

Durham E-Theses

Fast-ion deuterium-alpha measurements on MAST and MAST-U

ANDREW ROBERT JACKSON

How to cite:

JACKSON, ANDREW ROBERT (2023) Fast-ion deuterium-alpha measurements on MAST and MAST-U. Doctoral thesis, Durham University.

Use policy

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

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/14866/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way


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

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

**Fast-Ion Deuterium- α
Measurements on MAST and
MAST-U**

Andrew Robert Jackson

A thesis presented for the degree of
Doctor of Philosophy



Centre for Advanced Instrumentation
The University of Durham
United Kingdom
27th February 2023

Fast-Ion Deuterium- α Measurements on MAST and MAST-U

Andrew Robert Jackson

Abstract

Analysis and modelling of fast-ion deuterium- α (FIDA) spectroscopy data from both MAST and MAST-U was performed and is described, along with description of the setup of the diagnostic for the first experimental campaign on MAST-U. Analysis of archival MAST data featuring sawtooth crashes show that, while insufficient to distinguish between 3 different ways of modelling the effect of sawtooth crashes on the fast ions, can observe significant redistribution of the fast ions due to the crashes. FIDA data was used to add to a bank of evidence showing problems with the configuration of the neutral beam injection system, which was later confirmed. Analysis of the effect of a large fishbone instability on plasma rotation shows that the large deceleration of the plasma is unlikely to be due to forces induced by a radial inflow of bulk ions, itself induced by a radial outflow of fast ions. An attempt is made to utilise anomalous diffusion in order to examine the effect of Toroidal Alfvén Eigenmode (TAE)s and fishbones on the fast ions. Investigation into the effect of locked modes and long-lived mode (LLM)s show severe redistribution due to the locked mode, with correlations made between the dynamics of the locked mode and the FIDA emission. Investigation of the LLM show significant reduction of FIDA emission in core channels, suggesting localised redistribution. Future integration of multiple fast-ion diagnostics including the Solid State Neutral Particle Analyser will allow for a deeper understanding of fast ion/instability interaction.

Supervisors: Ray Sharples and Ken McClements

Acknowledgements

It has been a difficult road to the completion of the PhD, which wasn't helped by the global pandemic that happened in latter half of it. I would not have reached the end without the help and support from a large number of people. In particular, I'd like to thank Mum and Dad, who gave me the support to be able to get to the end. I'd like to thank Ken McClements, Clive Michaels and Ray Sharples for their supervision throughout my time at CCFE. I'd like to thank Asger Jacobsen who, during his time at CCFE, gave me the support and encouragement I needed at the time to get through my issues and made me feel like, "Yeah, maybe I can do this". I'd like to thank Yoshika Terada and Fabio Federici, who made their house a home for me (and without them I would have likely died from starvation and exposure). I'd like to thank Sam Gibson, who gave me support and helped me to view the PhD from a better perspective. And I'd like to thank Steven Thomas, David Rankin, Iain Cross and Scott Walkinshaw, who listened to me moan about the PhD throughout.

Contents

Declaration	vii
List of Figures	viii
List of Tables	xxiii
1 Introduction and Background	1
1.1 The motivation for fusion research	1
1.2 The physics of fusion energy	2
1.2.1 Fusion fuels	2
1.2.2 Self-sustaining reactions	5
1.3 Magnetic confinement fusion	6
1.3.1 Particle dynamics in a tokamak	6
1.3.2 Passing and trapped particles	9
1.3.3 Conventional vs. spherical tokamaks	10
1.3.4 MAST and MAST-U	11
1.4 Heating	12
1.4.1 Neutral beam injection	12
1.5 Fast Ions	13
1.6 Instabilities that affect fast ions	15
1.6.1 Inverse Landau damping	16

1.6.2	Toroidal Alfvén Eigenmodes	16
1.6.3	Kink instability, fishbones and sawteeth	17
1.7	Tokamak diagnostics	20
2	The principles of the FIDA diagnostic and its simulation	24
2.1	Physics of the FIDA diagnostic	24
2.1.1	The FIDA mechanism	24
2.1.2	Typical components of the observed FIDA diagnostic spectra	27
2.2	Forward modelling	29
2.2.1	TRANSP/NUBEAM	31
2.2.1.1	MAST-U TRANSP/NUBEAM limitations	33
2.2.2	FIDASIM	35
2.3	Estimating error in FIDASIM synthetic spectra using deterministic sampling	37
2.4	Aside: Solid State Neutral Particle Analyser	43
3	The FIDA Diagnostic: setup and preparation for the MAST-U experimental campaign	45
3.1	FIDA setup on MAST-U	45
3.1.1	Optical fibre setup	46
3.1.1.1	Tokamak fibres and viewing geometry	46
3.1.1.2	Patch panel	49
3.1.1.3	Spectrometer fibres	51
3.1.2	Spectrometer setup	52
3.1.3	Mask	54
3.1.4	Demagnification lenses	56
3.1.5	Camera	56
3.2	FIDA data calibration	57
3.3	Flux and Wavelength calibration	59
3.4	Determination of optical fibre transmission	60

3.4.1	Spectrometer fibres	60
3.4.1.1	A caveat	60
3.4.1.2	Calibration of spectrometer fibres	61
3.4.2	Calibration of tokamak fibres	62
3.5	Validation of fibre alignment	64
3.5.1	Back-lighting of fibres	64
3.5.2	Finding the coordinates of the illuminated spots using Calcam	66
3.5.3	Projection of light spots onto fibre head	69
3.5.4	Comparing tangency radii	73
3.6	Potential vertical FIDA on MAST-U	75
3.7	Conclusion	77
4	Analysis of sawtooth crashes during MAST operation	78
4.1	Sawtooth crash	78
4.2	Using FIDA to diagnose sawteeth	79
4.3	Sawtooth models	80
4.4	Analysis	83
4.4.1	Comparison of different sawtooth models with FIDA data	84
4.4.2	Detailed analysis and comparison with FIDA data	87
4.5	Passive FIDA	93
4.6	Conclusions	95
5	Analysis of beam fractions during first experimental campaign on MAST-U	97
5.1	SS beam fraction investigation	97
5.1.1	Diagnosis of unknown beam peak component	98
5.1.2	Derivation of beam fractions	101
5.1.2.1	Shot #44254	103
5.1.2.2	Shot #45483	104
5.2	Conclusion of beam fraction investigation	105

6	Analysis of MAST-U data	109
6.1	Investigation of effect of fishbone on plasma rotation and fast ions	109
6.1.1	Effect of fishbone instability upon plasma rotation	113
6.1.1.1	Using the derived model	118
6.1.1.2	Loss of momentum drive with loss of fast ions	121
6.1.2	Analysis with TRANSP and FIDASIM	122
6.1.2.1	Using Mirnov coil data to estimate anomalous dif- fusion	122
6.1.2.2	TRANSP and FIDASIM results	126
6.1.2.3	Conclusion	128
6.2	Effect of long-lived mode and locked mode on fast ions	129
6.2.1	TRANSP/FIDASIM analysis of LLM	134
6.3	Conclusion	140
7	Summary, conclusions and future work	142
7.1	The principles of the FIDA diagnostic and its simulation	142
7.2	The FIDA Diagnostic: setup and preparation for the MAST-U ex- perimental campaign	144
7.3	Analysis of sawtooth crashes during MAST operation	146
7.4	Analysis of beam fractions during first experimental campaign on MAST-U	148
7.5	Analysis of MAST-U data	149
7.6	Future of the FIDA diagnostic	151
	Bibliography	153

Declaration

The work in this thesis is based on research carried out at the Culham Centre for Fusion Energy, UK Atomic Energy Authority, Culham Science Centre, Abingdon, England. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is the sole work of the author unless referenced to the contrary in the text.

Some of the work presented in this thesis has been published in journals and conference proceedings - the relevant publications are listed below.

Publications

A. R. Jackson, A. S. Jacobsen, K. G. McClements *et al.*, “Diagnosing fast ion redistribution due to sawtooth instabilities using fast ion deuterium- α spectroscopy in the Mega Amp Spherical Tokamak”, *Nuclear Fusion*, vol. 60, p. 126 035, 2020, ISSN: 0029-5515. DOI: 10.1088/1741-4326/abb619

A. R. Jackson, A. S. Jacobsen, K. G. McClements *et al.*, “Diagnosing fast ion redistribution due to sawtooth instabilities using fast ion deuterium- α spectroscopy in the Mega Amp Spherical Tokamak”, in *47th EPS Conference on Plasma Physics, EPS 2021*, vol. 45A, Sitges, Spain, 2021, pp. 689–692, ISBN: 9781713837046

Copyright © 2022 by Andrew Robert Jackson.

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

List of Figures

1.1	Binding energy per nucleon vs. mass number. Any reaction that moves along the curve to a higher binding energy per nucleon releases energy. Note the unusually high binding energy per nucleon for ${}^4_2\text{He}$ compared to the nearby isotopes, which is advantageous for fusion energy. Note also the change from logarithmic to linear scales in the x-axis. Adapted from [4].	3
1.2	Fusion cross section vs. centre of mass energy for various potential fusion reactions. Adapted from [4].	4
1.3	Diagram of the coordinate systems used to describe physics in the torus with a simple circular cross section. Adapted from [6]	8
1.4	Comparison of the profile of an example conventional and spherical tokamak. Adapted from [14]	11
1.5	Block diagram of the generation of high energy neutrals by the neutral injection system.	13

1.6 An example of the kink instability in a column of plasma, which carries a current density \mathbf{J} , that produces an azimuthal field \mathbf{B}_θ . The twisting produces regions of stronger and weaker \mathbf{B}_θ . This results in a net $\mathbf{J} \times \mathbf{B}$ force in the direction of the plasma displacement, causing further displacement, and so the column is unstable. The magnetic tension of an applied \mathbf{B}_ϕ along the length of the column can provide a degree of stabilisation. Adapted from [6]. 18

1.7 Characteristic fishbone trace in Mirnov coil signal for fishbone mode appearing in shot #45026, extant between 0.346 s-0.35 s. The Mirnov coil signal is proportional to the rate of change of the azimuthal component of the magnetic field. 19

2.1 Weight functions (with units of $(\text{Ph}\cdot\text{cm})/(\text{s}\cdot\text{nm}\cdot\text{sr}\cdot\text{fast-ion}\cdot\text{dE}\cdot\text{dP})$) for the range of wavelengths accessible by the FIDA diagnostic, for shot #45091 at $R = 0.96$ m (a) and $R = 1.413$ m (b). (c) shows an example of a weight function for a potential vertical view (with line-of-sight origin below the beam), viewing $R = 1.413$ m. For this view the FIDA emission is primarily blue-shifted, giving an upper, rather than lower, limit to the accessible wavelengths. 27

2.2 A typical FIDA spectrum after background subtraction on MAST/MAST-U, as generated by FIDASIM (described further in Section 2.2.2), and convolved with the instrument function of the diagnostic to account for instrumental broadening. Once background subtraction has been performed, only the full, half, and third energy beam peaks, the halo component and the FIDA component should remain. The small secondary peaks seen on each of the beam peaks are an artefact of the convolution with the instrument function 29

2.3	<p>Synthetic neutron rates predicted by TRANSP for MAST pulse #29904. In (a), the predicted neutron rates are shown for each sigma-point, with the $q^{(i)}$ corresponding to those given in Equation 2.9. In (b) the resultant expectation value (given by Equation 2.7) and standard deviation (given by the square root of the variance given by Equation 2.8) is shown, along with a comparison with $q^{(5)}$, the simulation using the unmodified, "default" values of T_e and n_e.</p>	41
2.4	<p>Synthetic spectra predicted with FIDASIM, using the results of the five TRANSP runs. In each sub-figure, only the region of the spectrum which was dominated by FIDA emission is shown. (a) shows the synthetic spectra derived from each sigma-point $q^{(i)}$, corresponding to those given in Equation 2.9. (b) shows the calculated expectation value (given by Equation 2.7) and standard deviation (given by the square root of the expression in Equation 2.8) is shown, along with a comparison with $q^{(5)}$, the simulation using the unmodified, "default" values of T_e and n_e.</p>	41
2.5	<p>Percentage uncertainty for synthetic spectra as determined by deterministic sampling, for pulse #29904 averaged over 0.265 s-0.268 s. Each subplot shows a different channel, with (a) $R = 0.936$ m, (b) $R = 0.982$ m, (c) $R = 1.027$ m, (d) $R = 1.072$ m, (e) $R = 1.139$ m, (f) $R = 1.205$ m and, (g) $R = 1.269$ m. The plots focus on the region of the spectra dominated by FIDA emission, with the greyed out region showing where full energy beam peak emission dominates instead. The green region designates the region where the absolute value of the FIDA emission was above 1×10^{15} Ph s⁻¹ nm⁻¹ m⁻² sr⁻¹, the approximate value of the uncertainty on the experimental FIDA data at this time.</p>	42

3.1	Schematic diagram of the FIDA setup, which is described in detail in this section. Light from the tokamak enters fibres in the fibre head, which travels to the patch panel. Light travels through the spectrometer fibres chosen for examination into the spectrometer. The diffracted light passes through the mask, blocking the full energy beam peak, before being imaged onto the camera CCD chip via the de-magnification lenses. After the raw data is recorded it can be processed for analysis.	46
3.2	Top-down view of the spectral and beam geometry present on MAST-Upgrade (MAST-U) during the first experimental campaign. The extent of the passive views is shown with the blue cross hatch, the extent of the active views is bounded by the red dashed lines, with an example active channel shown in black. This intersects the SS NBI beam (which lies in the midplane and is directed towards increasing Y), with the SS beam centreline shown in dark green, and with the spatial extent of the beam shown in light green. Note that this spatial extent is an approximation, based on the horizontal width of the beam aperture (34 cm). The beam has a focal length and non-zero divergence, and will become wider towards increasing Y. Also shown is the SW beam centreline (magenta), but the SW beam is vertically displaced upwards by 0.65 m, ensuring no contamination of the inner channels. The SW beam of course also has a finite width, but is not shown here for figure clarity.	48
3.3	Close-up view of the ends of the FIDA optical fibres as mounted in the active and passive optical head assemblies, along with the Celeste-III and CXRS-U fibres that share the same head. Adapted from [73]. . . .	50

3.4	Patch panel as set up for the first MAST-U experimental campaign. Highlighted regions show the terminal fibre connectors for the tokamak fibres, with blue for Bundle 3, orange for Bundle 4, and green for Bundle 5. Fibre connectors not highlighted include those for Bundles 1 and 2, those for additionally calibration fibres for FIDA and a few for the BES diagnostic.	51
3.5	Schematic diagram of a Volume Phase Holographic (VPH) spectrometer. Adapted from [75].	52
3.6	Binned camera image of mask with white background. Note that the vertical gradient in brightness is just an artefact of the lab lighting utilised.	54
3.7	Example of mask fitting for a single channel, using white background test shot -15531, and plasma shot featuring beam peaks 44922. The two troughs present in the test shot are aligned with the full and half energy beam peaks. Magnitude of test shot has been scaled up to better show troughs.	55
3.8	Flowchart describing FIDA data retrieval process	58
3.9	Full frame CCD image of a subset of the light from FIDA fibres present in the active head. The labels at the top of the figure refer to which spectrometer fibre set is shown, and the numbers on the right to the specific spectrometer fibre number. The spectrometer fibre Set B are connected to Bundle 3, and Set C to Bundle 4. The tokamak fibres are arranged in "one on, two off", from the bottom working upwards are 3/4A1,3/4A4,... up to 3/4H3. It can be seen that tokamak fibre 4G1 (connected to spectrometer fibre C3), is notably dimmer than any other fibre. It is believed that that this tokamak fibre is broken.	63
3.10	An example of back-lighting of some FIDA and CXRS-U fibres, for the purpose of verifying fibre alignment. Here, 2 torches have been held to the optical fibre connectors for tokamak fibres 3C3 and 3C4 on the FIDA patch panel (which joins tokamak fibres to spectrometer fibres). .	66

3.11	Point pairs, specified by the user in Calcam, are used to calibrate a model that converts between a 3D CAD coordinate and a 2D image pixel coordinate. The pairs define the necessary translation and rotation to move from the CAD lab frame to the camera frame, and the parameters of the distortion correction and the effective focal length of the imaging system to project the 3D location to the projected image location. This can then be used to extract the 3D coordinate for an arbitrary pixel coordinate.	68
3.12	Unit vectors for lens frame in machine frame. Top row ((a), (b), (c)) for the active lens, bottom row ((d), (e), (f)) for the passive lens. (a) and (d) shows the projection of the unit vectors in the X-Y plane, (b) and (e) show the projection in the X-Z plane, and (c) and (f) show the projection in the Y-Z plane.	70
3.13	Top down diagram of the vectors discussed in this section, for (a) the active lens and (b) the passive lens. For each an example light spot was chosen to show $\overrightarrow{OS}_{\text{machine}}$ and $\overrightarrow{LS}_{\text{machine}}$. Note that the vectors all have a non-zero extent in the Z axis as well. The dashed black line represents the outer wall of the vacuum vessel. Both the active and passive lenses are present in ports, so are slightly recessed compared to the outer wall.	71
3.14	Projection of location of fibre locations on active and passive heads, compared to design locations.	72
3.15	Comparison of the tangency radii for the active and passive channels that could be observed.	74
3.16	Lines-of-sight (red) for a potential vertical view examining the same radii as the existing toroidal view, showing the X-Y plane (a) and the R-Z plane (b). The centrelines of the SS beam (green) and the SW beam (magenta) are shown, along with exemplar flux surfaces.	76

3.17	FIDASIM spectra generated based on Shot #45091 at $t = 0.35$ s for a potential vertical view that could be implemented for future campaigns. (a) shows all channels, (b) shows the individual spectral components for the specific channel at $R = 1.029$ m.	76
4.1	Sheet that defines the helical flux used by the Kadomtsev model to determine the effect of sawtooth crashes. Flux passes through this sheet in opposite direction on either side of the $q = 1$ surface. The opposite fluxes of equal magnitude are reconnected during the crash. Adapted from [5].	80
4.2	The progression in time of the newly formed magnetic island, due to the helical flux reconnection, displacing the original plasma core. In (i) the magnetic equilibrium is in a overall $q > 1$ state. In (ii), q drops below 1 in the core. In (iii) the magnetic surface on the right side of the core reconnects with the surfaces outside, connecting regions of equal and opposite helical flux. In (iv) and (v), the process continues, with the hot core being expelled, and the new cooler core forming. In (vi), the process is complete, with a new, cooler core formed, and the plasma in an overall $q > 1$ state. Adapted from [5].	81
4.3	The simulated fast-ion density in an azimuthal cross section of MAST before and after a sawtooth crash, with the flux surfaces shown in black.	82
4.4	Time traces of plasma parameters for the three pulses used in averaging. (a) NBI power, (b) neutron rate, (c) core electron temperature, (d) core electron density.	84
4.5	Soft X-ray (SXR) trace (for a line-of-sight with tangency radius 0.5285 m) showing a sawtooth crash for shots #29880 (blue), #29881 (orange) and #29882 (green).	85

4.6	Experimental and synthetic spectra showing only the spectral region where FIDA light can be observed at $R=1.205$ m in the plasma midplane. The predictions based on the three models are shown. Based on Section 2.3 the FIDASIM traces have an approximate error of 20 % . . .	85
4.7	Experimental and synthetic spectra showing only the region where FIDA light can be observed at $R=1.139$ m in the plasma midplane. The contribution to the spectrum from the full energy beam peak can be seen at the lower wavelengths, but the discrepancy between the experimental and synthetic spectra persists to higher wavelengths. Based on Section 2.3 the FIDASIM traces have an approximate error of 20 %	86
4.8	Spectral radiance across a number of FIDA channels. Note that the full energy beam peak appears at different locations depending on channel. The grey shaded area in plots (a)-(e) designates where the effect of the peak has a significant effect on the synthetic data, so direct comparisons between the observations and the models cannot be made here. Based on Section 2.3 the relative change in the FIDASIM traces have approximate absolute error of 30 percentage points (note that this is not a relative error, but is an absolute uncertainty in the calculated relative change).	88
4.9	Simulated fast ion distributions for $R = 1.205$ m in the device midplane using the Kadomtsev model. (a) Before crash, (b) after crash. The green contours represent the weight function for the FIDA diagnostic in the fast-ion dominated wavelength range and the spatial location of measurement.	91

4.10	The relative change in the fast-ion distribution function for $R = 1.205\text{m}$ in the device midplane. (a) Kadomtsev, (b) ergodic, (c) Porcelli. Areas in green represent phase-space regions with a fast-ion presence after the crash that had no presence before. Note that in some regions there is a positive change of greater than 100%. The black contours represent the weight function for the FIDA diagnostic in the fast-ion dominated wavelength range and the spatial location of measurement.	92
4.11	Relative change in the spectral radiance at $R = 1.205\text{m}$ in the device midplane, averaged over the indicated time windows. Note that Figure 4.11(a) examines the same data as Figure 4.8(e). Again, based on Section 2.3 the relative change in the FIDASIM traces have approximate absolute error of 30 percentage points (note that this is not a relative error, but is an absolute uncertainty in the calculated relative change).	94
4.12	(a) Available toroidally-viewing passive chords, with selected chord highlighted in red, and the NBI beamline to emphasise the passive nature of the view. Data from this chord is shown in Figure 4.12b. The black circles show the minimum and maximum extent of the plasma. (b) Spike and decay of passive signal seen during a sawtooth crash. The decaying part of the signal has been fitted with an exponential.	96
5.1	Spectra observed by the FIDA active view during the beam-into-gas phase in pulse #44254 at $R = 1.01\text{m}$ averaged over times 0.3705 s-0.3715 s. 4 peaks are observed in the redshifted D- α portion of the spectrum.	100
5.2	(a) Power and (b) flux fractions for Shot 44254, calculated across the FIDA channels	103
5.3	(a) Power and (b) flux fractions for Shot 45483, calculated across the FIDA channels.	105
5.4	Possible configurations of magnet polarities for MAST-U NBI ion source. Adapted from [103].	107

6.1	Time traces of plasma parameters of Shot #45026. (a) SS beam power, (b) neutron rate (note that the absolute calibration of the fission chamber is incorrect), (c) line integrated electron density, (d) D- α intensity.	110
6.2	(a) Time trace of the low-frequency Mirnov coil signal and (b) Spectrogram of low-frequency Mirnov coil time trace from shot #45026, showing the main resonant instabilities of the plasma. Specifically, TAEs dominate the early beam on portion of the shot, before transitioning to a series of fishbones. A number of these occur at the same time as significant drops in the neutron rate, with a emphasis on the fishbone that occurs at ~ 0.319 s (termed the "major" fishbone), at which time there is a very large drop in the neutron rate <i>and</i> plasma rotation velocity (discussed in section 6.1.1). At ~ 0.356 s a fishbone transitions into an LLM, which terminates with a sawtooth crash at 0.3883 s (seen as a feature across all shown frequencies, very narrow in time). Another sawtooth crash occurs at 0.409 s, before the beam terminates at 0.4287 s.	111
6.3	Wavelength integrated time traces of the FIDA emission in shot #45026. The FIDA emissivities are integrated between 660.52 nm-661.99 nm. Radii are (a) 1.029 m, (b) 1.097 m, (c) 1.164 m, (d) 1.23 m, (e) 1.293 m, and (f) 1.354 m. The times of the SS beam switch-on and switch-off are again indicated by red dashed and dashed-dotted lines.	112
6.4	Bulk plasma toroidal rotation derived from the CXRS diagnostic for shot #45026. The green trace shows the location of the magnetic axis over time, as defined by the archived EFIT run for the shot.	113

6.5	The (i) spectra, and (ii) global neutron rate and rotation velocity at $R = 1.04$ m for shots (a) 45026, (b) 45031, and (c) 45163, focussed specifically on the time of the analysed fishbone in each shot. In (a) and (c), there is a drop in the neutron rate during the fishbone, followed shortly after by a drop in the rotation rate. In (b), the effect appears to be smaller. There is a small drop in the rotation rate, and a flattening of the previously growing neutron rate.	115
6.6	Wavelength integrated (660.77 nm-661.52 nm) FIDA radial profiles before (blue) and after (red) the chirping phase of the selected fishbone in the shots (a) 45026, (b) 45031 and (c) 45163.	116
6.7	Rotation velocity profile at start and end of the fishbone (to within the time resolution of the diagnostic). Hashed region is ignored when averaging profiles.	119
6.8	Anomalous diffusion trace used in TRANSP modelling derived from binned RMS Mirnov coil signal, scaled in order to match the drop in neutron rate during the major fishbone, for a reasonable scaling of the fission chamber data.	124
6.9	Traces of the neutron rate, comparing the fission chamber data (as calibrated at the time of analysis) to TRANSP simulations with various levels of anomalous diffusion, and with none.	124
6.10	Measured neutron rate fitted to each of the 3 trial anomalous diffusion traces with a multiplicative scaling factor, derived using the minimisation of the sum of the squared residuals.	125
6.11	Wavelength-integrated time traces for the experimental FIDA data and FIDASIM data for Shot #45026. (a) is at $R = 1.029$ m, (b) is at $R = 1.293$ m. Based on the results of Section 2.3, an characteristic error of 20% error is assumed on the FIDASIM values. A full deterministic sampling analysis could be valuable here to determine a specific value. .	126

6.12	Experimental (blue) and FIDASIM (red) spectra before (a) and after (b) major fishbone at $R = 1.293$ m. Green region designates integration range used in Figure 6.11. Based on the results of Section 2.3, an characteristic error of 20 % error is assumed on the FIDASIM spectra.	127
6.13	Synthetic fast-ion distributions produced by TRANSP/NUBEAM before (a) and after (b) major fishbone at $R = 1.293$ m, along with the relative change (c). Lime contours denote the extent of the weight function, which is only for those wavelengths that had good agreement between the experiment and simulation. Therefore data within the contour can be utilised for analysis, whereas the data outside of the contour either cannot be seen by the diagnostic or has been shown in Figure 6.12 to not match experiment. Dark green region in (c) designates region where the fast ion distribution was 0 before the crash.	128
6.14	Time traces of plasma parameters of Shot #45091. (a) NBI power, (b) neutron rate, (c) line integrated electron density, (d) D- α intensity.	131
6.15	Amplitude of the saddle coil data (black) and neutron rate (green) vs. time focusing on the times just before and after the plasma bulk rotation goes to zero. The first drop in the neutron rate is due to a sawtooth, but subsequent drops are correlated with the behaviour of the locked mode. SS beam off time is shown in red, SW off time in blue.	131

- 6.16 (a) Time trace of the low-frequency Mirnov coil signal and (b) Spectrogram of low-frequency Mirnov coil time trace in shot #45091, showing the main resonant behaviour of the plasma. Specifically, after the SW and SS beam turn on (times of this are shown by the magenta and green vertical dotted lines respectively) TAEs are excited. Shortly thereafter, the SW beam erroneously turns off (magenta dot-dashed line) for a short duration while TAEs continue before it comes back on (blue dashed line). There is a very short period of relative quiescence before the excitation of fishbones that rapidly transition to an LLM. There is a sawtooth crash just after 0.5s, but this does not halt the the plasma rotation slowdown generated by the LLM. The bulk plasma comes to a halt at ~ 0.54 s. A locked mode, invisible to the Mirnov coil due to it being stationary in the machine frame, is excited here, which persists until ~ 0.59 s, after which it appears to reduce in magnitude or stops entirely. Both beams terminate soon after (SW - blue dot-dashed line, SS - green dot-dashed line). 132
- 6.17 Wavelength integrated (660.46 nm-661.93 nm) time traces of the FIDA data both for the full duration of the SS beam (a) and zoomed into the end of the shot (b) at (i) $R = 1.029$ m, (ii) $R = 1.23$ m, and (iii) $R = 1.354$ m. 133
- 6.18 Toroidal rotation profile from CXRS in shot #45091, focusing specifically on the post-sawtooth, locked mode period of the shot. After the sawtooth the plasma slows rapidly across the entire plasma radius, coming to a near stop at ~ 0.54 s, except for a spatially localised region at ~ 1.25 m – 1.3 m, which exhibits a counter-current rotation of ~ 20 m s⁻¹. 135

6.19	Electron temperature profiles from TS for the outboard region of the plasma in shot #45091, for times when the plasma rotation has gone to near zero (or negative, as shown in Figure 6.18). A plateauing of the temperature can be seen at ~ 1.3 m, indicative of the existence of a magnetic island (the locked mode).	135
6.20	Fast-ion slowing down time at each FIDA channel averaged across the quiescent plasma period 0.32 s-0.35 s for shot #45091. Error bars here are the standard deviation of the slowing down times over this period. .	136
6.21	Detail of Figure 6.16, showing the Mirnov signal (a) and spectrogram (b) zoomed into the fishbones and LLM specifically.	136
6.22	Radial profile of the FIDA and FIDASIM signals at $\lambda = 660.69$ nm ($E_{min} = 45.9$ keV) averaged over the times 0.347 s-0.35 s, the end of the quiescent period in shot #45091. Note that the channels in the range 0.77 m-0.96 m do not have usable experimental data at the selected wavelength due to being obscured by the beam peaks. While the general profile shapes in the FIDA (blue) and FIDASIM (red) agree, the absolute values do not. An additional scaling factor of 1.45 brings agreement around the peak values (green) but there remains a discrepancy in the edge channels, which requires further investigation to resolve. Based on the results of Section 2.3, an characteristic error of 20 % error is assumed on the FIDASIM values, but this does not full resolve the discrepancy.	137

6.23 The percentage change in both the FIDA and FIDASIM signal at $\lambda = 660.69 \text{ nm}$ ($E_{min} = 45.9 \text{ keV}$), going from the end of the quiescent duration (0.347 s-0.35 s) to subsequent time-slices throughout the fishbone/LLM portion of shot #45091. The time-slices for the "after" duration are (a) 0.366 s-0.369 s, (b) 0.377 s-0.380 s, (c) 0.383 s-0.386 s, (d) 0.390 s-0.393 s, (e) 0.397 s-0.4 s, (f) 0.407 s-0.41 s, (g) 0.437 s-0.44 s, (h) 0.467 s-0.47 s, and (i) 0.492 s-0.495 s. Based on the results of Section 2.3, an approximate absolute error of 30 percentage points is assumed on the FIDASIM values. A full deterministic sampling analysis would be useful to determine specific uncertainties. 138

List of Tables

3.1	Pairings of the tokamak fibres used for the active and passive FIDA, for the specific channels in the spectrometer.	50
3.2	Transmissivity of fibres from patch panel to spectrometer. These are found by placing the individual fibre inputs into the output of an integrating sphere and binning the counts found in the CCD images produced. These are then normalised to the channel of maximum count, and scaled by 0.65. The scaling is based on the transmissivity known from the MAST calibration, taking into account the likely degradation over the intra-MAST/MAST-U period. The channel used to normalise is highlighted in orange.	61
3.3	Integrating sphere luminance used for the various fibre calibrations. . .	64
3.4	Transmission ratios used during the first experimental campaign on MAST-U for the tokamak fibres.	65
3.5	Coordinates defining the position vectors of the active and passive lenses in machine coordinates, \vec{OL}_{active} and $\vec{OL}_{passive}$, relative to the origin (which is defined at the intersection of the centre column axis and geometric midplane). These are known from the MAST-U design documents.	69

3.6	Basis vectors of the lens coordinate system $\{uvw\}$ in terms of the basis vectors of the machine coordinate system $\{xyz\}$, for both the active and passive lenses. The values a , b and c are the coefficients for the machine basis vectors \hat{e}_x , \hat{e}_y and \hat{e}_z respectively.	70
5.1	Fit parameters generated for the specific channel shown in Figure 5.1. Note that the errors derived are solely the standard deviation error in the fitting.	100
5.2	Power and flux fractions for Shot #44254. The mean and standard deviation were evaluated from data across all FIDA channels.	104
5.3	Power and flux fractions for Shot #45483, taken as value from the outermost FIDA channel where the half and third energy peaks could be distinguished from each other, at $R = 1.293$ m.	105
6.1	Percentage change in the wavelength integrated FIDA signal due to the "major" fishbone across $R = 1.029$ m- 1.354 m.	113

Acronyms

ADAS Atomic Data and Analysis Structure. 101, 103

BES Beam Emission Spectroscopy. xii, 35, 51

CAD Computer Aided Design. xiii, 67, 68

CCD charge coupled device. xi, 46, 52, 56, 57, 60, 144

CCFE Culham Centre for Fusion Energy. 145

CX Charge Exchange. 45

CXRS charge exchange recombination spectroscopy. xi, xii, xvii, xx, 22, 24, 31, 34, 49, 50, 66, 98, 104, 113, 120, 135, 148

DS Deterministic Sampling. 37, 38, 142, 143

FIDA fast-ion deuterium- α . i, ix–xii, xv–xviii, xx–xxiv, 20, 22, 24–31, 33–36, 42–47, 50, 51, 53, 54, 57, 58, 60, 63, 66, 73–77, 79, 82–96, 98–108, 110, 112–114, 116, 119, 121–123, 126, 129, 131–142, 144–152

FWHM full width half maximum. 54

ICF inertial confinement fusion. 5

ICRH ion cyclotron resonance heating. 28

IDL Interactive Data Language. 57, 145

ITER International Thermonuclear Experimental Reactor. 2, 5

JET Joint European Torus. 4, 5, 106

LLM long-lived mode. i, xvii, xx–xxii, 19, 20, 109–111, 130–132, 134, 136–141, 151

MAST Mega Ampere Spherical Tokamak. i, ix, x, xxiii, 11, 12, 15, 18–21, 25, 27–30, 32–35, 38, 41, 46, 47, 51–53, 56, 57, 59–61, 69, 75, 77, 78, 139, 142, 144, 146

MAST-U MAST-Upgrade. i, ix, xi, xii, xvi, xxiii, 11–15, 18, 20–22, 24, 25, 27–35, 43–45, 47–54, 58–62, 64, 65, 67, 69, 75, 77, 78, 95, 96, 103, 105–107, 109, 111, 135, 142, 144–147, 149–152

MCF magnetic confinement fusion. 6

MHD magneto-hydrodynamic. 15, 19, 30, 120, 122, 129, 150

MSE Motional Stark Effect. 32, 33, 140, 146, 151

NBI neutral beam injection. xi, xvi, xix, 12, 16, 24, 27, 33, 34, 48, 55, 97, 102, 103, 105–108, 110, 131, 148

NPA Neutral Particle Analyser. 43, 122

RO Responsible Officer. 98

SSNPA Solid State Neutral Particle Analyser. 43, 44, 79, 96, 122, 140, 145, 147, 150, 152

TAE Toroidal Alfvén Eigenmode. i, xx, 17, 109, 123–125, 127, 130, 132, 150

TFTR Tokamak Fusion Test Reactor. 4

TS Thomson Scattering. xxi, 20, 32, 38, 135, 136

UCLA University of California, Los Angeles. 57

UKAEA UK Atomic Energy Authority. 49, 104

VPH Volume Phase Holographic. xii, 52, 53

Introduction and Background

1.1 The motivation for fusion research

With the rapid growth of developing countries and significant population rise, world energy demands are expected to increase by 28 % in the period between 2015 and 2040 alone [1], [2]. Rapidly increasing energy use raises environmental issues, with the threat of catastrophic climate change prompting investigation into low-carbon alternatives. Nuclear fusion energy has the potential to be a part of the solution to this looming crisis. Fusion energy has many attractive properties in regards to a sustainable future. Of the two components needed for the fuel mix deuterium (D) is found abundantly in water, the other, tritium (T) can be bred in fusion reactors, the fusion process itself creates no carbon dioxide or other greenhouse gases and the nuclear waste produced (mostly activated structural materials [3]) becomes safe to handle in a time on the order of a hundred years, as opposed to tens of thousands of years when dealing with the waste produced by fission power plants [3]. Additionally, a fusion reactor cannot undergo catastrophic phenomena akin to meltdown in a fission plant, making it inherently much safer. However, the route to producing net energy from fusion has proven extremely challenging. The conditions necessary for a self-sustaining reaction are extreme (discussed further in section 1.2.2) and generating these conditions is problematic from both a physics and materials standpoint. It is the goal of the international research effort to

devise methods and technologies to overcome these issues and deliver a device that produces net energy output, which is exemplified by the ongoing construction of the International Thermonuclear Experimental Reactor (ITER) device.

1.2 The physics of fusion energy

This section documents some of the physics principles that underpin fusion energy, including the selection of appropriate isotopes to use as a fuel mix, what is needed to achieve a self-sustaining reaction, how magnetic fields can be used to confine ions and how this is implemented practically for fusion energy in the form of tokamaks.

1.2.1 Fusion fuels

All nuclear energy ultimately derives from the liberation of energy from the mass defect found between nuclei and their constituent nucleons. Any reaction where the binding energy per nucleon increases in the product compared to the reactants releases energy. Figure 1.1 shows the binding energy per nucleon versus the isotope mass number.

Nearly all fusion energy research focuses on the fusion of various isotopes of hydrogen into helium. The most promising reaction for fusion energy is that of deuterium and tritium (DT):



Conservation of both energy and momentum determine how much of the released energy is carried by each product. This is inversely proportional to their mass, giving 3.5 MeV for the α particle and 14.1 MeV for the neutron. DT is considered promising for a few reasons: the energy released is high compared to many other fusion reactions and deuterium is found plentifully in nature (tritium is not, but can be bred within fusion reactors) [5]. For the purposes of energy generation




Figure 1.1: Binding energy per nucleon vs. mass number. Any reaction that moves along the curve to a higher binding energy per nucleon releases energy. Note the unusually high binding energy per nucleon for ${}^4_2\text{He}$ compared to the nearby isotopes, which is advantageous for fusion energy. Note also the change from logarithmic to linear scales in the x-axis. Adapted from [4].

the most important advantage is the large reaction cross section at relatively low energy. Figure 1.2 shows the cross section versus the centre of mass (COM) energy for a number of different reactions of interest to energy generation.

At lower energies, between 10 keV and 100 keV, the DT reaction has the highest cross section by around 2 orders of magnitude compared to the other reactions. The reaction rate for a unit volume of plasma will be dependent on the number densities of the two reactants and the product of the cross section of the reaction and the relative velocity of the nuclei averaged over their velocity distributions. This average is known as the *reactivity*, $\langle\sigma v\rangle$. For a fusion reactor to be effective the plasma must be kept hot enough to achieve a high $\langle\sigma v\rangle$, which is easiest to accomplish with DT, with temperatures around 10^8 K- 10^9 K [6].

Tritium is a radioactive isotope of hydrogen, with a half life of around 12.6 years [7]. This is problematic as a site using tritium is subject to strict regulations regarding how much is allowed to be present at one time, and any stockpile will be constantly decreasing over time. However, the neutron emitted in the DT reaction can, in conjunction with a suitable neutron multiplier like beryllium, be used to





Figure 1.2: Fusion cross section vs. centre of mass energy for various potential fusion reactions. Adapted from [4].

breed tritium via reaction with lithium [8]. This then allows for future power plants to maintain a stable and sufficient supply.

The majority of experimental reactors currently operational utilise pure deuterium fuel nearly exclusively. DT campaigns have been performed only on the now decommissioned Tokamak Fusion Test Reactor (TFTR) and on the Joint European Torus (JET) [9]. The DD reaction is [4]:



The first two reaction pathways occur with an almost 50:50 split, whereas the final channel has around a 10 000 times smaller cross section.

Other fuel mixes exist, some of which are *aneutronic*, such as proton-lithium and proton-boron. These eliminate problems of neutron irradiation, but are not currently thought to be feasible for energy generation due to low cross-sections at reasonable temperatures [10].

1.2.2 Self-sustaining reactions

A fusion power plant must, fundamentally, be able to extract more energy from the fusing plasma (from an economic standpoint, much more) than went into heating the plasma, which is parameterised as quality factor Q , the ratio of the fusion power produced by the plasma to the input heating power. A $Q \approx 1$ was achieved by JET during its 1997 DT campaign [9], and ITER is designed to achieve a $Q \approx 10$ [11]. Taking this to the logical conclusion, the ultimate goal is to achieve *ignition*, where the energy output of the plasma is enough to sustain the reaction even after removing all external sources of heating energy, and so $Q \rightarrow \infty$.

A mechanism that confines the extremely energetic α particles to the plasma would allow for self-heating, making ignition feasible. By analysing the balance between the self-heating and power loss due to imperfect confinement, it is possible to derive a criterion for the conditions necessary to achieve ignition. This is the fusion triple product (also known as the Lawson criterion) which is dependent on n , the number density of the reactants, τ_E , the energy confinement time, and T , the plasma temperature [6]. For a DT plasma between 10–20 keV this is:

$$nT\tau_E \gtrsim 3 \times 10^{21} \text{ m}^{-3} \text{ keV s}^{-1} \quad (1.3)$$

By examining the parameters of the criterion some types of confinement can be proposed. By achieving a high temperature and sufficiently high fuel density, the plasma can be allowed to be confined for only a short time. This is the principle behind inertial confinement fusion (ICF). Alternatively, confining a relatively low

density plasma for an extended period at a high temperature can still satisfy the criterion. This is the principle behind magnetic confinement fusion (MCF), which is the subject of the following section.

1.3 Magnetic confinement fusion

MCF is currently the most advanced method to achieving fusion energy available, with research beginning over 60 years ago [5]. This section covers some of the basic physics behind the use of magnetic fields to confine fusion plasmas, with a focus on the most popular type of MCF device, the tokamak.

