one of these is used at once). Novel flow patterns have been observed when two conditions are satisfied: 1) the use of either of the HFS midplane gas values for fuelling and 2) The plasma-wall gap (distance from the centre column to the separatrix, determined from EFIT) at the high field side midplane is $\lesssim 2.5$cm.

An example of the time evolution of C III flows in a discharge when these conditions are met, measured from the HM07 midplane view, is shown in Fig 6.7(a)-(e). For these measurements the fringes in the raw images were oriented horizontally, i.e. the direction of highest spatial resolution is horizontal. Image (a) is before the start of the high field side fuelling; the plasma is limited on the centre column and shows a flow pattern similar to that discussed in the previous section. Gas puffing starts after t = 94ms, around the same time that the plasma edge starts to lift away from the centre column. At the same time, poloidally narrow ‘stripes’ in both the flow and brightness profile start to become visible (image (b)). Between 0.127 - 0.143s (images (b) and (c)), a fast ($\sim$ 15km/s) counter-current toroidal flow is seen to develop in the upper half of the plasma, and clear ‘stripes’ of oppositely directed flow, i.e. neighbouring plasma regions rotating in opposite directions past each other, are seen in the lower half (image (c)). In the upper half although the fast toroidal rotation dominates, superimposed on this are poloidally localised increases and decreases in the line-of-sight flow speed following the same spatial pattern as the stripes in the bottom part of the plasma. This flow pattern persists with very little change between about 0.143 - 0.160s.

The white line in image (c) is an equilibrium magnetic field line on the LCFS, as reconstructed by EFIT, projected on to the image. This shows that the observed flow pattern is very well aligned to the magnetic field lines at the separatrix. The measured flows are therefore interpreted as the toroidal projection of parallel flows along the field lines, with the stripe pattern showing oppositely directed flows along neighbouring field lines. Looking closely at figure 6.7(c) there is also fine radial structure in the flows in the lower half of the plasma, with narrow radial layers flowing in different speeds or opposite directions. These 'barber-pole' flow patterns have been observed in 45 different plasma shots throughout the MAST campaign including a variety of plasma configurations, and HFS midplane gas fuelling and a small plasma-wall gap appear to be sufficient and necessary conditions for their observation. These patterns can be seen very weakly in some C II and He II data, but only clearly in C III; it is not clear whether this is due to different localisation of the emission or if the flows of different species are behaving differently. This is the first time flow patterns such as these have been observed in a tokamak plasma, and the phenomenon is worthy of further investigation in the future. The appearance of the images and conditions for the observation of these patterns suggest they may be related to the large, localised particle source provided by the gas puffing, and the magnetic topology at this time in the shot along with interaction of the plasma with the centre column.

As the plasma-wall gap increases, the pattern becomes less prominent: in image (d) at 177ms, only a single reversed flow stripe is clearly visible at around the same poloidal location as the gas injection. The lower half of the plasma appears to have very little toroidal rotation relative to the top, although it is in the same direction. In image (e) at 293ms, the entire inner edge of the plasma is clearly rotating in the same direction (counter-current), although the upper part of the plasma still shows faster rotation than the lower.

6.2.3 Divertor flows of multiple impurity species

Using a series of repeated plasma discharges, divertor flow data for each of the three measurable ion species has been obtained for comparable plasma conditions. The ability to study multiple impurity species in the same plasma has two main potential advantages: 1) Obtaining information from a larger spatial region than possible with a single impurity, due to different spatial localisation of emission from each species, and 2) In plasma regions where information is obtained in multiple species, it may be possible to use the data to investigate the influences of different driving mechanisms of the impurity flow e.g. electric field vs thermal or frictional forces, due to the different charge/mass ratios of the different ions. An illustration of the spatial localisation is shown in figure 6.8, where each colour represents the

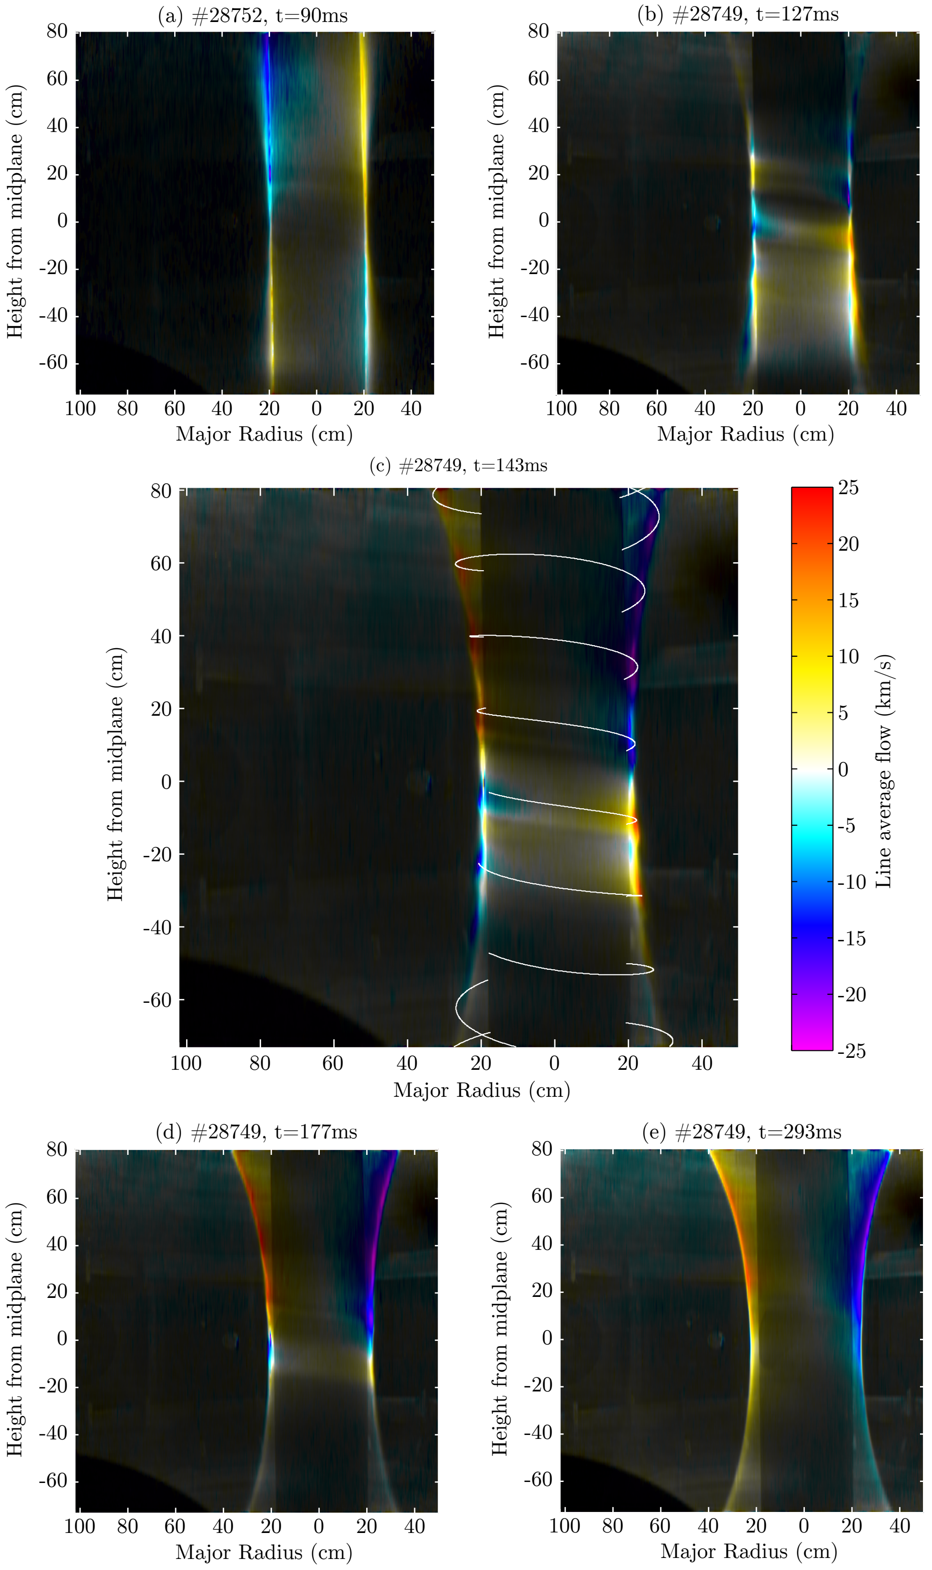

Figure 6.7: Typical time evolution, from (a) to (e), of C III flows in the high field side SOL in a DND plasma using high field side midplane fuelling. The exposure time was 10ms except for in (a), where it was 1ms. The white line superimposed on (c) is an equilibrium magnetic field line on the separatrix, calculated from EFIT and superimposed on the image using the camera projection model from section 5.2.3.

emission from a different impurity line in H-Mode. Each has been normalised to its maximum brightness, so the image does not provide a comparison of the relative brightness of the three lines but only the differences in spatial distribution. He II (red in fig. 6.8), emitting in the hottest temperature range of the three species, is brightest in a very narrow region around the LCFS in the divertor legs and around the LCFS above the X-point. C III (blue in fig. 6.8) can be seen localised further from the LCFS both on the outboard side and in the PFR, while C II emission, at the lowest temperatures, extends further out again on the outboard side and further into the PFR.

