CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	Inter-ELM evolution of the edge current density profile on the ASDEX upgrade tokamak
Authors	Dunne, Michael G.
Publication date	2013
Original Citation	Dunne, M. G. 2013. Inter-ELM evolution of the edge current density profile on the ASDEX upgrade tokamak. PhD Thesis, University College Cork.
Type of publication	Doctoral thesis
Rights	© 2013, Michael Dunne. - http://creativecommons.org/licenses/by-nc-nd/3.0/
Download date	2026-11-08 11:43:11
Item downloaded from	https://hdl.handle.net/10468/1409



UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Inter-ELM evolution of the edge current density profile on the ASDEX Upgrade tokamak

Michael G. Dunne



Department of Physics,
University College Cork

Supervisor: Dr. Patrick J. McCarthy
Head of department: Prof. John McInerney

Thesis submitted for the degree of Doctor of Philosophy to the
National University of Ireland, Cork
June 2013

Contents

1 Introduction **1**

1.1 The case for nuclear power . . . 2
1.2 Confinement . . . 4
1.3 Description of a typical tokamak device . . . 6
1.3.1 Poloidal cross section . . . 7
1.3.2 Plasma facing materials . . . 8
1.3.3 Plasma heating techniques . . . 10
1.3.4 Plasma position control . . . 10
1.4 Magnetohydrodynamics . . . 11
1.4.1 Plasma equilibrium . . . 11
1.5 Equilibrium reconstruction . . . 13
1.5.1 Predictive equilibrium calculation . . . 13
1.5.2 Function parameterisation . . . 14
1.5.3 Interpretive equilibrium reconstruction . . . 14
1.5.4 Fixed boundary solvers . . . 15
1.5.5 Resonant surfaces . . . 16
1.6 H-mode . . . 16
1.6.1 Plasma coordinate systems . . . 17
1.6.2 Edge Localised Modes - ELMs . . . 18
1.7 Focus of this thesis . . . 19

2 ELM phenomenology **21**

2.1 Pedestal evolution . . . 22
2.1.1 Gradient length evolution . . . 24
2.1.2 Comparison to ion measurements . . . 25
2.1.3 Evolution of pedestal in real space . . . 27
2.2 ELM mitigation techniques . . . 29
2.3 Pedestal relevant instabilities . . . 31
2.3.1 MHD modes . . . 31

2.3.2 Kinetic ballooning modes . . . 34
2.3.3 The ELM cycle . . . 35
2.4 Experimental evidence for MHD ELM description . . . 37

3 Current density analysis **41**
3.1 Diagnostics . . . 44
3.2 Pressure measurements . . . 46
3.2.1 ECE . . . 46
3.2.2 Thomson scattering . . . 47
3.2.3 Lithium beam emission spectroscopy . . . 47
3.2.4 Interferometry . . . 48
3.2.5 Charge-exchange recombination spectroscopy . . . 48
3.2.6 Ion density . . . 49
3.3 Magnetic measurements . . . 49
3.3.1 Current measurements . . . 50
3.4 Current density determination . . . 51
3.5 Example current density profiles . . . 52
3.5.1 Confidence band calculation . . . 54
3.6 Power scan . . . 55
3.7 Recovery from ELM crash . . . 56
3.7.1 Radial and temporal ELM resolved behaviour . . . 60
3.7.2 Normalised pressure gradient and current density evolution 62
3.8 Fuelling study . . . 63
3.9 Conclusions . . . 65

4 Sensitivity and accuracy of current density profiles and separatrix location **67**
4.1 Methodology . . . 67
4.2 Profile shifts . . . 68
4.3 Width scans . . . 70
4.3.1 Width scan with peak gradient location fixed . . . 71
4.3.2 Alternative width scan . . . 72
4.4 Interpretation of effects . . . 75
4.4.1 Comparison of errors . . . 77
4.5 Separatrix location . . . 78
4.5.1 Separatrix location using electron cyclotron forward modelling . . . 82
4.6 Conclusions . . . 84

5 Theoretical edge current density **87**
5.1 Model description . . . 87
5.2 Comparison to CLISTE . . . 91
5.2.1 Current diffusion boundary conditions . . . 93
5.3 Fuelling study . . . 94
5.4 Importance of T_i in current density recovery . . . 98