1.3.1 Particle dynamics in a tokamak

A charged particle of mass m moving through a magnetic and electric field experiences the Lorentz force [12], \mathbf{F}_L :

$$\mathbf{F}_L = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (1.4)$$

where q is the charge of the particle, \mathbf{E} is the electric field, \mathbf{v} is the particle velocity and \mathbf{B} is the magnetic field. Assuming no electric field and a straight, homogeneous magnetic field of strength B the particle's motion perpendicular to the magnetic field is restricted to a gyration around a field line with a gyro-frequency ω of:

$$\omega = \frac{qB}{m} \quad (1.5)$$

The particle's total velocity v can be decomposed into a component parallel to the magnetic field, v_{\parallel} , and a component perpendicular, v_{\perp} . The radius of the gyration of the particle around the field line, the Larmor radius r_L is given by:

$$r_L = \frac{v_{\perp}}{\omega} \quad (1.6)$$

The particle's parallel motion is unaffected. Total particle energy is conserved, and as long as the magnetic field varies little over the r_L the magnetic moment $\mu = mv_{\perp}^2/2B$ is an adiabatic invariant. Under these conditions, a particle moving from lower B to higher B will experience a force anti-parallel to the parallel velocity. Particles with insufficiently high ratio of v_{\parallel} to total v can be reflected due to this, with the effect called a *magnetic mirror*. The *pitch* parameter $p = v_{\parallel}/v$ can be defined to parameterise the population of particles that are trapped by the mirror effect. However, particles can be scattered in pitch due to particle collisions or waves, increasing particle pitch and allowing an end loss of particles. The principle behind the tokamak is to bend the field into a torus, which eliminates these end losses. This leads to the characteristic 'doughnut' shape. Figure 1.3 illustrates the coordinate systems used to describe this configuration in the simplest case of a torus with circular cross section. The ϕ and θ directions are known as the toroidal and poloidal directions respectively. Note also that the distance from the major axis R is the *major radius* (where R_0 is the distance from the major to minor axis) and the distance from the minor axis r is the *minor radius* (with a usually denoting the distance from the minor axis to the edge of the poloidal cross section of the plasma). The primary toroidal magnetic field, B_{ϕ} , can be generated using field coils wrapped poloidally around the torus.

However, the act of changing the geometry from a straight cylinder to a torus means that B_{ϕ} alone is no longer sufficient to confine the plasma. The toroidal geometry means that the field is no longer straight (self-evident) or homogeneous (the magnetic field strength is now inversely proportional to R). This introduces a *drift*, particle velocity components perpendicular to the field lines, specifically the curvature and ∇B drifts, which combined give a drift velocity \mathbf{v}_d :

$$\mathbf{v}_d = \frac{v_{\parallel}^2 + \frac{1}{2}v_{\perp}^2}{\omega} \frac{\mathbf{B} \times \nabla B}{B^2} \quad (1.7)$$

These act in opposite directions for positively and negatively charged particles as




Figure 1.3: Diagram of the coordinate systems used to describe physics in the torus with a simple circular cross section. Adapted from [6]

ω is opposite in direction, causing the ions and electrons to separate in opposite directions, both perpendicular to both \mathbf{B} and ∇B . The charge separation will induce an electric field \mathbf{E} , which in turn induces a drift perpendicular to both \mathbf{E} and \mathbf{B} called the $\mathbf{E} \times \mathbf{B}$ drift, given by:

$$\mathbf{v}_d = \frac{\mathbf{E} \times \mathbf{B}}{B^2} \quad (1.8)$$

This acts in the same direction for both electrons and ions, pushing both out to larger R , eventually causing a loss of confinement. To combat this it is necessary to generate a poloidal component to the magnetic field, B_θ . The resultant overall magnetic field spirals around the torus as a helix, which has the effect of cancelling the net drift experienced by the particles, resulting in effective plasma confinement [5]. The most common method to induce a B_θ is to drive a toroidal current through the plasma, which will then generate a poloidal magnetic field. The toroidal current can be inductively generated by ramping current through toroidally-orientated coils or a central solenoid. This creates a changing magnetic flux through the centre of

the torus, which in turn drives the toroidal current [5]. The resultant magnetic field has a structure of an axisymmetric infinite set of nested toroidal *flux surfaces*, to which the helical magnetic field lines are confined. The plasma pressure is a constant value on each surface, and it is useful to define the poloidal flux function ψ , the poloidal flux contained within a flux surface. The normalised poloidal flux $\psi_N = \psi/\psi_a$ where ψ_a is the value of the poloidal flux at the last closed flux surface, varies from 0 at the magnetic axis to 1 at the last closed flux surface, and is a useful "radial" coordinate to use as the flux surface cross-sections are rarely concentric circles.

A measure of how 'twisted' the magnetic field lines are on a particular flux surface is given by the safety factor, q , [5]. This is essentially the number of circuits taken by the field line toroidally for every circuit poloidally. Formally, if the change in toroidal angle needed for a field line to return to the same position in the poloidal plane is $\Delta\phi$ then:

$$q = \frac{\Delta\phi}{2\pi} \tag{1.9}$$

It can be shown in the limit of a large aspect ratio tokamak that q can also be expressed as $q = rB_\phi/R_0B_\theta$.

1.3.2 Passing and trapped particles

A toroidal configuration means that $|\mathbf{B}| \propto 1/R$. Following a field line in either direction from the outboard midplane, a particle will encounter an increasing magnitude of the magnetic field. This induces a magnetic mirror effect, meaning that a particle with insufficient v_{\parallel} will be reflected. These particles are known as *trapped* particles, and those that have sufficient v_{\parallel} to overcome the mirror force are known as *passing* particles. Neither type are confined exactly to a flux surface due to the effect of the curvature and ∇B drifts but the trapped orbits, known as *banana*

orbits due to the distinctive shape of the orbit in the poloidal projection, exhibit a larger deviation.

Coulomb collisions between unlike particles can cause cross-field transport of particles, as the change in particle velocity causes a change in the location of the particle guiding centre (collisions of like particles generate no net transport). The transport can be envisioned as a random walk, with a step size proportional to the particle's Larmor radius. However, the deviation of the orbit in a banana orbit is much wider than a Larmor radius, and forms a pathway for a random walk with a larger step size and so a higher diffusion rate. This is known as *neoclassical* transport. The cumulative effects of the curvature and ∇B drifts also causes banana orbits to precess in the toroidal direction.

1.3.3 Conventional vs. spherical tokamaks

One of the parameters that describe a tokamak is the *aspect ratio*, $A = R_0/a$. From the start of tokamak construction in the 1960s to the 1980s conventional tokamak design typically had aspect ratios of around 3 to 4 [5]. Around this time research into the effect of reducing aspect ratio revealed that doing so can greatly improve the magnetohydrodynamic (MHD) stability of the plasma. This led to the development of the *spherical tokamak* concept [13]. Spherical tokamaks also exhibit a marked increase in the value of the plasma β , which is the ratio of the mean plasma pressure, p , to the 'pressure' of the magnetic field used to confine it, given by [5]:

$$\beta = \frac{p}{(B^2/2\mu_0)} \quad (1.10)$$

where B is the mean magnetic field strength and μ_0 is the vacuum permeability. A higher β is desirable as it implies that less power is needed to generate a magnetic field to confine a given pressure, and so can be used as a measure of the efficiency of the confinement. However, increasing β in a tokamak with fixed minor radius and

toroidal magnetic field leads to an increase in the pressure gradient, which will almost invariably lead to the generation of instabilities that disrupt the confinement. It was found that by reducing the aspect ratio, it was possible to sustain a stable plasma at a much higher β value, as weaker field can be used. Figure 1.4 shows a representative example of the cross-sections of a conventional tokamak versus a spherical tokamak.




Figure 1.4: Comparison of the profile of an example conventional and spherical tokamak. Adapted from [14]

On a practical level, spherical tokamaks are smaller and less expensive to build (primarily due to needing smaller magnetic fields) [15]. The more compact design does however introduce some complications, as there is less space in and around the device for coils, heating and fuelling systems, and diagnostics.

1.3.4 MAST and MAST-U

The Mega Ampere Spherical Tokamak (MAST) was located at the Culham Centre for Fusion Energy (CCFE), began its experimental programme in December 1999 and was decommissioned in September 2013. Over this period a great deal of research was performed, including but not limited to the development of new plasma start-up techniques, the characterisation of edge localised modes (ELMs)[5]. Following the shutdown of MAST, its successor MAST-Upgrade (MAST-U) was constructed. This has improvements that include increased magnetic field strength, heating power and pulse length, together with a more versatile divertor and improved diagnostics. Construction and preparation for experiments concluded in 2021, with the first experimental campaign occurring in the latter half of the year. At the time of writing, preparations are being made to begin the second experi-

mental campaign in September/October 2022.

The following sections detail some essential background in regards to the generation and behaviour of fast ions in tokamak plasmas.

1.4 Heating

In order to get any thermonuclear yield and, in the future, to reach the point of ignition it is required that the plasma be heated to maximise the reactivity, as described in Section 1.2.1. On both MAST and MAST-U there were/are two methods of heating the plasma: Ohmic heating, and neutral beam injection (NBI). Ohmic heating is inherently generated from passing a current through a plasma with a finite resistivity. This is effective at relatively low temperature, but as plasma resistivity goes as $T_e^{-3/2}$ [16] it is ineffective at reaching the temperatures required for ignition. Neutral beams allow more fusion-relevant temperatures to be reached. In DT plasma, most of the heat should come from the generated and confined α -particles, but NBI is still required to drive current [17].

1.4.1 Neutral beam injection

Injected neutral particles can enter the plasma as they are not deflected by the magnetic fields used for confinement. Once they enter the plasma they can be ionised via impact ionisation from both ions and electrons, and via charge exchange with plasma ions. These newly generated high energy ions then thermalise via Coulomb collisions with the plasma electrons and ions, transferring their energy into the bulk plasma and thereby causing heating.

High energy neutrals can be generated via a neutral beam injection system [18]. A block diagram of the NBI is shown in Figure 1.5. These components will now be described in order.




Figure 1.5: Block diagram of the generation of high energy neutrals by the neutral injection system.

A source of ions is required for acceleration. On MAST-U this is present as a "bucket" source, which is essentially a volume of plasma, generated by heated wire filaments, confined by permanent magnets surrounding it on all sides but one. Ideally this plasma would contain solely mono-atomic deuterium nuclei, but in practice both di- and tri-atomic singly-ionised deuterium ions are present in small quantities. On the side of the source without the permanent magnets, instead there is a metallic grid with circular holes. Current passed through this accelerates the source ions. The accelerated ions pass through a flow of D_2 gas, which neutralises the beam ions via charge exchange, producing fast neutrals that travel onwards and slow ions. Additionally, the D_2^+ and D_3^+ ions will be dissociated and neutralised after acceleration, leading to populations of D neutrals with a half and a third of the energy of the D that were accelerated as individual nuclei. Not all of the beam ions will be neutralised, and some that are neutralised will be re-ionised via neutral impact. The remaining beam ions are deflected by magnets into a "beam dump", leaving a neutral beam to enter the plasma.

1.5 Fast Ions

In a tokamak plasma the velocity distribution of the majority of plasma particles can be described by a Maxwellian, $f_M(v)$ [12]:

$$f_M(v) = n \left(\frac{m}{2\pi k_B T} \right)^{3/2} \exp \left[-\frac{mv^2}{2k_B T} \right] \quad (1.11)$$

where v is the speed of the particles, n is the number density of the species in question, m is the species mass, $k_B = 1.380\,649 \times 10^{-23} \text{ J K}^{-1}$ is Boltzmann's

constant and T is the species temperature in kelvin. Due to the temperatures and energies used in tokamaks it is convenient to discuss temperatures in units of eV, related to the conventional temperature by $eT_{eV} = k_B T$, where $e = 1.602 \times 10^{-19}$ C is the elementary charge. Future references to temperature in this thesis will be in eV.

These particles are described as "thermal" as they are in thermal equilibrium, and hence have a Maxwellian distribution characterised by a single temperature. Often, there is also a population of supra-thermal ions (referred to hereafter as fast ions) with non-Maxwellian distributions (and so are not in thermal equilibrium). The highly energetic α -particles generated by the DT reaction, as described in section 1.2.2 are such a population. On MAST-U, utilising DD fuel, there is also a substantial supra-thermal ion population, namely the ions born from the injected beam neutrals. Effective heating of the plasma requires the good confinement of these particles so that they deposit the majority of their energy back into the thermal population via collisions. The highly energetic nature of this population leads to high velocities and large gyro-radii, leading to significantly different behaviour compared to the thermal population.

Fast ions deposit their energy through Coulomb collisions in the plasma. At high energies the fast ions primarily deposit energy into the bulk electrons, with little associated pitch angle scattering. As the particle slows down, it begins to heat the bulk ions more, along with significant scattering of the fast ion pitch angle. The energy, E_{crit} , at which the rate of loss of energy of the fast ions to electrons is equal to that of the bulk ions is [19]:

$$E_{\text{crit}} = 14.8 T_e \left(\frac{A^{3/2}}{n_e} \sum_j \frac{n_j Z_j^2}{A_j} \right)^{2/3} \quad (1.12)$$

where E_{crit} is in eV, and where T_e is the electron temperature (also in eV), A and A_j are the atomic weight of the fast ions and species j of the bulk ions respectively, n_e and n_j are the electron and bulk ion species j number density respectively,

and Z_j is the atomic charge of bulk ion species j . For a deuterium plasma with negligible impurities there is only one bulk ion species, and with a deuterium beam: $A = A_j = 2$, $Z_j = 1$, and $n_e \approx n_j$ (by quasi-neutrality). This reduces equation 1.12 to:

$$E_{\text{crit}} \approx 18.6 T_e \quad (1.13)$$

The characteristic time for a fast ion to thermalise, τ_{FI} , is given by [19]:

$$\tau_{FI} = \frac{\tau_{es}}{3} \ln \left(1 + \left(\frac{E_0}{E_{\text{crit}}} \right)^{3/2} \right) \quad (1.14)$$

where E_0 is the initial energy of the fast ion, and τ_{es} is the ion-electron slowing down time, which is given by [5]:

$$\tau_{es} = 6(2\pi^3)^{1/2} \frac{\epsilon_0^2 m_{FI} T_e^{3/2}}{m_e^{1/2} n_e e^4 \ln \Lambda} \quad (1.15)$$

where ϵ_0 is the vacuum permittivity, m_{FI} and m_e are the fast ion and electron mass respectively, e is the electron charge, and $\ln \Lambda$ is the Coulomb logarithm. For a deuterium plasma with $T_e > 10 \text{ eV} > (m_e/m_i)T_i$ where T_i is the ion temperature (and $m_e/m_i \approx (1/3662)$), the Coulomb logarithm for electron-ion collisions can be given as [20]:

$$\ln \Lambda = 24 - \ln \left(\frac{1}{T_e} \sqrt{\frac{n_e}{1 \times 10^6}} \right) \quad (1.16)$$

with n_e in units of m^{-3} . For both MAST and MAST-U typically $T_i \sim T_e \sim 1 \text{ keV}$, fulfilling the above criterion.

1.6 Instabilities that affect fast ions

Gradients in the fast-ion distribution constitutes a source of free energy in the plasma, and can excite magneto-hydrodynamic (MHD) modes present in the toka-

mak. The mechanism through which this occurs is *inverse Landau damping*.

1.6.1 Inverse Landau damping

Landau damping [21] is a process through which plasma oscillations may be damped collisionlessly. Linearisation of the Vlasov equation reveals a singularity at $v_p = \omega/k$, the phase velocity of the wave. Complex analysis of this gives a term that can represent either wave damping or wave growth, dependent on the sign of the gradient particle distribution function at v_p . This can be understood intuitively by considering that a particle travelling near v_p will experience a mostly steady electric field from the wave. Effective energy transfer can then occur between the wave and particle, with slower particles taking energy from the wave, and faster ones putting energy into it. In the case of the bulk ions and electrons, which have a Maxwellian distribution, the sign of the gradient will always be negative at v_p , so there are more particles slower than v_p than there are faster than it, leading to a net loss of energy from the wave to the particles and so the wave is damped. However, when a non-Maxwellian population exists like that of the fast ions derived from the NBI, the gradient of the distribution can be positive at v_p . More particles at velocities higher than v_p than below leads to a growth of the wave. This is known as inverse Landau damping, and can lead to resonant parts of the fast-ion distribution driving modes unstable.

1.6.2 Toroidal Alfvén Eigenmodes

It was shown by Alfvén [22] that a magnetised plasma can sustain waves. In a homogeneous plasma cylinder with a homogeneous field along the cylinder axis, shear Alfvén waves can propagate, similar to the plucking of a string under tension, with curving of the magnetic field lines providing the restoring tension and the plasma mass providing the inertia. The wave propagates along the direction of the magnetic field with a phase velocity of $V_A = B_0/\sqrt{\mu_0\rho}$, known as the Alfvén

speed, where B_0 is the strength of the unperturbed magnetic field, μ_0 is the vacuum permeability, and ρ is the plasma mass density. The wave has a dispersion relation relating the wave frequency ω to the parallel component of the wavevector, k_{\parallel} , $\omega = k_{\parallel} V_A$. However, in a non-homogeneous magnetised plasma, for example in a cylindrical plasma where the density is inversely proportional to the radial position, these waves experience significant *continuum damping* [23], where phase mixing of the waves moving at different speeds at different radii cause the waves to dissipate. In a tokamak, B_{ϕ} the toroidal component of the magnetic field, also varies as a function of radius, adding to the issue. Bending the cylinder into a torus adds a periodic constraint to the parallel wavenumber, as does adding a poloidal component to the magnetic field. This requires that $k_{\parallel} = (n - mq)/R$, where n and m are toroidal and poloidal mode numbers, and R is the radius from the centre of the torus. Waves that lie in this Alfvén continuum are damped strongly by the aforementioned continuum damping. However, the $B_{\phi} \propto 1/R$ dependence means that there is a coupling between consecutive poloidal modes (for the same n), leading to frequency gaps in the Alfvén continuum, within which can exist a weakly damped *Toroidal Alfvén Eigenmode (TAE)* of frequency $\omega_{\text{TAE}} = v_A/2qR$. These modes have their wave energy maximised on the flux surface with $q = (m + 1/2)/n$. Fast ions can exchange energy with TAEs, driving them unstable, if they satisfy the resonance condition $\omega_{\text{TAE}} - n\omega_{\phi} - l\omega_{\theta} = 0$ [24], where ω_{ϕ} and ω_{θ} are the toroidal and poloidal frequencies of the fast ion orbits, and l is an integer. TAEs can cause significant fast-ion redistribution [25].

1.6.3 Kink instability, fishbones and sawteeth

A fundamental instability that tokamak plasmas can be susceptible to is the kink instability [26]. A current carrying plasma cylinder generates an azimuthal magnetic field, \mathbf{B}_{θ} . Any displacement or twisting of the plasma causes an increase in the magnitude of \mathbf{B}_{θ} on the inside of the bend, and a decrease on the outside, shown in Figure 1.6. This produces a net $\mathbf{J} \times \mathbf{B}$ force that enhances any initial




Figure 1.6: An example of the kink instability in a column of plasma, which carries a current density \mathbf{J} , that produces an azimuthal field \mathbf{B}_θ . The twisting produces regions of stronger and weaker \mathbf{B}_θ . This results in a net $\mathbf{J} \times \mathbf{B}$ force in the direction of the plasma displacement, causing further displacement, and so the column is unstable. The magnetic tension of an applied \mathbf{B}_ϕ along the length of the column can provide a degree of stabilisation. Adapted from [6].

plasma displacement, rendering the column unstable. Bending a magnetic field line costs energy, therefore imposing a field \mathbf{B}_ϕ along the length of the column provides stabilisation against the instability.

By considering that the most unstable kink modes in a tokamak are those that bend the field lines the least, implying modes of the longest wavelength, the Kruskal-Shafranov limit [27] states that a toroidal plasma is stable to the kink instability as long as it satisfies $q > 1$. $q < 1$ near the plasma boundary can cause a dangerous *external* kink, which moves the entire plasma column. On MAST and MAST-U, most plasmas have a q_{95} , q at $\psi_N = 0.95$, of $\gtrsim 5$, which sufficiently prevents such disastrous occurrences. However, *internal* kink instabilities also exist, which do not move the plasma boundary. Specifically, a $m = n = 1$ mode can exist whenever a $q = 1$ surface exists within the plasma. A peaking of the plasma current on axis can cause a *sawtooth crash*, which rapidly redistributes plasma from the core towards the edge. These are not driven by fast ions, and in fact can be stabilised by their presence [28], but they can cause significant redistribution of the fast ions. These are described in greater detail in Chapter 4. Alternatively, there can be

a resonant exchange between the injected fast ions and the mode in a so-called *fishbone* instability, due to the characteristic trace it makes in Mirnov coil data (see Section 1.7), an example of which is shown in Figure 1.7.




Figure 1.7: Characteristic fishbone trace in Mirnov coil signal for fishbone mode appearing in shot #45026, extant between 0.346 s-0.35 s. The Mirnov coil signal is proportional to the rate of change of the azimuthal component of the magnetic field.

The primary hypothesis for the drive is the resonance between the toroidal precession of the banana orbits of trapped fast ions and the $n = m = 1$ mode. Fast ions can then drive the mode unstable via inverse Landau damping, due to a negative radial gradient. There is also evidence that passing fast ions can resonate via the toroidal rotation frequency [29]. This resonance can cause fast-ion redistribution to greater radii. It has been shown using the ideal MHD stability code MISHKA that for a typical MAST equilibrium the primary mode of a fishbone is maximised around $\rho \sim 0.35$ [30], with $\rho = \sqrt{\Phi_N}$, where Φ_N is the toroidal flux normalised to its boundary value.

After a fishbone occurs, it is often seen that the $n = m = 1$ mode can evolve into a saturated internal kink mode, which is called a long-lived mode (LLM) [31]. As the name suggests, these can persist for hundreds of milliseconds, during which the plasma bulk toroidal rotation profile is flattened due to an increase in the toroidal plasma viscosity due to a breaking of the toroidal symmetry and eddy currents generated in the vessel wall causing a counter-rotation torque [32] (effect

is discussed further below, in regards to locked modes). Such modes can cause significant fast ion losses. LLMs were fairly common on MAST [31] but appear to be less common on MAST-U. However lower density shots, which are more suitable for FIDA analysis due to greater penetration of the beam into the plasma therefore higher signal (see Chapter 2) are also suited to the presence of LLMs. This is because lower densities have a larger current penetration time, suppressing the appearance of sawtooth crashes and therefore allowing more LLMs to develop. Analysis in Section 6.2 discusses a LLM on MAST-U.

After the slowing down of the plasma it is not uncommon to observe *locked modes*. These are 2/1 modes, magnetic islands that due to the interaction of the island's field perturbation with the conductive wall of the tokamak, and with the *error fields*, small imperfections in the axisymmetry of the magnetic equilibrium [33], cause eddy currents that generate a $\mathbf{j} \times \mathbf{B}$ torque that drives the plasma rotation to zero. The magnetic island remains, now static in space, and the lack of the opposing torque allows the mode to grow. This can cause significant fast-ion redistribution. Care must be taken when diagnosing such modes as they require the use of a separate set of diagnostic coils to observe, which is discussed in Section 1.7.

1.7 Tokamak diagnostics

In order to fully diagnose a tokamak plasma shot, analysing the overall behaviour and specifying the instabilities, there must be a holistic approach to the tokamak diagnostics. Chapters 2 and 3 describe in detail the principals and experimental setup of the FIDA diagnostic. Here, the operation of some other crucial diagnostics are described.

Thomson Scattering (TS) [34] is a diagnostic technique used on both MAST and MAST-U [35]. It gives radially and temporally resolved measurements of electron density and temperature. Laser photons injected into the plasma scatter from the free electrons. The velocity of the free electrons causes the observed scattered

spectrum to be Doppler broadened compared to that injected. The electron temperature can then be determined from the width of the scattered spectrum, and the electron density from the absolute intensity.

Both MAST and MAST-U utilised a setup of 8 Nd:YAG lasers. Each laser has a repetition rate of only 30 Hz, but are programmed to fire in sequence, such that measurement frequencies of ≈ 250 Hz are achieved. The laser array can also be programmed to fire in a "burst mode" to allow for a very high time resolution ($\sim 1 \mu\text{s}$) for a short duration of time. No shots using this mode are analysed during this thesis.

The total neutron rate can be derived from the use of a *fission chamber*. An aluminium cylinder encases an ionisation chamber with electrodes coated in enriched $^{235}_{92}\text{U}$, which all-together is encased by a polythene neutron moderator [36]. Neutrons emitted by DD fusion within the vacuum vessel are slowed by the moderator to allow for a greater cross-section of fission with the uranium. Fission reactions generate high energy fission products that cause ionisation within the chamber, generating current pulses that can be detected. Appropriate calibrations and models [37] can then be applied to generate a global neutron rate for the discharge.

The plasma simulation code TRANSP/NUBEAM (described further in Section 2.2.1) can provide a synthetic neutron rate as a comparison with the experimental value. During the first experimental campaign on MAST-U it was found that there was a significant discrepancy between the experimental and synthetic values, which initially cast doubt on the accuracy of the calibration of the fission chamber. Improvements to the simulation reduced the discrepancy (discussed in Section 2.2.1.1 and Chapter 5), but a significant difference still remains. Further investigation will be required to address this issue in future experimental campaigns.

Mirnov coils are a crucial diagnostic for diagnosing instabilities. Coils of wire are situated around the tokamak vacuum vessel, with coils orientated in each orthogonal tokamak direction (radially, poloidally and toroidally) [38]. The voltage, V ,

induced in a coil is proportional to the rate of change of the of the magnetic field, $\dot{\mathbf{B}}$, along the normal to the plane of the coil, \mathbf{n} . For a coil with area A and N number of turns, $\dot{\mathbf{B}}$ can therefore be defined [34]:

$$\dot{\mathbf{B}} \cdot \mathbf{n} = \frac{V}{NA} \quad (1.17)$$

Integration of the Mirnov coil voltage gives the magnetic flux through the coil, and from this the magnetic field strength in the normal direction $\mathbf{B} \cdot \mathbf{n}$ can be determined. This is crucial information for equilibrium codes like EFIT++ [39] (described further in Section 2.2.1), constraining the boundary values of the field strength. During the work described in this thesis, the primary coils utilised are those orientated with the normal in the poloidal direction: hereafter this specific set of coils will be synecdochically referred to as the Mirnov coils. These coils are extremely useful at determining the specific type of resonant instability occurring during a shot, specifically by performing a spectrographic analysis of the coil voltage. Different instabilities have distinct signatures on a spectrogram, allowing confident determination of the specific type. Large coils with a normal primarily in the radial direction, known as the saddle coils, are also useful, specifically in the detection of locked modes, which are invisible in the poloidal Mirnov coil spectrogram due to the plasma rotation having gone to zero, so there is no varying component in the poloidal direction.

FIDA is described at length in this thesis, both the physical mechanisms behind it (Chapter 2) and the practical implementation on MAST-U (Chapter 3). FIDA is a specialised use of charge exchange recombination spectroscopy (CXRS) [40], where injected neutral atoms act as donors to plasma ions for charge exchange, allowing the emission of line radiation. On MAST-U FIDA examines Doppler-shifted D- α emission, but there is also a system examining impurity emission [41], specifically that of carbon, normally examining the C⁵⁺ emission line at 529 nm. By examining the broadening and overall Doppler shift of the line, the ion temperature and bulk

plasma rotation velocity can both be determined for broad range of localised radial positions.

The principles of the FIDA diagnostic and its simulation


In this chapter the principles of the FIDA diagnostic are presented, discussing the mechanisms of emission and detection, with a breakdown of the components that make up the captured spectra. The forward modelling of the captured spectra is described, with description of the analysis workflow using the codes TRANSP/NUBEAM and FIDASIM.

2.1 Physics of the FIDA diagnostic

2.1.1 The FIDA mechanism

FIDA is a specific form of charge exchange recombination spectroscopy (CXRS)[40]. In a typical MAST-U plasma (core electron temperature $T_e \sim 1$ keV) the overwhelming majority of thermal deuterons present are ions, and so there is no source of thermal neutrals with which the fast ions can charge exchange. In order to make spectral measurements, neutral deuterium atoms injected by the NBI system can undergo charge exchange with the population of fast ions existing in the plasma. These generate fast neutrals, which can be excited by electron collisions. Subsequent spontaneous de-excitation from principal quantum state $n = 3$ to $n = 2$

generates deuterium- α photons (equivalent to the Balmer- α transition in protium, but with a rest wavelength of 656.1 nm, compared to 656.3 nm)[42]–[44]:



where D^0 is a deuterium atom, D^+ is a deuterium ion, D^{0*} is a deuterium atom in an excited state, and γ is a D- α photon.

The emission of photons by fast ions, as part of a significantly supra-thermal population, will exhibit significant Doppler shift, generating a spectral component Doppler shifted away from the rest wavelength. This is proportional to the component of velocity of the fast neutrals (formerly fast ions) parallel to the line of sight of the diagnostic, v_{\parallel} , with the wavelength shift, $\Delta\lambda$ given by:

$$\Delta\lambda = \left(\frac{v_{\parallel}}{c}\right) \lambda_0 \quad (2.2)$$

where c is the speed of light, and λ_0 is the rest wavelength of the D- α transition.

On MAST-U, the diagnostic views are primarily along the direction of the neutral beam, ensuring that the FIDA emission is primarily red-shifted. The diagnostic is set up (see Section 3.1.2) to view only the red-shifted side of the spectrum. On MAST there was an additional view that viewed vertically down onto the beam, which observed a primarily blue shifted spectrum [45]. No such view *currently* exists on MAST-U, but there has been discussion about implementing a vertical view for a future campaign, likely not before the fourth experimental campaign on MAST-U at the earliest. This potential view is discussed further in Section 3.6.

The primary goal of FIDA spectroscopy is to diagnose the fast-ion distribution function, $f(E, p, R, Z)$, the density of fast ions in both phase and real space, where E is energy, $p = v_{B\parallel}/v$ is pitch (the ratio of the component of the velocity of the fast ion parallel to the magnetic field, $V_{B\parallel}$, to its full velocity, v), and R and Z are the radial and vertical coordinate in the poloidal cross-section. The FIDA diagnostic

measures signal at only 11 discrete intersections of the lines of sight and the neutral beam, which lies in the device midplane ($Z=0.0$ m). As the neutral beam has a finite width, the FIDA measurement is not a true "point" measurement, but comes from a somewhat cigar shaped region formed by the intersection. However, it is much more local than other "line-integrated" diagnostics. For the purposes of discussion, the FIDA channels will be discussed by the radius of the intersection of the line of sight with the centre of the beam (the device midplane), so it is easier to discuss the local velocity-space fast-ion distribution function $f(E, p)$ at these locations.

The relationship between the distribution function f and the FIDA signal from the charge exchange of beam neutrals and fast ions, $S_f(\lambda_1, \lambda_2, \phi)$, where λ_1 and λ_2 denote the wavelength limits of the observed signal and ϕ is the angle of the line of sight to the magnetic field, can be described by convolving the signal with a *weight function* $w(\lambda_1, \lambda_2, \phi, n_b, E, p)$ [44], [46], where n_b is the local neutral beam density:

$$S_f(\lambda_1, \lambda_2, \phi) = \int \int w(\lambda_1, \lambda_2, \phi, n_b, E, p) f(E, p) dE dp \quad (2.3)$$

The weight function can be thought of as the sensitivity of the diagnostic to different regions of velocity space [47], [48]. Due to the ϕ dependence different lines-of-sight have different regions of sensitivity. As a non-zero signal requires the presence of a neutral beam, the specific values taken by the weight function are dependent on the beam density, but the delineation between regions of phase space with zero and non-zero values should not vary significantly. This means that, in general, the sensitive (non-zero weight function values) and insensitive (weight function value of zero) regions are dependent on only the spectral and neutral beam geometry, so it is possible to compare phase-space coverage of specific lines of sight. Figure 2.1 shows an example of this variation, showing the weight function for Shot #45091 at $R = 0.96$ m (2.1(a)) and $R = 1.413$ m (2.1(b)), over the range of wavelengths accessible by each channel (determined by the location and extent of the full energy beam peak, which is described in Section 2.1.2) Additionally, Figure 2.1(c) shows

the weight function that would be found for the potential vertical view mentioned earlier, also observing the NBI beam at $R = 1.413$ m, from below. For this view, the FIDA emission is primarily blue-shifted, and the upper wavelength cutoff dependent on either the location and extent of the full energy beam peak or simply the cold D- α , dependent on the channel observed.

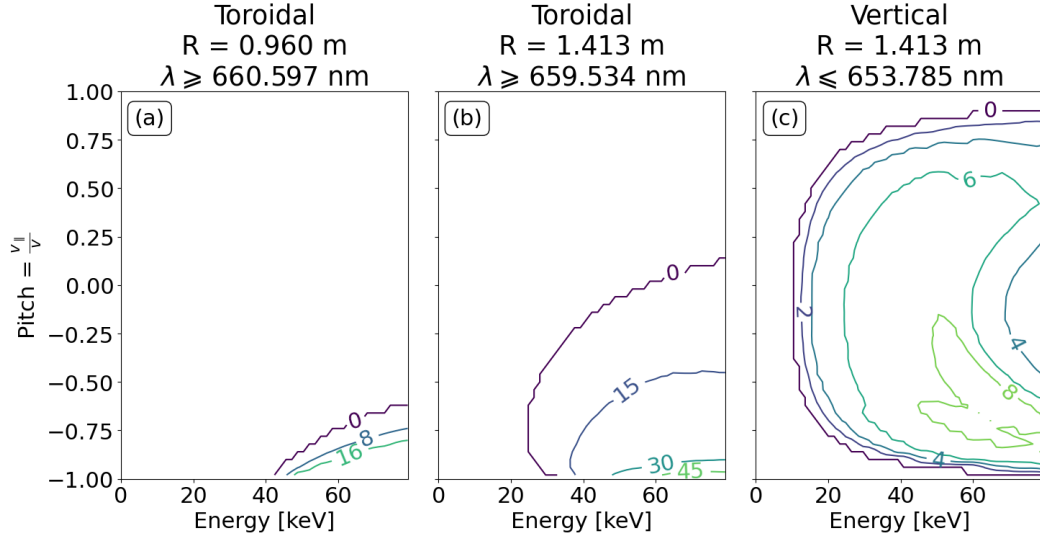


Figure 2.1: Weight functions (with units of $(\text{Ph}\cdot\text{cm})/(\text{s}\cdot\text{nm}\cdot\text{sr}\cdot\text{fast-ion}\cdot\text{dE}\cdot\text{dP})$) for the range of wavelengths accessible by the FIDA diagnostic, for shot #45091 at $R = 0.96$ m (a) and $R = 1.413$ m (b). (c) shows an example of a weight function for a potential vertical view (with line-of-sight origin below the beam), viewing $R = 1.413$ m. For this view the FIDA emission is primarily blue-shifted, giving an upper, rather than lower, limit to the accessible wavelengths.

Figure 2.1 shows that on MAST and MAST-U the toroidal views are primarily sensitive to passing particles, while the vertical view on MAST and the prospective vertical view on MAST-U are primarily sensitive to trapped particles.

2.1.2 Typical components of the observed FIDA diagnostic spectra

Unfortunately (for the purpose of FIDA analysis) the observed spectra obtained with the FIDA diagnostic does not contain only S_f , the FIDA emission. A typical measured spectrum, S is comprised as so [49]:

$$S = S_f + I_{cx} + S_{inj} + S_{halo} + b + i_{ncx} + s_{cold} + s_f \quad (2.4)$$

The components in uppercase represent *active* spectral components, present only when observing the neutral beam present in the plasma, while lowercase components are *passive*, and are present both when observing or not observing the neutral beam.

The active components include S_f , the the desired FIDA signal, and I_{cx} , impurity lines generated by charge exchange reactions with the injected neutrals. S_{inj} is emission directly from excited injected beam neutrals. As described in Section 1.4.1, the beam contains deuterium atoms at the beam energy, and also at a half and a third of the beam energy. This leads to 3 distinctive peaks, one for each energetic component, with the central wavelength determined by the speed of the atoms in each component. These spectral components will be referred to as the *full*, *half* or *third energy beam peak*. The S_{halo} component is generated by the charge exchange of beam neutrals with thermal ions. This also produces a "cloud" of thermal temperature neutrals around the neutral beam, called the *halo* neutrals. These form an additional population with which charge exchange can occur.

The passive components b , i_{inx} , s_{cold} and s_f are, respectively: bremsstrahlung; intrinsic impurity lines; emission near the rest wavelength from neutrals in the colder plasma edge; and emission from the charge exchange of fast ions with the colder edge neutrals [49]. It should be noted that all of the passive components except s_f are present irrespective of the presence of a neutral beam in the plasma at all. s_f requires a population of fast ions so, in the case of MAST and MAST-U, needs the presence of a beam. However, this beam can be one that is not directly observed by the diagnostic. Alternative supra-ohmic heating like ion cyclotron resonance heating (ICRH) [50] present on other tokamaks can produce fast-ions without the presence of beams, so would also generate a s_f component.

An example of a typical spectrum from the FIDA diagnostic is shown in Figure 2.2.

On MAST and MAST-U the majority of the FIDA light is redshifted, but most is obscured by the beam peaks, leaving only a small region of usable spectrum at the higher wavelengths.

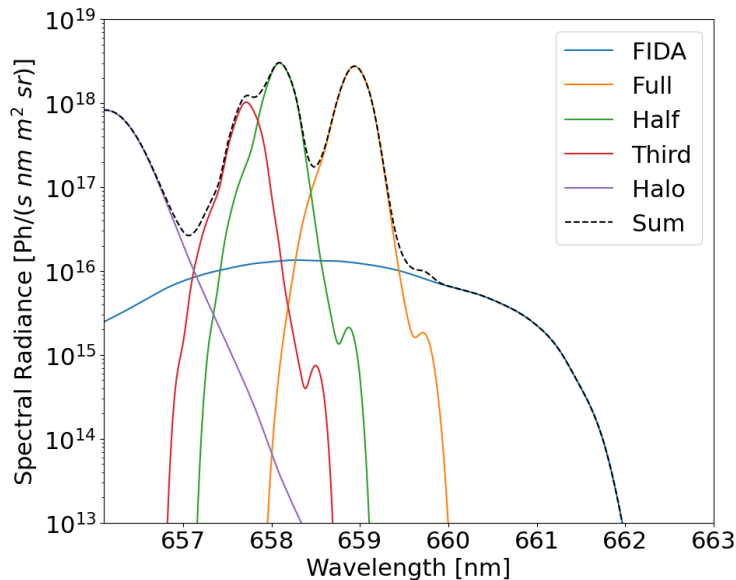


Figure 2.2: A typical FIDA spectrum after background subtraction on MAST/MAST-U, as generated by FIDASIM (described further in Section 2.2.2), and convolved with the instrument function of the diagnostic to account for instrumental broadening. Once background subtraction has been performed, only the full, half, and third energy beam peaks, the halo component and the FIDA component should remain. The small secondary peaks seen on each of the beam peaks are an artefact of the convolution with the instrument function

2.2 Forward modelling

Directly extracting the local fast-ion distribution from the FIDA signal requires the inversion of Equation 2.3, a highly non-trivial problem. It has been shown that by using multiple lines of sight observing the same spatial location (i.e. lines of sight from separate ports directed towards the same volume in space) it is possible to reconstruct the distribution function f [51]. Such tomographic inversions have been performed quite extensively on data from ASDEX Upgrade [46], [52], [53], where 3 to 5 viewpoints were available, while some work has been done on MAST data [54], utilising the toroidal and vertical views. However, in both cases external

constraints on the reconstructed distribution have to be imposed in order to reduce severe artifacting. On MAST-U, due to only having a single viewpoint onto the beam, tomographic inversion of FIDA data is not very suitable to interrogate the fast-ion distribution. If the prospective vertical view discussed in Section 3.6 is implemented in a future campaign it may be possible and valuable to examine tomographic inversions of MAST-U data.

In this thesis interrogation of the FIDA data on MAST and MAST-U was instead performed with *forward modelling*. A plasma transport code can be used to generate a synthetic fast-ion distribution, which is then given as input to the synthetic diagnostic code FIDASIM [55]–[57], which produces the predicted spectra for the FIDA diagnostic, and is described further in Section 2.2.2. Good agreement between experiment and synthetic data gives evidence towards the underlying synthetic fast-ion distribution being an accurate representation of reality, at least in the sensitive regions of velocity space. Disagreement can suggest the presence of non-classical plasma dynamics such as magneto-hydrodynamic (MHD) instabilities, which can then be investigated further.

In this thesis, the plasma transport code used to generate the synthetic fast-ion distribution was TRANSP/NUBEAM [58], [59], which is described in more detail in Section 2.2.1. Other codes could be used, for example, the Monte Carlo orbit-following code ASCOT [60] can produce a 3D synthetic fast-ion distribution $F(E, p, R, Z, \zeta)$, where ζ is the toroidal angle, compared to the 2D distribution produced by TRANSP. However ASCOT can only solve for steady-state plasmas, and TRANSP/NUBEAM to FIDASIM workflows are already well established compared to ASCOT, where such a workflow is only recently begun to be established. ASCOT is not used in the work in this thesis.

2.2.1 TRANSP/NUBEAM

TRANSP is a well established plasma transport code (having existed in some form since the 1980s [58]) that can be used to both analyse existing shot data (interpretive mode) and predict physical parameters of tokamak plasmas (predictive mode). In the course of this thesis TRANSP is used exclusively in interpretive mode. An accessible breakdown of the basic operation of TRANSP is given by Ongena *et al.* [61]. In summary, a wide variety of diagnostic data on parameters including, but not limited to, electron temperature and density from Thomson Scattering [35] and bulk ion temperature and plasma rotation velocity from CXRS [41], along with magnetic equilibria from the equilibrium reconstruction code EFIT++ [39] are given as input. The code treats the equilibria in 2D, but keeps the plasma parameters constant on each flux surface (1D), so the code is sometimes referred to as a "1.5D" code. The plasma parameters are mapped to flux surfaces, denoted by $\rho = \sqrt{\Phi_N}$, where Φ_N is the toroidal flux normalised to its boundary value. Then the particle continuity equations are solved, utilising the input data and additional data, such as the positions and energies of injected neutrals, from the code module NUBEAM. NUBEAM [59] is a Monte Carlo code that simulates and tracks the fast particles, from their injection as neutrals in the neutral beam, to their ionisation and subsequent evolution as fast ions within the plasma. This requires knowledge of the beam geometry and the power, energy and current fractions to be known as a function of time. The exact knowledge of the current fractions is potentially problematic for the first experimental campaign on MAST-U, which is described further in Section 2.2.1.1 and Chapter 5. With the solution of the continuity equation, the ion and electron energy balance equations can then be solved. TRANSP/NUBEAM produces particle diffusivities, thermal conductivities and, crucially for FIDA, spatially and phase-space resolved fast-ion distributions at requested times throughout the shot. The flux-surface mapped plasma parameters and the fast-ion distribution then contribute to the input to FIDASIM.