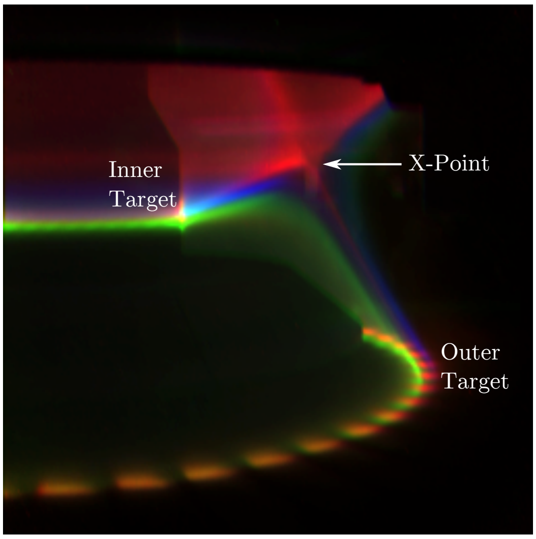

Figure 6.8: Distribution of emission from the C III (blue), C II (green) and He II (red) spectral lines in the MAST divertor during H-Mode, obtained with the CIS diagnostic in repeated LSND plasma discharges. Different localisation of the emission is clearly seen for each line, primarily due to the different temperatures at which the lines are brightest.

The discharge used for the multi-species divertor measurements was a LSND plasma with plasma current $I_p = 600$ kA and 1.2MW of NBI heating power. Plots of the plasma current, heating power, line-integrated density and midplane D_α brightness for the repeated shots used here are shown in figure 6.9. These are mostly very well matched, except for 29625 which has a delayed L-H transition and is significantly under-dense in the H-Mode phase, after delayed onset of the heating beam. This shot was used for comparisons in the L-Mode phase which is better matched, and 29628 was used for the H-Mode phase (but is over-dense in the L-Mode phase). A LSND plasma was used rather than DND for two reasons: the signal level in the divertor was found to be much larger in LSND plasmas, thus allowing higher frame rates, and the position of the X-point is lower in the vacuum vessel so the region near the X-point is more accessible to the diagnostic. This particular discharge was chosen since it had periods of L-Mode ($t < 0.33$ s), quiescent ELM-free H-Mode (approx. $0.33 < t < 0.38$s) and low frequency ELMing H-Mode (0.38s onwards) which allowed comparison in all 3 operational regimes. In the ELMy phase the inter-ELM spacing was large enough to accommodate the integration times for C II and C III, however good inter-ELM He II data was not obtained. The fringe pattern was rotated on the detector for these measurements such that the fringes were rotated anti-clockwise from the horizontal by 22.5°, in order to reduce demodulation artefacts on the high field side divertor leg.

Brightness and line-integrated flow images for each species, in the L-Mode and ELM-free H-Mode phases of the shot, are shown in figure 6.10. Each is integrated over 16ms from around the same time during the shot, with the exact times chosen so as to best match the density between shots. The white areas at the divertor targets for He II indicate regions with unusually low fringe contrast, indicating spectral contamination of the He II line as discussed in section 3.1, and therefore unusable data. The horizontal stripes at the inner target for C II are artefacts remaining due to the sharp emission feature at the strike point and rotated fringe pattern. For all three species, opposite signs of flow are observed near the inboard and outboard divertor targets: positive (away from the camera) at the inboard side and negative (towards the camera) at the outboard side. Considering the direction of the magnetic field lines, this is consistent with flow of all three impurity species along the field lines towards the divertor target plates, i.e. with the background plasma. This observation is consistent with C III CIS measurements in the DIII-D divertor [Howard et al., 2010a]. Between the L-mode and H-mode phases, C III and He II measurements show a clear increase in flow speed towards the targets by a factor of ∼1.5 - 2 and ∼2-3, respectively, while C II shows no significant change. Considering the spatial localisation of the emission, this may suggest an increase in flow towards the divertor which is greatest near the LCFS and weakens out into the far SOL and PFR. However, this interpretation does not take into account any change in localisation of the emission between L and H mode, nor the fact that the impurity flows may be behaving independently from both the background plasma and, to some extent, each other. It is, however, consistent with recent high speed imaging observations that the dynamics of the PFR plasma are not strongly dependent on the confinement regime. Tomographic inversion of the multi-species data would allow clearer interpretation, by providing unambiguous information about the spatial localisation of the recorded flow information for different species and confinement modes. This will be discussed in section 6.2.5.

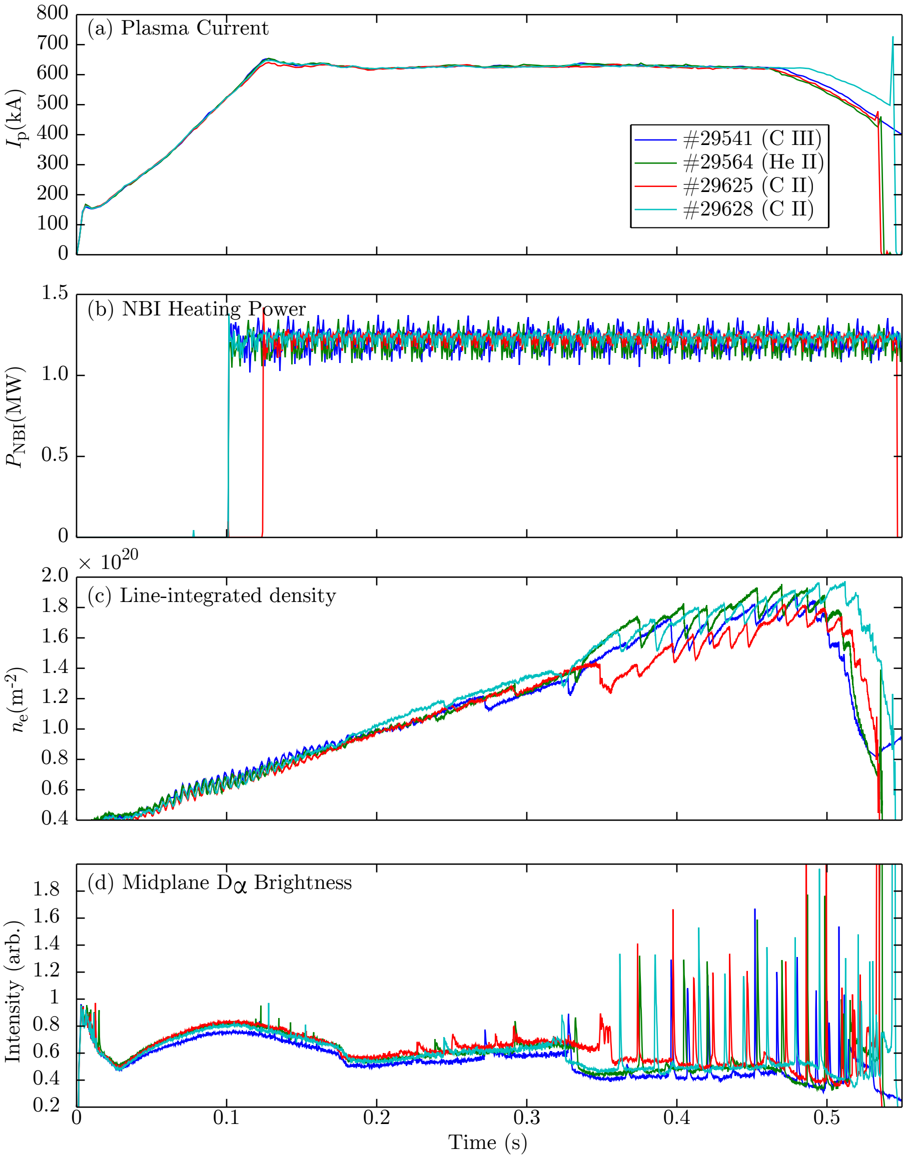

Figure 6.9: Summary of plasma shots used for multi-species divertor measurements: (a) Plasma current, (b) NBI heating power, (c) line integrated electron density and (d) midplane D_α brightness. Species names in the legend indicate which filter was used on the CIS diagnostic for each shot.