TRANSP does not simulate or predict the presence of plasma instabilities in general. The presence of such instabilities will normally cause the simulation to reflect reality less well. This itself can be used as a diagnostic tool, by comparing synthetic spectra from FIDASIM to experiment, deviations can tell us about the magnitude of the effect of the instability on the fast ions. There are a few limited ways in which TRANSP can model the effects of some instabilities. Chapter 4 extensively discusses the nature and utilisation of the sawtooth models that can be implemented in a TRANSP run in order to simulate the effect of sawtooth crashes on the plasma. More generally the effects of instabilities on the fast ions can be approximated by introducing an additional, user-specified *anomalous diffusion*, D_b , to the fast ions. This introduces an additional flux to the fast ions, Γ_{anom} , given by [62]:

$$\Gamma_{anom} = -D_b \nabla n_b \quad (2.5)$$

where n_b is the fast ion density. The diffusion can be given as a constant, or as a profile that can vary in time, radius and particle energy. An additional option can allow the code to determine an appropriate value by attempting to find the value that causes the predicted and measured neutron rates to match. On MAST-U the selection of appropriate anomalous diffusion values is complicated by diagnostic issues, described further in Section 2.2.1.1.

The use of TRANSP/NUBEAM on MAST and MAST-U is streamlined by the use of the OMFIT framework [63], which is a collection of modules that make the preparation of TRANSP/NUBEAM input easier. Firstly, a magnetic equilibrium is required. EFIT++ solves the Grad-Shafranov equation with imposed boundary conditions from various diagnostics to generate the equilibrium [64]. The diagnostics include magnetic coils, TS and the Motional Stark Effect (MSE) diagnostic. Including the data from more diagnostics constrains the reconstruction to produce more accurate EFIT++ equilibria. During and for a duration after the first experimental campaign on MAST-U only magnetics data was available,

meaning that the available equilibria are not as constrained as would be preferred. MSE data was taken during this campaign, and eventually will be made available to better constrain the reconstructions. However, this was not possible for the analysis of MAST-U data included in this thesis. With the equilibrium available, the plasma parameters are mapped to ρ , to which the user fits profiles. Along with NBI parameters and other quantities, all of this is submitted to TRANSP/NUBEAM, which can then run. The output is then used in FIDASIM, which is described in Section 2.2.2. However, there are some caveats to MAST-U TRANSP/NUBEAM runs that may effect the results, which are described below.

2.2.1.1 MAST-U TRANSP/NUBEAM limitations

During the setup of a TRANSP run within OMFIT we must specify the effective charge, Z_{eff} , which is defined [5]:

$$Z_{\text{eff}} = \frac{\sum_j n_j Z_j^2}{\sum_j n_j Z_j} \quad (2.6)$$

where where n_j and Z_j are the number density and atomic number of particle species j , respectively. In a completely pure deuterium plasma $Z_{\text{eff}} = Z = 1$. In reality, impurities will inevitably increase the value of Z_{eff} . During the first campaign there were no diagnostics capable of measuring Z_{eff} . In all TRANSP/NUBEAM runs for the first campaign on MAST-U Z_{eff} is set to 1.5, constant in both time and space, to take into account the presence of carbon, the primary impurity in MAST-U plasmas. On MAST there was an RGB camera diagnostic [65] which would take 2D images of the plasma emission, whereupon it is possible to infer the Z_{eff} profile in both time and space via the bremsstrahlung emission. The RGB camera diagnostic should be available for the second experimental campaign on MAST-U, resolving this issue in the future. In the context of synthetic FIDA spectra, produced from FIDASIM (discussed in Section 2.2.2) from the TRANSP output, the exact value of Z_{eff} is less important, as the population of fast ions observable with the fast-

ion deuterium- α (FIDA) diagnostic is well above the critical energy (discussed in Section 1.5) and so is primarily interacting with the electron population.

During the first campaign, no measurements of the edge neutral density were available, and is set to a default value in TRANSP/NUBEAM. This parameter influences the calculation of charge exchange losses of fast ions that gyrate slightly outside the last closed flux surface, so will affect the calculation of both active and passive FIDA near the edge.

NUBEAM must be provided with the beam energy, beam power, and flux fractions (fraction of injected deuterons at full, half and third energy). It was found that the SS NBI system had been configured incorrectly for the majority of the first experimental campaign on MAST-U (from mid-March 2021 onwards). This meant that the flux fractions were significantly different to what had been stated. Chapter 5 discusses this extensively, along with the use of FIDA and CXRS data to derive more accurate beam fractions. However, the analysis does not allow for a full retroactive calculation of accurate beam fractions for the different beam conditions of each shot. There will also be some effect on the overall beam power, although this is thought to be relatively minor. Overall, this potentially introduces error into the TRANSP/NUBEAM simulations.

It was previously established on MAST [66] that there was a discrepancy between the predicted and measured neutron rates, with the measured rate being $\sim 40\%$ lower than the predicted, with the source of the discrepancy hypothesised to be due to the guiding center approximation utilised within NUBEAM. Within this thesis, all synthetic neutron rates are corrected with a factor of 0.6. However, on MAST-U it was found during the first campaign that predicted neutron rates were significantly higher than those measured by the fission chamber, to a degree larger than that explained by [66]. Initially the predicted neutron rates were roughly four times higher than the measured rates, but investigation into the beam fractions utilised at the time (documented in Chapter 5) reduced this to roughly three times higher, which is still quite poor. The value of anomalous diffusion required to give

agreement is unreasonably high, suggesting that either there is a calibration issue with the fission chamber or there are as yet further unaccounted for errors in the TRANSP setup for MAST-U. Due to this finding appropriate values for anomalous diffusion in simulations is not trivial, but an attempt is made in Section 6.1.2.1.

2.2.2 FIDASIM

FIDASIM[55]–[57] is a synthetic diagnostic code that models expected radiance received by the FIDA diagnostic at a single time during a plasma discharge. The code can model not just the FIDA signal, but also the BES peaks and also the passive FIDA signal, among others.

FIDASIM takes in the electron density and temperature, ion temperature, Z effective (from which, along with the electron density, the ion density can be found), the cold neutral density and the various directional components of the bulk plasma rotation. The field parameters, as defined by EFIT++ are also input, along with the synthetic fast-ion distribution as derived by TRANSP/NUBEAM. All of these are mapped to a 2D Cartesian R - Z grid, referred to as the interpolation grid. As mentioned above FIDASIM can accept 3D data, varying in R, Z , and ϕ , which would then be defined onto a cylindrical grid, but only axisymmetric data is used in this thesis.

The beam power, energy and the current fractions at the simulation time are given, and are used to define the initial flux of each beam component (full, half and third energy, see Section 1.4.1).

The code utilises a Monte Carlo style simulation by following injected neutrals from the neutral beam. On both MAST and MAST-U the FIDA diagnostic observes the SS beam, so it is these injected particles that are tracked. However, on MAST as the SW beam was directed in the midplane, during dual beam shots there could be contamination of the inner channels by the beam peaks generated by the SW beam. On MAST-U as the SW beam has been raised out of the midplane

this is no longer an issue. In order to simulate the injected neutrals, the beam geometry must be known. This is the same geometry as used by NUBEAM, and includes the coordinates of the beam source grid and the unit vector defining the beam centreline, along with the shape (rectangular or circular) and dimensions of the source grid. Additionally, a beam focal length and divergence are defined. Marker particles are instantiated at random locations on the source grid, and are given velocity vectors with a direction dependent on the beam focal length. In a beam with zero divergence (not feasible in reality) all markers would travel through the focal point. FIDASIM introduces a small random component in the direction vector, dependent on the magnitude of the beam divergence.

A 3D grid, referred to as the beam grid is defined to have its x -axis along the beam centerline, is defined to encompass the extent of the beam inside the plasma volume, including the cloud of halo neutrals around the beam. The plasma and field parameters, and the fast-ion distribution function are defined for all of the cells. It is with these that the collisional-radiative model can be applied to the flux of markers from the beam. The model tracks the population of different energy levels of the markers as they undergoes various processes including electron and ion impact excitation and de-excitation, ionisation and crucially, charge-exchange between the injected neutral and the fast-ion population. It also does this for the derived halo neutrals as well.

The specific geometry of the FIDA system is given, with lens location, direction vectors and a line-of-sight radius defining each line-of-sight present as cylindrical volumes through the plasma. The calculated emission from the transition $n = 3$ to $n = 2$ in the intersection volume of the line-of-sight with the neutral beam, emitted along the direction of the line-of-sight then gives the synthetic FIDA spectral radiance returned by the code. This is also performed for other spectral components as shown in Figure 2.2.

2.3 Estimating error in FIDASIM synthetic spectra using deterministic sampling

The propagation of error from the input of a simulation to the output can be difficult due to the highly non-linear nature of the transformation of the inputs to outputs. This is especially true for TRANSP/NUBEAM simulations, due to the time needed to perform a single run (which is typically of the order of a day's worth of compute time, slightly more in real time). The standard, Monte Carlo, method of propagating uncertainties would be to randomly sample from the uncertainty distributions of the uncertain input variables then running the simulation, doing this many times with different samples (many thousands are not uncommon) [67], then determining the mean and variance of the output of each. This is not feasible with the TRANSP/NUBEAM to FIDASIM workflow. However, alternative methods have been developed to reduce the number of the required simulations while still giving reasonable estimates of the output uncertainty, one of which is known as Deterministic Sampling (DS). The use of DS has already been demonstrated in the case of TRANSP simulations by Sahlberg *et. al.* [68]. Additionally, a report by Sahlberg [69] gives an in-depth description of the technique and its effectiveness. A brief description of the technique based on [68], [69] will be given below.

Instead of sampling randomly, specific samples and sample weights are carefully chosen to represent the underlying statistical moments of the distributions of the input variables. These samples are known as *sigma-points*. Together the weights $w^{(i)}$ and sigma-points $q^{(i)}$, where i denotes the specific sigma-point, constitute the *ensemble*. Technically, the ensemble need only represent the first and second moments (mean and variance) to provide an estimate of the error, but including higher moments like skew and kurtosis generates more accurate values, especially for non symmetric uncertainty distributions, and would allow for deeper statistical investigation, if desired, of the output distribution. In theory, ensembles can be generated for many variables with different types of error distribution (and by en-

coding mixed moments correlations between variables can also be accounted for). However, increasing the number of variables increases the number sigma-points required, and increases the difficulty in selecting appropriate weights and sigma-points. The Sahlberg report [69] gives an in-depth description of the selection of sigma-points for multi-variable ensembles and how the weights may be algorithmically generated. The simulation is run multiple times, using each sigma point. The results of each, $S(q^{(i)})$ can then be combined along with the weights to generate the output expectation value, E , and variance σ^2 as follows [68]:

$$E = \sum_i w^{(i)} S(q^{(i)}) \quad (2.7)$$

$$\sigma^2 = \sum_i w^{(i)} \left(S(q^{(i)}) - E \right)^2 \quad (2.8)$$

To investigate how DS may be used with FIDASIM, a two variable case propagating the error in the TS generated electron temperature, T_e , and density n_e was performed for shot 29904 (on MAST), with the goal of generating a general characteristic uncertainty for the synthetic spectra. This shot was chosen as there was an extended period of quiescent plasma, approximately between 0.245 s-0.268 s, and a high-quality TRANSP run had already been constructed in the course of a previous investigation by Marco Cecconello *et. al.* [70]. For the FIDASIM run the time 0.265 s-0.268 s was examined, as this was near the end of the quiescent period.

For this example it was assumed that the error in each of the variables was uncorrelated, which is not strictly true as they derive from the same fitted data. Correlations can be encoded [69], but is difficult for time and space series data, and is beyond the scope of this thesis. Electron density and temperature were chosen as the fast-ion slowing down is primarily dependent on the interaction of the fast ions with the electron population (see 1.5), so would potentially have a large impact on the synthetic spectra. Extending the analysis to more TRANSP input variables, like ion temperature and plasma rotation, could be an avenue for

future investigation. It was assumed that the underlying uncertainty in each variable was Gaussian. This is significant, as the selection of sigma-points and weights is significantly simplified in this case (with non-symmetric distributions most difficult, then symmetric non-Gaussian, then most simplified is the Gaussian). For more accurate output statistics it is still important to encode the input skew and kurtosis, but for a Gaussian distribution this is always 0 and 3, respectively.

In this case a suitable ensemble can be generated with the Corner-method documented in [69], section 3.6.2.. Here, the entire T_e and n_e profiles are shifted up and down in different combinations, and simulations ran with the shifted profiles. This can be represented by a matrix, q , where each row represents a sigma-point, and each column the value of one of the variables being examined for that sigma-point:

$$q = \sqrt{3} \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ -1 & 1 \\ -1 & -1 \\ 0 & 0 \end{pmatrix} \Sigma + M \quad (2.9)$$

where:

$$\Sigma = \begin{pmatrix} \sigma_{T_e} & 0 \\ 0 & \sigma_{n_e} \end{pmatrix} \quad (2.10) \quad M = \begin{pmatrix} T_{e0} & n_{e0} \\ T_{e0} & n_{e0} \\ T_{e0} & n_{e0} \\ T_{e0} & n_{e0} \\ T_{e0} & n_{e0} \end{pmatrix} \quad (2.11)$$

where σ_{T_e} and σ_{n_e} are the uncertainty of T_e and n_e respectively, and T_{e0} and n_{e0} are the original T_e and n_e profiles respectively. The weights, w , for these sigma-points are [69]:

$$w = \left(\frac{1}{12} \quad \frac{1}{12} \quad \frac{1}{12} \quad \frac{1}{12} \quad \frac{2}{3} \right) \quad (2.12)$$

TRANSP takes in profiles fitted to the experimental data in OMFIT [63], so representative relative uncertainties are selected for both variables. For this investigation, these were selected by examining the relative uncertainty of the Thomson data during the quiescent period and over the spatial extent of the measurements. It was observed that over the greater part of the radial range, $\sim 0.5 \text{ m} - 1.1 \text{ m}$, the relative error of T_e and n_e did not typically exceed 10% and 5%, respectively. From $\sim 1.1 \text{ m} - 1.3 \text{ m}$ and $\sim 0.4 \text{ m} - 0.5 \text{ m}$, this increased for T_e to $\sim 12\%$, whereas for n_e it remained below 5%. It is noted that these are maximal values of the uncertainty, with a wide variability below that. For both T_e and n_e relative uncertainty rapidly rose for measurements at radii $R < 0.4 \text{ m}$ and $R > 1.3 \text{ m}$, as the absolute values of both measurements rapidly decrease at the edges of the plasma. It was decided to use a spatially uniform 10% uncertainty for the T_e measurement, and a spatially uniform 5% for the n_e measurement. In the future, introducing some spatial variability may be prudent for more detailed uncertainty analysis, but for generating a general characteristic uncertainty was deemed unnecessary.

With this, five TRANSP runs were prepared and ran, using the original run performed during [70] as a template, modified to reflect the sigma points as shown by Equation 2.9. As an example of the output the neutron rate generated in each run and the derived expectation value and standard deviation are shown in Figure 2.3.

Figure 2.3(b) shows the expectation value and standard deviation derived from the traces shown in Figure 2.3(a) and Equations 2.7 and 2.8. Additionally, the result of the TRANSP run with sigma-point $q^{(5)}$ (the original, "default", run) is shown as a comparison. In this case, there is a small difference between the default and expectation value, but well within the derived standard deviation.

The results of the five TRANSP runs, with the fast-ion output averaged over 0.265 s-0.268 s as stated above, are then propagated through FIDASIM. The results of each run for an example channel at $R = 1.072 \text{ m}$ are shown in Figure 2.4(a), with the calculated expectation value and standard deviation shown in Figure 2.4(b) (again with a comparison to the unmodified run, $q^{(5)}$).




Figure 2.3: Synthetic neutron rates predicted by TRANSP for MAST pulse #29904. In (a), the predicted neutron rates are shown for each sigma-point, with the $q^{(i)}$ corresponding to those given in Equation 2.9. In (b) the resultant expectation value (given by Equation 2.7) and standard deviation (given by the square root of the variance given by Equation 2.8) is shown, along with a comparison with $q^{(5)}$, the simulation using the unmodified, "default" values of T_e and n_e .




Figure 2.4: Synthetic spectra predicted with FIDASIM, using the results of the five TRANSP runs. In each sub-figure, only the region of the spectrum which was dominated by FIDA emission is shown. (a) shows the synthetic spectra derived from each sigma-point $q^{(i)}$, corresponding to those given in Equation 2.9. (b) shows the calculated expectation value (given by Equation 2.7) and standard deviation (given by the square root of the expression in Equation 2.8) is shown, along with a comparison with $q^{(5)}$, the simulation using the unmodified, "default" values of T_e and n_e .

In order to determine some characteristic uncertainties for the FIDASIM intensities, the percentage uncertainty of the synthetic spectra across multiple channels is examined. Figure 2.5 shows the percentage uncertainty as determined by deterministic sampling across a range of channels, from $R = 0.936$ m, the first channel outboard of the magnetic axis (which is located at approximately $R = 0.92$ m at

the time of the simulations, based on the archived EFIT data), to $R = 1.269$ m, the final channel with $R < 1.3$ m, beyond which the Thomson error largely diverges from that used in the sampling.

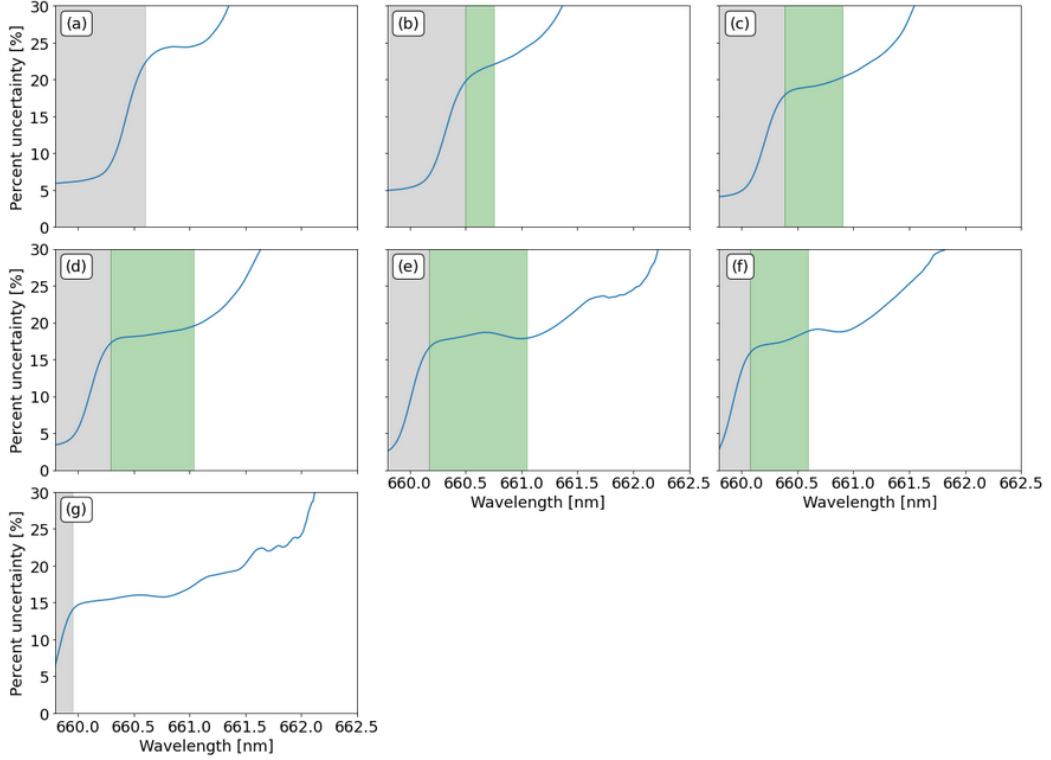


Figure 2.5: Percentage uncertainty for synthetic spectra as determined by deterministic sampling, for pulse #29904 averaged over 0.265 s-0.268 s. Each subplot shows a different channel, with (a) $R = 0.936$ m, (b) $R = 0.982$ m, (c) $R = 1.027$ m, (d) $R = 1.072$ m, (e) $R = 1.139$ m, (f) $R = 1.205$ m and, (g) $R = 1.269$ m. The plots focus on the region of the spectra dominated by FIDA emission, with the greyed out region showing where full energy beam peak emission dominates instead. The green region designates the region where the absolute value of the FIDA emission was above 1×10^{15} $\text{Ph s}^{-1} \text{ nm}^{-1} \text{ m}^{-2} \text{ sr}^{-1}$, the approximate value of the uncertainty on the experimental FIDA data at this time.

The position of the full energy beam peak is different in each channel, and the greyed out region denotes the wavelength range that is dominated by the beam peak emission. This should be ignored for the purpose of analysing the uncertainty in the FIDA data. In each channel, the percentage uncertainty asymptotically increases towards higher wavelength, which is a consequence of the FIDA radiance falling to zero. The focus is therefore directed towards the region of the spectra

with a "useful" level of spectral radiance. The error on the experimental FIDA data was $\sim 1 \times 10^{15} \text{ Ph s}^{-1} \text{ nm}^{-1} \text{ m}^{-2} \text{ sr}^{-1}$ at the time of the simulation, so the region of the synthetic spectra with spectral radiance above this value was classed as "useful", and in Figure 2.5 was highlighted in green. It can be seen that only channels with $R = 0.982 \text{ m} - 1.205 \text{ m}$ (Figures 2.5(b) - (f)) contained significant regions of "useful" spectra. The average value of the percentage uncertainty over the "useful" region for these channels was 21.1% , 19.2% , 18.5% , 18% , and 17.4%. A useful approximation for the uncertainty in the synthetic spectra in general would therefore be to assume a 20% error for the region dominated by FIDA emission. This value is adopted for all future comparisons of synthetic data to experimental in the rest of this thesis.

2.4 Aside: Solid State Neutral Particle Analyser

In the context of this thesis it is relevant to briefly describe the Solid State Neutral Particle Analyser (SSNPA) due to its highly complementary nature with FIDA. At the end of the charge exchange process described in Section 2.1.1, a fast neutral exists. Emission from quantum state $n = 3$ to $n = 2$ generates the photons detected by the FIDA diagnostic. The fast neutral is unconfined, and so leaves the plasma. The detection and analysis of this neutral is the focus of the Neutral Particle Analyser (NPA). In a conventional NPA [71] the neutral enters the diagnostic and is then re-ionised. Within the NPA, parallel electric and magnetic fields, perpendicular to direction of incoming ion, deflect the particle's motion based on its mass and energy. The ion then impinges on a detector, and the position can be used to infer the particle mass and energy. Such devices give good resolution of both mass and energy, but are large and so it is impractical to implement numerous channels into the plasma. The SSNPA [72] instead uses an array of photodiodes to detect the incoming neutral particle. On MAST-U 3 stacked arrays of 15 photodiodes view the plasma, allowing a radial profile to be discerned (with a 16th fully masked

channel to allow for background noise to be determined). Each array uses a filter foil on the photodiodes, with each array using a different thickness. The energy needed to penetrate the foil increases with thickness, and so the SSNPA is able to have a restricted form of energy discrimination of the incoming neutrals. On MAST-U the SSNPA lines-of-sight are directed much more radially compared to the mostly tangential views of the FIDA system. The sensitivity of the SSNPA is therefore much more towards trapped fast ions compared to the FIDA, making it highly complementary to the extant FIDA system. However, no SSNPA data is utilised within this thesis, primarily due to the system lines-of-sight radii being unsure until after the conclusion of the first experimental campaign, and the more subtle nature of the SSNPA analysis, as there is no "passive" SSNPA array present. This makes analysis more involved in order to discern the active signal. During the second campaign, it would be ideal to work with the SSNPA team to allow for a more complete analysis.

The FIDA Diagnostic: setup and preparation for the MAST-U experimental campaign

This chapter focuses on the practical parts of the diagnostic, describing the overall setup of the fibres, spectrometer and camera. The calibration procedures are briefly described, along with a summary of the data processing necessary to convert a raw signal in counts per pixel into a calibrated signal of *photons/s m² sr nm*. The work performed during the preparation for the first experimental campaign on MAST-U is described, specifically the calculation of the transmission of the optical fibres from the tokamak, and the validation of the fibre alignment. A brief discussion of the feasibility and benefits of a vertical view on MAST-U for a future campaign is given.

3.1 FIDA setup on MAST-U

The following sections describe the FIDA system on MAST-U, generally following the order of components that a photon emitted by a FIDA Charge Exchange (CX) reaction would follow through the various fibre optics, through the spectrometer,

to the charge coupled device (CCD) camera.

Figure 3.1 shows a schematic of the overall diagnostic system. Section 3.1.1 describes the fibre head that attaches the fibres to the tokamak and the fibres that join the tokamak to the patch panel (Section 3.1.1.2). The spectrometer fibres 3.1.1.3 join the patch panel to the spectrometer. The spectrometer (Section 3.1.2) diffracts the light, which passes through the mask (Section 3.1.3) to remove the full energy beam peak. The light passes through de-magnification lenses (Section 3.1.4) and is then imaged into the camera (Section 3.1.5). The raw data can then be processed (Section 3.2).




Figure 3.1: Schematic diagram of the FIDA setup, which is described in detail in this section. Light from the tokamak enters fibres in the fibre head, which travels to the patch panel. Light travels through the spectrometer fibres chosen for examination into the spectrometer. The diffracted light passes through the mask, blocking the full energy beam peak, before being imaged onto the camera CCD chip via the de-magnification lenses. After the raw data is recorded it can be processed for analysis.

3.1.1 Optical fibre setup

3.1.1.1 Tokamak fibres and viewing geometry

The spectrometer can accept a total of 22 optical fibres at any one time - 11 active and 11 corresponding passive channels.

On MAST 5 views were available, necessitating 5 separate sets of optical fibres. Hereafter each set of fibres will be designated Bundle X, where X is a number 1 through 5. On MAST Bundles 1 and 2 were used on active vertical and toroidally

displaced passive vertical views. These views are not present on MAST-U, and so these Bundles were not used during the first MAST-U experimental campaign.

On MAST-U 3 views were utilised, which correspond to Bundles 3 through 5. Bundle 3 corresponds to a toroidally directed active view of 32 fibres, of which a chosen subset of 11 fibres are selected for analysis. Bundle 4 consists of 32 fibres used with a vertically displaced passive view. This view was aligned and calibrated for fibre transmission (see sections 3.4 and 3.5) but was not utilised during the first MAST-U experimental campaign. Bundle 5 was the primary passive view, which consists of a toroidally displaced set of 32 fibres, of which the 11 corresponding to the active fibres are chosen (section 3.5.4 discusses further how this subset is chosen).

Figure 3.2 shows the top-down view of the FIDA viewing geometry available during the first experimental campaign on MAST-U. The red and blue hatched regions show the extent of the views available for the active and passive views respectively. It should be noted that the SS beam lies in the device midplane, while the SW beam is vertically displaced upwards by 0.65 m. This ensures that the active views only intersect the SS beam, and not the SW beam. On MAST, the SW beam lay in the midplane, which led to beam emission peaks from the SW beam contaminating the spectra when examining inner channels. Thankfully this is avoided on MAST-U.

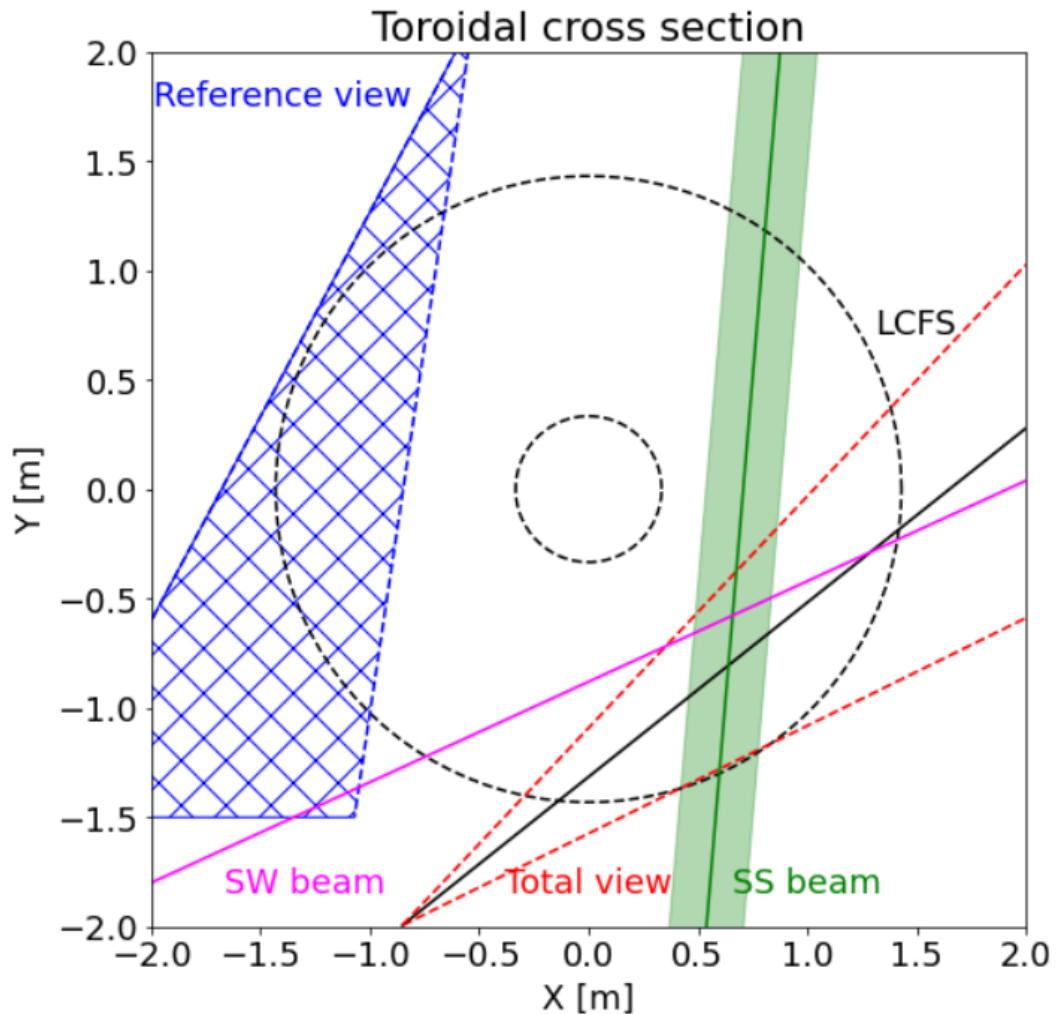


Figure 3.2: Top-down view of the spectral and beam geometry present on MAST-Upgrade (MAST-U) during the first experimental campaign. The extent of the passive views is shown with the blue cross hatch, the extent of the active views is bounded by the red dashed lines, with an example active channel shown in black. This intersects the SS NBI beam (which lies in the midplane and is directed towards increasing Y), with the SS beam centreline shown in dark green, and with the spatial extent of the beam shown in light green. Note that this spatial extent is an approximation, based on the horizontal width of the beam aperture (34 cm). The beam has a focal length and non-zero divergence, and will become wider towards increasing Y . Also shown is the SW beam centreline (magenta), but the SW beam is vertically displaced upwards by 0.65 m, ensuring no contamination of the inner channels. The SW beam of course also has a finite width, but is not shown here for figure clarity.

The fibres are mounted to the tokamak in two separate optics assemblies (assembled by Christopher Wade and Effy May-Smith of the UK Atomic Energy Authority (UKAEA)), referred to as 'heads', shared with the CXRS-Upgrade and Celeste III systems [73]. The 64 fibres of the active and vertically displaced passive views share one head, mounted on sector 7. As mentioned previously, the vertically displaced passive view was not utilised in the first MAST-U experimental campaign, so this head is referred hereafter as the 'active' head. The other head, mounted approximately 90° clockwise from the active head in sector 12, containing the 32 toroidally displaced passive views will be referred hereafter as the 'passive' head.


Figure 3.3 shows the arrangement of fibres within the active (3.3(a)) and passive (3.3(b)) heads.

During assembly, the fibres are further subdivided into bundles of 4, the edge fibres of which are denoted by the red and green dots in Figure 3.3. A individual fibre is therefore specified by three characters: 3,4 or 5 representing either the active, vertically displaced passive or toroidally displaced passive respectively; a letter, A through H inclusive, designating sub-bundle, and finally a number 1 through 4 inclusive to designate an individual fibre in the sub-bundle. For example, in figure 3.3(b) the rightmost FIDA fibre would be 5A1, moving to the left there would be 5A2, 5A3, 5A4, 5B1, and so on until the leftmost fibre, which would be 5H4.


The fibre optics leave the head, protected by a plastic conduit, and are fed through penetrations in the blockhouse wall into the optics laboratory. Here, they terminate in the patch panel, described in Section 3.1.1.2.

3.1.1.2 Patch panel

The patch panel consists of a metal panel with 191 fibre optic connectors. Figure 3.4 shows an image of the patch panel, with the fibre connectors of the termini of the 96 tokamak fibres highlighted, with blue for Bundle 3, orange for Bundle 4, and green for Bundle 5.



((a)) View of fibre ends as secured in the active fibre head. Note that the upper FIDA fibres are the active fibres, the lower are for the vertically displaced passive view not used in the MAST-U campaign. Adapted from [73].



((b)) View of fibre ends as secured in the passive fibre head. Adapted from [73].

Figure 3.3: Close-up view of the ends of the FIDA optical fibres as mounted in the active and passive optical head assemblies, along with the Celeste-III and CXRS-U fibres that share the same head. Adapted from [73].

	Channel Number										
	0	1	2	3	4	5	6	7	8	9	10
Active	3A1	3A4	3B3	3C2	3D1	3D4	3E3	3F2	3G1	3G4	3H3
Passive	5A1	5A3	5B2	5C1	5C4	5D2	5E2	5F1	5F4	5G3	5H2

Table 3.1: Pairings of the tokamak fibres used for the active and passive FIDA, for the specific channels in the spectrometer.

In theory, the presence of the patch panel allows for the tokamak channel selection to be changed between shots. However, the action of removing and re-inserting fibre optics into the panel can lead to the introduction of contaminants into the optical path, degrading the overall transmission. During the first campaign on MAST-U the selected fibers were not changed. Table 3.1 shows the specific tokamak fibres chosen for the first experimental campaign on MAST-U.

The methodology behind which active fibre is paired with which passive fibre is discussed in Section 3.5.4.




Figure 3.4: Patch panel as set up for the first MAST-U experimental campaign. Highlighted regions show the terminal fibre connectors for the tokamak fibres, with blue for Bundle 3, orange for Bundle 4, and green for Bundle 5. Fibre connectors not highlighted include those for Bundles 1 and 2, those for additionally calibration fibres for FIDA and a few for the BES diagnostic.

3.1.1.3 Spectrometer fibres

On MAST-U there are 2 sets of 12 fibres available to take light from the patch panel to the spectrometer, labelled B and C, 1 through 12. A third set of 12 are present, labelled A, but these were only utilised for experiments during the operation of MAST. The termini of these fibres are mounted within an aluminium block, arranged in 3 curved vertical columns. Curved columns are necessary to compensate for optical effects, described in more detail in Section 3.1.2, that induce a curvature in the spectrometer output from light that enters the spectrometer away from the optical midplane, which would cause the spectra from fibers not in the optical midplane to be offset, the effect of which would increase as the distance

from the midplane is increased. Choosing the correct input curvature compensates for this, ensuring that the spectra from each channel is (for the most part) aligned on the CCD. During the first MAST-U experimental campaign only fibres B1-B11 and C1-C11 were used, as there was not enough space on the camera's CCD chip for the 12 channel to be fully captured, however B12 and C12 were included in the calibration.

3.1.2 Spectrometer setup

The spectrometer setup is functionally unchanged since its operation during the experiments on MAST. A detailed account can be found in [74]. A brief summary will be given here.

The spectrometer used on MAST-U is a Kaiser Optical Systems *HoloSpec f/1.8* transmission grating spectrometer, where $f/1.8$ refers to the f-number of the system. This utilises a Volume Phase Holographic (VPH) grating of spatial frequency $\nu = 2144$ lines/cm. Spectrometers of this type are discussed thoroughly by R. Bell [75]. Figure 3.5 shows the schematic representation of the spectrometer.




Figure 3.5: Schematic diagram of a Volume Phase Holographic (VPH) spectrometer. Adapted from [75].

The associated grating equation for this type of spectrometer is [75]:

$$\lambda\nu = \cos(\gamma) [\sin(\theta_1 + \phi_1) + \sin(\theta_2 + \phi_2)] \quad (3.1)$$

where, in both Equation 3.1 and Figure 3.5, λ is wavelength, γ is the vertical angle from the optical axis at the lens to an off-axis ray, $\theta_1 = 40^\circ$ is the grating angle of incidence on the optical axis, $\theta_2 = 50^\circ$ the angle of diffraction at the central wavelength, ϕ_1 is the horizontal angle to the entrance slit, and ϕ_2 is horizontal angle to the diffracted light from the optical axis. Additionally, $f_1 = 82.5$ mm and $f_2 = 85$ mm are the focal lengths of the input and output lenses respectively, x_1 and x_2 are the horizontal distances to the ray in the entrance and image plane respectively, and y is the vertical distance to the input ray. The central wavelength of the this setup can be derived by setting $\gamma = \phi_1 = \phi_2 = 0$, and is 657.1 nm. The dependence on γ in Equation 3.1 introduces a curvature into the spectra output [75]. By using a curved input slit of the appropriate curvature, this effect is compensated for, resulting in a linear arrangement of the output spectra.

The low f-number of this system, compared to more traditional Czerny-Turner spectrometers which have typical f-numbers of f/5 to f/9 [75], means that there is a higher throughput, with more photons potentially captured from the coupled input fibre optics. Unlike Czerny-Turner spectrometers, transmission VPH grating spectrometers do not have a widely tuneable wavelength range, and the specific grating chosen determines the useful range of wavelengths. For FIDA this is no real disadvantage, as the FIDA spectrum is only observable over a small range of wavelengths. Within the spectrometer, before the grating, there is a holder for an optical filter. On MAST, due to having both vertical (primarily blue-shifted) and toroidal (primarily red-shifted) views available, there were 3 bandpass filters available. Two were used only on MAST, specifically one with the passband covering the blue-shifted side of the spectrum and a filter that passes both red and blue shifted light, for mixed mode operation. On MAST-U only the toroidal view was

used, so only a filter passing the red-shifted side of the spectrum was needed, with the centre of the passband at 660.6 nm and a full width half maximum (FWHM) of 7.6 nm. This effectively blocks the halo emission peak (see Section 2.1.2).

3.1.3 Mask

Spectral features observed by the FIDA diagnostic, such as the cold D- α peak and the beam emission peaks as discussed in 2.1.2, are significantly brighter than the light produced by the CX reactions between fast ions and beam neutrals that we are primarily interested in. On MAST-U the halo emission peak is blocked by the bandpass filter, but the beam emission peaks are not. Indeed, the full energy beam emission peak usually has a spectral radiance several orders of magnitude higher than the FIDA light itself. This can be problematic due to frame-transfer smearing, described further in Section 3.1.5. To combat this a mask can be placed into the intermediate image plane between the spectrometer optics and camera optics.

The mask consists of thin metal bars, orientated primarily in the vertical direction and on MAST-U a double bar mask is used to block the full and half energy peaks of the neutral beam. Figure 3.6 shows the mask against a white background on a binned camera image.




Figure 3.6: Binned camera image of mask with white background. Note that the vertical gradient in brightness is just an artefact of the lab lighting utilised.

As described in Section 2.1.2 the beam emission peaks exhibit lower amounts of Doppler shift as the examined channel moves further out in radius. The blocking bars are therefore angled to account for this. The overall position of the peaks are dependent on the beam energy, which can be varied shot-to-shot, but during typical experimental operation the beam energy varied between 63.2 keV-67.5 keV, which causes a variation in the wavelength position of the full energy peak of less than

0.15 nm. It is deemed acceptable to fit the bars to the peaks at any energy in this range, as any subsequent variation will not significantly degrade the effectiveness of the blocking.

The mask is mounted on a translation stage. To align the mask, an image of a spectrum of a typical discharge with NBI is taken, along with a test image with a white background place behind the mask. By translating the mask to match the troughs in the white background image with the correct peaks in the spectra, another white background shot can be taken, and confirmation of the fitting can be achieved. Figure 3.7 shows an example of this for a single channel using a test shot with white background, shot number -15531 (the negative simply denotes a test exposure of the FIDA camera, as opposed to a plasma shot), and a plasma shot with NBI, shot number 44922. The trough in the test shot is matched to the peak of the plasma shot, matching both the full and half energy peaks. This is repeated for each channel.

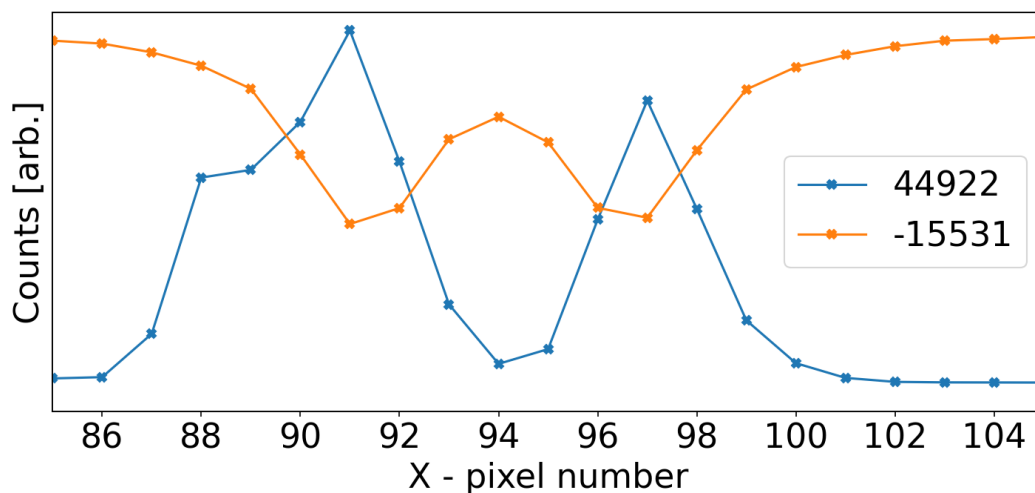


Figure 3.7: Example of mask fitting for a single channel, using white background test shot -15531, and plasma shot featuring beam peaks 44922. The two troughs present in the test shot are aligned with the full and half energy beam peaks. Magnitude of test shot has been scaled up to better show troughs.

At the start of the first experimental campaign the mask was not in place. Utilising the data from the unmasked shots and the procedure documented above, the mask was inserted. There was subsequent adjustment of the mask, which in the end gave

a reduction in the magnitude of the full energy beam peak. However, it is noticed that the masking is imperfect, especially when compared to select MAST data. It is thought that the mask may be slightly out of the image plane, resulting in a bad focusing of the mask onto the camera. Further investigation before the start of the second campaign to resolve this issue would be prudent.

3.1.4 Demagnification lenses

There is a final set of lenses that image the spectra and mask to the camera. The utilised lenses cause the 12th channel to be only partially captured on the CCD, so this channel is not utilised for data acquisition during experiments.

3.1.5 Camera

The camera setup is functionally unchanged since its operation during the experiments on MAST. A detailed account can be found in [74]. A brief summary will be given here.

A Photometrics Cascade 128+ 16 bit frame transfer CCD camera [76] is utilised as a detector. The absolute calibration of the CCD signal to convert from electrons to $photons/s\ nm\ m^2\ sr$ is documented further in Section 3.2.

The chip consists of two 128x132 pixel areas, with one area used for imaging and the other used for the read-out process. Of the 128x132 area only 128x128 pixels are actually available for imaging, with the remaining four pixel columns not exposed to light, and are used to determine the electronic offset (bias) of the CCD chip in each frame. After an exposure, the charge is shifted to the second, masked, 128x132 area from which it is subsequently read to memory. During this process light continues to be incident on the sensor, meaning that each row continues to accumulate charge, even as the charge is shifted into the unexposed area. This is known as frame-transfer smearing and is the primary reason the mask described in Section 3.1.3 is used, as the bright feature could cause significant charge to

accumulate in the short transfer window (approximately 80 ns per row). This could extend into the FIDA light regions of neighbouring channels, which is much less bright, and so could have a proportionally high impact. Within the spectra of a single channel the vertical direction contains no information, so sets of pixels in each column can be binned (summed) on a per channel basis. An example of a full frame image, before binning, can be seen in Figure 3.9, and a binned image shown in Figure 3.6 (these images are for different situations, but shows how the binning process changes the image). This further reduces smearing by making the transfer to the darkened area faster.

Fundamentally, the minimum uncertainty on the data acquired by this setup is determined by various noise sources in the detector setup. By the inherent quantised nature of light, the number of photons incident on the detector within some interval can be modelled by a Poisson distribution. This directly leads to the occurrence of shot noise. The shot noise is defined as the square root of the signal, which here is the number of electrons detected. A second type of noise is the readout noise, generated by the noise in the readout electronics. This is set to a constant value for the specific gain used. Finally, even on an unexposed pixel thermal energy can generate free electrons, which forms the 'dark current'. This noise can be reduced by cooling the CCD. The camera has built-in cooling that is able to keep the internals of the device around $-41\text{ }^{\circ}\text{C}$. At this temperature the dark current is negligible, and is not included in the determination of the signal noise.

3.2 FIDA data calibration

In order to process the raw data from the FIDA diagnostic into a useful form, extensive Interactive Data Language (IDL) code was written by Clive Michael of University of California, Los Angeles (UCLA), who is the current Responsible Officer (RO) of the diagnostic, and was the RO during the operation at MAST. The code was originally written for MAST, and some necessary additions were required

to allow it to process MAST-U data. A brief summary of the code is given below and Figure 3.8 shows the main steps taken by the code to process the data.

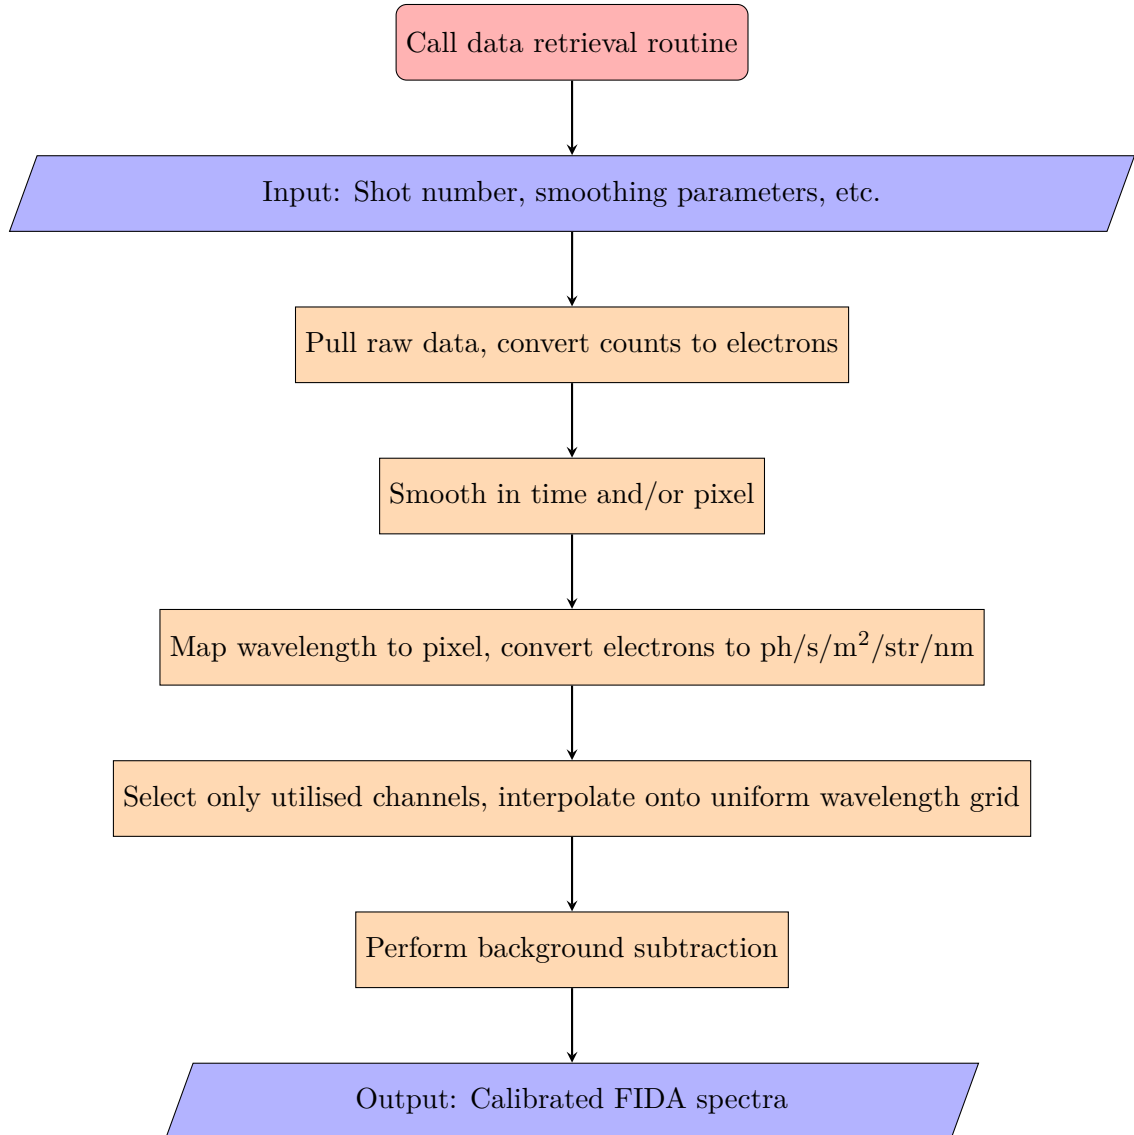


Figure 3.8: Flowchart describing FIDA data retrieval process

The user provides the shot number they want to load, along with parameters dictating any desired smoothing in time or pixel (wavelength). The default is no smoothing. The raw counts per pixel information is then retrieved from the database, and is converted to electrons per pixel by multiplying by the correct gain conversion factor. The smoothing is then performed, if requested. The mapping of wavelength to pixel for each channel is performed by simulating a synthetic camera

image, using the model described in Section 3.1.2 and by Equation 3.1, along with a calibration performed with an integrating sphere during the operation of MAST. At the same time the *etendue* of each channel is calculated. Etendue is a measure of the ability of an optical system to accept light, with units of $\text{m}^2 \text{sr}$. Dividing the signal by the wavelength bin size (range of wavelengths the pixel is capturing), the time step, the etendue and the transmission efficiencies of the various fibres converts the signal to $\text{photons/s/nm/m}^2/\text{sr}$. The transmission efficiencies were determined before the start of the first experimental campaign on MAST-U, the process of which is described in Section 3.4. The specific active and passive channels used are selected and paired, and the different channels interpolated onto the same wavelength grid, as due to the aforementioned distortion necessitating curved input slits the wavelength grid for each channel is very slightly different to each other. Then the final background subtraction can be performed, to give the 'difference' signal that, in theory, has eliminated all of the lower case components specified in Section 2.1.2.

3.3 Flux and Wavelength calibration

To have a fully functioning diagnostic both an absolute flux and wavelength calibration are needed. For MAST an integrating sphere was utilised to perform an absolute flux calibration, comparing the expected and received number of photons. As the spectrometer setup has not changed on MAST-U compared to MAST, this part of the calibration was not repeated. In the interim between devices, the optical fibres have been moved, removed and some repaired. For this reason, a full determination of the transmission efficiencies was performed, which is documented in Section 3.4. Due to a mechanical failure, the shutter that protects the active view window from the effects of deposition with the vessel was stuck partially open for some of the first MAST-U experimental campaign. This led to a further degradation of the overall transmission of the optical system due to the deposition

of material onto the optical window. Due to this and some additional factors first documented on MAST[45], an additional scaling factor of 2.35 was required to convert to absolute flux (an approximation found by combining the factor used in MAST [45], with an estimate of subsequent degradation of the window during the first campaign on MAST-U). Although this scaling factor is uncertain, it does not affect relative measurements made with the FIDA diagnostic.

The MAST-U wavelength calibration was performed by the RO, and used a hydrogen lamp. A single peak was observed after the light passed through the spectrometer, and the location on the CCD chip provided the offset for the simulation described in Section 3.2 to find the correct dispersion relation.

3.4 Determination of optical fibre transmission

In order to have an absolute calibration of the spectrometer the transmission of the optical fibres used must be known. The FIDA fibre optics ensemble consists of 2 separate sections that need to be calibrated separately: fibres from the fibre heads attached to the tokamak that terminate at the patch panel, and the fibres from the patch panel to the spectrometer.

3.4.1 Spectrometer fibres

3.4.1.1 A caveat

To perform a rigorous absolute calibration of the spectrometer fibres, as was performed on MAST[74], the spectrometer fibre output block would be placed at the output of an emitting integrating sphere, with the other ends of the fibres placed in a absolutely calibrated detecting integrating sphere. Integrating spheres are an optical device that can act as a evenly spatially distributed light source or as a detector of light emitted from a source in arbitrary directions. On MAST-U, only the emitting sphere was available. Instead the spectrometer-camera setup is

Fibre Set	Channel Number										
	0	1	2	3	4	5	6	7	8	9	10
Set B	0.604	0.623	0.635	0.65	0.629	0.629	0.619	0.548	0.586	0.539	0.512
Set C	0.549	0.611	0.617	0.593	0.648	0.65	0.63	0.62	0.584	0.544	0.56

Table 3.2: Transmissivity of fibres from patch panel to spectrometer. These are found by placing the individual fibre inputs into the output of an integrating sphere and binning the counts found in the CCD images produced. These are then normalised to the channel of maximum count, and scaled by 0.65. The scaling is based on the transmissivity known from the MAST calibration, taking into account the likely degradation over the intra-MAST/MAST-U period. The channel used to normalise is highlighted in orange.

used as a detector. This means that the following calibrations are inherently relative and not absolute, and are relative transmission values of the spectrometer fibre-spectrometer system, rather than just the fibres. In Section 3.4.1.2, some assumptions are made of the absolute transmission of the spectrometer fibres to generate an absolute calibration of the fibres.

3.4.1.2 Calibration of spectrometer fibres

The individual spectrometer fibre input ends are placed sequentially within the output of an integrating sphere, and full frame images are taken with the spectrometer camera. Additionally, a background image is taken, with all of the fibre inputs left in the ambient laboratory light conditions. This background is subtracted in the following analysis.

The counts for the specific channel in the image is binned, using the channel bins as described in Section 3.1.5. These are normalised to the channel bin with the highest counts, giving a relative calibration of the set of fibres.

To make an "absolute" calibration, it was decided to scale the the normalised calibration by 0.65, to match what was found during the calibration for MAST. It could be advantageous to verify this in the present day during a future inter-campaign shutdown. This gives the absolute transmissivities for the spectrometer fibres shown in Table 3.2.

3.4.2 Calibration of tokamak fibres

As described in Section 3.4.1.1, a lack of equipment meant that a direct absolute calibration of the tokamak fibres was not possible. However, with the "absolute" calibration given in Section 3.4.1.2 and Table 3.2 it is possible to get an estimate of the absolute transmission. The overall procedures to calibrate the active and the passive head are much the same, with the primary difference being that there are 2 bundles to calibrate in the active head (Bundles 3 and 4) compared to only 1 in the passive head (Bundle 5).

The head is detached from the tokamak, and affixed in such a way that the output of an integrating sphere is placed directly at the input to the fibres. Spectrometer fibre Set B is used exclusively with the active tokamak fibres, Bundle 3. Spectrometer fibre Set C is used for both passive views, Bundles 4 and 5. As the CCD camera can only fully fit the spectra of 2×11 fibres, 6 images are needed to calibrate the full complement of 96 tokamak fibres, 3 for each head. Fibres Sets B and C, 1-11 are utilised, with B12 and C12 not used. Tokamak fibres are connected to the spectrometer fibres in a sequence of "one on, two off", i.e. 3A1, 3A4,... until 3H3, which are connected to spectrometer fibre B11, B10,... and so on until B1. A full frame CCD image is taken, then the spectrometer fibres are shifted down one tokamak fibre, i.e 3A2 to B11, 3B1 to B10, and so on until 3H4 connects to B1, and another image taken. A final shift provides the final necessary image (for the active head). Note that in the final image, spectrometer fibre B1 is not connected. Figure 3.9 shows a full frame image of the first tokamak fibre configuration, with the designations of the spectrometer fibres labelled.

A noticeable feature of this image is the relative dimness of 4G1, which is spectrometer fibre C3. It is thought that this tokamak fibre is broken. This fibre was not used during the first MAST-U experimental campaign. A frame with no integrating sphere input is taken for background subtraction.

The spectra from each fibre are binned, and the ratio of the counts of the binned




Figure 3.9: Full frame CCD image of a subset of the light from FIDA fibres present in the active head. The labels at the top of the figure refer to which spectrometer fibre set is shown, and the numbers on the right to the specific spectrometer fibre number. The spectrometer fibre Set B are connected to Bundle 3, and Set C to Bundle 4. The tokamak fibres are arranged in "one on, two off", from the bottom working upwards are 3/4A1,3/4A4,... up to 3/4H3. It can be seen that tokamak fibre 4G1 (connected to spectrometer fibre C3), is notably dimmer than any other fibre. It is believed that that this tokamak fibre is broken.

tokamak fibre plus spectrometer fibre to the spectrometer fibre alone is found. This needs to be scaled to account for different integrating sphere luminance used during the tokamak fibre calibration versus the spectrometer fibre calibration. Table 3.3 shows the luminance recorded from the integrating sphere during each calibration. This gives scaling factors for each fibre ratio: Bundle 3 is 0.643; Bundle 4 is 0.646; and Bundle 5 is 0.528.

When this calibration procedure was performed for Bundle 5, it was noted that

	B	C	III	IV	V
Luminance [$k\text{ cd m}^{-2}$]	2.995	3.011	4.656	4.656	5.699

Table 3.3: Integrating sphere luminance used for the various fibre calibrations.

the transmission factors were significantly lower than expected, and much lower than those found for Bundle 3 and 4. The mean and standard deviation of the transmission factors for Bundle 3 was 0.614 ± 0.048 and for Bundle 4 was 0.579 ± 0.092 , but for Bundle 5 was 0.202 ± 0.018 . It is unknown what caused such a discrepancy. During the calibration weeks, the integrating sphere was utilised by a number of diagnostic teams, and is speculated that the either the luminance or sphere slit width was altered unbeknownst to the author. It was decided to scale the Bundle 5 transmission factors by 3. It would be prudent to repeat the fibre transmission during a future inter-campaign break to resolve this issue. Table 3.4 shows the final values for the tokamak fibre transmissivities used on the first experimental campaign on MAST-U.

3.5 Validation of fibre alignment

Once the fibre heads have been installed it is important to verify the alignment of the lines-of-sight, to ensure that they match the design to within some tolerance. An active fibre missing the neutral beam would of course be very problematic.

3.5.1 Back-lighting of fibres

To verify the alignment, back-lighting of the fibres was performed by the author and Clive Michael. A Canon 300D with EF-S 18 mm - 55 mm lens was affixed to an unoccupied (at the time of the alignment) port near the fibre head, directed towards the inner wall of the vessel opposite the head. For the active head the observation port was located in sector 7, for the passive head the port was in sector 12. Ideally, all of the fibres would have been back-lit from the patch panel simultaneously, but due to the layout of the patch panel and the space limitations

Specific fibre	Tokamak fibre bundle		
	3	4	5
A1	0.685	0.642	0.523
A2	0.688	0.694	0.67
A3	0.728	0.66	0.712
A4	0.671	0.655	0.706
B1	0.626	0.596	0.69
B2	0.714	0.621	0.702
B3	0.613	0.633	0.622
B4	0.629	0.584	0.636
C1	0.597	0.633	0.54
C2	0.634	0.628	0.609
C3	0.605	0.611	0.567
C4	0.593	0.609	0.618
D1	0.524	0.58	0.523
D2	0.623	0.549	0.512
D3	0.594	0.575	0.608
D4	0.585	0.624	0.532
E1	0.61	0.559	0.525
E2	0.598	0.549	0.605
E3	0.531	0.55	0.605
E4	0.555	0.595	0.575
F1	0.589	0.452	0.602
F2	0.651	0.569	0.616
F3	0.624	0.596	0.612
F4	0.612	0.608	0.557
G1	0.639	0.129	0.622
G2	0.603	0.538	0.604
G3	0.548	0.557	0.612
G4	0.587	0.565	0.655
H1	0.543	0.621	0.625
H2	0.583	0.548	0.637
H3	0.656	0.585	0.551
H4	0.609	0.6	0.664

Table 3.4: Transmission ratios used during the first experimental campaign on MAST-U for the tokamak fibres.

within the fibre cabinet two fibres at a time were illuminated by holding two torches directly to the fibre connectors. A long exposure (between 10 and 30 seconds) was then taken with the camera. The light from the illuminated fibres was visible on the inner tokamak wall. This process was repeated for all 96 fibres. However, due to having a restricted viewpoint, and the presence of various components within the tokamak, it was not possible to observe the illuminated spot from every fibre. Figure 3.10 shows an example of an image taken during this process, showing the spots of FIDA fibres 3C3 and 3C4.





Figure 3.10: An example of back-lighting of some FIDA and CXRS-U fibres, for the purpose of verifying fibre alignment. Here, 2 torches have been held to the optical fibre connectors for tokamak fibres 3C3 and 3C4 on the FIDA patch panel (which joins tokamak fibres to spectrometer fibres).

3.5.2 Finding the coordinates of the illuminated spots using Calcam


Calcam [77] is software developed for use in fusion diagnostics, that allows for the mapping of a two dimensional coordinate in a camera image to a three dimensional

coordinate in a Computer Aided Design (CAD) model of a tokamak.

A conversion model for both the active and passive must be calibrated. The model takes a 3D CAD model coordinate, converts to the camera frame by translation and rotation, then projects along the normal axis of the camera lens to generate a normalised pixel coordinate. This then must be corrected to account for the camera distortion present when using either rectilinear (used here) or fisheye lenses, then converted to an actual pixel coordinate by accounting for the effective focal length of the imaging system. In order to generate a Calcam calibration, images taken from the active and passive viewpoints are loaded into Calcam, along with a MAST-U CAD model. Prominent and defined features observable in both the CAD model and photograph are identified, and an associated pair of points are placed in the same 'location' in each. Well defined corners and holes are the easiest to utilise for this purpose. After defining multiple points spread across the calibration image, a fit can be determined that minimises the error when converting from 3D CAD coordinate to 2D pixel coordinate, and vice versa. From the calibration model fit, a wire-frame projection of the CAD model from the perspective of the virtual camera can be imposed on the real image, allowing for a visual inspection of the quality of the calibration. By specifying enough points accurately the mapping becomes well defined, allowing for an arbitrary pixel location in the images to be converted to the equivalent location in the tokamak. Figure 3.11 shows the calibrations used for the active (Figure 3.11(a)) and passive (Figure 3.11(b)) viewpoints. The centre of the fibre lights spots that were visible (i.e. not obscured by a feature within the vessel) were selected, and the corresponding CAD coordinates determined. These coordinates were then used to verify the alignment of the fibres mounted within the optical head, documented in Section 3.5.3.



((a)) Point pairs in Calcam for the active head calibrations.



((b)) Point pairs in Calcam for the passive head calibrations.

Figure 3.11: Point pairs, specified by the user in Calcam, are used to calibrate a model that converts between a 3D CAD coordinate and a 2D image pixel coordinate. The pairs define the necessary translation and rotation to move from the CAD lab frame to the camera frame, and the parameters of the distortion correction and the effective focal length of the imaging system to project the 3D location to the projected image location. This can then be used to extract the 3D coordinate for an arbitrary pixel coordinate.

3.5.3 Projection of light spots onto fibre head

In order to verify the alignment, the position of the fibre within the head is found by projecting the coordinate of the light spot back onto the head, which is compared with the design position. It is defined here that in the following analysis all coordinates are Cartesian unless otherwise stated, that the machine coordinates will utilise the notation $\{xyz\}$, while coordinates in the frame of the lens attached to the head will utilise $\{uvw\}$. The coordinates of the active and passive lenses (known from the MAST-U design documents) are shown in Table 3.5. The origin of the machine coordinates is located within the centre column of MAST-U, with $x = 0$ and $y = 0$ located on the column axis, and $z = 0$ defined at the geometric midplane of the device. Table 3.5 then defines the position vectors of the lenses in machine coordinates, \vec{OL}_{active} and $\vec{OL}_{passive}$ for the active and passive lenses respectively.

	x [m]	y [m]	z [m]
Active lens	-0.852	-1.994	0.302
Passive lens	-0.529	2.146	0.214

Table 3.5: Coordinates defining the position vectors of the active and passive lenses in machine coordinates, \vec{OL}_{active} and $\vec{OL}_{passive}$, relative to the origin (which is defined at the intersection of the centre column axis and geometric midplane). These are known from the MAST-U design documents.

The x and y coordinates for the active lens are similar on MAST-U to what they were on MAST. However, due to the presence of new features within the vacuum vessel that would obscure the old location, a periscope setup is used to avoid the new features, making the z coordinate on MAST-U greater than on MAST.

Also known from design specifications are the basis vectors for a coordinate system where the uv -plane is in the plane of the lens (and subsequently in the plane of the end of the fibres), with the w axis directed perpendicular, defined in terms of the machine basis vectors. These are shown in Table 3.6, with their projections in various frames shown in Figure 3.12.

Head	Lens basis vector	$\hat{\mathbf{e}}_{lens} = a \hat{\mathbf{e}}_x + b \hat{\mathbf{e}}_y + c \hat{\mathbf{e}}_z$		
		a	b	c
Active head	$\hat{\mathbf{e}}_u$	-0.56561629	0.8130726	0.13780843
	$\hat{\mathbf{e}}_v$	0.21222229	-0.01796995	0.97705618
	$\hat{\mathbf{e}}_w$	0.79689402	0.58188491	-0.16238803
Passive head	$\hat{\mathbf{e}}_u$	0.94489881	-0.31090941	0.10247721
	$\hat{\mathbf{e}}_v$	-0.13278402	-0.07787374	0.98808101
	$\hat{\mathbf{e}}_w$	-0.30085627	-0.94670507	-0.11504357

Table 3.6: Basis vectors of the lens coordinate system $\{uvw\}$ in terms of the basis vectors of the machine coordinate system $\{xyz\}$, for both the active and passive lenses. The values a , b and c are the coefficients for the machine basis vectors $\hat{\mathbf{e}}_x$, $\hat{\mathbf{e}}_y$ and $\hat{\mathbf{e}}_z$ respectively.

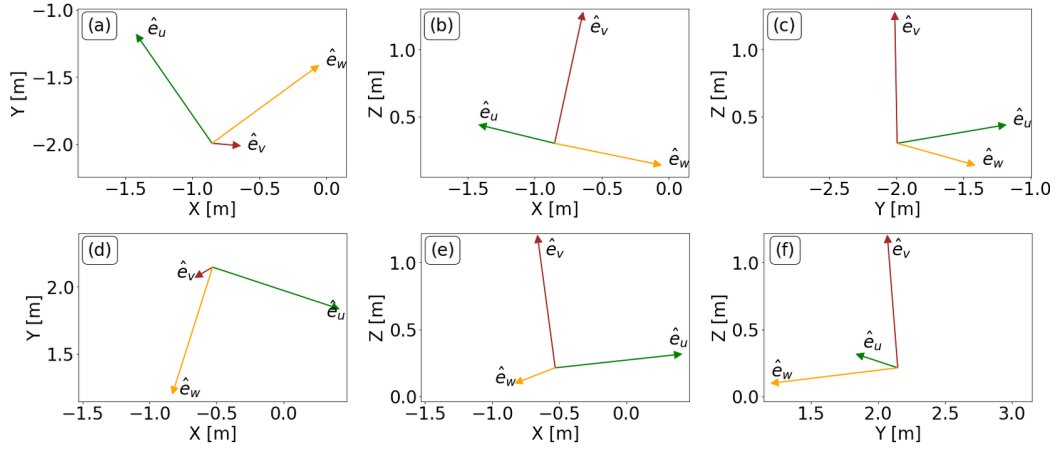


Figure 3.12: Unit vectors for lens frame in machine frame. Top row ((a), (b), (c)) for the active lens, bottom row ((d), (e), (f)) for the passive lens. (a) and (d) shows the projection of the unit vectors in the X-Y plane, (b) and (e) show the projection in the X-Z plane, and (c) and (f) show the projection in the Y-Z plane.

A simple vector subtraction, $\vec{L}\hat{S}_{machine} = \vec{O}\hat{S}_{machine} - \vec{O}\hat{L}_{machine}$, generates the vector from the lens to the spot, $\vec{L}\hat{S}_{machine}$, where $\vec{O}\hat{S}_{machine}$ and $\vec{O}\hat{L}_{machine}$ are the position vectors of the spot and the lens in machine coordinates, respectively. Figure 3.13 shows examples of the vectors used here for the active (Figure 3.13(a)) and passive (Figure 3.13(b)) lenses.

This can then be transformed into the lens coordinates by:

$$\vec{L}\hat{S}_{lens} = \underline{\underline{R}} \vec{L}\hat{S}_{machine} \quad (3.2)$$

where $\underline{\underline{R}}$ is a rotation matrix defined:

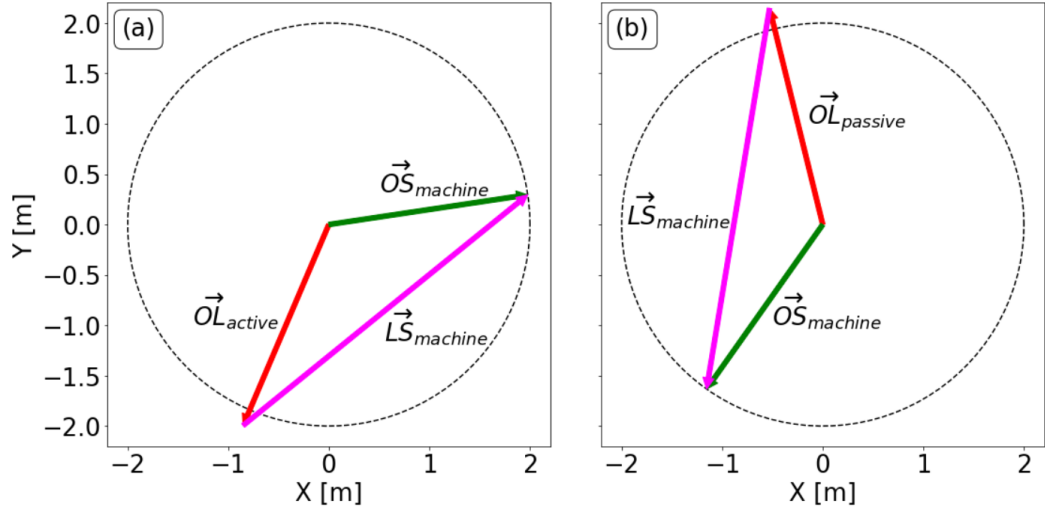


Figure 3.13: Top down diagram of the vectors discussed in this section, for (a) the active lens and (b) the passive lens. For each an example light spot was chosen to show $\vec{O}\vec{S}_{\text{machine}}$ and $\vec{L}\vec{S}_{\text{machine}}$. Note that the vectors all have a non-zero extent in the Z axis as well. The dashed black line represents the outer wall of the vacuum vessel. Both the active and passive lenses are present in ports, so are slightly recessed compared to the outer wall.

$$\begin{pmatrix} a_u & b_u & c_u \\ a_v & b_v & c_v \\ a_w & b_w & c_w \end{pmatrix} \quad (3.3)$$

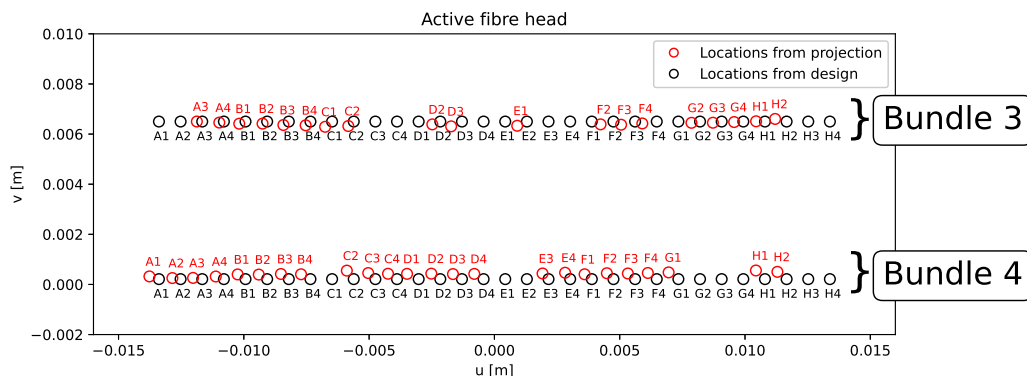
where a , b and c are the coefficients of the machine unit vectors \hat{e}_x , \hat{e}_y and \hat{e}_z respectively, with the subscript u , v or w denoting which lens unit vector the coefficient belongs to. These are defined in Table 3.6.

$\vec{L}\vec{S}_{\text{lens}}$ is a position vector from the new origin of the lens to the spot. The lens fibres all lie at $w = 72$ mm. Using the Cartesian form of a line:

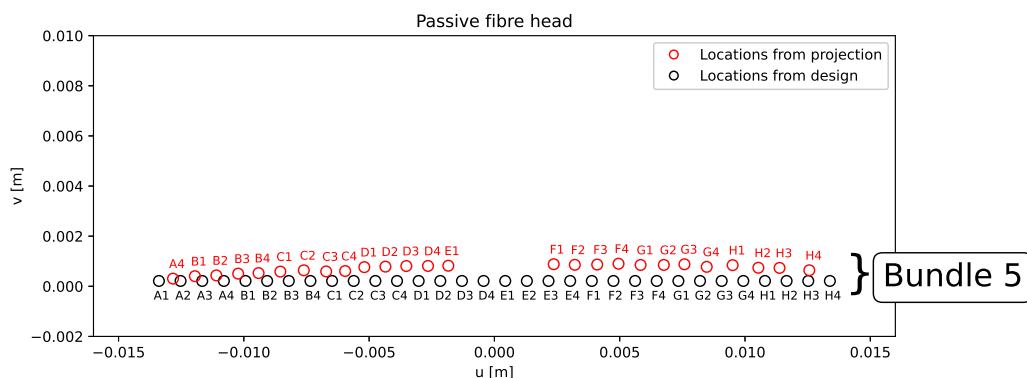
$$\frac{u - u_1}{\alpha} = \frac{v - v_1}{\beta} = \frac{w - w_1}{\gamma} \quad (3.4)$$

where u_1 , v_1 and w_1 are the coordinates of a point on the line (which here can just be $[0, 0, 0]$ as it is a position vector through the origin) and α , β and γ are the coefficients of the components of the vector $\vec{L}\vec{S}_{\text{lens}}$, the u and v coordinates of the projected spot can be determined.

Utilising the above method, we can generate the projected positions for the fibres and compare them to their design positions. It should be re-stated that only some of the light spots were not obscured within the tokamak, so projections could not be found for every fibre. Figure 3.14 shows the results, with Figure 3.14(a) showing the active head and Figure 3.14(b) showing the passive head.



((a)) Projection of location of active fibre locations on head, compared to design locations.



((b)) Projection of location of active fibre locations on head, compared to design locations.

Figure 3.14: Projection of location of fibre locations on active and passive heads, compared to design locations.

It can be seen that the agreement between the observed fibres and design predictions for the active head is fairly good. The inherent ambiguity of precisely selecting the centre of a diffuse light spot on a surface that is not normal to the direction of the incident ray means that 1-to-1 matches are difficult to obtain. For the passive, there is a larger discrepancy. However, it was known from the that due to a bowing of an aluminium plate in the passive head [73] that the row of fibres had an unavoidable arcing, which is reflected in the projected positions. However,

the is simply a reality of experimental design. The vertical displacement (in the \hat{e}_v direction) translates to a small displacement primarily in the machine frame Z direction, which should not have a large effect on the tangency radius of the specific channels. The horizontal displacement can be remedied by the specific selection of active-passive fibre pairs, documented further in section 3.5.4.

3.5.4 Comparing tangency radii

The 11 active fibres are chosen to have a reasonably regular spacing in examined radius, defined by the intersection of the line-of-sight and the neutral beam, spanning most of the range between the magnetic axis and last closed flux surface of a typical plasma. The 11 passive fibres are chosen to best match the tangency radii of the active channels.

The tangency radius of a channel can be found by finding the point along the line of sight that most closely approaches the central axis of the machine. The channel will form a tangent to a flux surface at this radius (assuming toroidal symmetry). It is very important to note that this is **not** the same as the radius at which a particular active channel is observing a FIDA signal, which is defined by the intersection of the channel and the neutral beam. However, the tangency radii are important to match active channels with appropriate passive channels, as channels with similar tangency radii will, under the condition of toroidal symmetry, experience similar plasma conditions along their length.

To generate the tangency radii, the Z-axis is ignored, and the point on the chord closest to (0, 0) is found. This is done by finding the projection of the vector between the lens and the machine origin on to the direction of the line of sight, using the information about the line vectors derived in section 3.5.3. The point found will be the closest to (0, 0) on the chord, and the radius for this point derived.

The default configuration is to pair up active fibres with their correspondingly designated passive fibre: 3A1 with 5A1, 3A2 with 5A2, etc. However, this must be

verified and changed as needed. As mentioned in section 3.5.2, only some of the light spots of both the active and passive channels were visible, due to being obscured from where the backlighting photographs were taken by components within the vacuum vessel. Therefore the tangency radii for only the subset of fibres whose backlit light spots could be observed was available for comparison. Based on a lack of complete data, some element of judgement is needed. The resultant tangency radii for the active and passive channels are shown in figure 3.15:




Figure 3.15: Comparison of the tangency radii for the active and passive channels that could be observed.

It can be observed that for channels with lower tangency radii, the default pairings work quite well. However, for channels with high tangency radii, the comparison becomes worse. It is important to note that it is for the high tangency radii channels that the correct pairing is most important, as these channels have lower active FIDA emission (due to less fast ions), so the effect of the passive subtrac-

tion becomes proportionally more important. With the information presented, the decision was made by the FIDA RO to shift the passive fibres down one position (pairing fibre 3H1 with 5H2 for example). However, some results from MAST-U (discussed in Chapter 6) suggest that the edge channels continue to have poor background subtraction. This issue should be thoroughly investigated in a future inter-campaign diagnostic calibration period.

3.6 Potential vertical FIDA on MAST-U

It is possible that a vertical FIDA view could be implemented for a future campaign on MAST-U. As shown in Figure 2.1 vertical views allow for sensitivity to the presence of trapped fast ions, which would be beneficial in studies of instabilities primarily driven by trapped fast ions, like fishbones, and would be complimentary to the phase-space observed by the current toroidally directed view. If the vertical view's lines of sight were designed to intersect the beam at the same radii as the toroidal view, this would open the possibility of performing velocity-space tomography, as was performed on MAST [54].

A possible location for the vertical active view's lens (subject to future decisions regarding the utilisation of MAST-U ports) is at machine coordinates $\mathbf{x} = 1.02\text{ m}$, $\mathbf{y} = -1.02\text{ m}$, $\mathbf{z} = 0.696\text{ m}$). The lines-of-sight would be directed downwards towards the SS beam. Assuming the choice of the same radii of observation as the existing toroidal view, the SW beam does not intersect any of the lines-of-sight. Figure 3.16 shows the lines-of-sight of the view.

Such lines-of-sight observe FIDA light that is primarily blue-shifted. A large advantage of such an orientation is that the Doppler shift of the beam peaks is much less, leaving much more of the FIDA emission exposed for analysis. Figure 3.17 shows what such a view would observe (with the same intersection radii as the existing view), using Shot #45091 at $t = 0.35\text{ s}$ as an example.

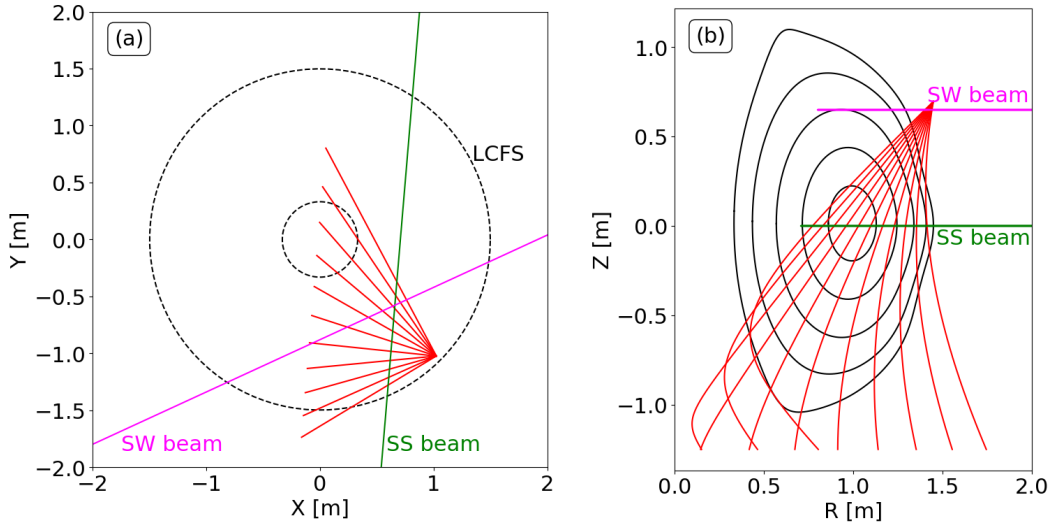


Figure 3.16: Lines-of-sight (red) for a potential vertical view examining the same radii as the existing toroidal view, showing the X-Y plane (a) and the R-Z plane (b). The centrelines of the SS beam (green) and the SW beam (magenta) are shown, along with exemplar flux surfaces.

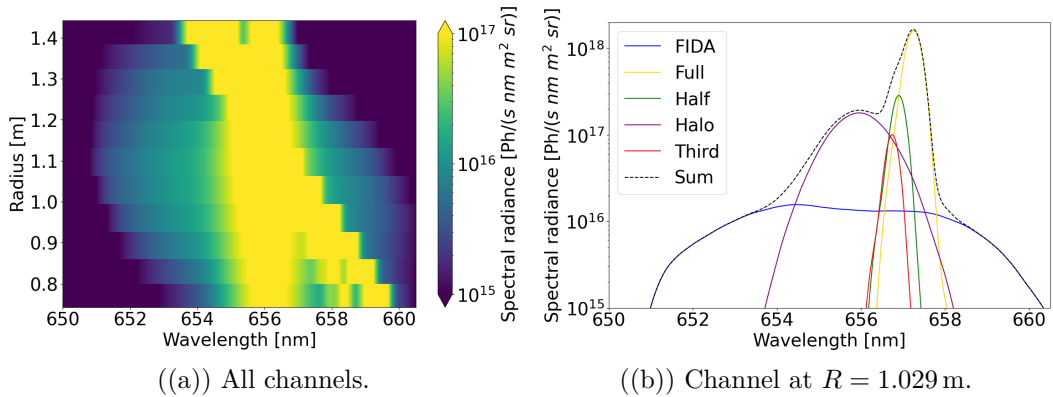


Figure 3.17: FIDASIM spectra generated based on Shot #45091 at $t = 0.35$ s for a potential vertical view that could be implemented for future campaigns. (a) shows all channels, (b) shows the individual spectral components for the specific channel at $R = 1.029$ m.