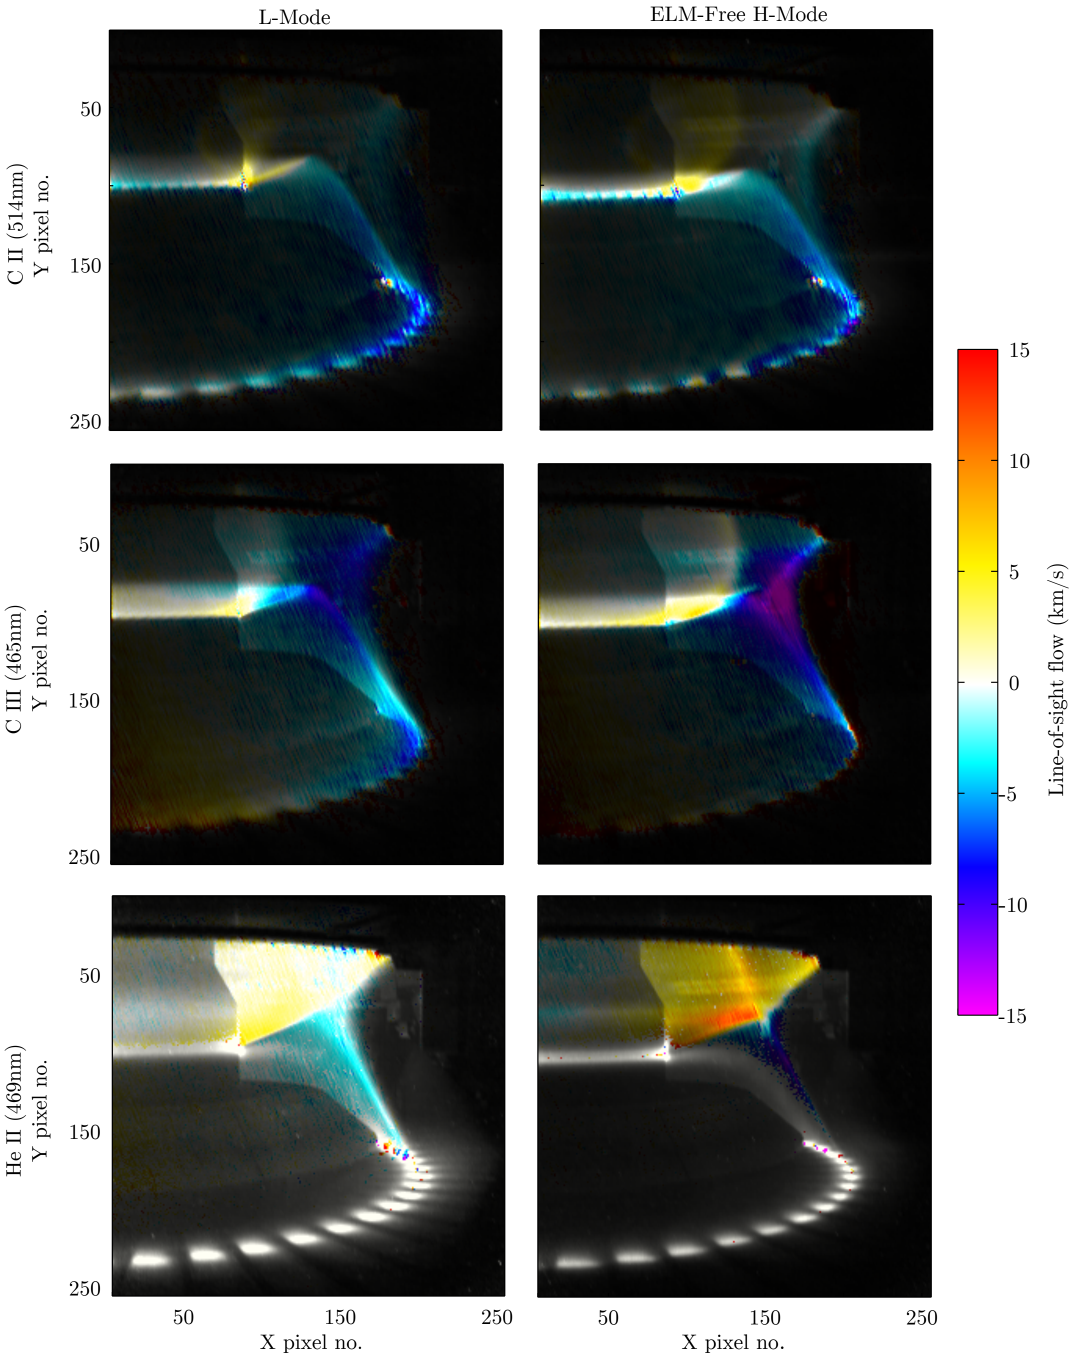

Figure 6.10: Line-average brightness and flow images for C II (top), C III (middle) and He II (bottom) in repeated LSND plasma shots, in L-Mode (left) and ELM-free H-Mode (right) phases of the discharge.

6.2.4 Resonant Magnetic Perturbations

Mitigation of the high transient particle and heat loads to divertor surfaces from ELMs is an area of much current interest and concern, due its potential impact on divertor operating lifetimes on future high power devices. One of the leading methods for achieving this is the application of Resonant magnetic perturbations (RMPs). These are small perturbations to the equilibrium magnetic field which are resonant, in the sense of having the same helicity as, field lines on a particular flux surface in the plasma edge, and are applied by a set of coils outside the plasma. This causes ergodisation of the magnetic field around the resonant flux surface and destruction of the usual closed flux surface geometry, changing the confinement characteristics. The idea behind this technique is to reduce the edge pressure gradient below the value which would trigger an ELM, replacing the undesirable large bursts of outward transport with a steady state, lower level of transport along the ergodised field lines. Complete suppression of ELMs has been demonstrated using this technique on DIII-D [Evans et al., 2004a], whilst ELM mitigation (increase in the ELM frequency with a corresponding reduction in the peak heat flux per ELM) has been demonstrated on JET [Liang et al., 2007], ASDEX-Upgade [Suttrop et al., 2011] and MAST [Kirk et al., 2011]. This ELM mitigation is more difficult to explain than the complete suppression on DIII-D with the above picture, and all of the effects of RMPs on tokamak plasmas are not yet well understood.

A recent experimental observation is the appearance of 'lobe' structures in the divertor in the presence of applied RMPs, near the X-Point and caused by the modification of the magnetic topology. These structures were first predicted theoretically [Evans et al., 2004b], before being directly observed with visible imaging on MAST [Kirk et al., 2012] and extreme UV / Soft X-Ray imaging on DIII-D [Shafer et al., 2012]. A photograph of such lobes on MAST (viewed in C III light) is shown in figure 6.11; these are very fine spatial structures with individual lobes having widths as small as 3 - 4 mm [Harrison et al., 2014]. Three dimensional fluid modelling with the EMC3-EIRENE code package has predicted the appearance of novel helical SOL flow patterns associated with RMPs on multiple machines (for example see Lunt et al. [2012]), due to the modified connection length to the divertor targets along the perturbed field lines. In the vicinity of the lobes these are predicted to take the form of strongly counter-flowing channels following a similar spatial structure to the visible lobes. Due to the spatially complex and 3D nature of these structures, the capabilities of the coherence imaging diagnostic are uniquely suited for seeking experimental evidence of this phenomenon. Measurements of spatial flow patterns near the X-Point associated with RMPs have not previously been attempted.

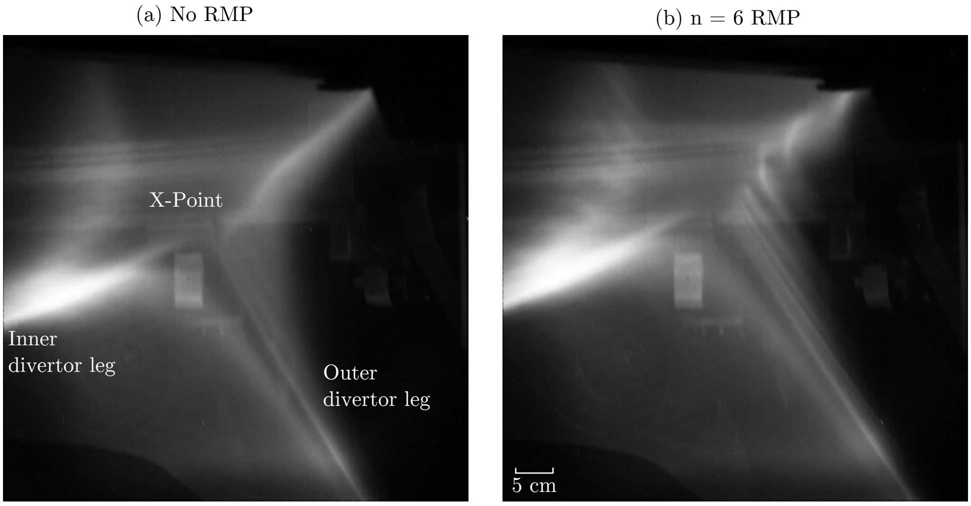

Figure 6.11: C III emission in the region of the X-point (a) without applied RMPs and (b) with applied RMPs. With the RMPs on, lobe structures protruding from the plasma edge are clearly visible. The scale bar in (b) shows the scale at the tangency plane.

To attempt to best resolve the fine spatial structure associated with the X-point lobes, the zoom capability of the CIS instrument was used to obtain a narrowed field of view of around 14 degrees ($f_1 \approx 50$mm) around the lower X-Point, from the usual HL07 divertor viewing port. This gives a FWHM of the spatial response to flows of around 1.6cm perpendicular to the fringes, and the fringe pattern was rotated such that the maximum spatial resolution was approximately aligned across the lobes. The main disadvantage of this viewing configuration was that there were no radial sight-lines in the view for phase offset calibration. The flow measurements referred to in this section are therefore not absolutely calibrated, and should only be compared within any given plasma shot. RMPs flow data were only obtained for C III, since the longer exposure times required for the other species were longer than the inter-ELM time with the RMPs on.

Brightness and $\bar{v}$ images for the X-point region in a LSND plasma in H-Mode, with and without an applied n = 4 RMP, are shown in figure 6.12. Modification of the flows can be clearly seen in the image with the applied RMP compared to the image without, and considerable spatial structure in the flow is seen in the vicinity of the lobes. The peak-to-peak amplitude of this flow structure is up to ~ 8 km/s near the separatrix. While clear modification of the flows by the RMPs is observed, the distinct spatial structure seen in the simulations is not clear from the recorded images. This could be due to a number of effects, including line integration and viewing geometry effects due to the 3D lobe and flow structures; the C III ions ex-

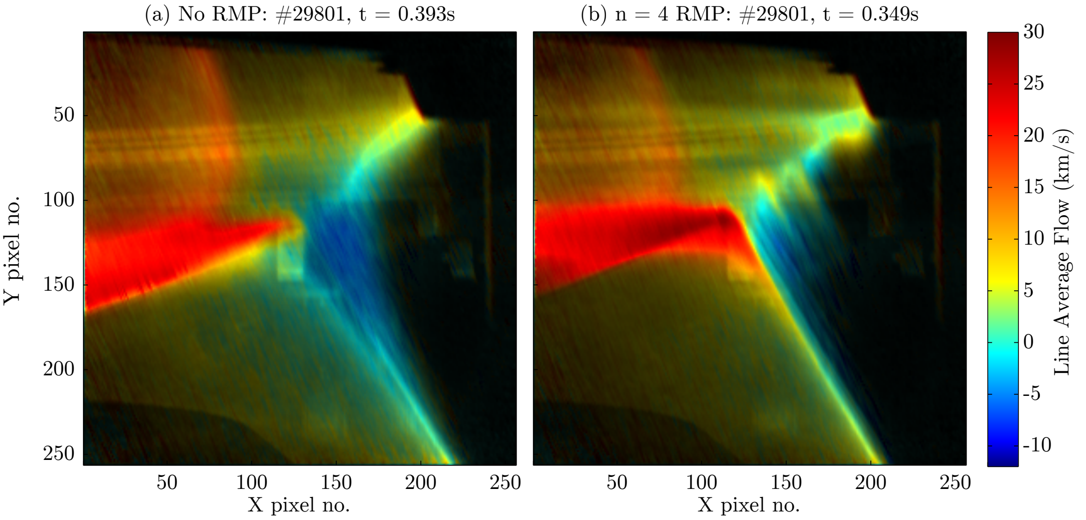

Figure 6.12: Line integrated brightness and flow images with (right) and without (left) an applied n = 4 RMP in an H-Mode LSND plasma. The calibration offset is unknown for these measurements, hence the flow speeds are only correct to within a constant offset. The bright rectangular feature close to the X-point is reflected light from a structural component supporting the P4 coil. Reflections from the P3 coil can also be seen through the plasma.

periencing different flow to the background plasma; and the localisation of the C III emission. Comparison with EMC3-EIRENE results for the relevant MAST divertor conditions is required to assess how the data compare to modelling predictions, and would be a very valuable avenue for future work. Due to the inherently 3D nature of the lobe structures, tomographic inversion is not applicable to these results and the forward modelling approach must be used instead. Extending the existing forward modelling code for use with non-axisymmetric plasma and magnetic field information would be required in order to do this.