Figure 3.17(a) shows the spectra across all channels, with Figure 3.17(b) showing a specific example of the spectrum at $R = 1.029$ m, to show how the individual spectral components look. It can be seen that there is a large, unobscured region of blue-shifted FIDA light across a majority of the channels. Shot #45091 had a beam energy of 63.4 keV, and for the channel shown in Figure 3.17(b) emission from the injection energy down to a minimum energy ≈ 20 keV, which is a significantly larger range than that observable by the toroidal view, even in the toroidal outer

channels which have the smallest Doppler shift of the beam peaks. Such a view would require the careful selection of a passive view and a port at some toroidal angle displacement, chosen to ensure no intersection with the SW beam, would be ideal.

3.7 Conclusion

While the calculation of the transmission of the tokamak fibres was mostly successful, the ambiguity regarding the Bundle 5 measurements recommends a future set of measurements to more accurately determine the transmission. Ideally, a full absolute calibration utilising a pre-calibrated detector to examine the transmission of both the tokamak and spectrometer fibres could be performed to confirm the overall calibration, although the current calibration is sufficient. The backlighting confirmed that the overall positioning of the fibres was good, and the calculation of the tangency radii allowed the selection of the correct pairing of active and passive fibres. However, it can be seen that the pairing is worse at larger radii. The relatively small signal seen in edge channels means that this data is often not useful for analysis anyway. A vertical FIDA view would be quite desirable, allowing a filling out of the currently unobserved parts of the phase-space, and perhaps allowing tomographic inversion studies in the future. However, this view is still just a prospective examination of the advantages of such a system, and has to contend with one of the unfortunate realities of working on a spherical tokamak, specifically a lack of space for diagnostic ports. Such a system would need to commandeer two appropriate ports, one for the active view and one for the passive. As mentioned, this was done on MAST, but different priorities for experiments on MAST-U may make such a system infeasible for implementation for the near future.

Analysis of sawtooth crashes during MAST operation

This work was published in [78] and presented in [79]. This chapter exclusively focuses on the analysis of archival data taken on MAST, which was taken significantly prior to the practical work described in Chapter 3, which described work performed for setup for the campaign on MAST-U. The author had no involvement in the setup or running of the diagnostic when it was utilised on MAST.

4.1 Sawtooth crash

Sawtooth crashes are internal kinks with a dominant harmonic of $m = n = 1$, which are characterised by a sudden flattening of temperature and density profiles in the core of the plasma [80], and a small increase outside, defining an *inversion* radius at the transitional radius. The sawtooth instability and its interaction with the fast ion population is of particular interest to the development of ITER plasmas, as crashes can degrade the quality of the fast ion confinement or more generally cause redistribution in both configuration and velocity space, in particular an expulsion of fast ions from the sawtooth region [81]. ITER plasmas will rely on good fast ion confinement as the primary source of heating comes from the fusion α -particles.

The presence of fast ions can act to stabilise these internal kinks, delaying the onset of the instability. However, this delay can cause the mode amplitude to grow further, therefore inducing a more violent crash when it finally does occur. [82], [83].

4.2 Using FIDA to diagnose sawteeth

FIDA can be used to study the effect of sawteeth on fast ion confinement. The first such studies were carried out on the DIII-D conventional tokamak [25], [84]. Later the technique was applied to the investigation of fast ion behaviour during sawteeth in the ASDEX Upgrade conventional tokamak [53], and the inverse problem of determining sawtooth-induced changes in fast ion distributions from FIDA spectra has also been addressed in this device [85]. More recently, FIDA has been used in conjunction with a SSNPA to study the effects of sawteeth on fast ions in the NSTX spherical tokamak [86].

The MAST data investigated in this chapter have been previously examined by a number of researchers [54], [87], [88]. The study presented here extends the work carried out by Ceconello *et. al.* [88] and utilises TRANSP/NUBEAM data generated in the course of that study. Specifically, the author has extended the data analysis to include FIDA data, with FIDASIM synthetic spectra.

The structure of the rest of this chapter is as follows: in section 4.3 the models used to describe the effect of sawtooth crashes on the plasma are discussed; in subsection 4.4.1 the different models are compared using the experimental and synthetic FIDA data; in subsection 4.4.2 further analysis regarding the redistribution of the fast ions in the experiment is described and in section 4.5 some preliminary results regarding the analysis of passive FIDA emission are presented.

4.3 Sawtooth models

The Kadomtsev model [89] characterises the sawtooth crash as a process of rapid magnetic reconnection. The model specifies a helical sheet, defined as having the magnetic axis and $q = 1$ field line within it, shown in Figure 4.1.




Figure 4.1: Sheet that defines the helical flux used by the Kadomtsev model to determine the effect of sawtooth crashes. Flux passes through this sheet in opposite direction on either side of the $q = 1$ surface. The opposite fluxes of equal magnitude are reconnected during the crash. Adapted from [5].

The magnetic flux passing through this sheet, termed helical flux, is directed in opposite directions on either side of the $q = 1$ surface. The model states equal and opposite helical flux on either side of the $q = 1$ surface progressively reconnects, forming a magnetic island that grows to form the new magnetic core of the plasma, a process shown in Figure 4.2.

The new core is colder and less dense than the original hot core, which is dispersed towards greater minor radii. In this model complete reconnection occurs - directly after the crash $q \geq 1$ everywhere in the plasma. Experimentally this model has been shown to be not applicable to all plasma pulses. Porcelli *et al.* [90] describe a number of pulses, some of which follow the Kadomtsev model and several that do not. The authors then propose an alternative model (referred hereafter as the Porcelli model) which allows for incomplete reconnection, meaning that $q(r = 0) < 1$. Here the formation of the magnetic island begins as in the Kadomtsev model but




Figure 4.2: The progression in time of the newly formed magnetic island, due to the helical flux reconnection, displacing the original plasma core. In (i) the magnetic equilibrium is in an overall $q > 1$ state. In (ii), q drops below 1 in the core. In (iii) the magnetic surface on the right side of the core reconnects with the surfaces outside, connecting regions of equal and opposite helical flux. In (iv) and (v), the process continues, with the hot core being expelled, and the new cooler core forming. In (vi), the process is complete, with a new, cooler core formed, and the plasma in an overall $q > 1$ state. Adapted from [5].

is interrupted by magnetic turbulence once the island has grown to some threshold size. The Kadomtsev model moves plasma at $q = 1$ pre-sawtooth to the magnetic axis after the sawtooth, and other regions are mixed with helical flux-matching as described above. This has a flattening effect on the temperature and density profiles, but can cause some previous structure to be retained. Alternatively, ergodic mixing can be utilised, ignoring the requirement of flux-matching and redistributing the plasma in such a way that the temperature and density profiles are completely flattened within the mixing region. Henceforth the model implementing this complete flattening of profiles is referred to as the ergodic model. The Porcelli model specifies two regions that are mixed separately, based on how large the magnetic island is allowed to grow before it is interrupted. Once the new locations for the NUBEAM markers have been determined, the post-crash equilibrium combined with assumed (by the TRANSP module KDSAW, that governs the sawtooth models) conservation of magnetic moment and velocity parallel to the magnetic field determine the change in the energy and velocity of the markers [91]. It should be




Figure 4.3: The simulated fast-ion density in an azimuthal cross section of MAST before and after a sawtooth crash, with the flux surfaces shown in black.

noted that while conservation of magnetic moment is often justifiable (as it is an adiabatic invariant) there is no immediately obvious physical reason for assuming conservation of the parallel velocity. TRANSP/NUBEAM implements the models in the KDSAW module, and a full explanation of the behaviour of the models is given in the online documentation [92]. Figure 4.3 shows synthetic spatial distributions of fast ions on MAST (generated by TRANSP/NUBEAM) before (Figure 4.3(a)) and after (Figure 4.3(b)) a sawtooth crash.

The three models (Kadomtsev, ergodic and Porcelli) were used in TRANSP/NUBEAM runs previously analysed in [88] to provide simulated fast-ion distributions before and after each sawtooth in pulse #29880. These were then utilised by the author, with FIDASIM, to generate synthetic spectra, which are compared to the experimental FIDA data.

4.4 Analysis

The data used in this study derive from a series of three pulses (#29880, #29881, #29882) which were designed to be as similar as possible. Figure 4.4 shows the time evolution of each pulse in a variety of parameters. The sawtooth crashes are clearly seen in the neutron rate and core electron temperature, but for these pulses the electron density is largely unaffected by the instability. To analyse the FIDA data, time slices were taken before and after the third sawtooth in the overall sequence of sawteeth (which occurs between $t = 0.25\text{ s} - 0.26\text{ s}$ in all three shots). The EFIT equilibrium reconstruction yields a q -profile greater than 1 at the time of the first sawtooth, and is not low enough after the second sawtooth to trigger the Porcelli model. A q substantially less than 1 is required to trigger the 3 sawtooth models in TRANSP/NUBEAM, a condition that is met by the time of the 3rd sawtooth in the sequence. Additionally, there is a secular downward trend in the experimental FIDA radiance over the course of the pulses, so the third sawtooth provides the best available signal-to-noise ratio. Figure 4.5 shows a soft x-ray (SXR) tangential camera trace, for a line-of-sight with a tangency radius of 0.5285 m, before, during and after the selected single crash present in all three pulses. The SXR camera is uncalibrated, and the signal has units of volts. The soft X-ray emission can be regarded as a proxy for electron temperature.

The three pulses were performed with the same setup and it was observed that, in the sequence of crashes that occurs in all of the pulses, the selected crash is separated in time by a maximum of 7 ms between the pulses. This is significantly less than both the energy confinement time and the fast-ion slowing down time at around the time of the selected crash. From this it is concluded that the pulses are repeatable and that it is appropriate to average across pulses (after aligning the selected crashes in time) as well as averaging across time, thereby improving the signal-to-noise ratio. This technique was also used in the previous investigations [87], [88].




Figure 4.4: Time traces of plasma parameters for the three pulses used in averaging. (a) NBI power, (b) neutron rate, (c) core electron temperature, (d) core electron density.

4.4.1 Comparison of different sawtooth models with FIDA data

In the work of Ceconello *et. al.* [88] neutron camera data were insufficient to determine which of the three sawtooth crash models used (Kadomtsev, ergodic and Porcelli) provided the best fit to the data. It was speculated that FIDA data may help in this regard. Figure 4.6 shows a before- and after-crash timeslice of the FIDA data, averaged over 3 ms (1-4 ms before crash, 1-4 ms after the crash) and across the three pulses (with the crash times aligned), along with synthetic spectra generated using each of the sawtooth models mentioned above, for a channel observing the plasma at a radius of 1.205 m in the plasma midplane. The corresponding plasma minor radius ρ (defined as $\rho = \sqrt{\psi_N}$) is 0.569, 0.565 and 0.58 before the crash and 0.573, 0.57 and 0.584 after the crash for the Kadomtsev, ergodic and Porcelli models respectively. The synthetic data were only generated for pulse #29880. It is appropriate to do this for the reasons given at the beginning of this section, i.e. that the plasma conditions are closely replicated in the three pulses that are




Figure 4.5: Soft X-ray (SXR) trace (for a line-of-sight with tangency radius 0.5285 m) showing a sawtooth crash for shots #29880 (blue), #29881 (orange) and #29882 (green).

analysed.




Figure 4.6: Experimental and synthetic spectra showing only the spectral region where FIDA light can be observed at $R=1.205$ m in the plasma midplane. The predictions based on the three models are shown. Based on Section 2.3 the FIDASIM traces have an approximate error of 20 %.

Before the crash there is a small difference between the predicted spectra. As mentioned previously, due to the equilibrium reconstruction used, none of the models trigger at the time of the first sawtooth, and only the Kadomtsev and ergodic model

trigger at the second sawtooth. Therefore a small difference in the models before the crash is observed. Drops in spectral radiance are observed in both the experimental and synthetic data for the channel shown in Figure 4.6 and for channels closer to the core, with channels further out showing small to no drop. The absolute values for the spectral radiance are reproduced well by the TRANSP/NUBEAM and FIDASIM modelling, for the channel shown in Figure 4.6. However, this degree of agreement is only seen in the channel at $R = 1.205$ m; other channels show serious deviation between experiment and prediction, an example of which is given in Figure 4.7. This figure shows the same region of wavelength space as Figure 4.6 but for a channel at $R=1.139$ m (corresponding to pre-crash normalised minor radii $\rho = 0.419, 0.414, 0.435$, and post-crash radii $\rho = 0.425, 0.42, 0.44$ for the Kadomtsev, ergodic, and Porcelli models respectively) in the plasma midplane.




Figure 4.7: Experimental and synthetic spectra showing only the region where FIDA light can be observed at $R=1.139$ m in the plasma midplane. The contribution to the spectrum from the full energy beam peak can be seen at the lower wavelengths, but the discrepancy between the experimental and synthetic spectra persists to higher wavelengths. Based on Section 2.3 the FIDASIM traces have an approximate error of 20%.

This discrepancy in the absolute magnitude, possibly due to calibration factors, is discussed further in section 4.4.2. An EFIT equilibrium with a lower core q would lead to a larger sawtoothing region, and so more fast ions would be expelled

by the crash. This would imply a lowering of the FIDASIM spectrum within the sawtooth region, and an increase outside it, potentially giving better agreement between the absolute magnitudes of the experimental and synthetic spectra. In all channels, over most of the usable wavelength space, the 3 sawtooth models produce only a small difference in the predicted spectra compared to the overall change due to the sawtooth. The errors shown are based on the combination of photon shot noise and the CCD read noise for both the active and passive channels used (dark current noise is neglected due to the cooling of the CCD chip). It must be emphasised that this represents the minimum possible error on the FIDA data and that the true error may be larger, especially considering the possible problems due to uncertain calibration, discussed in subsection 4.4.2. Due to this uncertainty and the size of the error, no clear determination of the suitability of one model over the others can be made. It is concluded that FIDA measurements cannot resolve the issue of which model best represents the data [88].

4.4.2 Detailed analysis and comparison with FIDA data

While in the previous section it was shown that FIDA is not sensitive enough to distinguish between the 3 sawtooth models, and that there was significant deviation of the experiment and prediction, by eliminating a source of calibration error it is shown in this section that good agreement between the forward-modelled fast-ion redistribution due to the sawtooth crash and the FIDA diagnostic measurements can be achieved within the experimental error.

Only some of the FIDA channels contain usable spectral data. Channels close to the plasma edge view plasma regions containing relatively few fast ions and those close to the plasma core view plasma regions containing few neutrals with which the fast ions can charge exchange. Both of these effects lead to a low level of FIDA radiance. Of the channels with a usable level of FIDA radiance it was observed that only a single channel, observing the plasma at a major radius of 1.205 m, consistently showed agreement between the absolute experimental values and the

FIDASIM predictions. However, it has been shown previously that a scaling factor is required in order to model the experimental beam peaks correctly [45]. Applying that factor here would lead to larger disagreement between the model and data, and would have varying effects on the agreement in the other channels. Due to the presence of the optical mask in the FIDA instrument (see Section 3.1.3) it is not possible to use the comparison of the experimental and synthetic beam peaks here as a test of the calibration and scaling validity. Investigating relative changes in the spectra between pre- and post-crash reduces or eliminates some of the effects of absolute calibration errors. Figure 4.8 shows the relative change (difference in signal after the crash vs before, divided by the signal before) in the observed spectral radiance across the fast-ion dominated portion of the spectrum, and across channels that have an observable level of FIDA radiance averaged using the same methods described for Figure 4.6 above.

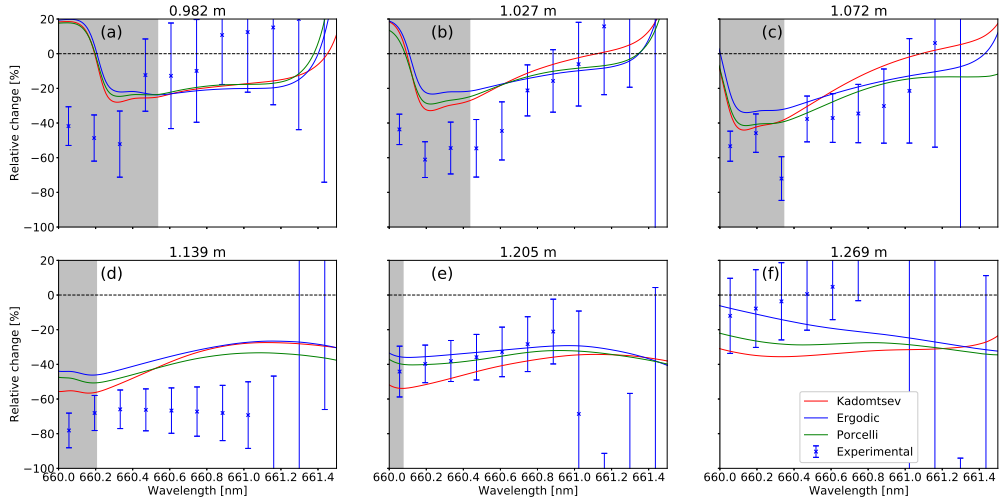


Figure 4.8: Spectral radiance across a number of FIDA channels. Note that the full energy beam peak appears at different locations depending on channel. The grey shaded area in plots (a)-(e) designates where the effect of the peak has a significant effect on the synthetic data, so direct comparisons between the observations and the models cannot be made here. Based on Section 2.3 the relative change in the FIDASIM traces have approximate **absolute** error of 30 percentage points (note that this is not a relative error, but is an absolute uncertainty in the calculated relative change).

As stated previously, the location of the beam peaks in wavelength space are de-

pendent on the relative orientations of the line-of-sight to the NBI beamline, so they appear at different locations in wavelength for the different FIDA channels, which can reduce the usable areas of the spectrum. The sawtooth crash affects the beam peaks due to changes in the beam deposition profile. Unusable parts of the spectra due to this effect are shown by the greyed out portions in figure 4.8. These were determined by calculating the relative change in the synthetic spectra before and after the sawtooth crash firstly using all FIDASIM spectral components and again with only the FIDA component. A region of the wavelength range was designated as unusable if the difference between these two calculations was greater than 0.5 percentage points in terms of the vertical scale used in Figure 4.8. The figure of 0.5 percentage points is somewhat arbitrary, but was chosen to ensure that sawtooth-induced changes in the beam peaks (which were not caused by fast ion redistribution) did not have a significant effect on the overall change in the measured intensity. Note as well that the error on the experimental data increases significantly at wavelengths above 661.0 nm, as the absolute signal experiences a rapid drop-off. It can be seen that in the valid region the experimental and synthetic results are comparable for the most part except at 1.139 m (Figure 4.8(d)), where there is significant disparity between them.

Through examination of Thomson scattering data (by observing the radius within which a drop in the electron temperature is observed, and outside an increase is observed), it can be determined that the outboard radius of the inversion region in the midplane is approximately 1.16 m (corresponding to pre-crash normalised minor radii $\rho = 0.467, 0.462, 0.481$, and post-crash radii $\rho = 0.472, 0.468, 0.486$ for the Kadomtsev, ergodic and Porcelli models respectively), slightly higher than the tangency radius corresponding to Figure 4.8(d) and slightly lower than that corresponding to Figure 4.8(e). Comparing these two plots, it can be seen that the relative change in emission following the sawtooth crash is considerably larger in Figure 4.8(d) than it is in Figure 4.8(e). This in itself is unsurprising, since the effects of the sawtooth on the emission are expected to become smaller as the line

of sight moves away from the sawtooth region. The fact that an effect is seen at all in Figure 4.8(e), despite the tangency radius lying outside the sawtooth region, can be explained by the fact that the fast ion orbit widths and Larmor radii are typically larger than the difference between the tangency radius of this spectrometer chord (1.205 m) and the sawtooth inversion radius (1.16 m).

When the drops in FIDA emission with the three sawtooth models are compared, it is evident from Figure 4.8(d) that all of them predict a drop in emission that is smaller than the measured value. The estimated FIDASIM error ameliorates this somewhat, but the shown channel is still anomalous compared to the others shown. In making these comparisons, it should be noted that while finite Larmor radius corrections are taken into account in NUBEAM for the purpose of evaluating neutralisation probabilities and collision rates, they are not taken into account in terms of the local magnetic field used to track particle orbits, since a guiding centre approximation is used [59]. This introduces an error in the synthetic spectra for tangency radii lying close to the boundary of the sawtooth region. The discrepancy between measurements and modelling in Figure 4.8(d) could perhaps be due to a pitch angle (i.e. Larmor radius) dependence of the effects of the sawtooth on the fast ion population that is not captured by any of the models, although further investigation would be needed to confirm this. The use of ASCOT or another full-orbit code could be used here to test this hypothesis.

Figure 4.8(d) shows the largest drop in the observed FIDA radiance across all valid channels, along with the largest drop in the ergodic and Porcelli models. Figure 4.8(e) shows the largest drop in the Kadomtsev model, with similar drops in the other models to Figure 4.8(d). This is somewhat counter-intuitive: it may be thought that channels corresponding to locations closer to the core would have larger drops in radiance, since the core is inside the sawtooth region and contains more fast ions prior to the crash. The drop in radiance in the channel used in Figure 4.8(e) and the lesser drops in other channels can be correlated with the change in the fast-ion distributions predicted by TRANSP/NUBEAM due to the

good correspondence between the observed and predicted values.

Figure 4.9 shows the computed fast-ion distributions at a major radius of 1.205 m in the midplane of the tokamak (the Kadomtsev model is used here, as an example). Pitch here is defined as the ratio of the component of the marker's velocity parallel to the magnetic field to the full velocity. On MAST during conventional operation the toroidal component of the magnetic field, B_ϕ , was directed clockwise as seen from above and current anti-clockwise, so particles with negative pitch are travelling in the co-current direction. The overlaid contours show the extent of the weight function for the FIDA diagnostic [45], [93], specifically in the chosen spatial location and over the range of wavelengths that are dominated by the presence of the fast ions. For comparison with Figure 4.8(e) the relative change in the fast-ion distributions across the three models is shown in Figure 4.10.




Figure 4.9: Simulated fast ion distributions for $R = 1.205\text{m}$ in the device midplane using the Kadomtsev model. (a) Before crash, (b) after crash. The green contours represent the weight function for the FIDA diagnostic in the fast-ion dominated wavelength range and the spatial location of measurement.

It can be observed across all of the models in Figure 4.10 that there is a significant

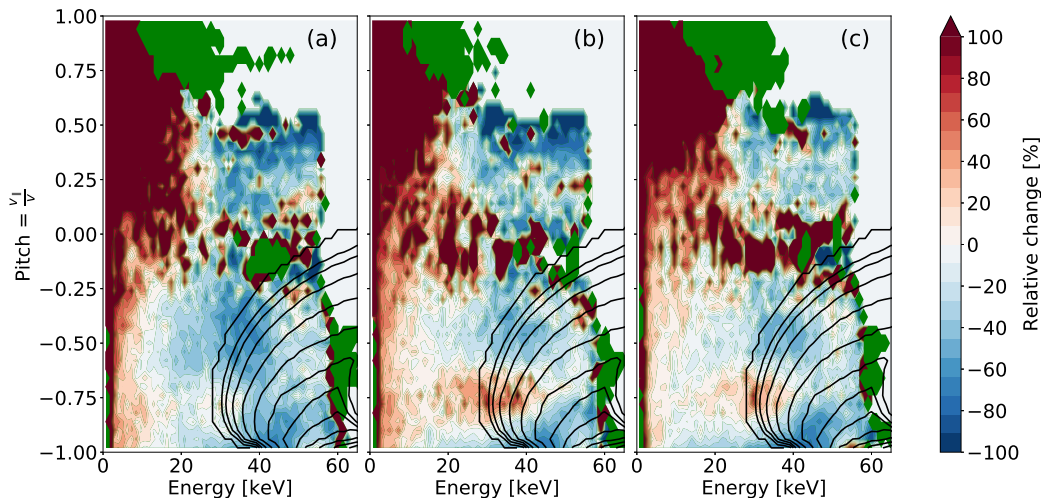


Figure 4.10: The relative change in the fast-ion distribution function for $R = 1.205\text{m}$ in the device midplane. (a) Kadomtsev, (b) ergodic, (c) Porcelli. Areas in green represent phase-space regions with a fast-ion presence after the crash that had no presence before. Note that in some regions there is a positive change of greater than 100%. The black contours represent the weight function for the FIDA diagnostic in the fast-ion dominated wavelength range and the spatial location of measurement.

loss of fast ions (up to 50%) from energies between 45-55 keV across pitches between -0.25 and -1 , with the Kadomtsev (Figure 4.10(a)) model showing additional losses down to 30 keV and the ergodic (Figure 4.10(b)) and Porcelli (Figure 4.10(c)) models showing an increase in the fast ions around a pitch of -0.75 across energies 25-40 keV. Additionally there is an increase in fast ions with energies less than 20 keV across positive pitch values. The higher energy particles are expelled from the sawtooth region during the crash and lower energy particles from larger radii are mixed in. These lower energy fast ions will have undergone significant pitch angle scattering. The apparent preference for positive pitch angle after the crash across the models in Figure 4.10 is likely to be an artefact of the normalisation. It can be seen that the weight function primarily intersects the region of significant loss, which is reflected in the decrease in FIDA radiance shown in Figure 4.8, with the smaller regions of increase within the intersection present in the ergodic and Porcelli models leading to the smaller decrease present in those models compared

to the Kadomtsev.

The shape of the weight function depends on the angle between the line-of-sight and the magnetic field, and hence is channel-dependent. In the other channels the weight function primarily intersects the population of newly injected fast ions around 55 keV, so in these channels a smaller change in FIDA radiance due to redistribution of the fast-ion population by the sawtooth is observed. It should be emphasised when examining synthetic fast-ion distributions that any regions of the phase-space outside of the overlap with the weight function cannot be probed with the FIDA diagnostic, so cannot be compared with experimental values.

The recovery of the fast-ion distribution after the crash can be observed in Figure 4.11, where the spectra move towards pre-crash conditions as the time window averaged for the 'after' portion of the data moves further from the crash. The total time interval covered in this figure is comparable to the fast-ion slowing down time, which is the timescale on which it could be expected for the fast-ion distribution to recover following a sawtooth crash. Figure 4.4(b) shows an overall degradation of the fast-ion confinement, which could contribute to the incomplete recovery of the spectra. There is a significant disagreement between the experimental and synthetic data in Figure 4.11(e). While a definitive explanation for this is not yet known, examination of the D- α signal for one of the three pulses (#29880) in the period after the third sawtooth shows a significant increase around the start of this averaging window. This could be a consequence of more edge neutrals permeating into the hot plasma, which would be likely to generate an increase in the passive FIDA signal, thereby affecting the passive subtraction. Such an effect was not taken into account in the FIDASIM calculation for these pulses.

4.5 Passive FIDA

As described in section 2.1.2 there is a FIDA spectral component generated by charge exchange reactions between fast ions and neutrals present in the colder

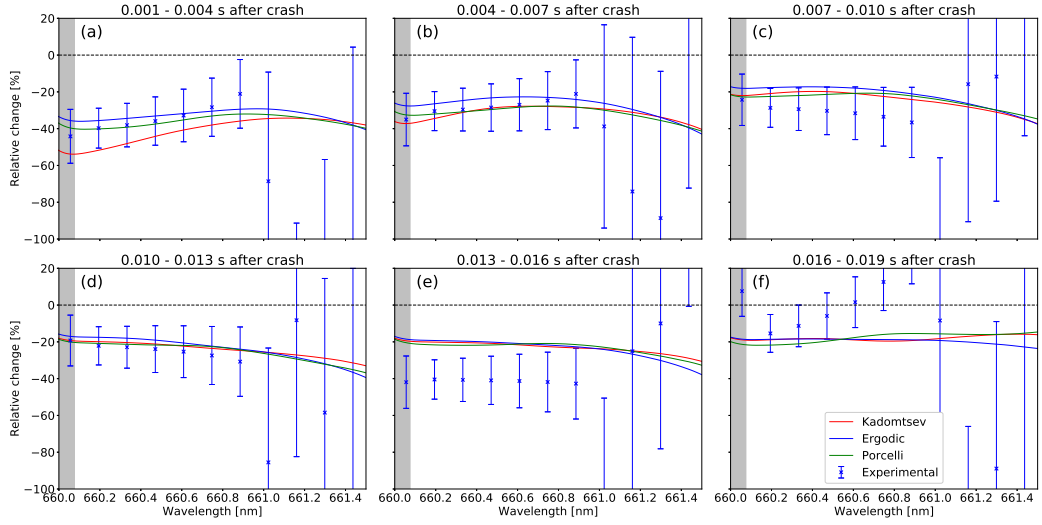


Figure 4.11: Relative change in the spectral radiance at $R = 1.205\text{m}$ in the device midplane, averaged over the indicated time windows. Note that Figure 4.11(a) examines the same data as Figure 4.8(e). Again, based on Section 2.3 the relative change in the FIDASIM traces have approximate **absolute** error of 30 percentage points (note that this is not a relative error, but is an absolute uncertainty in the calculated relative change).

edge region of the plasma, rather than the beam neutrals. In this section, the effect of the sawtooth crashes (in #29880) on the passive emission observed by a toroidally directed passive view on MAST is examined. Examination of the passive component of the FIDA spectrum in response to sawtooth crash has been examined on a conventional tokamak by Bolte *et. al.* [94]. The following shows a brief first examination on a spherical tokamak.

Figure 4.12(a) shows a plan view of the MAST plasma, with the black circles denoting the minimum and maximum extent of the plasma at the chosen time. The fan of chords available to the passive FIDA view are shown in grey, with an example channel shown in red. Figure 4.12(b) shows the integrated radiance, I , for this channel for a single sawtooth crash. A decaying exponential of the form:

$$I(t) = ae^{-\frac{(t-t_0)}{\tau}} + c \quad (4.1)$$

where t is time, t_0 is a time offset, and a , c and τ are constants, is used to describe the radiance temporal decay. The decay time τ is fitted to the decaying part of

the signal. It is observed that there is a large spike in integrated radiance across all channels that intersect the plasma, with decay times between 0.3-0.5 ms. This feature is not observed in data integrated in regions free from FIDA light, indicating that background fluctuations are not the cause. Therefore the data suggests that there is a sudden significant increase in the population of passing fast ions followed by a rapid decrease. It is likely that this rapid decrease is due primarily to charge exchange between these passing fast ions and edge neutrals rather than direct fast ion orbit losses. A typical neutral density, n_0 , in the MAST plasma edge is a few times $1 \times 10^{16} \text{ m}^{-3}$ [95]: using the velocity, v , and charge exchange cross-section σ for a 40 keV deuteron [5], the charge-exchange time is estimated to be $\tau_{CX} = 1/(vn_0\sigma) \sim 0.3 \text{ ms}$, consistent with the decay time in Figure 4.12(b). Bolte *et. al.* [94] have carried out a detailed examination of passive FIDA emission on the DIII-D tokamak, where they also observe the preference for the ejection of passing particles from the core to the edge. Additionally, they are able to determine, with an accurately calibrated diagnostic and the use of passive FIDASIM simulations, that approximately 1% of the fast-ion population is expelled into the edge region. Looking to the future, MAST-U has a fast-ion loss detector, which is expected to provide direct information on sawtooth-induced fast-ion losses.

4.6 Conclusions

Based on forward modelling performed with the FIDASIM code in conjunction with experimental FIDA data it is not possible to distinguish between complete (Kadomtsev and ergodically mixed Kadomtsev) and incomplete (Porcelli) reconnection models for the sawtooth crashes, due to the level of noise present in the FIDA data and the relatively small differences seen in the FIDASIM data between these models. While issues regarding the calibration of the diagnostic mean that there are significant difficulties in analysing the absolute FIDA radiances, examining the relative change in the spectra and synthetic fast-ion distribution resulting

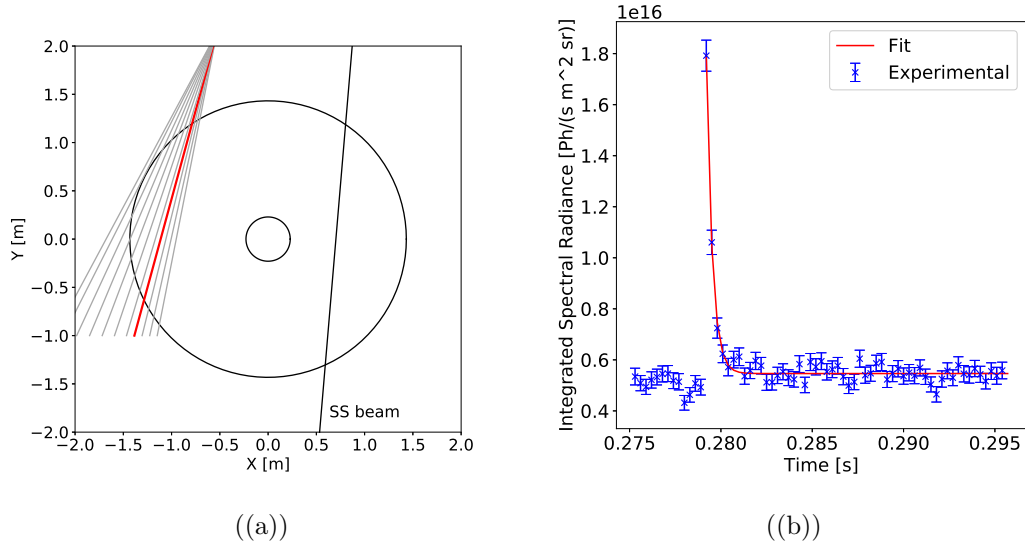


Figure 4.12: (a) Available toroidally-viewing passive chords, with selected chord highlighted in red, and the NBI beamline to emphasise the passive nature of the view. Data from this chord is shown in Figure 4.12b. The black circles show the minimum and maximum extent of the plasma. (b) Spike and decay of passive signal seen during a sawtooth crash. The decaying part of the signal has been fitted with an exponential.

from the sawtooth crashes allows for further investigation into the fast-ion redistribution. Large drops in the measured spectral radiance are reflected in the synthetic spectra, which indicates significant depletion of the fast-ion phase-space at energies around 40 keV, which progressively recovers over the inter-sawtooth period. Passive FIDA data show that sawtooth crashes lead to a sudden increase and subsequent rapid decay of passing fast ions in the edge region of the plasma. It is concluded that additional fast ion diagnostics, such as the SSNPA installed on MAST-U, are likely to be required in order to discriminate experimentally between alternative models of fast ion redistribution and loss resulting from sawtooth crashes.

Analysis of beam fractions during first experimental campaign on MAST-U

5.1 SS beam fraction investigation

NBI commissioning was ongoing during the early part of the experimental campaign. As part of the commissioning, dedicated beam-into-gas shots were performed. Beam-into-gas shots allow for the examination of the beam performance, as the only spectral components that can be observed are the beam peaks, generated by collisional excitation of the beam atoms. Additionally, there is much less beam attenuation in beam-into-gas shots compared to beam-into-plasma. These shots can be used to examine the beam condition, and verify the flux fraction (equivalently called the current fractions in some of the literature) of the three energy components of the neutral beam. However, part of the injection system was replaced after these shots, making an analysis from before the replacement less useful. During a later shot, #44254, the SS beam erroneously continued to operate after the termination of the plasma, allowing a serendipitous extra examination of the beam in a beam-into-gas situation. Beam-into-plasma shots can also be used

as long as background subtraction to eliminate impurity lines is performed. In general, FIDA cannot be used for this purpose due to the presence of the mask obscuring the beam peaks, but in the early part of the campaign the mask was not installed. Additionally, the mask was removed for the last day of the campaign. This allows the analysis of both a (relatively) early campaign shot and an analysis of the last plasma shot of the campaign (#45483).

The motivation for this investigation was that, for #44254, four beam peaks were observed, instead of the usual three, suggesting there was a contamination of the beam. Section 5.1.1 documents the diagnosis of potential candidates for the contamination using FIDA spectrometer data. Additionally, it was noted that the relative magnitudes of the beam peaks were not what was expected. At the end of the campaign, for shot #45483, it was also noted by the RO of the CXRS diagnostic that the data from the new Celeste-4 system also gave calculated SS beam fractions significantly different to those provided in the shot database (which was approximately Full = 0.78, Half = 0.16 and Third = 0.06, for both shots). Section 5.1.2 documents the derivation of the specific fraction values for the early and late campaign shots using FIDA data, as a confirmation of the discrepancy. The SW beam is not observed by the FIDA diagnostic, and is not analysed in any way.

5.1.1 Diagnosis of unknown beam peak component

The Doppler shift from the D- α rest wavelength of the full, half and third energy beam peaks is dependent on the component of the particle velocity parallel to the FIDA line of sight. The wavelength positions of the peaks, λ_s where s refers to the specific singly-charged species being accelerated (mono-, di- or tri-atomic deuterium), can therefore be derived from knowledge of the direction vectors of the lines of sight, \mathbf{a}_{los} and the neutral beam, \mathbf{a}_{nbi} , along with the beam energy, E_{nbi} :

$$\lambda_s = \lambda_0 \left[1 + \frac{1}{c} \left(\sqrt{\frac{2E_{\text{nbi}}}{m_s}} \frac{\mathbf{a}_{\text{los}} \cdot \mathbf{a}_{\text{nbi}}}{|\mathbf{a}_{\text{los}}||\mathbf{a}_{\text{nbi}}|} \right) \right] \quad (5.1)$$

where m_s is the mass of the species in question. From Equation 5.1 with the average beam energy $E_{nbi} = 51.7 \text{ keV}$ the location of the beam peaks in each channel in wavelength space can be calculated. When this is done, it can be observed that there is a fourth peak unaccounted for by the mono-, di- and tri-atomic deuterium. Information on the species that is causing this peak can be obtained by examination of the FIDA spectra.

The non-background subtracted spectrum was used for the fitting for shot #44254, as the passive data did not extend to include the unknown peak. This is acceptable for the beam-into-gas portion of the shot, as there are no impurity peaks. The fit must now include a parameter D , to match the spectrally invariant background light detected. The spectrum in each channel was fitted using a function $f(\lambda)$ consisting of the sum of four Gaussians of different amplitudes $a_{u,t,h,f}$ (where the subscript denotes the peak: u = unknown; t = third; h = half; f = full) and wavelength position $\lambda_{u,t,h,f}$, but the same width c_λ , and a constant background level, D :

$$f(\lambda) = a_u \exp \left[\frac{-(\lambda - \lambda_u)^2}{2c_\lambda^2} \right] + a_t \exp \left[\frac{-(\lambda - \lambda_t)^2}{2c_\lambda^2} \right] + a_h \exp \left[\frac{-(\lambda - \lambda_h)^2}{2c_\lambda^2} \right] + a_f \exp \left[\frac{-(\lambda - \lambda_f)^2}{2c_\lambda^2} \right] + D \quad (5.2)$$

Figure 5.1 shows an example FIDA spectrum from a single channel in this shot at time 0.3705 s-0.3715 s, after the termination of the plasma but before the beam shutoff, along with the derived fit:

The parameters derived to generate the fit in Figure 5.1 are shown in Table 5.1.

Rearranging Equation 5.1 for the mass allows an estimate of the mass of the unknown accelerated molecule. The mean value across all of the channels is 18.6 ± 0.4 u, where u is the atomic mass unit.

Some inferences can be made about the possible molecular composition. Ammonia (H_3N - molecular mass 17.03 u), water (H_2O - 18.015 u), heavy water (D_2O -

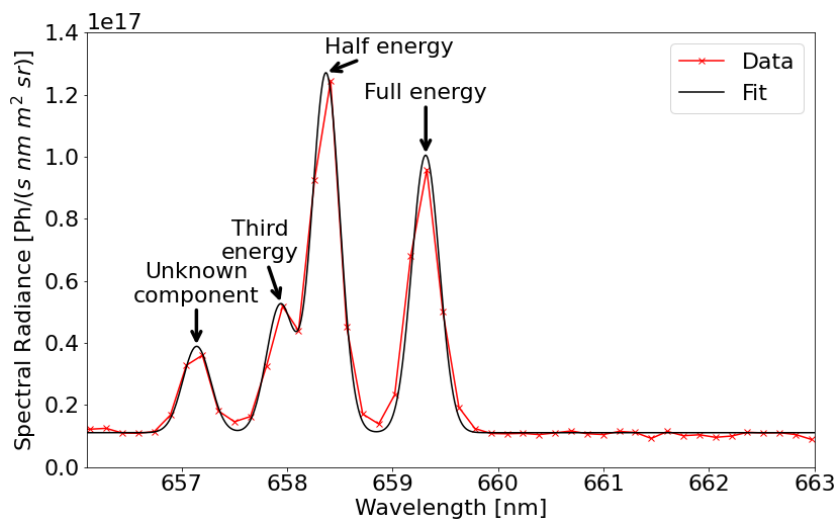


Figure 5.1: Spectra observed by the FIDA active view during the beam-into-gas phase in pulse #44254 at $R = 1.01$ m averaged over times 0.3705 s-0.3715 s. 4 peaks are observed in the redshifted D- α portion of the spectrum.

Fit parameter	Value	Error
a_u	2.78×10^{16}	0.22×10^{16}
a_t	4.1×10^{16}	0.34×10^{16}
a_h	1.16×10^{17}	0.26×10^{16}
a_f	8.94×10^{16}	0.27×10^{16}
λ_u	657.14	0.02
λ_t	657.94	0.02
λ_h	658.37	0.005
λ_f	659.31	0.006
c_λ	0.13	0.003
D	1.11×10^{16}	0.035×10^{15}

Table 5.1: Fit parameters generated for the specific channel shown in Figure 5.1. Note that the errors derived are solely the standard deviation error in the fitting.

20.027 u), semiheavy water (DHO - 19.02 u) and deuterated ammonia (ND₃ - 20.049 u) could all be potential candidates. The results suggest that water or semiheavy water are the most likely cause of the fourth peak, but the relatively low wavelength resolution of the FIDA spectrometer may obscure small changes in the wavelength position of the peak. Water is known to be a problem component on other beams [96], but usually occurs in very small proportions. Section 5.1.2.1 documents the calculated beam fractions for all of the spectral components.

5.1.2 Derivation of beam fractions

Using the fitted FIDA spectra and electron temperature and density data from Thomson scattering, an estimate of the beam power fractions can be made. D- α beam emission is a collisional-radiative process, determined by excitation into $n = 3$ and de-excitation from higher states into $n = 3$, and the subsequent radiative de-excitation into $n = 2$, determined by the Einstein rate coefficients for spontaneous de-excitation. The Atomic Data and Analysis Structure (ADAS) [97] is a suite of codes and data pertaining to radiative phenomena in plasmas. From ADAS the ADF22 data set contains the effective beam emission coefficients, k , for a range of beam energies, electron densities and temperatures, as well as the degree of plasma impurity, for a variety of plasma and beam ions. Utilising these, the intensity of the spectral line of a specific energy fraction of the beam, ϵ_F , where F denotes the specific fraction (full, half or third energy), can be given by [98]:

$$\epsilon_F = \int_S k(E_F, n_e, T_e) n_e n_b^F ds \quad (5.3)$$

where n_e is the electron density, T_e is the electron temperature, n_b^F is the density of the specific beam component. E_F is E_{nbi} , $E_{nbi}/2$ or $E_{nbi}/3$ for full, half and third energy components respectively. These are integrated over S , the length of the intersection of the line-of-sight of the diagnostic and the neutral beam.

The power injected by a specific fraction, P_F can then be estimated by:

$$P_F = \int_A n_b^F v_F E_F dA = \sqrt{\frac{2}{m_D}} E_F^{3/2} \int_A n_b^F dA \quad (5.4)$$

where $v_F = \sqrt{2E_F/m_D}$ is the velocity of the injected deuterium atom, m_D is the deuterium mass, and A is the cross-sectional area of the beam.

For the calculated approximation the beam is assumed to have a radius $r_{nbi} = 0.1$ m and all quantities are assumed to be constant across the beam cross-section.

The FIDA lines-of-sight are designed to be roughly aligned tangential to the flux surfaces at the NBI intersection. This allows the approximation of constancy across the line-of-sight/beam intersection, and across the beam cross section. With these assumptions, Equation 5.3 reduces to:

$$\epsilon_F = k_F n_e n_b^F (2 r_{nbi} / \sin \theta) \quad (5.5)$$

where k_F is the value of k for the energy E_F , and θ is the angle between the beam and the FIDA line-of-sight. Equation 5.4 reduces to:

$$P_F = \sqrt{\frac{2}{m_D}} E_F^{3/2} n_b^F \pi r_{nbi}^2 \quad (5.6)$$

It follows from Equations 5.5 and 5.6 that the beam power fractions p_F are then given by:

$$p_F = \frac{\frac{\epsilon_F}{X^{3/2} k_F}}{\frac{\epsilon_{full}}{k_{full}} + \frac{\epsilon_{half}}{2^{3/2} k_{half}} + \frac{\epsilon_{third}}{3^{3/2} k_{third}}} \quad (5.7)$$

where $X = \{1, 2, 3\}$ if the fraction is full, half and third respectively. The elimination of the explicit dependence on n_e , n_b^F , θ and r_{nbi} makes the p_F calculation more robust than the calculation of the absolute values of P_F . Additionally, p_F is dependent only on the relative values of each of the ϵ_F and k_F components to each other, so the calculation is robust to an incorrect absolute calibration of the FIDA diagnostic. Power fractions and flux fractions, f_F , can be converted to each other by:

$$f_F = \frac{p_F X}{p_{full} + 2p_{half} + 3p_{third}} \quad (5.8)$$

$$p_F = \frac{(1/X)f_F}{f_{full} + (1/2)f_{half} + (1/3)f_{third}} \quad (5.9)$$

The ADAS rate coefficient k is specific to beam-into-plasma shots. Due to relatively low beam energy for shot #44254 the unknown and third energy peaks are contaminated in the plasma phase by the cold D- α emission and the carbon impurity lines, which are poorly subtracted in this case. Data on the cross-section of Balmer- α excitation of a hydrogen beam injected into hydrogen gas for a variety of beam energies [99], [100] can be used for an equivalent deuterium setup by taking the values at a half of the deuterium beam energy. With the assumption of $k_F^{gas} = \sigma_F v_F$, where σ_F is the cross-section for D- α emission for the specific beam component, the analysis can be extended roughly to beam-into-gas. Equation 5.7 can be extended to include the "unknown" fraction by making an assumption about the molecule that is accelerated. Based on the results from Section 5.1.1 water is assumed for this component.

5.1.2.1 Shot #44254

Figure 5.2 shows the calculated power and flux fractions across all channels. Due to the lack of beam attenuation in beam-into-gas shots, the mean and standard deviation of the calculated values can be made across the full complement of calculated fractions, given in Table 5.2.

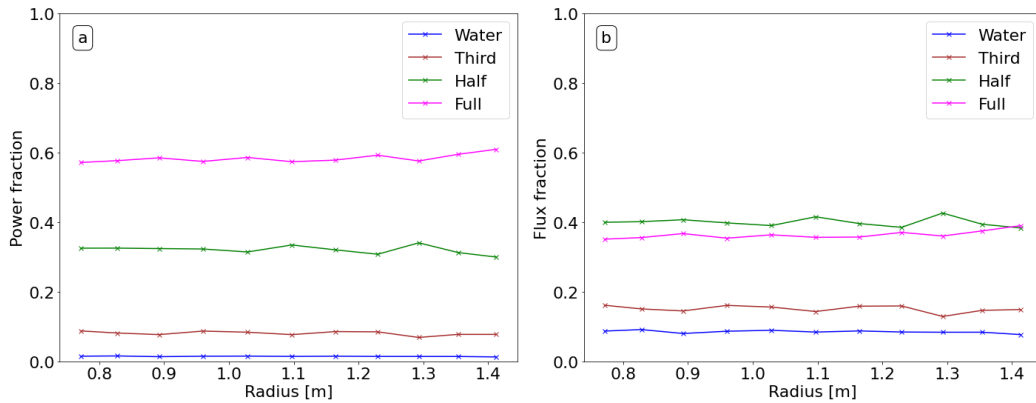


Figure 5.2: (a) Power and (b) flux fractions for Shot 44254, calculated across the FIDA channels

It was observed that the half and full values calculated from the FIDA spectrometer are significantly different from those predicted by the NBI team on MAST-U. Their

Component	Power fraction		Flux fraction	
	Value	Std. dev.	Value	Std. dev.
Full	0.584	0.011	0.364	0.011
Half	0.321	0.011	0.4	0.012
Third	0.081	0.005	0.151	0.009
Water	0.014	0.001	0.085	0.004

Table 5.2: Power and flux fractions for Shot #44254. The mean and standard deviation were evaluated from data across all FIDA channels.

predicted flux fractions were: Full - 0.743, Half - 0.125, Third - 0.132. There was no associated water prediction as it is an unwanted contaminant. It was suspected that the anomalies could be due to these measurements being taken relatively early in the operation of the new PINI, and that the fractions might become closer to the predicted values later in the campaign. As mentioned previously, for the majority of the campaign the mask had been installed on the FIDA spectrometer, so in general no beam fraction measurements could be made using it.

5.1.2.2 Shot #45483

The beam power fractions for Shot #45483 were initially estimated by Christopher Wade of UKAEA utilising the Celeste 4.1 spectrometer, part of the CXRS system. These estimates were: Full - 63 %, Half - 31 %, Third - 6 %. Due to being a beam-into-plasma shot, beam attenuation is more pronounced than for the beam-into-gas. Figure 5.3 shows the power (5.3(a)) and flux (5.3(b)) fractions calculated for shot #45483.

The half and third components are attenuated more than the full, causing the full fraction to rise as the beam continues to penetrate the plasma. The two outermost FIDA channels, at $R = 1.354\text{m}$ and $R = 1.413\text{m}$ could not be used as, due to the relatively low wavelength resolution of the FIDA spectrometer, the half and third energy peaks were not distinguishable from each other in these channels. The most accurate value possible for the "true" values of the fractions is taken from the outermost channel that could distinguish the half and third peaks, at $R = 1.293\text{m}$,

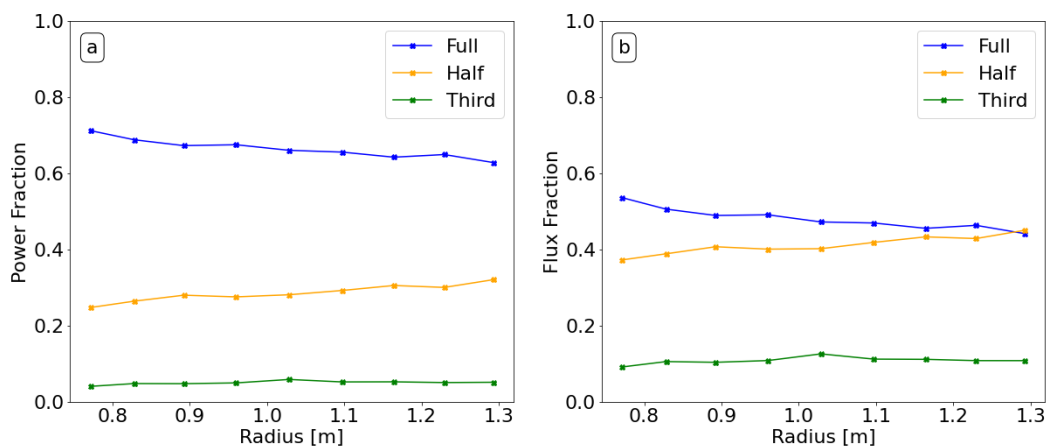


Figure 5.3: (a) Power and (b) flux fractions for Shot 45483, calculated across the FIDA channels.

Component	Power fraction	Flux fraction
Full	0.63	0.44
Half	0.32	0.45
Third	0.05	0.11

Table 5.3: Power and flux fractions for Shot #45483, taken as value from the outermost FIDA channel where the half and third energy peaks could be distinguished from each other, at $R = 1.293$ m.

shown in Table 5.3.

While the values shown in Tables 5.2 and 5.3 are quite different, they are much closer to each other than either is to the original values stated in the shot database.

5.2 Conclusion of beam fraction investigation

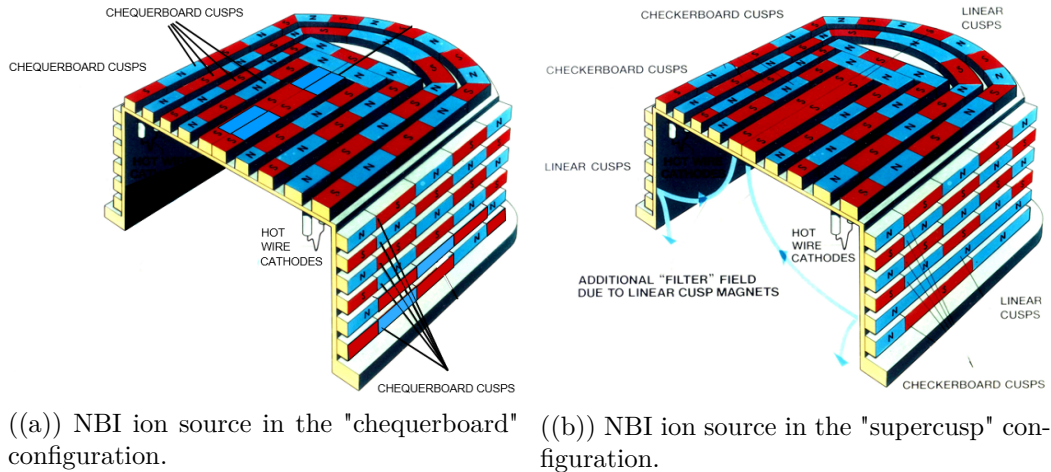
The above analysis makes a number of assumptions, but the fractional calculations make a number of calibration issues irrelevant, and the results are only dependent on relative intensities observed. The values for the flux fractions calculated are strikingly different than those predicted by the current NBI model used by the NBI group on MAST-U, which is used in TRANSP by default. For shot #45483 the original prediction for flux fractions gave Full = 0.78, Half = 0.16 and Third = 0.06. Such a significant difference in the ratios additionally calls into question the value used by the beam group for the overall injected power, which is itself a

function of the fractions. Together, these issues cause severe doubt in regards to the accuracy of any MAST-U FIDASIM simulations made with the "known" values.

While the uncertainty in the overall beam power injected means that a direct FIDA/FIDASIM comparison of the absolute values of the spectral radiance of the beam peaks is inherently suspect, the relative magnitudes of the simulated peaks compared to those of the experimental peaks can allow a comparison of the different fractions used. Looking at $R = 1.293$ m, the relative magnitude of the experimental peaks, using the full energy peak as a reference, is: 1:1.06:0.27. For the "original" fractions, the simulated ratio is: 1:0.24:0.12, and for the "new" fractions it is: 1:1.14:0.46. The new fractions give significantly better relative heights for the beam peaks than the original fractions do.

A complicating factor in regards to the beam fractions is the effect of a varying neutraliser gas flow rate between shots, as beam fractions can vary with this in a non-linear way. The total injected power and the beam fractions are normally determined assuming the maximum utilised neutraliser flow of 12 mbar L s^{-1} . For shots #44254 and #45483 the neutraliser flow rate was 10 mbar L s^{-1} . It has been shown on the JET NBI system that the full energy fraction is relatively unaffected by varying neutraliser gas flows, while both half and third are somewhat more affected [101]. As the MAST-U FIDA system is almost entirely looking at fast-ions injected at the full energy, varying flow rates on this scale are not a hugely concerning issue.

After the issue of the beam fractions was raised an investigation of the SS PINI was performed. Figure 5.4 show the two possible configurations for the magnets that surround the NBI ion source on MAST-U, specifically *chequerboard* (Figure 5.4(a)), where the magnets are arranged alternating north-south around the body and on the backplate, and *supercusp* (Figure 5.4(b)) where there are portions of non-alternating magnets on the back and sides (both configurations were previously used on JET [102]).



((a)) NBI ion source in the "chequerboard" configuration. ((b)) NBI ion source in the "supercusp" configuration.

Figure 5.4: Possible configurations of magnet polarities for MAST-U NBI ion source. Adapted from [103].

Chequerboard gives higher injected power but a lower full energy fraction compared to supercusp. It had been assumed that the ion source was in supercusp configuration. However, inspection of the ion source revealed that the magnets on the ion source backplate were in checkerboard configuration, whereas those on the body were in supercusp. This is not a standard configuration, so the behaviour is not known precisely. Based on internal discussions the author has had with members of the NBI team, it is believed that the behaviour of the PINI will have been closer to chequerboard than supercusp, which aligns with the lower beam fractions diagnosed by the Celeste 4.1 and FIDA systems.

In conclusion, there is a reasonable degree of confidence that the true full energy beam fraction in the south beam is significantly lower than that originally predicted, due to an error in the magnetic configuration of the NBI ion source. This has consequences for the calculation of the total injected power. It is unknown, at the time of writing, to what degree corrected beam fractions and power can be determined retroactively for the first experimental campaign on MAST-U and, if it is possible, on what timescale such corrections can be made. This has significant consequences for the generation of accurate TRANSP and FIDASIM runs. It was decided that for MAST-U data analysis performed in the remainder of this thesis,

the flux fractions of 44:43:13 (as determined by the spectra from Celeste-4 system, which had higher wavelength resolution than the FIDA system, so would be able to get better fits) would be used, with the total power being given by the already predicted value (as no more accurate values are available at the time of writing). Until a formal review by the NBI team is complete, this is the best that can be done. It is inevitable that this will introduce an unknown degree of uncertainty into both the TRANSP and subsequent FIDASIM simulations. When (and if) better values are available, it would be prudent to re-evaluate the analysis.

Analysis of MAST-U data

This chapter documents analysis performed on data taken during the first experimental campaign on MAST-U. Section 6.1 describes analysis of Shot #45026, specifically investigating the effect of a large fishbone instability on the plasma rotation, and using anomalous diffusion to better investigate the effect of the fishbone on the fast ions. Section 6.2 focuses on Shot #45091, probing the effects of a long-lived mode (LLM) and locked mode on the fast ions.

6.1 Investigation of effect of fishbone on plasma rotation and fast ions

Shot #45026 was performed as part of the MU01-FPP-01 experiment, investigating neutral beam current drive and fast ion slowing down. The pulse was heated with only the SS beam for a duration of 0.327s, with an average power of 1.49 MW and average primary energy of 65.7 keV. Figure 6.1 shows the time evolution of some plasma parameters for the shot. Examination of these in conjunction with the Mirnov coil spectrogram, shown in Figure 6.2, reveal some interesting features of the shot.

The early phase of SS beam injection is dominated by TAEs, with frequencies of ~ 100 kHz, comparable to the frequency $\omega_{TAE}/2\pi$ given in Section 1.6.2. There




Figure 6.1: Time traces of plasma parameters of Shot #45026. (a) SS beam power, (b) neutron rate (note that the absolute calibration of the fission chamber is incorrect), (c) line integrated electron density, (d) D- α intensity.

were then several prominent fishbones, with associated drops in neutron rate (as seen in Figure 6.1(b)). This period ends with the development of an LLM, which itself is terminated by a sawtooth crash. Another sawtooth crash follows shortly after, then the SS NBI beam terminates.

The behaviour of the wavelength-integrated FIDA data over the duration when the SS beam is on, across a range of radii, is shown in Figure 6.3. The fishbones have little effect on emission in channels outside of this radius range.

After a transient spike at SS beam turn on, the emission exhibits a fairly flat trend in the inner channels (1.029 m-1.164 m) and a slow secular increase in the outer channels (1.23 m-1.354 m), up until approximately 0.25 s, pre-fishbones. It can be seen that the fishbone at ~ 0.28 s and the "major" fishbone at ~ 0.319 s cause significant drops in the FIDA data across the majority of the examined channels, especially the "major" fishbone. Table 6.1 shows the percentage change in the wavelength integrated FIDA signal due to the "major" fishbone, representing strong evidence that a significant fraction of high energy fast ions are being re-




Figure 6.2: (a) Time trace of the low-frequency Mirnov coil signal and (b) Spectrogram of low-frequency Mirnov coil time trace from shot #45026, showing the main resonant instabilities of the plasma. Specifically, TAEs dominate the early beam on portion of the shot, before transitioning to a series of fishbones. A number of these occur at the same time as significant drops in the neutron rate, with a emphasis on the fishbone that occurs at ~ 0.319 s (termed the "major" fishbone), at which time there is a very large drop in the neutron rate *and* plasma rotation velocity (discussed in section 6.1.1). At ~ 0.356 s a fishbone transitions into an LLM, which terminates with a sawtooth crash at 0.3883s (seen as a feature across all shown frequencies, very narrow in time). Another sawtooth crash occurs at 0.409s, before the beam terminates at 0.4287s.

distributed by the fishbone across a large range of radii. The large drops during the major fishbone at large radii are at first glance surprising, as fishbones are core localised. However, on MAST-U, due to a relatively low plasma current and magnetic field, the particles will experience large orbit excursion, as well as having large Larmor orbits. This makes it possible for fast-ions detected at the larger radii

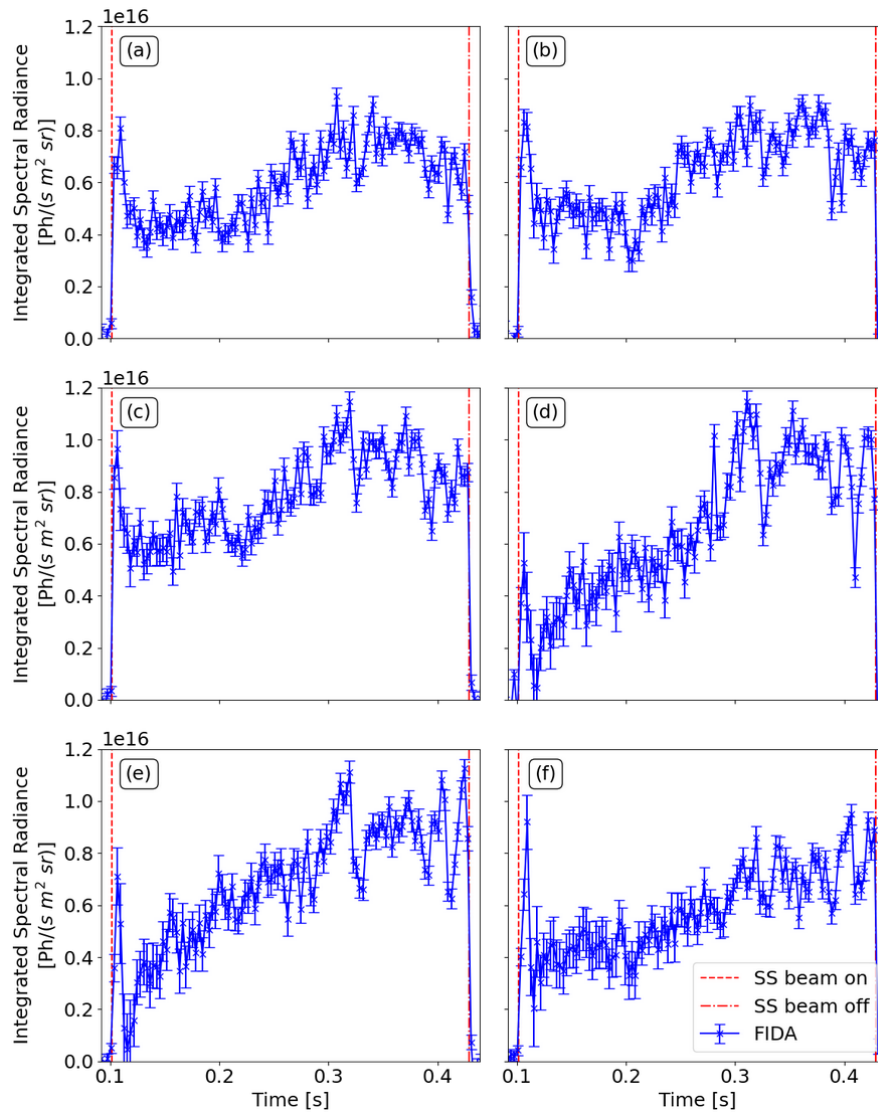


Figure 6.3: Wavelength integrated time traces of the FIDA emission in shot # 45026. The FIDA emissivities are integrated between 660.52 nm-661.99 nm. Radii are (a) 1.029 m, (b) 1.097 m, (c) 1.164 m, (d) 1.23 m, (e) 1.293 m, and (f) 1.354 m. The times of the SS beam switch-on and switch-off are again indicated by red dashed and dashed-dotted lines.

to be influenced by more core localised effects.

Radius [m]	Percentage change	Error [percentage points]
1.029	-16.6	7.1
1.097	-25.2	6.3
1.164	-33.8	4.7
1.23	-42.1	5.8
1.293	-32.6	5.5
1.354	-24.7	7.2

Table 6.1: Percentage change in the wavelength integrated FIDA signal due to the "major" fishbone across $R = 1.029\text{ m}-1.354\text{ m}$.

6.1.1 Effect of fishbone instability upon plasma rotation

Figure 6.4 shows the bulk plasma toroidal rotation as observed by the CXRS diagnostic.




Figure 6.4: Bulk plasma toroidal rotation derived from the CXRS diagnostic for shot #45026. The green trace shows the location of the magnetic axis over time, as defined by the archived EFIT run for the shot.

Over the duration of SS beam injection, the velocity of the plasma in the $\sim 0.9\text{ m} - 1.1\text{ m}$ range increases until $\sim 0.23\text{ s}$, when it reaches a plateau. At $\sim 0.3\text{ s}$ there is another increase in velocity, followed by a sharp crash at 0.325 s , which coincides with the end of the chirping phase of the major fishbone. As discussed previously there is a significant drop in the FIDA emission over the course of the chirping

phase of the fishbone. Similar behaviour was observed during two other shots from MU01, 45031 and 45163. Specifically, a fishbone is observed that coincides with a drop (or in the case of 45031, a reduction of the rate of growth) in neutron rate and drop in the FIDA emission, followed by a drop in the plasma rotation rate. The shots are not perfectly analogous, but share these features. In shot 45031, the fishbone occurs over the period $\sim 0.245\text{ s} - 0.25\text{ s}$, and in shot 45163 occurs over $\sim 0.322\text{ s} - 0.327\text{ s}$.

Figure 6.5 shows the behaviour of the neutron rate and a central slice of the rotation velocity during the selected fishbone in each shot and Figure 6.6 shows the FIDA emission versus R before and after the selected fishbone in each shot integrated over 660.77 nm-661.52 nm, a wavelength range that is sensitive to fast-ions with $E_{min} \geq 47.4\text{ keV}$.

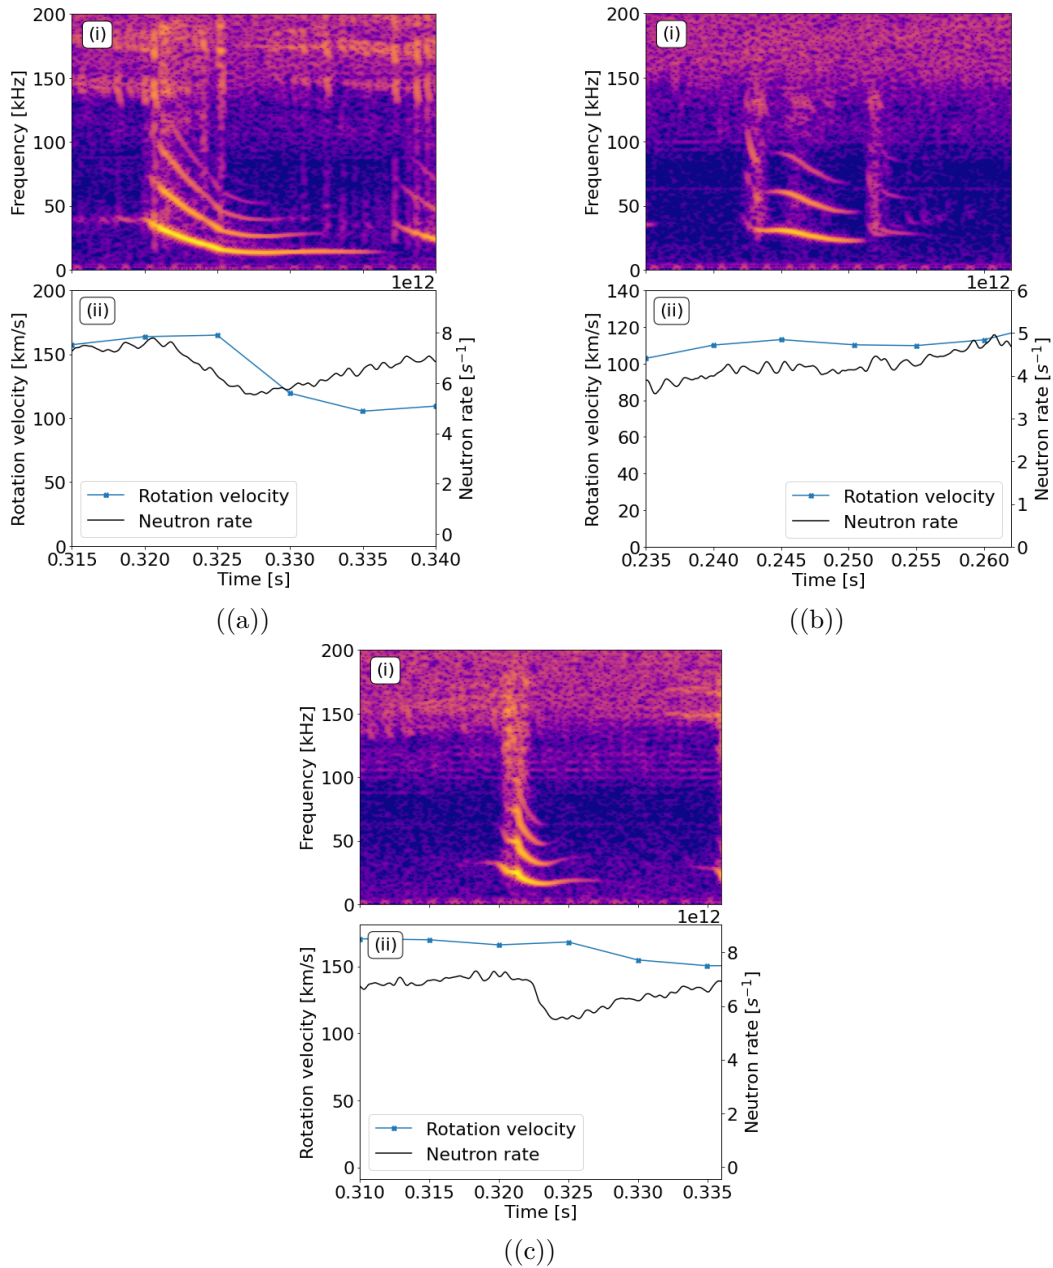


Figure 6.5: The (i) spectra, and (ii) global neutron rate and rotation velocity at $R = 1.04$ m for shots (a) 45026, (b) 45031, and (c) 45163, focussed specifically on the time of the analysed fishbone in each shot. In (a) and (c), there is a drop in the neutron rate during the fishbone, followed shortly after by a drop in the rotation rate. In (b), the effect appears to be smaller. There is a small drop in the rotation rate, and a flattening of the previously growing neutron rate.




Figure 6.6: Wavelength integrated (660.77 nm-661.52 nm) FIDA radial profiles before (blue) and after (red) the chirping phase of the selected fishbone in the shots (a) 45026, (b) 45031 and (c) 45163.

It was decided to investigate whether there may be a casual connection between the depletion of the fast ions (shown by the drop in the FIDA emission and the global neutron rate) and the drop in the rotation velocity. Previous studies have shown a relationship between the radial transport and loss of fast ions and toroidal deceleration. A. Thyagaraja *et. al.* [104] developed a model describing this effect. Succinctly, a radial current of fast ions must then produce an inward radial return current of bulk ions in order to preserve quasi-neutrality. This return current then crosses on the poloidal magnetic field to produce a $\mathbf{j} \times \mathbf{B}$ force in the counter-plasma current direction, acting to decelerate the bulk plasma, since the toroidal rotation is normally in the plasma current direction. The paper employs equation of motion of the bulk ions in the toroidal direction that includes a momentum injection rate per unit volume (from the neutral beam), ρ_{NBI} :

$$m_i n_i \frac{dv_\zeta}{dt} = \rho_{NBI} - m_i n_i \nu_\zeta v_\zeta - j_r(t) B_\theta \quad (6.1)$$

where v_ζ is the toroidal velocity of the bulk ions, m_i is the mass of a bulk ion, n_i is the number density of the bulk ions, ν_ζ is the toroidal momentum dissipation rate (assumed to be $1/\tau_E$, where τ_E is the energy confinement time). The NBI starts at $t = 0$, at which time $v_\zeta = 0$, and the fishbone that causes the radial outflow of fast ions appears at $t = t_0$. The return current j_r is assumed to have the form:

$$j_r(t) = \begin{cases} 0, & t < t_0 \\ j_0 e^{-\gamma(t-t_0)}, & t \geq t_0 \end{cases} \quad (6.2)$$

where the damping constant $\gamma = 1/\tau_{inst} > 0$, where τ_{inst} is a time constant associated with the instability causing the expulsion of fast ions. Equation 6.1 can be re-arranged into a form solvable with the integration factor $I(t) = e^{\nu_\zeta t}$, which we can integrate over the range $t = 0$ to $t = t'$:

$$m_i n_i \int_0^{t'} \frac{d}{dt} (e^{\nu_\zeta t} v_\zeta) dt = \rho_{NBI} \int_0^{t'} e^{\nu_\zeta t} dt - B_\vartheta \int_0^{t'} j_r(t) e^{\nu_\zeta t} dt \quad (6.3)$$

The integrals on the left hand side and in the first term on the right hand side can be evaluated trivially (with the initial condition that $v_\zeta(0) = 0$).

As j_r is 0 for $t < t_0$, the solution for $t < t_0$ is:

$$v_\zeta = \frac{\rho_{NBI}}{m_i n_i \nu_\zeta} (1 - e^{-\nu_\zeta t}), t < t_0 \quad (6.4)$$

For $t \geq t_0$ the second term on the right hand side of 6.3 becomes, for $\nu_\zeta \neq \gamma$:

$$j_0 e^{\gamma t_0} \int_{t_0}^{t'} e^{(\nu_\zeta - \gamma)t} dt = \frac{j_0 e^{\gamma t_0}}{(\nu_\zeta - \gamma)} (e^{(\nu_\zeta - \gamma)t'} - e^{(\nu_\zeta - \gamma)t_0}) \quad (6.5)$$

When $\nu_\zeta = \gamma$, the integral reduces to:

$$j_0 e^{\gamma t_0} \int_{t_0}^{t'} dt = j_0 e^{\gamma t_0} (t' - t_0) \quad (6.6)$$

Bringing all the cases together gives (replacing t' with t):

$$v_{\zeta} = \begin{cases} \frac{\rho_{NBI}}{m_i n_i \nu_{\zeta}} (1 - e^{-\nu_{\zeta} t}), & t < t_0 \\ \frac{\rho_{NBI}}{m_i n_i \nu_{\zeta}} (1 - e^{-\nu_{\zeta} t}) + \frac{B_{\vartheta} j_0}{(\gamma - \nu_{\zeta}) m_i n_i} (e^{-\gamma(t-t_0)} - e^{-\nu_{\zeta}(t-t_0)}), & t \geq t_0, \nu_{\zeta} \neq \gamma \\ \frac{\rho_{NBI}}{m_i n_i \nu_{\zeta}} (1 - e^{-\nu_{\zeta} t}) - \frac{B_{\vartheta} j_0}{m_i n_i} (t - t_0) e^{-\gamma(t-t_0)}, & t \geq t_0, \nu_{\zeta} = \gamma \end{cases} \quad (6.7)$$

Note that at time $t = t_0$ and $t = \infty$ the $t \geq t_0, \nu_{\zeta} \neq \gamma$ and $t \geq t_0, \nu_{\zeta} = \gamma$ branches of the solution return to the $t < t_0$ solution, ensuring that v_{ζ} is continuous in time, even when the current perturbation is not. Note also that for any $t \geq t_0$ the second term in each of the $t \geq t_0, \nu_{\zeta} \neq \gamma$ and $t \geq t_0, \nu_{\zeta} = \gamma$ branches is always negative, irrespective of the value of either ν_{ζ} or γ . Therefore it can be said that for a outward radial current of fast ions, which generates a return bulk ion current to maintain quasi-neutrality, a force is produced on the bulk plasma that acts to slow the rotation velocity for a duration governed by either ν_{ζ} or γ , whichever is smaller.

6.1.1.1 Using the derived model

A "worked example" of the analysis will be given for shot 45026 (as it has the most distinct drop in rotation velocity), but the results of the analysis for all three shots (45026, 45031, and 45163) will be discussed below.

In principal, an estimate of the return current could be found by estimating the likely density of fast ions in the plasma core, along with assumptions about the fraction of fast ions expelled by the instability and the timescale over which this occurs. However, these values are highly uncertain, and so an alternative approach is adopted assuming the model in the previous section is applicable, and solving the second case of equation 6.7 for j_0 :

$$j_0 = \frac{\gamma - \nu_\zeta}{B_\theta (e^{-\gamma(t-t_0)} - e^{-\nu_\zeta(t-t_0)})} \left(m_i n_i v_\zeta - \frac{\rho_{NBI}}{\nu_\zeta} (1 - e^{-\nu_\zeta t}) \right) \quad (6.8)$$

It should be noted that the model does not directly account for spatial variation of the plasma, and it can be observed that at the magnetic axis, B_θ goes to 0, which would imply an undefined value for j_0 . Instead, to use the model it is necessary to average all profiles over the region within which there is significant deceleration. This is the core region, bounded by limits defined by the first radii at which the change in the velocity of the plasma is less than 5%, moving from the magnetic axis to both higher and lower major radius, shown in Figure 6.7. The values of parameters before the crash are used.




Figure 6.7: Rotation velocity profile at start and end of the fishbone (to within the time resolution of the diagnostic). Hashed region is ignored when averaging profiles.

For this calculation it is assumed that $Z_{eff} \approx 1$, therefore by quasi-neutrality $n_i \approx n_e$. Averaged over the region of interest this gives $n_i = 1.56 \times 10^{19} \text{ m}^{-3}$. As discussed above, ν_ζ is assumed to be approximately equal to the inverse of the energy confinement time. The stored plasma energy is estimated from EFIT, giving $\nu_\zeta = 55.9 \text{ s}^{-1}$. The decay rate of the return current γ is assumed to be equal to $1/\tau_{instability}$, where $\tau_{instability}$ is the timescale of the instability. The drop in the FIDA emission occurs during the chirping-down period of the fishbone, which

occurs over 0.008 s, giving $\gamma = 125 \text{ s}^{-1}$. B_{θ} is estimated from EFIT and averaged over the region of interest, giving $B_{\theta} = 0.0442 \text{ T}$. The rotation at $t = 0.016 \text{ s}$, after both the chirping and non-chirping phases of the fishbone, can be obtained from the CXRS data. Averaging over the region of interest gives $v_{\zeta} = 96.9 \text{ km s}^{-1}$

As we take ρ_{NBI} to be a constant, the first case of Equation 6.7 can be used to generate a simple empirical estimate of this parameter. Early in the beam-on time of the plasma ($t \ll 1/\nu_{\zeta}$) we can approximate $e^{-\nu_{\zeta}t}$ with a Maclaurin expansion and neglect terms of order t^2 and above:

$$v_{\zeta} \simeq \frac{\rho_{NBI}}{m_i n_i \nu_{\zeta}} \left[1 - \left(1 - \nu_{\zeta}t + \frac{1}{2}\nu_{\zeta}^2 t^2 - \dots \right) \right] \simeq \frac{\rho_{NBI}t}{m_i n_i} \quad (6.9)$$

Therefore:

$$\rho_{NBI} = m_i n_i \frac{v_{\zeta}}{t} \quad (6.10)$$

This gives an estimation of $\rho_{NBI} = 0.03 \text{ kg m}^{-2} \text{ s}^{-2}$.

Together, the parameter values listed above imply a value of j_0 returned by Equation 6.8 that is negative, which is unphysical. A value of $9 \times \rho_{NBI} = 0.28 \text{ kg m}^{-2} \text{ s}^{-2}$ is required for a positive j_0 . For shot 45031, the observed drop in rotation is not significant enough to properly define a region over which we can average profiles to repeat the above analysis, but it is possible for shot 45163, and similarly gives an unphysical result for the value of j_0 . This suggests that it is unlikely that the drop of rotation was due to a counter-rotation force being applied due to a return current of bulk ions caused by an outward flow of fast ions. This leaves two other possibilities: that the rotation drop was due to a reduction in the direct fast-ion torque or the presence of a counter-current MHD ($\mathbf{j} \times \mathbf{B}$) torque associated with the fishbone. We now discuss the first of those two possibilities, with reference to shot 45026. The second possibility is more difficult to quantify because Mirnov coil measurements do not provide a direct measurement of the mode amplitude in the

plasma, which we would need to evaluate the ($\mathbf{j} \times \mathbf{B}$) torque associated with the mode.

6.1.1.2 Loss of momentum drive with loss of fast ions

The ρ_{NBI} term in Equation 6.1 is dependent on the input beam ions exchanging momentum through collisions. It is possible that during the instability, the population of fast ions exchanging momentum with the bulk ions is expelled, along with the fast ion population that is contributing to the FIDA light, causing the observed drop in FIDA signal. These populations are not necessarily the same. FIDA light is only observable at wavelengths above the location of the beam emission peaks. FIDA emission at a particular wavelength corresponds to a minimum energy, E_{min} , (due to Doppler shift) of:

$$E_{\text{min}} = \frac{m_i}{2} \left(c \left(\frac{\lambda - \lambda_0}{\lambda_0} \right) \right)^2 \quad (6.11)$$

where $m_i = 3.34 \times 10^{-27}$ kg is the mass of deuterium, c is the speed of light, $\lambda_0 = 656.1$ nm is the rest wavelength of the deuterium- α transition, and λ is the wavelength of the emission in question.

Due to the presence of the beam peaks, observation of lower energy fast ions is limited. This is potentially problematic as at higher energies fast ions mainly undergo electron drag, so experience little pitch angle scattering or ion drag and impart negligible momentum to the bulk ions. Only after the fast ions have slowed down to the critical energy E_{crit} do they begin to significantly interact with the bulk ions.