6.2.5 Tomographically Inverted Data

Data from several shots have been tomographically inverted to obtain 2D profiles of C III parallel flow and emissivity in the divertor, covering the region from just above the lower X-point to the divertor targets, using the methods described in section 4.6. Inverted profiles based on the C III data presented in section 6.2.3 are shown in figure 6.13, for the L-Mode, ELM-free H-Mode and ELMy H-Mode discharge phases (the ELMy H-Mode data is from the period between ELMs). Here the sign convention for the flow is that positive parallel flows have a toroidal component in the same direction as the plasma current. Each profile was generated from a

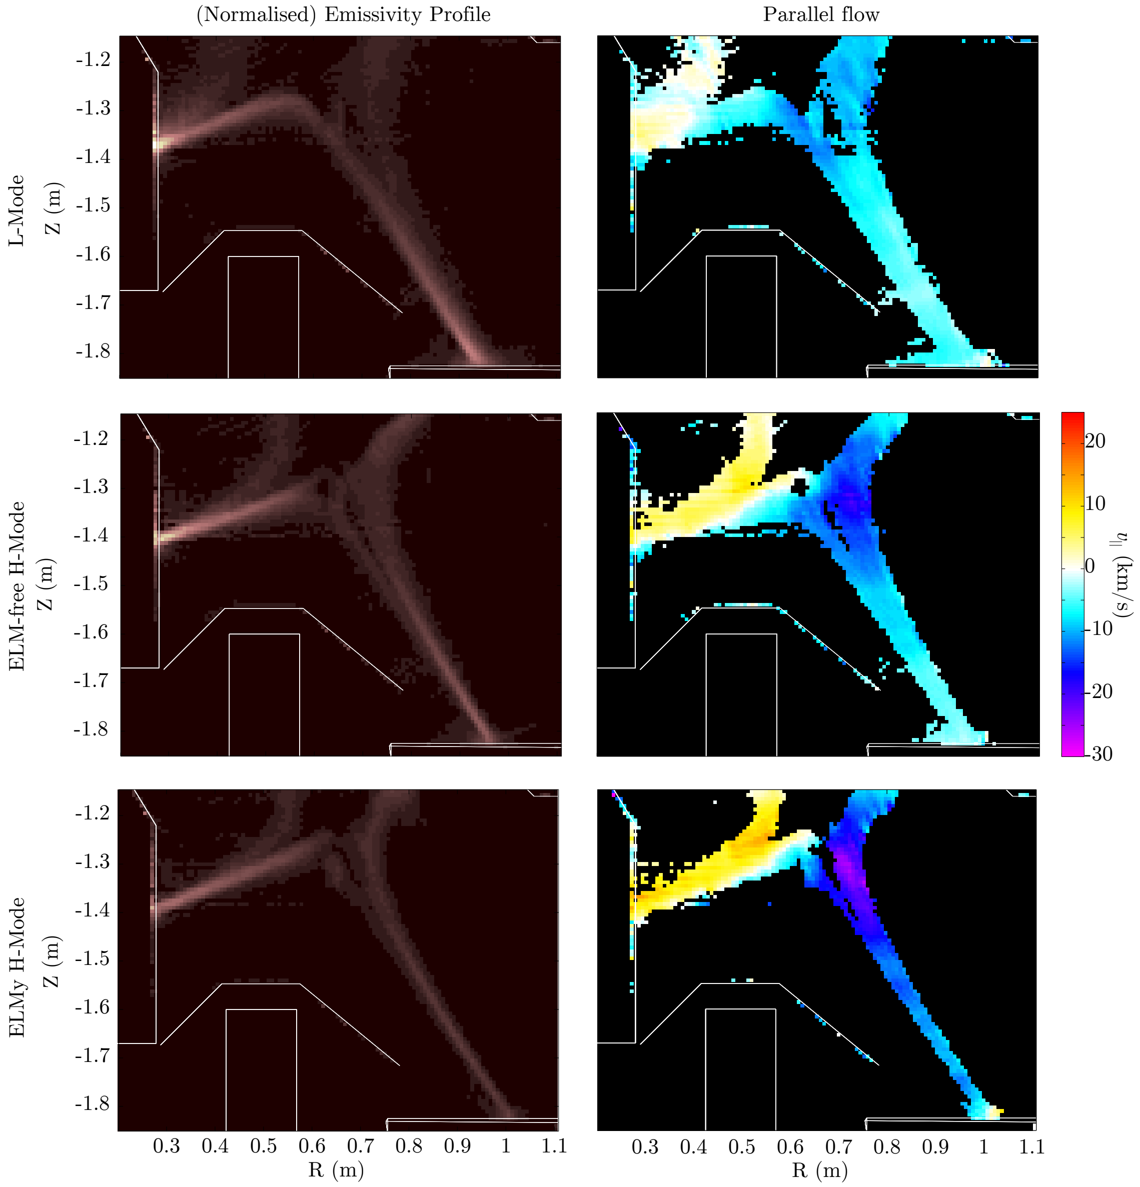

Figure 6.13: Tomographically inverted emissivity (left) and C III parallel flow (right) profiles for periods of different operation modes in shot #29541: L-Mode (top), ELM-Free H-Mode (centre) and ELMy H-mode (bottom).

single 1ms exposure. The general features of these images are the same as shown in the line-integrated data: parallel flow towards both divertor targets and increased flow speeds going from L-Mode → ELM-free H-Mode → ELMy H-Mode. Note the stripe pattern of positive and negative flows at the high field side in the L-Mode reconstruction (at $Z \geq -1.3$m): this is an artefact which was propagated through the inversion from the line integrated images, and is a known behaviour of the FFT demodulation in the presence of reflections from the poloidal field coils. There is also a rectangular artefact due to reflections from a supporting bracket of these coils (at $R \sim 0.65$m, $Z \sim -1.3$m), particularly noticeable on the ELMy H-Mode inversion. Reconstruction artefacts such as these, and those due to response matrix errors caused by line-of-sight calibration errors and some sight-lines being poorly constrained by the data, appear to be the dominant source of error in the inverted profiles (the effect of noise in the line-integrated data on the inversions was investigated numerically with a Monte-Carlo technique and found to be typically $\ll 1$km/s). These types of reconstruction artefacts can typically be readily identified since they correspond to known artefacts in the input line integrated data, and/or appear as static features in the images which do not respond to changing plasma behaviour.

In conjunction with suitable 2D SOL modelling, these profiles and similar inversions for C II and He II could in principle be used to investigate the physics of the flow differences seen between the different species. In practise, however, the C II and He II images are less amenable to inversion, and profiles for these impurities could not be reliably produced with sufficient quality using the inversion method described in the previous chapter. One of the main difficulties with these inversions is the geometry of the outboard divertor PFCs on MAST, consisting of a set of toroidally discrete tiles separated by relatively large gaps. For C II and He II the emission is peaked close to the target plates, which leads to a toroidally periodic, rather than toroidally symmetric emission pattern in the images. This can be seen in figure 6.10, where the C II and He II measurements show a much stronger toroidally periodic pattern than C III. Since toroidal symmetry is assumed in the formulation of the tomography problem, this causes artefacts in the inversions. In future it may be possible to implement an inversion geometry with toroidal periodicity matched to that of the target plates, in order to reduce this effect.

6.2.6 Summary

The MAST CIS flow diagnostic was operated successfully throughout the experimental campaign in May - September 2013. Using integration times between 1 -

16ms, data was obtained for all three impurity species and flow image noise levels were as low as < 1km/s. The phase calibration offset was successfully monitored using radial sight-lines, on both inter- and intra-shot timescales. This revealed initial problems with vibrations causing frame-to-frame changes in the flow calibration offset, however this problem was alleviated by re-designing the mounting scheme for the instrument. In future, a more robust mechanical design taking into account the vibration environment could improve this aspect of the calibration stability. On longer timescales, variation of the calibration offset from shot-to-shot were up to 16km/s over the course of a week. Since these slower drifts are likely to be dominated by thermal effects, improved thermal stabilisation is required to reduce this effect. Using phase shape calibrations performed before and after a week long period of operation, the phase shape calibration was found to be stable to within 3km/s over the whole frame. Although this has not caused significant problems with the present measurements, an in-situ calibration procedure is desirable to monitor the phase shape calibration more accurately. Comparisons between CIS and dispersive Doppler spectroscopy measurements using the ECELESTE diagnostic show reasonable agreement in the flow differences between different discharges, however the absolute calibration of the instrument could not be properly benchmarked due to absolute calibration difficulties with both instruments on these observations.