The maximal value of T_e during the shot is ~ 1.2 keV, in the plasma core. From Equation 1.12 under the approximation for a pure deuterium plasma, using a deuterium beam, this gives an $E_{\text{crit}} = 22.32$ keV. By re-arranging Equation 6.11 it can be found that the corresponding red-shifted D- α wavelength is 659.3 nm. In this

shot, the location of the primary energy beam peak varies over channel between 658.9 nm-660.4 nm. Due to the finite width of the beam peaks wavelengths of 659.3 nm and below are inaccessible to any channel, so FIDA cannot provide direct measurements of the population that primarily drive the momentum transfer to the bulk ions. Assuming that the drop in rotation during the major fishbone is caused by a reduction in the momentum transfer from the beam to the plasma, it is implied that the fishbone is causing an expulsion of both the relatively low energy fast ions that are the major part of the rotation drive and an expulsion of the higher energy fast ions that can be observed by the FIDA diagnostic. The current SSNPA diagnostic has a minimum observable energy "cutoff" at around 25 keV. This is not a hard limit, but foil transmission decreases quite rapidly below this energy. The crude energy resolution of the SSNPA means that the energies below the critical energy are not well resolved, compared to that achievable to a conventional NPA. An alternative FIDA view (like the proposed view in Section 3.6) would also potentially allow access to energies below the critical energy, but would depend on the specific lines of sight chosen to ensure sufficient sensitivity to these lower energies and that the beam peaks are not obscuring the features. However, the previously mentioned possibility of a torque associated with the MHD activity of the fishbone itself cannot be eliminated without further analysis.

6.1.2 Analysis with TRANSP and FIDASIM

6.1.2.1 Using Mirnov coil data to estimate anomalous diffusion

The presence of significant MHD activity throughout shot #45026 means that any TRANSP runs performed need some degree of anomalous fast-ion diffusion to describe aspects of the plasma behaviour. Additionally, the two sawtooth crashes near the end of the beam-on time would necessitate the use of the sawtooth models, as discussed in Chapter 4. However, at the time of analysis, only the magnetics constrained EFIT was available. At 0.3883s, the time of the first sawtooth, the

EFIT derived value of q_0 was 0.99. Similar to the difficulties mentioned in Section 4.4, this appears to not be low enough to cause the sawtooth model to trigger correctly. For this reason analysis of the sawteeth in this shot was not performed. Upon the availability of better constrained EFIT++ runs, it may be possible to re-visit such an analysis.

For the TAE and fishbone part of the shot, it was decided to perform modelling using a time-varying anomalous diffusion provided to TRANSP. Previous studies [105]–[110] found that there is some evidence that the global neutron rate, measured by the fission chamber, is proportional to the root mean square (RMS) amplitude of the Mirnov coil signal during TAEs and fishbones. As an initial investigation it was decided to assume that the overall anomalous diffusivity is proportional to the RMS amplitude of the Mirnov coil to construct a time-trace of the overall anomalous diffusion. This is a fairly simple model, accounting for time-variation, but not examining radial variation (which would be expected in the case of the fishbone, as it is a core-localised mode) or variation in the pitch or energy of the affected particles (again, the fishbone is primarily thought to affect trapped particles). A refined model including these dependencies would be a natural extension of this work that could be performed in the future.

The Mirnov signal utilised is from the coils specifically filtered to examine relatively low (≤ 200 kHz) frequencies. The signal is additionally low-pass filtered with a 200 kHz cutoff to further remove any high-frequency noise, then binned on 0.003 s intervals to match the utilised FIDA data timescale, then the RMS value is found. This allows a trace in time on the the averaged FIDA timescale. In order to scale this trace such that it could be used for the values of the diffusion, the scaling factor was manually set such that the drop in the neutron rate during the "major" fishbone was matched. Due to the consistent overprediction of the value of the neutron rate by TRANSP compared with the fission chamber measurements as mentioned in Section 2.2.1.1, there was a degree of ambiguity in scaling of the fission chamber measurements as well. It was reasoned that there was a figurative tension




Figure 6.8: Anomalous diffusion trace used in TRANSP modelling derived from binned RMS Mirnov coil signal, scaled in order to match the drop in neutron rate during the major fishbone, for a reasonable scaling of the fission chamber data.




Figure 6.9: Traces of the neutron rate, comparing the fission chamber data (as calibrated at the time of analysis) to TRANSP simulations with various levels of anomalous diffusion, and with none.

between scaling the diffusion up, to lower the neutron rate to match experimental values, and scaling the fission chamber fluxes up, to match the predicted values. Scaling the experimental values up brings the neutron rate during the TAEs closer to the prediction, therefore requiring a smaller anomalous diffusion, but would also increase the absolute size of the drop in the neutron rate during the "major" fishbone, which would require a bigger anomalous diffusion in order to match. By scaling the anomalous diffusion behaviour of the TAEs and fishbone together, the parameter space of the possible scaling of both the fission chamber and diffusion could be restricted, and the number of free parameters in the model reduced. Figure 6.8 shows the constructed trace of the anomalous diffusion, which was used to perform a high fidelity TRANSP run. Note that this analysis was restricted to examining only up to directly after the major fishbone.

To probe the parameter space described above, TRANSP runs were also performed with double and half that shown in Figure 6.8. The measured neutron rate was then fit to each of the synthetic traces using a multiplicative scaling factor a i.e. $a \times$ measured neutron rate using a minimisation of the sum of the squared residuals. The results of this investigation are shown in Figure 6.10.




Figure 6.10: Measured neutron rate fitted to each of the 3 trial anomalous diffusion traces with a multiplicative scaling factor, derived using the minimisation of the sum of the squared residuals.

The original trace required $a = 1.861$, the one with halved anomalous diffusivity $a = 2.179$, and the one with doubled anomalous diffusivity $a = 1.471$. The coefficient of determination R^2 allows the goodness of fit of each to be parameterised. The original trace gave $R^2 = 0.964$, the halved $R^2 = 0.975$, and the doubled $R^2 = 0.913$, suggesting that the half trace gives a slightly better fit. However, the original scaling of the anomalous diffusion was chosen to replicate the drop at the major fishbone. The relative drop of the measured neutron rate during the major fishbone (which is independent of the scaling factor) is 27.1%. The drops using the original, half and doubled traces respectively are 26.2%, 14.2% and 45.4%. It was decided to continue to utilise the original trace, as it best represents the behaviour at the fishbone, while reasonably reproducing the TAE behaviour. A more detailed examination of the parameter space could allow for a more exact determination of the scaling, which could be a goal of future work.

6.1.2.2 TRANSP and FIDASIM results

A series of times during the shot were simulated using the TRANSP -> FIDASIM workflow. It was noticed that channels towards the core did not give good agreement between experiment and simulation, while channels further towards the edge gave better agreement, but only for the more Doppler shifted (therefore higher minimum energy) part of the observable FIDA emission. Figure 6.11 shows time traces of the FIDA and FIDASIM emission, integrated over 660.52 nm-662.0 nm, which corresponds to a range of minimum energies of $\sim 42.5 \text{ keV} - 76 \text{ keV}$ (with some degree of uncertainty due to instrumental broadening). It should be noted that this range encompasses only fast ions that originated from the full energy component of the beam.




Figure 6.11: Wavelength-integrated time traces for the experimental FIDA data and FIDASIM data for Shot #45026. (a) is at $R = 1.029 \text{ m}$, (b) is at $R = 1.293 \text{ m}$. Based on the results of Section 2.3, an characteristic error of 20 % error is assumed on the FIDASIM values. A full deterministic sampling analysis could be valuable here to determine a specific value.

For Figure 6.11(a) at $R = 1.029 \text{ m}$ there is reasonable agreement at earlier times, with divergence occurring at $\sim 0.25 \text{ s}$, after which the experimental measurement increases, whereas the FIDASIM radiance has a slow decrease. Additionally, the major fishbone causes a noticeable drop in the experimental data, but this is not reflected by even a relative drop in the synthetic trace. The simple radius- and

energy-invariant anomalous diffusivity is not sufficient to replicate the behaviour. In contrast, in Figure 6.11(b) at $R = 1.293$ m there is good agreement in both the magnitude and behaviour over time of the FIDASIM compared to the experiment. There is a general increase in signal from $\sim 2.5 \times 10^{15} \text{ Ph s}^{-1} \text{ m}^{-2} \text{ sr}^{-1}$ to $\sim 8 \times 10^{15} \text{ Ph s}^{-1} \text{ m}^{-2} \text{ sr}^{-1}$ across the TAE phase of the shot. At the first, small fishbone at $t \sim 0.283$ s, there is a drop in the synthetic signal, but this does not match the magnitude of the drop. However, the drop in signal due to the major fishbone is modelled fairly well. Figure 6.12 shows the experimental and FIDASIM spectra at $R = 1.293$ m before (a) and after (b) the major fishbone.




Figure 6.12: Experimental (blue) and FIDASIM (red) spectra before (a) and after (b) major fishbone at $R = 1.293$ m. Green region designates integration range used in Figure 6.11. Based on the results of Section 2.3, a characteristic error of 20 % error is assumed on the FIDASIM spectra.

The drop in spectral radiance of the experimental data at the higher wavelengths ($\lambda \geq 660.5$ nm, $E_{min} \gtrsim 42.6$ keV) is matched well by the synthetic trace, but at lower wavelengths neither before or after is matched well. A more sophisticated model of the anomalous diffusion, taking into account variation of the behaviour with energy could allow this discrepancy to be probed further. There could be some effect due to the uncertainty of the beam fractions as described in Chapter 5. While the maximum wavelength at which there would be a contribution from the half energy (32.9 keV) purely based on the Doppler effects is 660 nm, instrumental broadening could cause some contribution at the lower wavelengths presented here. There could also be problems with the background subtraction, perhaps due to

improperly chosen active-passive channel pairings (see Section 3.5.4). Further investigation is required.

Figure 6.13 shows the synthetic fast-ion distributions produced by TRANSP/NUBEAM, with a lime green contour denoting the extent weight function for the part of the spectrum that had good experiment/simulation agreement overlaid.




Figure 6.13: Synthetic fast-ion distributions produced by TRANSP/NUBEAM before (a) and after (b) major fishbone at $R = 1.293$ m, along with the relative change (c). Lime contours denote the extent of the weight function, which is only for those wavelengths that had good agreement between the experiment and simulation. Therefore data within the contour can be utilised for analysis, whereas the data outside of the contour either cannot be seen by the diagnostic or has been shown in Figure 6.12 to not match experiment. Dark green region in (c) designates region where the fast ion distribution was 0 before the crash.

Within the lime green contour we can have some degree of trust in the distribution. Figure 6.13(c) shows how the the major fishbone has caused a large depletion of the high energy fast ions, with a drop of $\sim 44\% - 50\%$ across the pitch range $-0.5 - -1.0$, although it must be remembered that the basic anomalous diffusion model used was pitch invariant.

6.1.2.3 Conclusion

It is unlikely that the sharp drop in rotation observed during the major fishbone in shot #45026 is due to a $\mathbf{j} \times \mathbf{B}$ force generated by a bulk ion return current in response to a outflowing fast ion current, as the magnitude of the current needed is

larger than that likely to arise from expulsion of fast ions from the fishbone region. Investigating the alternative hypothesis of a sudden reduction of the fast ions most responsible for momentum transfer, those below the critical energy, is not possible with FIDA due to the relevant wavelength ranges being obscured by the presence of the beam peaks. A third possibility, that there is a counter-current MHD torque associated with the major fishbone, cannot be ruled out.

We have used a scaled and smoothed Mirnov coil trace as a proxy for the anomalous diffusion time evolution. Comparing this model with experiment is complicated by uncertainty in the absolute values of the neutron rate from the fission chamber, but we have endeavoured to circumvent this issue by matching to experiment the relative drop in the neutron rate before and after the major fishbone. Using this to generate synthetic spectra, it is found that one of the FIDA channels gives good agreement between synthetic and experimental data. More sophisticated anomalous diffusion profiles, varying in both radii and phase-space, may provide more comprehensive agreement between modelling and experiment. For the matching channel, it was additionally found that only the higher energy part of the observable spectra gave a good match. Further investigation is required to properly diagnose this discrepancy. For the matching part of the spectra, using the regions of the phase-space matched by experiment and simulation show a drop of $\sim 50\%$ of the fast ions is observed after the fishbone.

6.2 Effect of long-lived mode and locked mode on fast ions

Shot #45091 was also performed as part of the MU01-FPP-01 experiment. Figure 6.14 shows some plasma parameters for the shot. The pulse was heated with both beams, with the SS beam operating over 0.1681 s-0.6153 s, with an average power of 1.47 MW and average primary energy of 63.4 keV. The SW beam operated over 0.1535 s-0.2105 s before experiencing an unintended drop-out that lasted over the

period 0.2105 s-0.278 s, after which it resumed until its termination at 0.5973 s. The two "on" portions of the SW beam time had very similar average powers and primary energies of 1.71 MW and 72.5 keV for the first portion, and 1.74 MW and 72.6 keV for the second portion. The Mirnov coil signal and spectrogram, shown in Figure 6.16, shows that after the onset of the SS beam TAEs are excited (as discussed with shot #45026 in Section 6.1, the mode exists within a range broadly consistent with with the estimated $\omega_{\text{TAE}}/2\pi$ using the plasma parameters for the time of the mode). These persist until ~ 0.32 s, after which there is a very brief period of quiescence until ~ 0.35 s.

Subsequently there are a few fishbones that rapidly transition to an LLM. Just after 0.506 s a sawtooth crash occurs, evidenced by a large drop in the neutron rate (Figure 6.14(b)), a very high amplitude spike in the Mirnov coil signal (Figure 6.16(a)) which corresponds to a very short duration broad-band feature seen in Figure 6.16(b), and a reversal of the downward trend of the value of q_0 given by the magnetics-constrained EFIT equilibrium, increasing from ≈ 0.8 just before this time to ≈ 1 in the time after it. After the sawtooth, an increase in density (Figure 6.14(c)) coupled with a reduction in the D- α intensity (Figure 6.14(d)) points towards a transition from L to H mode. There are some studies [111], [112] showing evidence of a causal link between sawtooth crashes and L to H-mode transitions, specifically that the sawtooth generates a very rapid heat flux outwards from the core to the plasma edge, which stimulates a mode transition. However, this is still not fully understood and is an area undergoing active research.

During the H-mode, the LLM continues, falling in frequency as the plasma rotation slows down, which drops to approximately 0 km s^{-1} at ~ 0.54 s, which is reflected in the frequency of the LLM going to 0 kHz. The plasma appears to drop out of H-mode around this time as well. The amplitude of the integrated saddle coil signal, shown in Figure 6.15, implies that a locked mode evolves here after this, which persists until shortly before the SW beam terminates. The SS beam terminates shortly after.




Figure 6.14: Time traces of plasma parameters of Shot #45091. (a) NBI power, (b) neutron rate, (c) line integrated electron density, (d) D- α intensity.




Figure 6.15: Amplitude of the saddle coil data (black) and neutron rate (green) vs. time focusing on the times just before and after the plasma bulk rotation goes to zero. The first drop in the neutron rate is due to a sawtooth, but subsequent drops are correlated with the behaviour of the locked mode. SS beam off time is shown in red, SW off time in blue.

The behaviour in time of the FIDA data for a representative sample of channels across the plasma radius is given in Figure 6.17, both for the full duration of the SS beam on (6.17(a)) and zoomed to better show the behaviour during the sawtooth and locked mode period (6.17(b)).

A drop in the FIDA emission can be observed in the outermost channels (of which Figure 6.17(a)(iii) is an example) during the duration of the SW beam drop-out, which corresponds to a small drop in the neutron rate (Figure 6.14(c)). This drop is not observed in core or mid-range channels. As the SW beam injects far from the midplane (the beamline is centered on machine coordinate $Z = 0.65$ m), near the plasma edge, we would expect that the fast ions injected by the SW beam to be found nearer the plasma edge, which is borne out by the FIDA data. It is not immediately apparent if there is an effect on the fast ions from the LLM starting at ~ 0.38 s. An analysis of these measurements utilising FIDASIM is discussed in 6.2.1.






Figure 6.16: (a) Time trace of the low-frequency Mirnov coil signal and (b) Spectrogram of low-frequency Mirnov coil time trace in shot #45091, showing the main resonant behaviour of the plasma. Specifically, after the SW and SS beam turn on (times of this are shown by the magenta and green vertical dotted lines respectively) TAEs are excited. Shortly thereafter, the SW beam erroneously turns off (magenta dot-dashed line) for a short duration while TAEs continue before it comes back on (blue dashed line). There is a very short period of relative quiescence before the excitation of fishbones that rapidly transition to an LLM. There is a sawtooth crash just after 0.5s, but this does not halt the the plasma rotation slowdown generated by the LLM. The bulk plasma comes to a halt at ~ 0.54 s. A locked mode, invisible to the Mirnov coil due to it being stationary in the machine frame, is excited here, which persists until ~ 0.59 s, after which it appears to reduce in magnitude or stops entirely. Both beams terminate soon after (SW - blue dot-dashed line, SS - green dot-dashed line).

At ~ 0.5 s the sawtooth crash occurs, which is accompanied by an immediate drop in the integrated FIDA emission of $\sim 50\%$ across the mid-range channels (Figure 6.17(b)(ii)), with only small apparent drops in the in the core and outer channel (Figures 6.17(b)(i) and (iii)). Due to the lack of a MSE-constrained EFIT++ at the time of analysis it was decided that there was only limited value in examining the



(a) Time trace of wavelength integrated FIDA data for the full duration of the SS beam. SS beam on and off times shown by dotted and dash-dotted red lines respectively. The first SW beam off time is shown by a brown dash-dotted line, the second SW on time and the final off time are shown by green dotted and dash-dotted lines respectively.



(b) Time trace of wavelength integrated FIDA data for only the end of the shot, showing in more detail the effect of the sawtooth crash and the locked mode. The SS and SW beam off times are shown by red and green dash-dotted lines respectively.

Figure 6.17: Wavelength integrated (660.46 nm-661.93 nm) time traces of the FIDA data both for the full duration of the SS beam (a) and zoomed into the end of the shot (b) at (i) $R = 1.029$ m, (ii) $R = 1.23$ m, and (iii) $R = 1.354$ m.

sawtooth with the TRANSP/NUBEAM sawtooth models. The locked mode that develops after the sawtooth appears to cause significant fast-ion redistribution, as shown in Figure 6.15, where it can be observed that significant drops in the neutron rate that occur at approximately 0.54 s, 0.555 s and 0.577 s roughly correlate with stationary points of saddle coil amplitude. Across all of the FIDA channels from $R = 1.029$ m outwards there is an overall decrease of the FIDA emission as the saddle coil amplitude rises. In the mid-range and outer channels the FIDA signal becomes negative during the first peaking of the saddle coil amplitude at 0.555 s,

which indicates a problem with the background subtraction. Locked modes are a persistent and toroidally localised feature, so a breaking of the axisymmetry is therefore expected and could explain the negative values since the subtraction of passive FIDA emission is based on the assumption of axisymmetry. If the magnetic island associated with the locked mode were in the line-of-sight of the outer passive channels, the bulge in the plasma at this location due to the instability could cause the overall length of the intersection of the line-of-sight and the plasma to increase, causing an apparent increase in the passive emission. Care must be taken when analysing such FIDA data, and due to these problems not too much can be said. However, the overall drop in neutron rate does suggest a real loss of fast ions. The small sharp drops in the neutron rate during this time correlate with smaller drops in the FIDA emission, especially in the outer channels. However, the brief rally in the FIDA emission before the termination of the SW beam apparent in Figure 6.17(b) remains unexplained.

While the overall drop in bulk rotation to zero is correlated with the LLM it can be observed, in Figure 6.18, that upon the advent of the locked mode a highly localised region of the radius, $\approx R = 1.25 \text{ m} - 1.35 \text{ m}$, experiences a significant counter-current velocity of $\sim -20 \text{ km s}^{-1}$. The localisation corresponds to a flattening in the electron temperature over the locked mode duration, shown in Figure 6.19, which is indicative of the location of a magnetic island, suggesting that the counter-current rotation is caused by the locked mode. Locked modes can exert a counter-rotation torque on plasmas [33], but there could be some effect due to the expulsion of fast ions causing a bulk return current that exerts a torque, previously discussed in Section 6.1.1 and in [104]. Further investigation would be required to precisely determine the causes of this reverse rotation.

6.2.1 TRANSP/FIDASIM analysis of LLM

In order to better understand the effect of the LLM on the fast ions, analysis with TRANSP/NUBEAM and FIDASIM is prudent. As previously stated, the

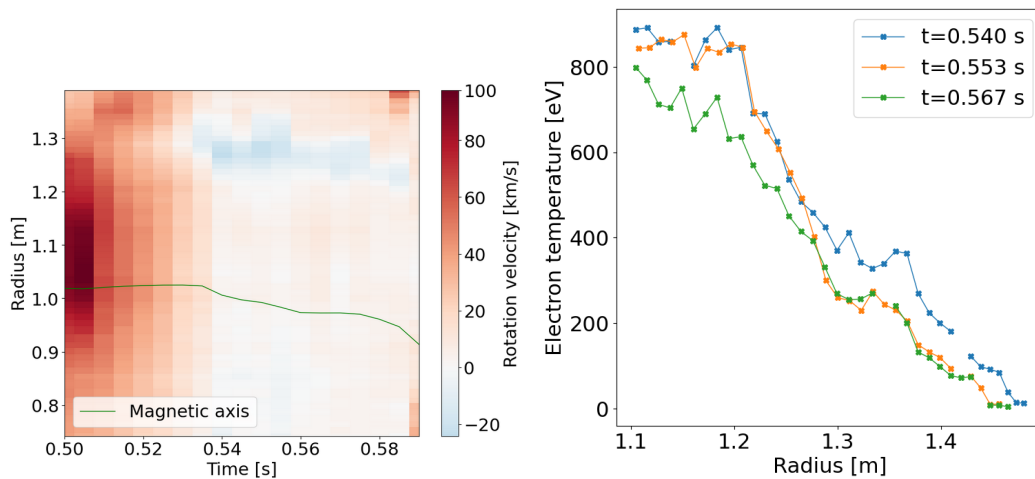


Figure 6.18: Toroidal rotation profile from CXRS in shot #45091, focusing specifically on the post-sawtooth, locked mode period of the shot. After the sawtooth the plasma slows rapidly across the entire plasma radius, coming to a near stop at ~ 0.54 s, except for a spatially localised region at ~ 1.25 m – 1.3 m, which exhibits a counter-current rotation of ~ 20 ms^{-1} .

Figure 6.19: Electron temperature profiles from TS for the outboard region of the plasma in shot #45091, for times when the plasma rotation has gone to near zero (or negative, as shown in Figure 6.18). A plateauing of the temperature can be seen at ~ 1.3 m, indicative of the existence of a magnetic island (the locked mode).

presence of instabilities not directly simulated by TRANSP/NUBEAM requires the use of something like anomalous diffusion to account for the redistribution caused by them. However, comparisons between quiescent and non-quiescent portions of a shot allow for a comparative analysis, and statements can be made about the effects of the instabilities on the fast ions. SS beam shots on MAST-U are extremely susceptible to instabilities, and finding a shot with a significant duration of quiescent plasma is difficult. In this shot there is a brief period of quiescent plasma between 0.32 s-0.35 s (see Figure 6.16). Using Equations 1.14 and 1.16 an estimate of the fast-ion slowing down time can be found for the plasma at each FIDA channel radius, based on TS data. Due to the nature of the laser setup for the TS, the precise radius of inspection by the diagnostic varies slightly in time, cycling positions every 5 measurements in time. However, at no time is any FIDA channel more than 0.0052 m from a TS measurement radius, and the maximum separation between neighbouring TS radii is 0.012 m. It is therefore deemed acceptable to




Figure 6.20: Fast-ion slowing down time at each FIDA channel averaged across the quiescent plasma period 0.32 s–0.35 s for shot #45091. Error bars here are the standard deviation of the slowing down times over this period.




Figure 6.21: Detail of Figure 6.16, showing the Mirnov signal (a) and spectrogram (b) zoomed into the fishbones and LLM specifically.

linearly interpolate the TS radii to the FIDA channel radii. Figure 6.20 shows the calculated slowing down time averaged over the quiescent period for the different FIDA channels

It can be seen that the slowing down time is at a maximum in the core region (~ 1.0 m) with a value of ≈ 0.07 s, which then drops rapidly across the more edge region. While it is true that the brevity of the quiescent duration means that the full steady state fast-ion distribution will not be entirely reestablished, the higher energy portion of the distribution will be reestablished much sooner than a full slowing down time (see [113], particularly Figure 1). For beam ion energies that are large compared to the bulk ion temperature, it can be shown [5] that the energy of a beam ion after time t , $E_b(t)$, injected at energy E_{b0} is:

$$E_b = E_{b0} \left(e^{-3t/\tau_{se}} - \left(\frac{E_{\text{crit}}}{E_{b0}} \right)^{3/2} \left(1 - e^{-3t/\tau_{se}} \right) \right)^{2/3} \quad (6.12)$$

Ions born at the start of the quiescent phase will have energy ~ 30 keV by the end (for the core region, ions in regions nearer the edge will have lower energy), so

with the FIDA analysis already restricted to only the > 40 keV high energy portion of the distribution (due to the beam peaks obscuring the low energy part of the spectrum), it is therefore justified to consider the FIDASIM synthetic spectra (for this energy range) to be valid near the end of the quiescent phase.




Figure 6.22: Radial profile of the FIDA and FIDASIM signals at $\lambda = 660.69$ nm ($E_{min} = 45.9$ keV) averaged over the times 0.347 s-0.35 s, the end of the quiescent period in shot #45091. Note that the channels in the range 0.77 m-0.96 m do not have usable experimental data at the selected wavelength due to being obscured by the beam peaks. While the general profile shapes in the FIDA (blue) and FIDASIM (red) agree, the absolute values do not. An additional scaling factor of 1.45 brings agreement around the peak values (green) but there remains a discrepancy in the edge channels, which requires further investigation to resolve. Based on the results of Section 2.3, a characteristic error of 20% error is assumed on the FIDASIM values, but this does not fully resolve the discrepancy.

Figure 6.21 shows the Mirnov coil signal and spectrogram, zoomed in to show the fishbone/LLM period in more detail. Within this time range a number of FIDASIM synthetic spectra have been generated to allow comparison with the experimental FIDA.

Figure 6.22 shows the radial profile of the FIDA and FIDASIM signals for $\lambda = 660.69$ nm, for a time slice at the very end of the quiescent period, 0.347 s – 0.35 s. Based on the arguments given above, there should be agreement between the ab-

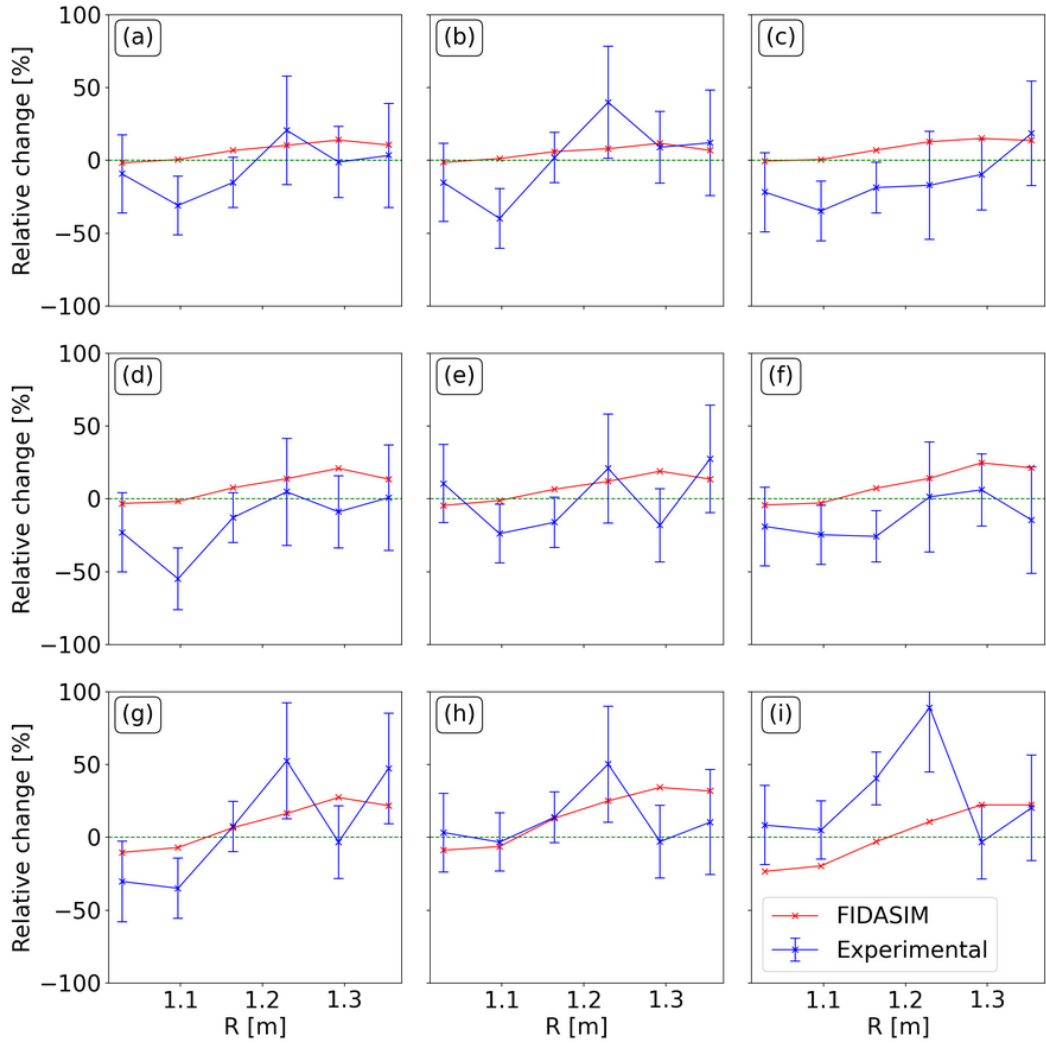


Figure 6.23: The percentage change in both the FIDA and FIDASIM signal at $\lambda = 660.69$ nm ($E_{min} = 45.9$ keV), going from the end of the quiescent duration (0.347 s-0.35 s) to subsequent time-slices throughout the fishbone/LLM portion of shot #45091. The time-slices for the "after" duration are (a) 0.366 s-0.369 s, (b) 0.377 s-0.380 s, (c) 0.383 s-0.386 s, (d) 0.390 s-0.393 s, (e) 0.397 s-0.4 s, (f) 0.407 s-0.41 s, (g) 0.437 s-0.44 s, (h) 0.467 s-0.47 s, and (i) 0.492 s-0.495 s. Based on the results of Section 2.3, an approximate absolute error of 30 percentage points is assumed on the FIDASIM values. A full deterministic sampling analysis would be useful to determine specific uncertainties.

solute values of the FIDA and FIDASIM intensities, but it can be seen that while the profile shapes are similar in the mid-range channels the absolute values are not. A scale factor of 1.45 brings these into good agreement. However, there remains a significant disagreement in the edge channels, which may be due to incorrect background subtraction. Due to these disagreements between synthetic

and experimental spectra at the lower part of the observable FIDA wavelength range, and to see more clearly the effect on the FIDA emission of instabilities, the relative change in both the synthetic and experimental signal going from the quiescent period into the period with fishbones and the LLM is examined. Figure 6.23 shows the relative change across a number of time-slices for a specifically chosen wavelength at $\lambda = 660.69$ nm, which corresponds with $E_{min} = 45.9$ keV. Towards higher wavelengths, the signal-to-noise ratio of the experimental signal decreases, and so the measurements in this range are of limited use in examining the relative change. Lower wavelengths restrict the radial range that can be examined.

Figures 6.23(a) and (b) show time-slices during and just after the fishbone phase, with (b) at the beginning of the LLM. Initially the LLM has a frequency around ~ 23 kHz, which decreases across the duration analysed until it has ~ 10 kHz during (e). There appears to be some decrease in emission in the channels at 1.097 m and 1.164 m during (a), with only the former in the case of (b). A localised depletion of fast-ions near the core corresponding roughly with a modelled fishbone mode structure was observed previously on MAST [30]. Near the start of the LLM, during (c), there appears to be a decrease in emission over a wider range of radii, with the aforementioned channels 1.097 m and 1.164 m showing a drop that is outside the uncertainty of the measurement. Across the early part of the LLM, (d)-(f), the behaviour is similar, again with particularly large drops in the channels at 1.097 m and 1.164 m. This provides evidence that the LLM is causing strong redistribution of passing fast ions away from the core channels. In the latter half of the LLM, figures (g)-(h), the emission moves back towards a pre-LLM level, before exceeding it at mid-radii just before the sawtooth (see (i)). The lessening of redistribution seen in (g) and (h) could be due in part in the reduction of the amplitude of the LLM and also due to the mode frequency dropping. As this occurs it is possible that the region of phase-space in resonance with the mode changes, moving from an observable to an un-observable region.

With the FIDA lines-of-sight primarily observing passing fast ions, lines-of-sight

with a phase-space coverage focused more on the trapped particles, like the prospective view described in Section 3.6, could help illuminate this issue. The SSNPA diagnostic has more perpendicular views and so a phase-space coverage that is more towards the trapped particles, and the integration of its data into this analysis could also shed some light. During the time range observed in (i), the plasma rotation is experiencing a brief increase after a long decrease over the duration of the LLM, before the sawtooth crash and locked mode cause a very rapid drop to zero. While the mode amplitude from the Mirnov coil (Figure 6.21) does not appear to significantly change, an increase in rotation implies a reduction of any decelerating effects of the LLM (a torque is still being provided by the beams at this time). Further investigation is required to fully diagnose this feature.

6.3 Conclusion

The effects of a locked mode and an LLM on the fast ion population in shot #45091 have been investigated. Later drops in the FIDA emission are correlated with a sawtooth crash and changes in the amplitude of a saddle coil signal, indicating the presence of a locked mode. A deeper analysis of the sawtooth crash can be performed when MSE-constrained EFIT++ runs become available. The locked mode appears to cause large amounts of fast-ion redistribution, with the caveat that background subtraction can become complicated due to the breaking of the axisymmetry arising from the presence of the mode. The locked mode is observed to drive the overall plasma rotation to zero, with the magnetic island itself undergoing significant rotation counter to the initial plasma rotation direction.

The LLM behaviour has been investigated by utilising FIDASIM. A time slice at the end of a quiescent phase before the development of the fishbones and LLM is modelled on the basis of classical slowing-down, assuming that the high-energy population of the fast ions is in a classical slowing down steady state, and therefore the FIDA and FIDASIM radiances should match when the fast ions in

TRANSP/NUBEAM are treated classically. Using this as a point of reference, the behaviour of the experimental and synthetic spectra over the duration of the LLM is examined. In the early part of the LLM, drops in the FIDA emission can be observed in the more core-localised channels, with the drops reducing over the course of the LLM until just before the sawtooth, where the emission is greater compared to the pre-LLM emission. Indirectly, the increase in plasma rotation during this time implies a lessening of the effect of the LLM, which may be a component in the explanation of this behaviour.

Summary, conclusions and future work

The work presented in this thesis has documented the commissioning and exploitation of a fast-ion deuterium- α (FIDA) spectroscopic diagnostic, examining the behaviour of "fast" (supra-thermal) ions on the MAST and MAST-U tokamaks. In a future burning plasma deuterium-tritium device, fast ions (the α -particle DT reaction product) will constitute the majority source of plasma heating, so must be well confined. Understanding how fast ions affect and are affected by plasma instabilities is therefore an essential avenue of investigation on the road to the development of nuclear fusion as a source of energy.

In this chapter the results from the main chapters are summarised, with discussion of possible future directions for this research.

7.1 The principles of the FIDA diagnostic and its simulation

Within Chapter 2 an attempt to use the technique of Deterministic Sampling (DS) to propagate errors in TRANSP input profiles to the synthetic spectra produced by FIDASIM, as an extension of the use of the technique with TRANSP alone as

performed by Sahlberg *et. al.* [68]. The analysis looked at two variables, electron density and temperature, with an assumption that the error on both variables was Gaussian. This allowed for a simple ensemble that could capture the first four moments of the output distribution (although only the first two, mean and variance, were examined in this thesis). The technique was able to produce, with relatively few simulations, an estimate of the uncertainty on the FIDASIM synthetic spectra. In this thesis a single shot is examined to generate a characteristic error. In future investigations, it could be beneficial to produce a systemic workflow for performing DS on the TRANSP and FIDASIM simulations of every analysed shot, to allow for even more accurate predictions of the uncertainty. It would be possible to create an OMFIT module that allows the creation of a TRANSP run as normal, but allow the user to also specify the sigma-point ensemble. OMFIT could then systematically vary the input profiles and submit the TRANSP runs with the combination of inputs as specified by the ensemble, greatly streamlining the process of the creation of the samples needed. The analysis could also be extended to incorporate more variables. However, as discussed briefly in Chapter 2, increasing the number of variables causes the number of sigma-points needed to properly encode the underlying uncertainty to rise sharply. For example, adding a third Gaussian variable would have increased the four-moment ensemble size from five sigma-points to nine [69]. Incorporating variables whose uncertainty is not Gaussian (for example, Z_{eff} , which by definition cannot take a value below 1, so cannot be described by a simple Gaussian [68]) would require even more sigma-points. This is a compounded problem when the technique is propagated to FIDASIM, as each full set of sigma-points must be ran for every time in the shot that is analysed. DS is a powerful technique, but requires the user to think carefully about what is required for the specific analysis ahead of the simulation step, to prevent running huge numbers of time-intensive simulations.

7.2 The FIDA Diagnostic: setup and preparation for the MAST-U experimental campaign

In Chapter 3 an overall description of the setup of the FIDA diagnostic, along with its commissioning for MAST-U, was given. While most of the setup including the spectrometer and camera were the same as was used for the MAST campaigns, it was still required to ensure that the spectra were properly captured on the CCD. With the installation of new tokamak fibres, plus the usage of a new fibre head and the completely new periscope within the vessel to allow the lines-of-sight to avoid in-vessel components that were not present in MAST, it was required to re-determine the tokamak fibre transmission and the alignment of the fibres.

Overall, the determination of tokamak fibre transmission was successful, with some important caveats. The tokamak fibre transmission was derived relatively, as the transmission of spectrometer fibres was not found, but instead values from MAST were utilised. A full transmission analysis using an absolutely calibrated detector should be considered before a future experimental campaign, but the values calculated for Bundles 3 and 4 are reasonable. The calculated transmission of Bundle 5 is distinctly lower than the others. A scaling factor of 3 brought the bundle as a whole reasonably in to line with the transmissions calculated for the other bundles, and subsequent analysis on MAST-U data in conjunction with FIDASIM has not shown egregious discrepancies, suggesting that instrumental or procedural error is responsible. It would however be prudent to re-perform the procedure for the determination of the transmission on at least Bundle 5 during some inter-campaign period in the near future.

Overall, the fibre alignment is good, with excellent alignment of the active head and slightly poorer alignment of the passive head. The alignment issues are simply part of the reality of experimentation and the limitations in the construction of the passive head. It is unlikely that the passive head will be reconstructed at

any point in the near future, so this is simply something that has to be dealt with elsewhere. The determination of the tangency radii was performed using the positions as defined by Calcam, and verifies that for core and mid-range channels the tangency radii match-ups are good even with the poorer alignment. However, the outer channels have larger mismatches. The most outer channels show poor background subtraction and it is possible that this could be due to poor tangency radius matching. A further alignment of the outer channels is recommended to resolve this.

Currently the data calibration code, converting the raw data in counts/pixel to photons/s/nm/m²/sr, is entirely written in IDL. This is a proprietary programming language, requiring the purchase of a licence for its use. Additionally, code documentation is sparse to non-existent, making the on-boarding of new FIDA users difficult and time-consuming. It is the opinion of the author that a full re-write in the open-source programming language Python would be highly advantageous, making the code portable without the inconvenience and expense of licensing. The re-write should comment the code in full to allow new users to easily understand the purpose of every script and function, along with full documentation. The code and documentation should be made available on the CCFE GitLab, allowing internal and external collaborators to access FIDA data much more easily.

A vertical FIDA view, giving a much more perpendicular set of lines-of-sight to the beam would be highly advantageous for fast ion analysis. Such views would provide data on the regions of phase-space containing trapped fast ions, mostly invisible to the existing views. There is a degree of redundancy here when considering the SSNPA, which also has mostly perpendicular views onto the neutral beam, however the SSNPA has only discrete and coarse energy resolution, whereas a FIDA system would be able to observe a wide range of energies, especially when considering the much more advantageous placement of the beam peaks in wavelength-space compared to that of the existing tangential FIDA system. However, the opportunity to enact such a view may be slim, given the priorities of MAST-U regarding the

use of a novel divertor [114]. The diagnostic "space" on the vessel is highly sought after (a compounded issue on spherical tokamaks compared to conventional), and fast-ion studies may be de-emphasised compared to divertor studies.

7.3 Analysis of sawtooth crashes during MAST operation

The work in Chapter 4 was performed on archived MAST data from 2013. Specifically, the behaviour of fast ions in response to sawtooth crashes were examined. TRANSP/NUBEAM simulations were performed by Cecconello *et. al.* [88] to examine the effect of sawteeth on neutron camera data. TRANSP/NUBEAM was run with 3 different models simulating the effect of sawteeth on the fast ions:

- Kadomtsev, a model featuring complete reconnection of the core field after the sawtooth (ie. $q \geq 1$ everywhere after crash), where regions inside and outside of the $q = 1$ surface that had the same helical flux were mixed
- Ergodic Kadomtsev, also features complete reconnection but core region is mixed without regards to helical flux matching
- Porcelli, featuring incomplete reconnection (q can be < 1 after crash).

It was determined in [88] that neutron camera data were insufficient to distinguish the differences in behaviour of the 3 models. This work was therefore extended in this chapter to utilise FIDA data. After averaging the data over multiple sawteeth and multiple repeat shots, experimental and synthetic FIDA spectra were found to have poor absolute agreement. With archived data it is difficult to know the specifics of the diagnostic setup, so identifying calibration errors is mostly intractable. To resolve this, shots containing sawteeth from MAST-U campaigns should be analysed in the same way as was performed in this chapter to generate a comparison. Until MSE constrained EFIT++ runs are available, this is not advised as

the accuracy of the sawtooth models is dependent on accurate q data, specifically the location of the $q = 1$ surface and the value to which q_{\min} drops before the crash. Magnetics only EFIT runs are much less likely to reflect the reality of the equilibrium in core regions.