In plasma measurements on MAST, the CIS diagnostic demonstrated strong capabilities in revealing spatially complex flow patterns which would be very difficult to diagnose and interpret using traditional dispersive systems. In limited plasmas, impurity flows consistent with flow towards the limiter in the SOL and co-current rotation of the confined plasma were observed. Strong responses of flows to high field side gas puff fuelling were also observed for the first time, most notably in the form of field-aligned, counter-rotating patterns around the centre column under certain conditions. In the divertor, data was obtained for multiple impurity species in different plasma conditions, suggesting a sudden increase in flow towards the divertor targets in H-Mode which is greatest at the strike point and less prominent in the far SOL and PFR. The CIS diagnostic was also used to perform the first measurements of divertor flow perturbations due to the application of RMPs for ELM control. Tomographically inverted data was successfully obtained for C III in a variety of shots, however toroidally periodic emission profiles from C II and He II prevented the same high quality of inversion for these species. Interpretation of the physics behind these more complex plasma observations is still ongoing, and in general will require comparisons with sophisticated modelling using codes such as EDGE2D, SOLPS and EMC3.

Chapter 7

Conclusions & Further Work

This thesis has presented the development of, and first results from, a Doppler coherence imaging spectroscopy (CIS) diagnostic for measuring impurity ion flows in the Scrape-off layer (SOL) and divertor of the Mega-Amp Spherical Tokamak (MAST). Coherence imaging spectroscopy is a technique based on narrow-band Fourier transform spectroscopy which allows low order spectral information, such as Doppler shifts and line widths, to be captured with 2D spatial resolution over wide, continuous fields of view using an imaging interferometer. The primary advantages of this technique over other flow diagnostics (Mach probes, impurity plume imaging, multi-chord Doppler spectroscopy) are the very large amount of spatial information obtained, and the simple relationship between the measured interferogram properties and spatially varying impurity ion flows.

In chapter 2, the principles and theory of the coherence imaging technique were reviewed. Impurity flows are measured using the Doppler shift of emission lines from partially ionised impurities in the cool SOL and divertor plasma. By interfering the light in a given spectral line with a phase delayed version of itself, the complex degree of coherence of the emission is measured. As was illustrated in figure 2.2, measurements of the interferogram properties (brightness, fringe phase, fringe contrast) at a single value of interferometer time delay are sufficient to recover low order spectral information about the emission line. It was shown that the measured fringe phase is related to the average impurity ion flow along the diagnostic's line of sight, weighted by the local impurity emissivity. This straightforward relationship combined with the large amount of spatial information captured enables tomographic inversion of the data, to obtain localised measurements of flows. However, for multiplet spectral lines the measured phase will be affected by any changes in the component line ratios, or by spectral contamination by continuum or nearby line emission. This effect cannot be separated from a real Doppler shift, hence carefully chosen narrow-band filters must be used to select for the spectral line of interest, and minimise any contamination.

Existing implementations of CIS diagnostics were also reviewed in chapter 2, all based on the polarisation interferometer concept but using different methods to perform the necessary small scan of the interferometer delay τ (e.g. time multiplexing and spatial multiplexing methods). The spatial heterodyne concept (described in section 2.5.3 and briefly illustrated in figure 2.7) was chosen for the MAST system, due to the simple 'static' instrument design and the fact that the time resolution is limited only by the detector framing rate.

7.1 Instrument Design

Chapter 3 presented in detail the design of a coherence imaging diagnostic for MAST. Preliminary spectral measurements revealed the best candidate spectral lines were the C III 465nm triplet, C II 514nm multiplet and He II 468nm multiplet. High resolution spectral measurements (see fig. 3.1) revealed some contamination of the C III line by nearby O II line emission, however numerical modelling suggested the effect on flow measurements would be around 1km/s, compared with the range of $\sim 5 - 30$ km/s usually measured for this species. The He II line showed unacceptable contamination near the divertor targets, making measurements in this region impractical for this spectral line.

The instrument design was based on the use of commercial camera lenses (supplied by Sigma imaging), a high speed CMOS camera (Photron APX-RS) and custom birefringent optics (supplied by Australian National University) in an imaging polarisation interferometer. The choice of lenses and spectral line selection filters were made based on simplified numerical modelling of the imaging system, based on limited optical data provided by the lens manufacturer. The choices made were a compromise between maximising collecting power, minimising vignetting, and controlling the blue-shift of the filter pass-bands for light at the edge of the field of view passing obliquely through the filter. Three different birefringent 'delay plates' were specified, to obtain three different interferometer phase delay values optimised for different measurements. The choice of these delay plates was based on calculations of the expected fringe contrast for the different multiplet lines (fig. 3.8), in order to maximise the fringe contrast. A summary of the component specifications of the diagnostic is given in Appendix A.

7.2 Data Analysis Methods

In Chapter 4 procedures were developed to extract brightness and flow measurements from the raw CIS data. As an input to this development to provide realistic test data, forward modelling code was developed to generate simulated diagnostic images (e.g. figure 4.2) from the output of OSM plasma simulations [Harrison, 2010; Lisgo et al., 2005] for MAST. Images created using this code were used throughout the data analysis testing and development, and could potentially also be employed to compare plasma modelling results to the experimental data.

The analysis of MAST CIS data begins by removal of bright spots and streaks in the images due to D-D fusion neutrons in NBI heated discharges, before the underlying brightness image is extracted using frequency domain filtering and boxcar smoothing perpendicular to the fringes. The brightness extraction was found to reproduce input test images to within $\sim 5\%$. The extracted brightness is then factored out of the raw image, and the flow information (fringe phase) is extracted using the analytic signal representation [Gabor, 1946; Lawrence Marple Jr., 1999] of the resulting signal. Windowing in Fourier space and apodisation of the input data around sharp jumps in brightness are used to reduce artefacts due to sharp spatial features and image noise, at the cost of smoothing the spatial response to flows.

Pixel-based tomographic inversion using the SART algorithm [Andersen and Kak, 1984] was implemented in MATLAB and tested with simulated divertor data, including integrated testing with the interferogram demodulation. The results show good quantitative agreement between the plasma profiles used to generate the test data and the recovered profiles from the simulated measurements. However, for noisy data the flow inversion shows some ringing artefacts at sharp image edges, and larger errors for parts of the plasma profile with low emissivity. These must be kept in mind when applying the techniques to real data, and inverted flow profiles can only be obtained from parts of the plasma which are emitting sufficiently brightly. The inversion testing also did not include sources of error which can become important in the real data, such as line of sight registration errors and artefacts due to reflections in the line integrated images.

7.3 Instrument testing & calibration

Detailed characterisation of both the individual diagnostic components and the integrated system were performed, and the results were presented in Chapter 5. From measurements with a calibrated integrating sphere, the Photron APX-RS camera was found to exhibit integral nonlinearity of around 4%, which was corrected for using a look-up-table based on these measurements. The conversion gain of the camera was measured to be $(23.3 \pm 0.2)e^-/ADU$ using mean-variance curves, and the read noise was estimated from dark frames as $(37 \pm 0.4)e^-$. The dark signal was found to drift with ambient temperature by around 1.8ADU/°C, which was sufficient to prevent accurate calibration of the interferometer's characteristic contrast, and therefore to prevent ion temperature measurements using the CIS instrument. The line selection band pass filters were also found to produce some unexpected structure in flat field images (see fig. 5.5), which cannot be fully accounted for without more detailed measurements of the spatial and angular dependence of the filter transmission. The absolute light sensitivity and vignetting of the assembled diagnostic were measured and were within approximately 20% of the design calculations.

The interferometer group delay, which sets the proportionality between Doppler shift and measured interferometer fringe phase, was calibrated by measuring the rate of fringe phase change with wavelength tuning of a tuneable laser source. This is the first time this type of calibration has been performed for a coherence imaging instrument. The calibrated group delay values were found to agree very well with the design calculations for the thinner two delay plates, however differed by about 5% for the thickest. This did not have any impact on the baseline flow measurement configurations. Measurements of the interferometer's intrinsic contrast showed the contrast was highest at the edges of the field of view, contrary to what might be naïvely expected. Measurements of non-uniformities in the interferometer optics provided a qualitative explanation for this behaviour. Measurements of the calibration stability in the lab confirmed the need for active temperature stabilisation of the birefringent interferometer optics, and the need for monitoring of the calibration during plasma operations.

To calibrate the system for measurements on MAST, procedures were developed based on using a combination of a spectral calibration lamp and light from the plasma itself, and were presented in section 5.3.2. Calibration of the diagnostic sight-lines through the plasma was performed with a two-stage calibration process, separating the calibration of the imaging properties of the instrument and its position and orientation on the tokamak, with both stages based on existing calibration codes.

7.4 CIS flow measurements on MAST

The CIS diagnostic was operated successfully throughout the May - September 2013 experimental campaign on MAST, and these results were presented in Chapter 6.

Using integration times between 1 - 16ms, data was obtained for all 3 impurity species and flow image noise levels were as low as < 1km/s (typical measured flow amplitudes were $\sim 5-30$ km/s). The phase calibration offset was successfully monitored using radial sight-lines, on both inter- and intra-shot timescales. This revealed initial problems with vibrations causing frame-to-frame changes in the flow calibration offset, however this problem was alleviated by re-designing the mounting scheme for the instrument. In future, a more robust mechanical design taking into account the vibration environment could improve this aspect of the calibration stability. On longer timescales, variation of the calibration offset from shot-to-shot were up to 16km/s over the course of a week. Since these slower drifts are likely dominated by thermal effects, improved thermal stabilisation would be required to reduce this effect. Using phase shape calibrations performed before and after a week long period of operation, the phase shape calibration was found to be stable to within 3km/s over the whole frame. Although this has not caused significant problems with the present measurements, an in-situ calibration procedure is desirable to monitor the phase shape calibration more accurately. Comparisons between CIS and dispersive Doppler spectroscopy measurements using the ECELESTE diagnostic showed reasonable agreement in the flow difference between different discharges, however the absolute calibration of the instrument could not be properly benchmarked due to absolute calibration difficulties with both instruments during these measurements.

In plasma measurements on MAST, the CIS diagnostic demonstrated strong capabilities to reveal spatially complex flow patterns, which would be very difficult to diagnose and interpret using traditional dispersive systems. In limited plasmas, impurity flows consistent with flow towards the limiter in the SOL and co-current rotation of the confined plasma were observed. Strong responses of flows to high field side gas puff fuelling were also observed, most notably in the form of field-aligned, counter-rotating patterns around the centre column under certain conditions (figure 6.7). In the divertor, data were obtained for multiple impurity species in different plasma conditions, suggesting a sudden increase in flow towards the divertor targets in H-Mode which is greatest at the strike point and less prominent in the far SOL and PFR (figure 6.10). The CIS diagnostic was also used to perform the first measurements of divertor flow spatial structures associated with the application of RMPs for ELM control (fig. 6.12). Tomographically inverted data was successfully obtained for C III in a variety of shots, however toroidally periodic emission profiles from C II and He II prevented the same high quality of inversion for these species. Further work is also required to properly quantify the errors on the tomographically inverted data. Interpretation of the physics behind the more complex plasma observations in this thesis is ongoing, and in part will require comparisons with sophisticated numerical modelling, e.g. using codes such as EDGE2D, SOLPS and EMC3.

7.5 Further Work

The results presented in this thesis suggest a number of avenues for future continuation of this work. First, several possible improvements to the instrument design have been identified which would benefit a future iteration of the diagnostic. Probably most important amongst these is improvement of the calibration stability and procedures. This would remove the need to frequently move the diagnostic to the lab for recalibration, and allow absolutely calibrated flow measurements when using any plasma view, which is not possible with the current system. This would involve 3 areas of work: 1) Improvement to the mechanical design of the instrument to be more robust against temperature changes and vibration; 2) Better thermal stabilisation, or athermalisation, of the birefringent optics, and 3) implementation of an in-situ calibration source and integrating sphere, with appropriate coupling of the calibration light into the instrument. With regard to (2) and (3), recent work on Doppler coherence imaging systems on the DIII-D tokamak has started to address these issues [Allen, S. private communication, 2014] with techniques which could be applicable to MAST.

One of the chief problems with data quality from the MAST instrument was the presence of artefacts due to sharp image features which are not well handled by the FFT based demodulation. From the results of modelling performed in chapter 5, these could be essentially eliminated if the brightness image could be obtained without the spatial smoothing due to removal of the interference fringes. This could be achieved with a dual-detector system design: using a polarising beamsplitter instead of the first polariser in the interferometer to direct light which is currently rejected to a second detector, without any interferometer optics. By using the resulting ‘full resolution’ image in the analysis of the interferogram, substantially better flow image quality could be obtained for scenes with sharp intensity features. An alternative use for the extra channel would be to construct a dual wavelength system for measurements of two spectral lines simultaneously, with essentially no loss of light in either channel compared to a single channel system.

With the installation of the novel ‘Super-X’ divertor [Katramados et al., 2011; Valanju et al., 2009] as part of the current upgrade to MAST, over the coming few years there will be great interest in diagnosing low temperature, dense divertor plasmas. The coherence imaging technique has the potential to be a very powerful tool for this task, particularly in novel divertor geometries such as Super-X, and the scope of coherence imaging on MAST-Upgrade could be broadened to include plasma properties other than flows. For example, electron density imaging using stark broadening of deuterium lines has been demonstrated using coherence imaging on a linear plasma device[Lischtschenko et al., 2010], however has not yet been implemented on a tokamak. Another interesting area wh

In addition to continued development of the diagnostic itself, large amounts of flow imaging data were obtained during this work but analysis and interpretation of this data, in the context of SOL and divertor physics, was largely outside of the project's scope. Continued analysis and interpretation of the flow data, by comparison with suitable 2-3D modelling and analytical calculations where available, would be a natural extension of this work. Due to the large amounts of data, it would be interesting to develop methods to analyse the data set as a whole and see if any general trends in flow behaviour can be identified.

Appendix A

Component Specifications & Data

Table A.1 summarises the specifications of the MAST CIS diagnostic components. Opto-mechanical and supporting components, some of which were custom manufactured for the MAST system, have not been included.

Component	**Manufacturer/Supplier**	**Description**
Camera	Photron	**FASTCAM APX-RS Model 250K** 1024x1024 px, 10bpp (Custom CMOS) Pixel size: 17μm square Frame rate (full frame): $\leq$ 3kHz Quantum Efficiency: $\sim 40\%, 400 - 700$ nm On-board frame storage: 2.5GiB (2048 full frames)
C III filter[2]	Andover / LOT Oriel	λ =(464.7 +0.5/-0)nm, $\Delta_{\text{FWHM}} = (3.0 \pm 0.5)$nm, $N^* = 1.45$
C II filter[2]		λ =(512.3 +0.5/-0)nm, $\Delta_{\text{FWHM}} =$ (5.0 +0.5/-0)nm, $N^* = 2.05$
He II filter[2]		λ =(468.65 +0.4/-1)nm, $\Delta_{\text{FWHM}} = (2.0 \pm 0.5)$nm, $N^* = 1.45$
Lens l_1	Sigma Imaging	Sigma **17-70mm f/2.8-4.5 DC Macro**, F-Mount
Lens l_2		Sigma **105mm f/2.8 EX DG**, F-Mount
Lens l_3		Sigma **150mm f/2.8 EX DG**, F-Mount
α-BBO Delay Plate	CLaser Photonics / ANU[1]	$\varnothing$ = 30 mm, AR coated 400 - 700nm, L = 4.6 mm
α-BBO Delay Plate		$\varnothing$ = 30 mm, AR coated 400 - 700nm, L = 6.5mm
α-BBO Delay Plate		$\varnothing$ = 30 mm, AR coated 400 - 700nm, L = 9.8 mm
α-BBO Savart Polariscope		$\varnothing$ = 30 mm, AR coated 400 - 700nm, L = 2.2mm
α-BBO Savart Polariscope		$\varnothing$ = 30 mm, AR coated 400 - 700nm, L = 4.0mm
Polarisers	Newport / ANU[1]	Newport model **20LP-VIS**, $\varnothing$ = 30.5 mm, Extinction ratio = 25000 : 1
Temperature Controller	Andover Corp	Andover Model **101FRDC00-CTRL** Regulation range 30°C - 60°C Regularion accuracy ±0.2°C
Temperature Cell	Andover Corp / ANU[1]	Andover Model **101FRDC00-50** $\varnothing_{\text{inside}} = 50$ mm, $\varnothing_{\text{outside}} = 90$ mm

Table A.1: Summary of component specifications for the MAST CIS diagnostic. [1]Components were modified at Australian National University, i.e. were re-mounted in custm optical mounts or modified to accept the custom mounts. [2]All filters were from the Andover / LOT Oriel 'Semi-custom' filter range, and had $\varnothing$ = 50mm and were mounted in M52 threaded rings. Tolerances quoted on the filter parameters are the manufacturing tolerances of the ordered filters.

Bibliography

Andersen, A. and A. Kak
1984. Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrasonic imaging*, 6(1):81–94.

Asakura, N.
2007. Understanding the SOL flow in L-mode plasma on divertor tokamaks, and its influence on the plasma transport. *Journal of Nuclear Materials*, 363-365:41–51.

Asakura, N., H. Kawashima, N. Ohno, G. Matsunaga, T. Nakano, and N. Oyama
2008. ELM propagation in the low- and high-field-side scrape-off layer of the JT-60U tokamak. *Journal of Physics: Conference Series*, 123:012009.

Baillard, X., a. Gauguet, S. Bize, P. Lemonde, P. Laurent, a. Clairon, and P. Rosenbusch
2006. Interference-filter-stabilized external-cavity diode lasers. *Optics Communications*, 266(2):609–613.

Boswell, C. J., J. L. Terry, B. Lipschultz, and J. Stillerman
2001. Applications of visible CCD cameras on the Alcator C-Mod tokamak. *Review of Scientific Instruments*, 72(1):935.

Brown, D.
1966. Decentering distortion of lenses. *Photogrammetric Engineering*, 32(3):444–462.

Bussard, R.
1991. Some physics considerations of magnetic inertial-electrostatic confinement: a new concept for spherical converging-flow fusion. *Fusion Technology*, 19(2):273–293.

Censor, Y.

1983. Finite series-expansion reconstruction methods. *Proceedings of the IEEE*, 71(3):409–419.

Chung, J.
2004. *Time resolved 2D Doppler imaging of ion dynamics*. PhD thesis, University of Greifswald.

Chung, J., R. König, J. Howard, M. Otte, and T. Klinger
2005. Time resolved coherence-imaging spectrometer on WEGA stellarator. *Plasma Physics and Controlled Fusion*, 47(6):919–940.

Chung, K.-S.
2012. Mach probes. *Plasma Sources Science and Technology*, 21(6):063001.

Conrady, A. E.
1919. Decentred Lens-Systems. *Monthly Notices of the Royal Astronomical Society*, 79:384–390.

Cox, M. and MAST Team
1999. The mega amp spherical tokamak. *Fusion Engineering and Design*, 46:397–404.

Darke, A., R. Hayward, G. Counsell, and K. Hawkins
2005. The MAST improved divertor. *Fusion Engineering and Design*, 75-79:285–289.

Eimerl, D., L. Davis, S. Velsko, E. K. Graham, and a. Zalkin
1987. Optical, mechanical, and thermal properties of barium borate. *Journal of Applied Physics*, 62(5):1968.

Elder, J., P. Stangeby, D. Whyte, S. Allen, a. McLean, J. Boedo, B. Bray, N. Brooks, M. Fenstermacher, M. Groth, C. Lasnier, S. Lisgo, D. Rudakov, W. Wampler, J. Watkins, and W. West
2005. OEDGE modeling of 13C deposition in the inner divertor of DIII-D. *Journal of Nuclear Materials*, 337-339:79–83.

Evans, T., R. Moyer, P. Thomas, J. Watkins, T. Osborne, J. Boedo, E. Doyle, M. Fenstermacher, K. Finken, R. Groebner, M. Groth, J. Harris, R. La Haye, C. Lasnier, S. Masuzaki, N. Ohyabu, D. Pretty, T. Rhodes, H. Reimerdes, D. Rudakov, M. Schaffer, G. Wang, and L. Zeng
2004a. Suppression of Large Edge-Localized Modes in High-Confinement DIII-

D Plasmas with a Stochastic Magnetic Boundary. *Physical Review Letters*, 92(23):235003.