Examining the relative change of the experimental and synthetic FIDA spectra, thereby removing sources of absolute error, produced much more consistent matching. The produced synthetic spectra matched experiment over most of the plasma radius. Mismatching of the data near the inversion radius was observed. While it is possible that limitations of the TRANSP/NUBEAM (in reference to how finite Larmor radius effects are dealt with) could contribute, a more in-depth analysis of this in particular would be necessary to narrow down the cause, perhaps using the ASCOT full orbit code. Again, performing this type of analysis on MAST-U data would be advantageous in order to rule out calibration issues.

The main conclusion is that, whether looking at absolute or relative values, the separation of the different synthetic spectra in the observable spectral region is small compared to the experimental error. It is therefore not possible, in these shots, to distinguish the Kadomtsev, Ergodic Kadomtsev and Porcelli models using FIDA data. On MAST-U, with the installation of the SSNPA, it would be a worthy avenue of investigation to repeat the analysis for sawtoothing shots, and integrate both synthetic FIDA spectra *and* synthetic SSNPA data. If the sawteeth are preferentially affecting particles in hitherto unobserved regions of phase-space, it may be possible to distinguish the effects using the SSNPA.

Examination of the passive FIDA spectra suggests that the sawtooth crashes expel passing fast ions from the core towards the edge, where they are then depleted due to charge-exchange losses (as suggested by the timescale of the loss). This result seems to reflect results found on DIII-D [94]. The obvious evolution of this analysis would be to integrate SSNPA data to examine the unobserved regions of phase-space, the fast ion loss detector to confirm the idea of charge-exchange versus direct losses, and to utilise the capabilities of more recent versions of FIDASIM to

generate synthetic passive spectra, which can then be compared to experiment before and after a sawtooth crash.

7.4 Analysis of beam fractions during first experimental campaign on MAST-U

In Chapter 5 the beam fraction ratios in the south beam were investigated using the unmasked FIDA spectra. Analysis early in the campaign had revealed incongruities between the stated and derived beam fractions of the SS beam. However, when the mask was put in place it was not possible to examine the beam fractions with FIDA data, and it was hoped that as beam commissioning advanced the issues regarding beam fractions would be resolved. After the conclusion of the campaign, the CXRS team then obtained evidence of the previously-assumed beam fractions being incorrect for the last (good) shot of the campaign. With this prompting, the analysis of the early shot was revisited and the CXRS team's result for the final shot was reproduced instead utilising the FIDA spectra, as (serendipitously) the FIDA mask had been removed for that shot. Taken all-together, this provided compelling evidence that the beam fractions had been assumed incorrectly, which significantly affects TRANSP/NUBEAM simulations (with knock-on effects on FIDASIM). Investigation of the SS NBI system by the NBI team revealed that the beam ion source had been configured incorrectly, and so was producing a significantly lower full energy fraction (and subsequently higher half and third fractions) than had been originally derived.

While "reasonable" beam fraction values could be derived, the mis-configuration means that it is unlikely to be possible to retroactively generate exact fraction information. This means that all TRANSP/NUBEAM and FIDASIM runs generated previously for the first campaign have a known but unquantifiable source of error (and indeed, all runs that had been performed before the mis-configuration was identified had to be re-run for the results presented in this thesis). Based on the

arguments presented at the end of the chapter, it was thought that the estimated current fractions (44:43:13) were reasonable for use in analysis, but were still subject to uncertainty. However, the mis-configuration has been remedied since the end of the first experimental campaign, and the SS beam ion source is now fully in a supercusp configuration. This should give current fractions closer to 80:15:5 for the second campaign, greatly increasing the proportion of full energy deuterons injected. This should improve the FIDA signal, as due to the beam peaks we are restricted to only examining the full energy component. Additionally, the calculations of the precise values of the fractions is likely to be much better, so TRANSP/NUBEAM (and consequently FIDASIM) runs will better reflect reality. A higher proportion of full-energy particles being injected may also lead to a greater drive of fast-particle driven instabilities, so stronger mode activity may be observed.

7.5 Analysis of MAST-U data

In Chapter 6 the analysis of data taken during the first experimental campaign on MAST-U was described. Shot #45026 featured a prominent fishbone instability that correlated with a sharp drop in plasma rotation. Similar, but less prominent examples were found in shots #45031 and #45163. An initial hypothesis based on a slightly modified version of a model presented in [104] was used to investigate the action of the fishbone in expelling fast ions from the core towards the plasma edge than that likely to have been produced in the experiment a return bulk current, generating in turn a $\mathbf{j} \times \mathbf{B}$ force opposing the direction of plasma rotation. Utilising the basic model it was concluded that the magnitude of fast ion current would have to be significantly higher than that likely to have been produced in the experiment in order to generate sufficient return bulk current that the opposing force could slow the plasma to the degree observed. An alternative hypothesis, that the fishbone removed the fast ions responsible for momenta transfer to the bulk, was proposed

but could not be tested with FIDA data due to spectral region corresponding to the fast ions of relevant energies being obscured due to the presence of the beam peaks. Further analysis would also be required to eliminate the possibility of a MHD torque associated with the fishbone. In future campaigns, the SSNPA should aid greatly in the diagnosis of fishbones, owing to its phase-space coverage focusing on the trapped regions. While the "slow" fast ions, those below the critical energy that are mostly responsible for the rotation drive, are not observable by either FIDA or the SSNPA (which can observe minimum energies ≈ 25 keV) for typical MAST-U plasmas, using both diagnostics in tandem should help the diagnosis of potential cases of the bulk return current generating plasma slowdown. A potential vertical FIDA view could help probe these lower energies, but would be dependent on the specific line of sight configuration chosen.

It was investigated whether an appropriate anomalous diffusion profile could be generated by utilising the signal from the Mirnov coils. This analysis was complicated by problems regarding the absolute calibration of the fission chamber, but an attempt was made by scaling the derived anomalous diffusion time trace by different scaling factors, then comparing the resultant synthetic neutron trace to the experiment. By matching the drop in the neutron rate to the major fishbone and the TAE portion simultaneously, a reasonable prediction for the anomalous diffusion can be found. However, the resultant synthetic FIDA spectra do not conform well to the experimental observations except in a single channel at $R = 1.293$ m. The radius, energy and pitch-invariant anomalous fast ion diffusivity is a simple model, and the matching may improve by varying the diffusivity in these dimensions as well. Constraints on the variation should be used, informed by knowledge of the physics, for example, the "kick" model proposed by M Podesta *et al.* [115]. Observations of experimental and synthetic mismatch towards the lower end of the observable FIDA range point towards a undetermined source of error, with problematic background subtraction being a potential cause. Further investigation, including revisiting the active-passive channel pairings, is advised.

Shot #45091 featured a prominent LLM following a short quiescent period of the plasma. The LLM correlated with a period of toroidal rotation slowdown, before the onset of a sawtooth crash and a locked mode. The affect of the locked mode on the FIDA radiance was examined. It was found that the evolution of the locked mode (as measured by a saddle coil) could be roughly correlated to drops in the neutron rate and FIDA emission, with the caveat that the locked mode could break the axisymmetry of the plasma, and so degrade the FIDA background subtraction.

An analysis of the effect of the LLM on the fast ion population was performed utilising FIDASIM synthetic spectra. It was argued that a short duration of quiescent plasma was sufficient to allow the high-energy part of the synthetic fast ion distribution to be modelled based on classical slowing down, to provide a comparison point to which time-slices throughout the LLM could be compared. Statistically significant drops are observed in core-localised FIDA channels during the earlier parts of the LLM, which reduce as the LLM continues. Before the sawtooth a short rise in the toroidal rotation is observed, which corresponds to an increase in the experimental FIDA emission. The increase in rotation may imply a lessening of the effect of the LLM on the rotation, which may also give an increase in the FIDA emission due to reduced redistribution of the fast ions.

7.6 Future of the FIDA diagnostic

While the analysis of data from the first experimental campaign on MAST-U was inhibited by a variety of problems as described above and throughout the thesis, the future use of the FIDA diagnostic looks bright. Primarily, the fixing of the beam misconfiguration means that shots from the second campaign and beyond will have a much higher full energy fast ion population, hopefully giving better signal in the observable FIDA region. Additionally, the fidelity of TRANSP/NUBEAM and FIDASIM simulations should be better with more accurate beam knowledge. MSE-constrained EFIT++ should be available for the second campaign, and should be

made available retroactively for the first campaign, allowing for more confidence in the synthetic results and deeper analysis using, for example, the TRANSP sawtooth models. The scope of fast-ion analysis will be greatly improved in the second campaign due to having access to the fully commissioned SSNPA and fast ion loss detector. The SSNPA in particular is highly complementary to FIDA, and utilising data from both diagnostics will allow for analysis over a wide swath of the fast-ion pitch range. Combining FIDA and SSNPA data with neutron camera measurements should extend further the coverage of the fast-ion phase space.

There are still unresolved problems however. The consistent mismatching between FIDA and FIDASIM in edge channels and in specific areas of the wavelength space suggest that a review of the active-passive channel pairings would be prudent. The masking issues should be investigated further, as a well placed mask should nearly entirely eliminate the emission from the full energy beam peak. While not directly affecting the FIDA diagnostic, the continuing issues regarding the fission chamber absolute calibration make analysis using anomalous diffusion very difficult, which greatly restricts the level of analysis that can be performed with FIDASIM, as anomalous diffusion is the primary way to deal with effects of instabilities on the fast ions within a simulation.

Once these issues are fully resolved, the integration of FIDA data with the other available fast ion diagnostics will provide MAST-U with a powerful, holistic fast ion diagnostic suite. With the improvements to the forward modelling described, much deeper analysis into instability/fast ion interaction will be possible.

Bibliography

- [1] EIA, “International Energy Outlook 2017 Overview”, *U.S. Energy Information Administration*, vol. IEO2017, no. 2017, p. 143, 2017, ISSN: 0163660X. DOI: [www.eia.gov/forecasts/ieo/pdf/0484\(2016\).pdf](http://www.eia.gov/forecasts/ieo/pdf/0484(2016).pdf). [Online]. Available: [https://www.eia.gov/outlooks/ieo/pdf/0484\(2017\).pdf](https://www.eia.gov/outlooks/ieo/pdf/0484(2017).pdf).
- [2] UN DESA, “World Population Prospects The 2017 Revision Key Findings and Advance Tables”, *World Population Prospects The 2017*, pp. 1–46, 2017, ISSN: 10187081. DOI: 10.1017/CB09781107415324.004. [Online]. Available: https://esa.un.org/unpd/wpp/Publications/Files/WPP2017_KeyFindings.pdf.
- [3] W. Gulden, S. Ciattaglia, V. Massaut *et al.*, “Main safety issues at the transition from ITER to fusion power plants”, *Nuclear Fusion*, vol. 47, no. 9, pp. 1391–1398, 2007, ISSN: 00295515. DOI: 10.1088/0029-5515/47/9/C01.
- [4] S. Atzeni and J. Meyer-ter-Vehn, *Physics of Inertial Confinement Fusion and High Energy Density in Matter*. Oxford: Clarendon Press, 2003.
- [5] J. Wesson, *Tokamaks*, 4th. Oxford: Oxford University Press, 2011.
- [6] W. M. Stacey, *Fusion Plasma Physics*, 2nd. Verlag: Wiley-VCH, 2012, ISBN: 9783527407422.

- [7] B. Brunelli and H. Knoepfel, *Safety, environmental impact, and economic prospects of nuclear fusion*, B. Brunelli and H. Knoepfel, Eds. New York and London: Plenum Press, 1990, ISBN: 0306435241.
- [8] M. Rubel, “Fusion Neutrons: Tritium Breeding and Impact on Wall Materials and Components of Diagnostic Systems”, *Journal of Fusion Energy*, vol. 38, no. 3-4, pp. 315–329, 2019, ISSN: 01640313. DOI: 10.1007/s10894-018-0182-1. [Online]. Available: <https://doi.org/10.1007/s10894-018-0182-1>.
- [9] M. Keilhacker, A. Gibson, C. Gormezano *et al.*, “High fusion performance from deuterium-tritium plasmas in JET”, *Nuclear Fusion*, vol. 39, no. 2, pp. 209–234, 1999, ISSN: 00295515. DOI: 10.1088/0029-5515/39/2/306.
- [10] A. A. Harms, D. R. Kingdon, K. F. Schoepf *et al.*, *Principles of Fusion Energy: An Introduction to Fusion Energy for Students of Science and Engineering*. Singapore: World Scientific, 2000.
- [11] ITER Physics Basis Editors, ITER Physics Expert Group Chairs and Co-Chairs, ITER Joint Central Team and Physics Integration Unit *et al.*, “ITER Physics Basis - Chapter 1: Overview and summary”, *Nuclear Fusion*, vol. 39, no. 12, pp. 2137–2174, 1999, ISSN: 0029-5515. DOI: 10.1088/0029-5515/39/12/301.
- [12] F. F. Chen, *Introduction to Plasma Physics and Controlled Fusion*, 3rd. Cham: Springer, 2016, ISBN: 978-3-319-22308-7. DOI: 10.1088/1751-8113/44/8/085201.
- [13] A. Sykes, “Physics of spherical tokamaks”, *Technical Physics*, vol. 44, no. 9, pp. 1047–1053, 1999, ISSN: 1063-7842. DOI: 10.1134/1.1259468.
- [14] CCFE, *Spherical Tokamaks*. [Online]. Available: <http://www.ccfе.ac.uk/ST.aspx>.

- [15] A. Sykes, A. Costley, C. Windsor *et al.*, “Compact fusion energy based on the spherical tokamak”, *Nuclear Fusion*, vol. 58, no. 1, p. 016 039, 2018, ISSN: 0029-5515. DOI: 10.1088/1741-4326/aa8c8d. [Online]. Available: <http://stacks.iop.org/0029-5515/58/i=1/a=016039?key=crossref.c82a427510f4e4a225bad04c82a0ab6b>.
- [16] L. Spitzer and R. Härm, “Transport phenomena in a completely ionized gas”, *Physical Review*, vol. 89, no. 5, pp. 977–981, 1953, ISSN: 0031899X. DOI: 10.1103/PhysRev.89.977.
- [17] T. Ohkawa, “New methods of driving plasma current in fusion devices”, *Nuclear Fusion*, vol. 10, no. 2, pp. 185–188, 1970, ISSN: 17414326. DOI: 10.1088/0029-5515/10/2/012.
- [18] D. Stork, “Neutral beam heating and current drive systems”, *Fusion Engineering and Design*, vol. 14, no. 1-2, pp. 111–133, 1991, ISSN: 09203796. DOI: 10.1016/0920-3796(91)90238-L.
- [19] T. H. Stix, “Heating of toroidal plasmas by neutral injection”, *Plasma Physics*, vol. 14, no. 4, pp. 367–384, 1972, ISSN: 00321028. DOI: 10.1088/0032-1028/14/4/002.
- [20] A. S. Richardson, “NRL Plasma Formulary”, Naval Research Laboratory, Washington, DC, Tech. Rep., 2019.
- [21] L. Landau, “On the vibrations of the electronic plasma”, *Zh. Eksp. Teor. Fiz.*, vol. 10, no. 1, 1946.
- [22] H. Alfvén, “Existence of Electromagnetic-Hydrodynamic Waves”, *Nature*, vol. 150, pp. 405–406, 1942, ISSN: 00280836.
- [23] W. W. Heidbrink, “Basic physics of Alfvén instabilities driven by energetic particles in toroidally confined plasmas”, *Physics of Plasmas*, vol. 15, no. 5, 2008, ISSN: 1070664X. DOI: 10.1063/1.2838239.

- [24] S. E. Sharapov, “Toroidal alfvén eigenmodes and fast particles in tokamaks”, *Fusion Science and Technology*, vol. 61, no. 2 T, pp. 104–112, 2012, ISSN: 15361055. DOI: 10.13182/FST12-A13497.
- [25] D. C. Pace, R. K. Fisher, M. García-Müoz *et al.*, “Transport of energetic ions due to sawteeth, Alfvén eigenmodes and microturbulence”, *Nuclear Fusion*, vol. 51, no. 4, p. 043 012, 2011, ISSN: 00295515. DOI: 10.1088/0029-5515/51/4/043012.
- [26] J. A. Wesson, “Hydromagnetic stability of tokamaks”, *Nuclear Fusion*, vol. 18, no. 1, pp. 87–132, 1978, ISSN: 17414326. DOI: 10.1088/0029-5515/18/1/010.
- [27] M. Kruskal and J. L. Tuck, “The Instability of a Pinched Fluid with a Longitudinal Magnetic Field”, in *Proceedings of the Royal Society of London. Series A, Mathematical and Physical*, vol. 245, 1958, pp. 222–237. [Online]. Available: <http://www.jstor.org/stable/2414705>.
- [28] F. Porcelli, “Fast particle stabilisation”, *Plasma Physics and Controlled Fusion*, vol. 33, no. 13, pp. 1601–1620, 1991, ISSN: 07413335. DOI: 10.1088/0741-3335/33/13/009.
- [29] R. Betti and P. Freidberg, “Destabilization of the Internal Kink by Energetic Circulating Ions”, *Physical Review Letters*, vol. 70, no. 22, pp. 3428–3430, 1993.
- [30] O. M. Jones, C. A. Michael, K. G. McClements *et al.*, “Fast-ion deuterium alpha spectroscopic observations of the effects of fishbones in the Mega-Ampere Spherical Tokamak”, *Plasma Physics and Controlled Fusion*, vol. 55, no. 8, 2013, ISSN: 07413335. DOI: 10.1088/0741-3335/55/8/085009.
- [31] I. T. Chapman, M. D. Hua, S. D. Pinches *et al.*, “Saturated ideal modes in advanced tokamak regimes in MAST”, *Nuclear Fusion*, vol. 50, no. 4, 2010, ISSN: 00295515. DOI: 10.1088/0029-5515/50/4/045007.

- [32] K. C. Shaing, “Magnetohydrodynamic-activity-induced toroidal momentum dissipation in collisionless regimes in tokamaks”, *Physics of Plasmas*, vol. 10, no. 5 I, pp. 1443–1448, 2003, ISSN: 1070664X. DOI: 10.1063/1.1567285.
- [33] D. F. Howell, T. C. Hender and G. Cunningham, “Locked mode thresholds on the MAST spherical tokamak”, *Nuclear Fusion*, vol. 47, no. 9, pp. 1336–1340, 2007, ISSN: 00295515. DOI: 10.1088/0029-5515/47/9/034.
- [34] I. H. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. Cambridge: Cambridge University Press, 2002, ISBN: 9780521803892. DOI: 10.1017/cbo9780521803892.
- [35] K. J. Gibson, N. Barratt, I. Chapman *et al.*, “New physics capabilities from the upgraded Thomson scattering diagnostic on MAST”, *Plasma Physics and Controlled Fusion*, vol. 52, no. 12, 2010, ISSN: 07413335. DOI: 10.1088/0741-3335/52/12/124041.
- [36] K. Stammers and M. J. Loughlin, “The calibration of the MAST neutron yield monitors”, *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 562, no. 1, pp. 521–530, 2006, ISSN: 01689002. DOI: 10.1016/j.nima.2006.03.012.
- [37] M. Kropík, “Neutron flux measurement utilizing the Campbell technique”, in *Proceedings of the 2000 International Conference on Nuclear Energy in Central Europe*, Slovenia, 2000.
- [38] M. J. Hole, L. C. Appel and R. Martin, “A high resolution Mirnov array for the mega ampere spherical tokamak”, *Review of Scientific Instruments*, vol. 80, no. 12, 2009, ISSN: 00346748. DOI: 10.1063/1.3272713.
- [39] L. C. Appel, M. K. Bevir and M. J. Walsh, “Equilibrium reconstruction in the START tokamak”, *Nuclear Fusion*, vol. 41, no. 2, pp. 169–180, 2001, ISSN: 00295515. DOI: 10.1088/0029-5515/41/2/303.

- [40] R. C. Isler, “An overview of charge-exchange spectroscopy as a plasma diagnostic”, *Plasma Physics and Controlled Fusion*, vol. 36, no. 2, pp. 171–208, 1994, ISSN: 13616587. DOI: 10.1088/0741-3335/36/2/001.
- [41] N. J. Conway, P. G. Carolan, J. McCone *et al.*, “High-throughput charge exchange recombination spectroscopy system on MAST”, *Review of Scientific Instruments*, vol. 77, no. 10, pp. 1–4, 2006, ISSN: 00346748. DOI: 10.1063/1.2354309.
- [42] A. Kramida, Y. Ralchenko, J. Reader *et al.*, *NIST Atomic Spectra Database (version 5.9)*, 2021. DOI: <https://dx.doi.org/10.18434/T4W30F>. [Online]. Available: <https://physics.nist.gov/asd>.
- [43] Y. Luo, W. W. Heidbrink, K. H. Burrell *et al.*, “Measurement of the $D\alpha$ spectrum produced by fast ions in DIII-D”, *Review of Scientific Instruments*, vol. 78, no. 3, p. 033 505, 2007, ISSN: 0034-6748. DOI: 10.1063/1.2712806. [Online]. Available: <http://aip.scitation.org/doi/10.1063/1.2712806>.
- [44] W. W. Heidbrink, “Fast-ion $D\alpha$ measurements of the fast-ion distribution”, *Review of Scientific Instruments*, vol. 81, no. 10, p. 10D727, 2010, ISSN: 00346748. DOI: 10.1063/1.3478739.
- [45] C. A. Michael, N. Conway, B. Crowley *et al.*, “Dual view FIDA measurements on MAST”, *Plasma Physics and Controlled Fusion*, vol. 55, no. 9, p. 095 007, 2013, ISSN: 07413335. DOI: 10.1088/0741-3335/55/9/095007.
- [46] A. S. Jacobsen, “Methods to determine fast-ion distribution functions from multi-diagnostic measurements”, Ph.D. dissertation, 2015.
- [47] M. Salewski, B. Geiger, D. Moseev *et al.*, “On velocity-space sensitivity of fast-ion D-alpha spectroscopy”, *Plasma Physics and Controlled Fusion*, vol. 56, no. 10, p. 105 005, 2014, ISSN: 13616587. DOI: 10.1088/0741-3335/56/10/105005.

- [48] A. S. Jacobsen, “Methods to determine fast-ion distribution functions from multi-diagnostic measurements”, Ph.D. dissertation, Technical University of Denmark, 2015.
- [49] W. W. Heidbrink, Y. Luo, C. M. Muscatello *et al.*, “A new fast-ion $D\alpha$ diagnostic for DIII-D”, *Review of Scientific Instruments*, vol. 79, no. 10, 2008, ISSN: 00346748. DOI: 10.1063/1.2956828.
- [50] D. G. Swanson, “Radio frequency heating in the ion-cyclotron range of frequencies”, *Physics of Fluids*, vol. 28, no. 9, pp. 2645–2677, 1985, ISSN: 10706631. DOI: 10.1063/1.865224.
- [51] M. Salewski, B. Geiger, S. K. Nielsen *et al.*, “Tomography of fast-ion velocity-space distributions from synthetic CTS and FIDA measurements”, *Nuclear Fusion*, vol. 52, no. 10, 2012, ISSN: 00295515. DOI: 10.1088/0029-5515/52/10/103008.
- [52] M. Salewski, B. Geiger, S. K. Nielsen *et al.*, “Combination of fast-ion diagnostics in velocity-space tomographies”, *Nuclear Fusion*, vol. 53, no. 6, 2013, ISSN: 00295515. DOI: 10.1088/0029-5515/53/6/063019.
- [53] M. Weiland, B. Geiger, A. S. Jacobsen *et al.*, “Enhancement of the FIDA diagnostic at ASDEX Upgrade for velocity space tomography”, *Plasma Physics and Controlled Fusion*, vol. 58, no. 2, p. 025 012, 2016, ISSN: 13616587. DOI: 10.1088/0741-3335/58/2/025012.
- [54] B. Madsen, M. Salewski, J. Huang *et al.*, “Velocity-space tomography using prior information at MAST”, *Review of Scientific Instruments*, vol. 89, no. 10, pp. 1–6, 2018, ISSN: 10897623. DOI: 10.1063/1.5035498. [Online]. Available: <http://dx.doi.org/10.1063/1.5035498>.
- [55] W. W. Heidbrink, D. Liu, Y. Luo *et al.*, “A code that simulates fast-ion $D\alpha$ and neutral particle measurements”, *Communications in Computational Physics*, vol. 10, no. 3, pp. 716–741, 2011, ISSN: 18152406. DOI: 10.4208/cicp.190810.080211a.

- [56] B. Geiger, L. Stagner, W. W. Heidbrink *et al.*, “Progress in modelling fast-ion D-alpha spectra and neutral particle analyzer fluxes using FIDASIM”, *Plasma Physics and Controlled Fusion*, vol. 62, no. 10, 2020, ISSN: 13616587. DOI: 10.1088/1361-6587/aba8d7.
- [57] L. E. Stagner, B. Geiger and W. W. Heidbrink, *FIDASIM: A Neutral Beam and Fast-ion Diagnostic Modeling Suite*. DOI: 10.5281/zenodo.1341369. [Online]. Available: <https://doi.org/10.5281/zenodo.1341369>.
- [58] R. J. Hawryluk, “An empirical approach to tokamak transport”, in *Physics of Plasmas Close to Thermonuclear Conditions*, 1981, pp. 19–46.
- [59] A. Pankin, D. McCune, R. Andre *et al.*, “The tokamak Monte Carlo fast ion module NUBEAM in the national transport code collaboration library”, *Computer Physics Communications*, vol. 159, no. 3, pp. 157–184, 2004, ISSN: 00104655. DOI: 10.1016/j.cpc.2003.11.002.
- [60] E. Hirvijoki, O. Asunta, T. Koskela *et al.*, “ASCOT: Solving the kinetic equation of minority particle species in tokamak plasmas”, *Computer Physics Communications*, vol. 185, no. 4, pp. 1310–1321, 2014, ISSN: 00104655. DOI: 10.1016/j.cpc.2014.01.014. [Online]. Available: <http://dx.doi.org/10.1016/j.cpc.2014.01.014>.
- [61] J. P. H. E. Ongena, I. Voitsekhovitch, M. Evrard *et al.*, “Numerical Transport Codes”, *Transactions of Fusion Science and Technology*, vol. 57, pp. 381–390, 2010.
- [62] M. Vachharajani and PPPL, *TRANSP Help*. [Online]. Available: <https://w3.pppl.gov/~pshare/help/transp.htm>.
- [63] O. Meneghini, S. P. Smith, L. L. Lao *et al.*, “Integrated modeling applications for tokamak experiments with OMFIT”, *Nuclear Fusion*, vol. 55, no. 8, 2015, ISSN: 17414326. DOI: 10.1088/0029-5515/55/8/083008.

- [64] L. L. Lao, H. S. John, R. D. Stambaugh *et al.*, “Reconstruction of current profile parameters and plasma shapes in tokamaks”, *Nuclear Fusion*, vol. 25, no. 11, pp. 1611–1622, 1985, ISSN: 17414326. DOI: 10.1088/0029-5515/25/11/007.
- [65] A. Patel, P. G. Carolan, N. J. Conway *et al.*, “Versatile multiwavelength imaging diagnostic in the MAST spherical tokamak”, *Review of Scientific Instruments*, vol. 75, no. 10 II, pp. 4145–4148, 2004, ISSN: 00346748. DOI: 10.1063/1.1789608.
- [66] M. Cecconello, W. Boeglin, D. Keeling *et al.*, “Discrepancy between estimated and measured fusion product rates on MAST using TRANSP/NUBEAM”, *Nuclear Fusion*, vol. 59, no. 1, 2019, ISSN: 17414326. DOI: 10.1088/1741-4326/aaea19.
- [67] A. M. d. Veen and M. G. Cox, “Getting started with uncertainty evaluation using the Monte Carlo method in R”, *Accreditation and Quality Assurance*, vol. 26, no. 3, pp. 129–141, 2021, ISSN: 14320517. DOI: 10.1007/s00769-021-01469-5. [Online]. Available: <https://doi.org/10.1007/s00769-021-01469-5>.
- [68] A. Sahlberg, C. Hellesen, J. Eriksson *et al.*, “Propagating transport-code input parameter uncertainties with deterministic sampling”, *Plasma Physics and Controlled Fusion*, vol. 60, no. 12, p. 125 010, Dec. 2018, ISSN: 0741-3335. DOI: 10.1088/1361-6587/aae80b. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6587/aae80b>.
- [69] A. Sahlberg, “Ensemble for deterministic sampling with positive weights”, Swedish Radiation Safety Authority, Tech. Rep., 2016.
- [70] M. Cecconello, W. Boeglin, D. Keeling *et al.*, “Discrepancy between estimated and measured fusion product rates on MAST using TRANSP/NUBEAM”, *Nuclear Fusion*, vol. 59, no. 1, p. 016 006, 2018, ISSN: 0029-5515. DOI: 10.1088/1741-4326/aaea19.

- [71] S. S. Medley and A. L. Roquemore, “Construction and operation of parallel electric and magnetic field spectrometers for mass/energy resolved multi-ion charge exchange diagnostics on the Tokamak Fusion Test Reactor”, *Review of Scientific Instruments*, vol. 69, no. 7, pp. 2651–2661, 1998, ISSN: 00346748. DOI: 10.1063/1.1148994.
- [72] D. Liu, W. W. Heidbrink, K. Tritz *et al.*, “Design of solid state neutral particle analyzer array for National Spherical Torus Experiment-Upgrade”, *Review of Scientific Instruments*, vol. 85, no. 11, 2014, ISSN: 10897623. DOI: 10.1063/1.4889913. [Online]. Available: <http://dx.doi.org/10.1063/1.4889913>.
- [73] C. Wade, “CXRS Fibre Bundle Assembly Report”, UKAEA, Tech. Rep. August, 2020.
- [74] E. Braeken, “Assembly and calibration of FIDA diagnostic on the MAST tokamak”, Eindhoven University of Technology, Eindhoven, Tech. Rep., 2011.
- [75] R. E. Bell, “Exploiting a transmission grating spectrometer”, *Review of Scientific Instruments*, vol. 75, no. 10 II, pp. 4158–4161, 2004, ISSN: 00346748. DOI: 10.1063/1.1787601.
- [76] Photometrics, *Photometrics Cascade 128+, PM-DS-002 Rev C1*, 2010. [Online]. Available: <https://www.photometrics.com/wp-content/uploads/2019/10/Cascade-128-Datasheet.pdf>.
- [77] S. Silburn, J. Harrison, T. Farley *et al.*, *Calcam*, 2020. DOI: 10.5281/zenodo.4698377. [Online]. Available: <https://github.com/euratom-software/calcam>.
- [78] A. R. Jackson, A. S. Jacobsen, K. G. McClements *et al.*, “Diagnosing fast ion redistribution due to sawtooth instabilities using fast ion deuterium- α spectroscopy in the Mega Amp Spherical Tokamak”, *Nuclear Fusion*, vol. 60, p. 126 035, 2020, ISSN: 0029-5515. DOI: 10.1088/1741-4326/abb619.

- [79] A. R. Jackson, A. S. Jacobsen, K. G. McClements *et al.*, “Diagnosing fast ion redistribution due to sawtooth instabilities using fast ion deuterium- α spectroscopy in the Mega Amp Spherical Tokamak”, in *47th EPS Conference on Plasma Physics, EPS 2021*, vol. 45A, Sitges, Spain, 2021, pp. 689–692, ISBN: 9781713837046.
- [80] S. von Goeler, W. Stodiek and N. Sauthoff, “Studies of Internal Disruptions and $m=1$ Oscillations in Tokamak Discharges with Soft—X-Ray Techniques”, *Physical Review Letters*, vol. 33, no. 20, pp. 1201–1203, 1974, ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.33.1201. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.33.1201>.
- [81] F. Wising, D. Anderson and M. Lisak, “Modelling of fast ion redistribution due to sawteeth in neutral beam heated plasmas”, *Plasma Physics and Controlled Fusion*, vol. 34, no. 5, pp. 853–862, 1992, ISSN: 13616587. DOI: 10.1088/0741-3335/34/5/013.
- [82] I. T. Chapman, S. D. Pinches, J. P. Graves *et al.*, “The physics of sawtooth stabilization”, *Plasma Physics and Controlled Fusion*, vol. 49, no. 12 B, B385–B394, 2007, ISSN: 07413335. DOI: 10.1088/0741-3335/49/12B/S35.
- [83] I. T. Chapman, R. J. Buttery, S. Coda *et al.*, “Empirical scaling of sawtooth period for onset of neoclassical tearing modes”, *Nuclear Fusion*, vol. 50, no. 10, p. 102 001, 2010, ISSN: 00295515. DOI: 10.1088/0029-5515/50/10/102001.
- [84] C. M. Muscatello, W. W. Heidbrink, Y. I. Kolesnichenko *et al.*, “Velocity-space studies of fast-ion transport at a sawtooth crash in neutral-beam heated plasmas”, *Plasma Physics and Controlled Fusion*, vol. 54, no. 2, p. 025 006, 2012, ISSN: 07413335. DOI: 10.1088/0741-3335/54/2/025006.
- [85] M. Salewski, B. Geiger, A. S. Jacobsen *et al.*, “High-definition velocity-space tomography of fast-ion dynamics”, *Nuclear Fusion*, vol. 56, no. 10, p. 106 024, 2016, ISSN: 17414326. DOI: 10.1088/0029-5515/56/10/106024.

- [86] D. Liu, W. W. Heidbrink, M. Podestà *et al.*, “Effect of sawtooth crashes on fast ion distribution in NSTX-U”, *Nuclear Fusion*, vol. 58, no. 8, p. 082 028, 2018, ISSN: 17414326. DOI: 10.1088/1741-4326/aac64f.
- [87] M. Cecconello, O. M. Jones, W. U. Boeglin *et al.*, “Energetic ion behaviour in MAST”, *Plasma Physics and Controlled Fusion*, vol. 57, no. 1, p. 014 006, 2015, ISSN: 13616587. DOI: 10.1088/0741-3335/57/1/014006.
- [88] M. Cecconello, A. Sperduti and the MAST team, “Study of the effect of sawteeth on fast ions and neutron emission in MAST using a neutron camera”, *Plasma Physics and Controlled Fusion*, vol. 60, no. 5, p. 055 008, 2018, ISSN: 13616587. DOI: 10.1088/1361-6587/aab6cc.
- [89] B. B. Kadomtsev, “On the tearing instability in tokamaks”, *Fizika Plazmy*, vol. 1, no. 5, pp. 710–714, 1975.
- [90] F. Porcelli, D. Boucher and M. Rosenbluth, “Model for the sawtooth period and amplitude”, *Plasma Physics and Controlled Fusion*, vol. 38, pp. 2163–2186, 1996.
- [91] M. Podesta, *Private communication*, 2019.
- [92] D. McCune, *NTCC module – kdsaw – Kadomtsev Sawtooth Model with Porcelli Upgrade*. 2005. [Online]. Available: <https://w3.pppl.gov/ntcc/KDSAW/DOC>.
- [93] W. W. Heidbrink, Y. Luo, K. H. Burrell *et al.*, “Measurements of fast-ion acceleration at cyclotron harmonics using Balmer-alpha spectroscopy”, *Plasma Physics and Controlled Fusion*, vol. 49, no. 9, pp. 1457–1475, 2007, ISSN: 0741-3335. DOI: 10.1088/0741-3335/49/9/008.
- [94] N. G. Bolte, W. W. Heidbrink, D. Pace *et al.*, “Measurement and simulation of passive fast-ion D-alpha emission from the DIII-D tokamak”, *Nuclear Fusion*, vol. 56, no. 11, p. 112 023, 2016, ISSN: 17414326. DOI: 10.1088/0029-5515/56/11/112023.

- [95] K. G. McClements, K. Tani, R. J. Akers *et al.*, “The effects of resonant magnetic perturbations and charge-exchange reactions on fast ion confinement and neutron emission in the Mega Amp Spherical Tokamak”, *Plasma Physics and Controlled Fusion*, vol. 60, no. 9, p. 095005, 2018, ISSN: 13616587. DOI: 10.1088/1361-6587/aad252.
- [96] J. F. Bonnal, G. Bracco, C. Breton *et al.*, “Doppler spectroscopy on high power neutral beams”, *Journal of Physics D: Applied Physics*, vol. 15, no. 5, pp. 805–822, 1982, ISSN: 00223727. DOI: 10.1088/0022-3727/15/5/011.
- [97] P. Summers H, *The ADAS User Manual, version 2.6*, 2004. [Online]. Available: <http://www.adas.ac.uk>.
- [98] S. Henderson, “Impurity Transport Studies on MAST”, Ph.D. dissertation, 2014.
- [99] I. D. Williams, J. Geddes and H. B. Gilbody, “Balmer α emission in collisions of H, H⁺, H²⁺ and H³⁺ with H₂”, *Journal of Physics B: Atomic and Molecular Physics*, vol. 15, no. 9, pp. 1377–1389, 1982, ISSN: 00223700. DOI: 10.1088/0022-3700/15/9/014.
- [100] I. D. Williams, J. Geddes and H. B. Gilbody, “Balmer α emission in H₂-H₂ collisions”, *Journal of Physics B: Atomic and Molecular Physics*, vol. 16, p. L765, 1983.
- [101] R. McAdams, A. Sparkes, A. Ash *et al.*, “Preparation for the next JET tritium campaign: Performance of the EP2 PINIs with grid gas delivery”, *Fusion Engineering and Design*, vol. 96-97, pp. 527–531, 2015, ISSN: 09203796. DOI: 10.1016/j.fusengdes.2015.01.043.
- [102] D. Ćirić, A. D. Ash, B. Crowley *et al.*, “Performance of upgraded JET neutral beam injectors”, *Fusion Engineering and Design*, vol. 86, no. 6-8, pp. 509–512, 2011, ISSN: 09203796. DOI: 10.1016/j.fusengdes.2010.11.035.
- [103] D. Ćirić, *Internal presentation: Neutral Beam Heating Overview*, 2014.

- [104] A. Thyagaraja, F. Schwander and K. G. McClements, “Rotation driven by fast ions in tokamaks”, *Physics of Plasmas*, vol. 14, no. 11, pp. 1–16, 2007, ISSN: 1070664X. DOI: 10.1063/1.2801716.
- [105] O. M. Jones, “Experimental fast-ion transport studies on the Mega-Amp Spherical Tokamak”, Ph.D. dissertation, University of Durham, 2015.
- [106] H. H. Duong, W. W. Heidbrink, E. J. Strait *et al.*, “Loss of energetic beam ions during TAE instabilities”, *Nuclear Fusion*, vol. 33, no. 5, pp. 749–765, 1993, ISSN: 00295515. DOI: 10.1088/0029-5515/33/5/I06.
- [107] W. W. Heidbrink, “Beam-driven chirping instability in DIII-D”, *Plasma Physics and Controlled Fusion*, vol. 37, no. 9, pp. 937–949, 1995, ISSN: 07413335. DOI: 10.1088/0741-3335/37/9/002.
- [108] E. D. Fredrickson, R. E. Bell, D. S. Darrow *et al.*, “Collective fast ion instability-induced losses in National Spherical Tokamak Experiment”, *Physics of Plasmas*, vol. 13, no. 5, 2006, ISSN: 1070664X. DOI: 10.1063/1.2178788.
- [109] M. Turnyanskiy, C. D. Challis, R. J. Akers *et al.*, “Measurement and control of the fast ion redistribution on MAST”, *Nuclear Fusion*, vol. 53, no. 5, 2013, ISSN: 00295515. DOI: 10.1088/0029-5515/53/5/053016.
- [110] O. M. Jones, M. Cecconello, K. G. McClements *et al.*, “Measurements and modelling of fast-ion redistribution due to resonant MHD instabilities in MAST”, *Plasma Physics and Controlled Fusion*, vol. 57, no. 12, 2015, ISSN: 13616587. DOI: 10.1088/0741-3335/57/12/125009.
- [111] Y. R. Martin and TCV team, “Synchronization of L-mode to H-mode transitions on the sawtooth cycle in Ohmic TCV plasmas”, *Plasma Physics and Controlled Fusion*, vol. 46, no. 5 SUPPL. A, A77–A85, 2004, ISSN: 07413335. DOI: 10.1088/0741-3335/46/5A/008.

- [112] L. M. Shao, G. S. Xu, N. Yan *et al.*, “L-mode to H-mode transition triggered by sawtooth-induced heat flux in EAST”, *Physics Letters, Section A: General, Atomic and Solid State Physics*, vol. 384, no. 9, 2020, ISSN: 03759601. DOI: 10.1016/j.physleta.2019.126184.
- [113] K. G. McClements, R. O. Dendy and A. Gondhalekar, “Modelling Neutral Particle Analyzer Measurements of Alpha-Particle Distributions in JET and TFTR”, in *Fifth IAEA Technical Committee Meeting on Alpha Particles in Fusion Research*, J. Jacquinet, B. E. Keen and G. J. Sadler, Eds., Abingdon, UK, 1997, P.Tu.02.
- [114] G. Fishpool, J. Canik, G. Cunningham *et al.*, “MAST-upgrade divertor facility and assessing performance of long-legged divertors”, *Journal of Nuclear Materials*, vol. 438, no. SUPPL, S356–S359, 2013, ISSN: 00223115. DOI: 10.1016/j.jnucmat.2013.01.067. [Online]. Available: <http://dx.doi.org/10.1016/j.jnucmat.2013.01.067>.
- [115] M. Podestà, M. Gorelenkova and R. B. White, “A reduced fast ion transport model for the tokamak transport code TRANSP”, *Plasma Physics and Controlled Fusion*, vol. 56, no. 5, 2014, ISSN: 13616587. DOI: 10.1088/0741-3335/56/5/055003.

Colophon

This thesis is based on a template developed by Matthew Townson and Andrew Reeves. It was typeset with L^AT_EX 2_ε. It was created using the *memoir* package, maintained by Lars Madsen, with the *madsen* chapter style. The font used is Latin Modern, derived from fonts designed by Donald E. Kunith.