Evans, T. E., R. K. W. Roeder, J. a. Carter, and B. I. Rapoport
2004b. Homoclinic tangles, bifurcations and edge stochasticity in diverted tokamaks. *Contributions to Plasma Physics*, 44(13):235–240.

Flores, A. and L. Gradinarsky
2000. Sensing atmospheric structure: Tropospheric tomographic results of the smallscale GPS campaign at the Onsala Space Observatory. *Earth, planets and space*, 52:941–945.

Françon, M. and S. Mallick
1971. *Polarization Interferometers: Applications in Microscopy and Macroscopy*, Wiley series in pure and applied optics. Wiley Interscience.

Gabor, D.
1946. Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, 93(26):429–441.

Gafert, J., K. Behringer, D. Coster, C. Dorn, K. Hirsch, M. Niethammer, and U. Schumacher
1997. First experimental determination of ion flow velocities and temperatures in the ASDEX Upgrade divertor. *Plasma Physics and Controlled Fusion*, 39:1981–1995.

Gangadhara, S. and B. LaBombard
2004. Impurity plume experiments in the edge plasma of the Alcator C-Mod tokamak. *Plasma Physics and Controlled Fusion*, 46(10):1617–1646.

Goodman, J.
2000. *Statistical Optics*. Wiley Interscience.

Ham, C.
2012. *External Internship: a feasibility study for Coherent [sic] Imaging*. Unpublished report, Eindhoven University of Technology.

Harlander, J., C. Englert, D. Babcock, and F. Roesler
2010. Design and laboratory tests of a Doppler Asymmetric Spatial Heterodyne (DASH) interferometer for upper atmospheric wind and temperature observations. *Optics Express*, 18(25):7297–7307.

Harrison, J.
2010. *Characterisation of Detached Plasmas on the MAST Tokamak*. Ph.d thesis, University of York.

Harrison, J., A. Kirk, I. Chapman, P. Cahyna, Y. Liu, E. Nardon, and A. Thornton
2014. Characteristics of X-point lobe structures in single-null discharges on MAST. *Nuclear Fusion*, 54(6):064015.

Hecht, E.
2002. *Optics*, fourth edition. Addison-Wesley.

Heikkila, J. and O. Silvén
1997. A four-step camera calibration procedure with implicit image correction. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1997. Proceedings.*, Pp. 1106 – 1112.

Howard, J.
2002. Electro-optically modulated polarizing Fourier-transform spectrometer for plasma spectroscopy applications. *Applied optics*, 41(1):197–208.

Howard, J.
2006. High-speed high-resolution plasma spectroscopy using spatial-multiplex coherence imaging techniques (invited). *Review of Scientific Instruments*, 77(10):10F111.

Howard, J.
2010. Coherence imaging spectro-polarimetry for magnetic fusion diagnostics. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(14):144010.

Howard, J., A. Diallo, M. Creese, S. Allen, R. Ellis, W. Meyer, M. Fenstermacher, G. D. Porter, N. H. Brooks, M. E. Van Zeeland, and R. L. Boivin
2011. Coherence Imaging of Flows in the DIII-D Divertor. *Contributions to Plasma Physics*, 51(2-3):194–200.

Howard, J., A. Diallo, M. Creese, B. D. Blackwell, S. L. Allen, R. M. Ellis, G. D. Porter, W. Meyer, M. E. Fenstermacher, N. H. Brooks, M. E. Van Zeeland, and R. L. Boivin
2010a. Doppler coherence imaging and tomography of flows in tokamak plasmas (invited). *Review of scientific instruments*, 81(10):10E528.

Howard, J., R. Jaspers, O. Lischtschenko, E. Delabie, and J. Chung
2010b. Imaging charge exchange recombination spectroscopy on the TEXTOR tokamak. *Plasma Physics and Controlled Fusion*, 52(12):125002.

Howard, J., C. Michael, F. Glass, and a. D. Cheetham
2001. Optical coherence techniques for plasma spectroscopy (invited). *Review of Scientific Instruments*, 72(1):888.

Hutchinson, I.
2005. *Principles of Plasma Diagnostics*, second edition. Cambridge University Press.

IAEA
2013. *Energy, Electricity and Nuclear Power Estimates for the Period Up to 2050*, number 1 in Reference Data Series. Vienna: IAEA.

Isler, R. C., N. H. Brooks, W. P. West, a. W. Leonard, G. R. McKee, and G. D. Porter
1999. Spectroscopic analysis of normal and reversed ion flows in the DIII-D divertor. *Physics of Plasmas*, 6(2):541.

Jenkins, F. A. and H. E. White
1981. *Fundamentals of Optics*, fourth edition. McGraw-Hill International.

Karmida, A., Y. Ralchenko, and NIST ASD Team
2013. NIST Atomic Spectra Database (ver. 5.1), [Online]. Available: http://physics.nist.gov/asd.

Kato, K.
1986. Second-harmonic generation to 2048 Åin -Ba2O4. *IEEE Journal of QUantum Electronics*, 22(7):1013–1014.

Katramados, I., G. Fishpool, M. Fursdon, G. Whitfield, V. Thompson, and H. Meyer
2011. MAST upgrade closed pumped divertor design and analysis. *Fusion Engineering and Design*, 86(9-11):1595–1598.

Keilhacker, M., A. Gibson, C. Gormezano, P. J. Lomas, P. R. Thomas, M. L. Watkins, P. Andrew, B. Balet, D. Borba, C. D. Challis, I. Coffey, G. A. Cottrell, H. P. L. D. Esch, N. Deliyanakis, A. Fasoli, C. W. Gowers, H. Y. Guo, G. T. A. Huysmans, T. T. C. Jones, W. Kerner, S. E. Sharapov, G. Sips, P. Smeulders,

B. J. D. Tubbing, M. G. V. Hellermann, D. J. Ward, and J. E. T. Team
1999. High fusion performance from deuterium-tritium plasmas in JET. *Nuclear Fusion*, 39(2):209–234.

Kirk, A., J. Harrison, Y. Liu, E. Nardon, I. T. Chapman, and P. Denner
2012. Observation of Lobes near the X Point in Resonant Magnetic Perturbation Experiments on MAST. *Physical Review Letters*, 108(25):255003.

Kirk, A., Y. Liu, E. Nardon, P. Tamain, P. Cahyna, I. Chapman, P. Denner, H. Meyer, S. Mordijck, and D. Temple
2011. Magnetic perturbation experiments on MAST L- and H-mode plasmas using internal coils. *Plasma Physics and Controlled Fusion*, 53(6):065011.

Konjević, N., A. Lesage, J. Fuhr, and W. Wiese
2002. Experimental Stark widths and shifts for spectral lines of neutral and ionized atoms. *Journal of physical and chemical reference data*, 31(3):819–927.

Koubiti, M., T. Nakano, L. Godbert-Mouret, Y. Marandet, J. Rosato, and R. Stamm
2011. Diagnostics of JT-60U divertor plasmas by StarkDoppler broadening of carbon spectral lines. *Journal of Nuclear Materials*, 415(1):S1151–S1154.

Laberge, M.
2007. An Acoustically Driven Magnetized Target Fusion Reactor. *Journal of Fusion Energy*, 27(1-2):65–68.

LaBombard, B., J. Rice, a.E Hubbard, J. Hughes, M. Greenwald, J. Irby, Y. Lin, B. Lipschultz, E. Marmar, C. Pitcher, N. Smick, S. Wolfe, S. Wukitch, and t. A. Group
2004. Transport-driven Scrape-Off-Layer flows and the boundary conditions imposed at the magnetic separatrix in a tokamak plasma. *Nuclear Fusion*, 44(10):1047–1066.

Lang, P., a. Loarte, G. Saibene, L. Baylor, M. Becoulet, M. Cavinato, S. Clement-Lorenzo, E. Daly, T. Evans, M. Fenstermacher, Y. Gribov, L. Horton, C. Lowry, Y. Martin, O. Neubauer, N. Oyama, M. Schaffer, D. Stork, W. Suttrop, P. Thomas, M. Tran, H. Wilson, a. Kavin, and O. Schmitz
2013. ELM control strategies and tools: status and potential for ITER. *Nuclear Fusion*, 53(4):043004.

Lawrence Marple Jr., S.
1999. Computing the discrete-time analytic signal via FFT. *IEEE Transactions on Signal Processing*, 47(9):2600.

Leggate, H., S. Lisgo, and J. Harrison
2013. Impurity transport studies using fast imaging of injected carbon on the MAST tokamak. In *41st EPS Conference on Plasma Physics, European Conference Abstracts*, P. P5.036. European Physical Society.

Lewis, E. N., P. J. Treado, R. C. Reeder, G. M. Story, a. E. Dowrey, C. Marcott, and I. W. Levin
1995. Fourier transform spectroscopic imaging using an infrared focal-plane array detector. *Analytical chemistry*, 67(19):3377–81.

Liang, Y., H. Koslowski, P. Thomas, E. Nardon, B. Alper, P. Andrew, Y. Andrew, G. Arnoux, Y. Baranov, M. Bécoulet, M. Beurskens, T. Biewer, M. Bigi, K. Crombe, E. De La Luna, P. de Vries, W. Fundamenski, S. Gerasimov, C. Giroud, M. Gryaznevich, N. Hawkes, S. Hotchin, D. Howell, S. Jachmich, V. Kiptily, L. Moreira, V. Parail, S. Pinches, E. Rachlew, and O. Zimmermann
2007. Active Control of Type-I Edge-Localized Modes with n=1 Perturbation Fields in the JET Tokamak. *Physical Review Letters*, 98(26):265004.

Lischtschenko, O., K. Bystrov, G. De Temmerman, J. Howard, R. J. E. Jaspers, and R. König
2010. Density measurements using coherence imaging spectroscopy based on Stark broadening. *Review of scientific instruments*, 81(10):10E521.

Lisgo, S., P. Börner, G. Counsell, J. Dowling, a. Kirk, R. Scannell, M. OMullane, and D. Reiter
2009. Interpretation of spatially resolved helium line ratios on MAST. *Journal of Nuclear Materials*, 390-391:1078–1080.

Lisgo, S., P. Stangeby, J. Elder, J. Boedo, B. Bray, N. Brooks, M. Fenstermacher, M. Groth, D. Reiter, D. Rudakov, J. Watkins, W. West, and D. Whyte
2005. Re-construction of detached divertor plasma conditions in DIII-D using spectroscopic and probe data. *Journal of Nuclear Materials*, 337-339:256–260.

Lunt, T., Y. Feng, M. Bernert, a. Herrmann, P. de Marné, R. McDermott, H. Müller, S. Potzel, T. Pütterich, S. Rathgeber, W. Suttrop, E. Viezzer, E. Wolfrum, and M. Willensdorfer
2012. First EMC3-Eirene simulations of the impact of the edge magnetic perturbations at ASDEX Upgrade compared with the experiment. *Nuclear Fusion*, 52(5):054013.

Marr, K., B. Lipschultz, B. LaBombard, and J. Terry
2005. Spectroscopic measurements of plasma flow in the SOL in C-Mod. *Journal of Nuclear Materials*, 337-339:286–290.

McLean, A., J. Elder, P. Stangeby, S. Allen, J. Boedo, N. Brooks, M. Fenstermacher, M. Groth, S. Lisgo, A. Nagy, D. Rudakov, W. Wampler, J. Watkins, W. West, and D. Whyte
2005. DIVIMP modeling of the toroidally symmetrical injection of 13CH4 into the upper SOL of DIII-D. *Journal of Nuclear Materials*, 337-339:124–128.

Meade, D.
2010. 50 Years of Fusion Research. *Nuclear Fusion*, 50(1):014004.

Michael, C. A., J. Howard, and B. D. Blackwell
2001. The MOSS camera on H-1NF. *Review of Scientific Instruments*, 72(1):1034.

Mirnov, S.
2009. Plasma-wall interactions and plasma behaviour in fusion devices with liquid lithium plasma facing components. *Journal of Nuclear Materials*, 390-391:876–885.

Morgan, T.
2011. *Measurement of MAST edge ion temperatures and velocities*. Phd., University of York.

Odstrčil, M., J. Mlyná, V. Weinzettl, P. Háček, T. Odstrčil, G. Verdoolaege, M. Berta, T. Szabolics, and A. Bencze
2014. Plasma tomographic reconstruction from tangentially viewing camera with background subtraction. *Review of scientific instruments*, 85(1):013509.

Okamoto, T., S. Kawata, and S. Minami
1984. Fourier transform spectrometer with a self-scanning photodiode array. *Applied optics*, 23(2):269.

Padgett, M. J. and a. R. Harvey
1995. A static Fourier-transform spectrometer based on Wollaston prisms. *Review of Scientific Instruments*, 66(4):2807.

Paige, C. C. and M. a. Saunders
1982. LSQR : An Algorithm for Sparse Linear Equations and Sparse Least Squares. *ACM Transactions on Mathematical Software*, 8(1):43–71.

Patella, D. and S. Patella
2009. Geophysical Tomography in Engineering Geological Applications: A Mini-Review with Examples. *Open Geology Journal*, 3:30–38.

Peng, Y. and D. Strickler
1986. Features of spherical torus plasmas. *Nuclear Fusion*, 26(6):769–777.

Persky, M. J.
1995. A review of spaceborne infrared Fourier transform spectrometers for remote sensing. *Review of Scientific Instruments*, 66(10):4763.

Piras, F., S. Coda, I. Furno, J.-M. Moret, R. A. Pitts, O. Sauter, B. Tal, G. Turri, A. Bencze, B. P. Duval, F. Felici, A. Pochelon, and C. Zucca
2009. Snowflake divertor plasmas on TCV. *Plasma Physics and Controlled Fusion*, 51(5):055009.

Romanelli, F., P. Barabaschi, D. Borba, G. Federici, L. Horton, R. Neu, D. Stork, and H. Zohm
2012. *Fusion Electricity: A roadmap to the realisation of fusion energy*. European Fusion Development Agreement, EFDA.

Shafer, M., E. Unterberg, D. Orlov, D. Battaglia, T. Evans, J. Harris, D. Hillis, R. Maingi, R. Moyer, R. Nazikian, and a. Wingen
2012. Experimental imaging of separatrix splitting on DIII-D. *Nuclear Fusion*, 52(12):122001.

Shimada, M., D. Campbell, V. Mukhovatov, M. Fujiwara, N. Kirneva, K. Lackner, M. Nagami, V. Pustovitov, N. Uckan, J. Wesley, N. Asakura, A. Costley, A. Donné, E. Doyle, A. Fasoli, C. Gormezano, Y. Gribov, O. Gruber, T. Hender, W. Houlberg, S. Ide, Y. Kamada, A. Leonard, B. Lipschultz, A. Loarte, K. Miyamoto, T. Osborne, A. Polevoi, and A. Sips
2007. Progress in the ITER physics basis. *Nuclear Fusion*, 47(6):S1–S17.

Smith, B., B. Laubscher, B. Cooke, P. LaDelfe, J. Harlander, J. Howard, and S. Milligan
1999. IRISHS: the infrared imaging spatial heterodyne spectrometer: a new pushbroom Fourier transform ultraspectral imager with no moving parts. In *Proc. SPIE 3698, Infrared Technology and Applications XXV*, volume 3698, Pp. 501–509.

Sobelman, I. I.
1979. *Atomic Spectra and Radiative Transitions*. Springer-Verlag.

Soukhanovskii, V. a., R. E. Bell, a. Diallo, S. Gerhardt, S. Kaye, E. Kolemen, B. P. LeBlanc, a. G. McLean, J. E. Menard, S. F. Paul, M. Podesta, R. Raman, T. D. Rognlien, a. L. Roquemore, D. D. Ryutov, F. Scotti, M. V. Umansky, D. Battaglia, M. G. Bell, D. a. Gates, R. Kaita, R. Maingi, D. Mueller, and S. a. Sabbagh
2012. Snowflake divertor configuration studies in National Spherical Torus Experiment. *Physics of Plasmas*, 19(8):082504.

Soukhanovskii, V. a., D. W. Johnson, R. Kaita, and a. L. Roquemore
2006. Electron density measurements in the National Spherical Torus Experiment detached divertor region using Stark broadening of deuterium infrared Paschen emission lines. *Review of Scientific Instruments*, 77(10):10F127.

Stammers, K. and M. Loughlin
2006. The calibration of the MAST neutron yield monitors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 562(1):521–530.

Steel, W.
1983. *Interferometry*, second edition. Cambridge University Press.

Summers, H. P.
2004. The ADAS User Manual, version 2.6.

Suttrop, W., T. Eich, J. C. Fuchs, S. Günter, a. Janzer, a. Herrmann, a. Kallenbach, P. T. Lang, T. Lunt, M. Maraschek, R. M. McDermott, a. Mlynek, T. Pütterich, M. Rott, T. Vierle, E. Wolfrum, Q. Yu, I. Zammuto, and H. Zohm
2011. First Observation of Edge Localized Modes Mitigation with Resonant and Nonresonant Magnetic Perturbations in ASDEX Upgrade. *Physical Review Letters*, 106(22):225004.

Sykes, A., R. Akers, and L. Appel
1997. High-performance of the START spherical tokamak. *Plasma Physics and Controlled Fusion*, 39:247–260.

Tal, B., B. Labit, D. Nagy, R. Chavan, B. Duval, and G. Veres
2013. Plasma radiation dynamics with the upgraded Absolute Extreme Ultraviolet tomographical system in the Tokamak à Configuration Variable. *Review of scientific instruments*, 84(12):123508.

Terry, P.
2000. Suppression of turbulence and transport by sheared flow. *Reviews of Modern Physics*, 72(1).

Timberlake, J.
1983. Photography of impurity injection into PLT plasmas. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 1(2):841.

Tucker, A. J.
1998. Computerized Ionospheric Tomography. *Johns Hopkins APL Technical Digest*, 19(1):66–71.

U.S. Energy Administration
2014. International Energy Statistics.

Valanju, P. M., M. Kotschenreuther, S. M. Mahajan, and J. Canik
2009. Super-X divertors and high power density fusion devices. *Physics of Plasmas*, 16(5):056110.

Veiras, F. E., L. I. Perez, and M. T. Garea
2010. Phase shift formulas in uniaxial media: an application to waveplates. *Applied optics*, 49(15):2769–77.

Wagner, F., G. Becker, and K. Behringer
1982. Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak. *Physical Review Letters*, 49(19):1408–1412.

Weber, T. R., S. L. Allen, and J. Howard
2012. C-III flow measurements with a coherence imaging spectrometer. *Review of Scientific Instruments*, 83(10):10E102.

Willis, M.
2000. *Algebraic Reconstruction Algorithms for Remote Sensing Image Enhancement.* Msc, Brigham Young University.

Wu, L., C. Zhang, and B. Zhao
2007. Analysis of the lateral displacement and optical path difference in wide-field-of-view polarization interference imaging spectrometer. *Optics Communications*, 273(1):67–73.

Zaniol, B., R. C. Isler, N. H. Brooks, W. P. West, and R. E. Olson
2001. Measurements of C V flows from thermal charge-exchange excitation in divertor plasmas. *Physics of Plasmas*, 8(10):4386.

Zhang, C., B. Xiangli, and B. Zhao

2002. A static polarization imaging spectrometer based on a Savart polariscope. *Optics communications*, 203(1-2):21–26.

Zhang, Z.

1999. Flexible camera calibration by viewing a plane from unknown orientations. *The Proceedings of the Seventh IEEE International Conference on Computer Vision, 1999.*, 1:666 – 673.