6 Analysis of Type-I ELMy H-mode database **101**
6.1 Input data . . . 101
6.1.1 Kinetic profiles . . . 102
6.1.2 Magnetic data . . . 102
6.1.3 Fitting . . . 103
6.2 Time resolved equilibrium uncertainty . . . 104
6.3 Output current density . . . 106
6.3.1 Dependence on pressure gradient . . . 106
6.3.2 Theoretical current density . . . 109
6.4 ELM frequency dependence . . . 111
6.5 Further observations of ELM frequency variation . . . 113
6.6 Separatrix movement at ELM crash . . . 116
6.7 Magnetic shaping . . . 117

7 ELM mitigation scenarios **121**
7.1 Type-II ELMs . . . 122
7.2 N_2 seeded ELMs . . . 126
7.3 External perturbations . . . 132
7.3.1 Comparison with other regimes . . . 137
7.4 Conclusions . . . 142

8 Summary and outlook **145**
8.0.1 Future work . . . 147

Bibliography **149**

Appendices **161**

A Formulae for neoclassical current density calculation **161**

Acronyms **163**

Frequently used parameters **165**

List of Figures

1.1 $n\tau_E$ as a function of temperature 4
1.2 Principal tokamak components 7
1.3 Sample AUG magnetic geometry 9
1.4 Comparison of electron temperature profiles before and after L-H transition 17

2.1 Effect of ELM crash on plasma parameters 22
2.2 Pedestal top behaviour at an ELM crash 23
2.3 Pedestal gradient evolution relative to an ELM crash 24
2.4 ELM resolved electron temperature and density gradient lengths . 25
2.5 Pre-ELM edge ion temperature profile 26
2.6 Comparison of ELM resolved T_i and T_e profile evolution 27
2.7 Real space T_e and n_e evolution 28
2.8 ELM size vs. ELM frequency for AUG 2012 experimental campaign 30
2.9 B-coil layout at AUG 32
2.10 Cartoon ELM stability cycle 33
2.11 EPED model of ELM cycle 36
2.12 Stability diagram for AUG discharge #23417 38
2.13 Spectrogram using Mirnov coil data during discharge #27963 . . . 39

3.1 L/H mode pressure and j comparison 42
3.2 Magnetic field line pitch angle deflection due to ELM crash 44
3.3 Plasma cross section showing diagnostics 45
3.4 Fast ion contribution to total pressure for #23221 52
3.5 Pressure and current density profile for AUG #23221 53
3.6 Input profiles for high power phase of #23223 55
3.7 Pressure and current density fits during #23223 56
3.8 Evolution of ∇p and j over an ELM cycle 57
3.9 Change of separatrix and magnetic axis relative to ELM 58
3.10 Divertor and plasma currents during the ELM cycle 59

3.11 Surface plot of LFS j from discharge #23221 60
3.12 ELM synchronised peak current density gradient 61
3.13 Evolution of normalised pressure gradient and current density relative to ELM phases 62
3.14 j comparison in fuelling study 64

4.1 Magnetic errors vs. pressure profile shift for discharge #23223 . . 69
4.2 Electron temperature profiles for discharge #23223 69
4.3 Input profiles for high power phase of pedestal width scan 71
4.4 Magnetic errors vs. pedestal width change for discharge #23223 . 72
4.5 T_e and j profiles for modified width scan. 73
4.6 Magnetic errors for modified width scan in discharge #23223 . . . 73
4.7 Combined magnetic errors for modified width scan in discharge #23223 74
4.8 Strikepoint location in (R,z)(a) and divertor coordinate s(b) for discharge #23223 76
4.9 CLISTE and pressure data confidence interval comparison 78
4.10 Statistical deviation in magnetic equilibrium separatrix as compared to n_e and function parameterisation. 80
4.11 Modified fitting of separatrix location distribution showing better agreement 81
4.12 Raw ECE vs. ECFM showing separatrix location. 82
4.13 Relative separatrix movement from T_e, n_e and EQH. 84

5.1 Rough sketch of the bootstrap mechanism 88
5.2 Pre-ELM $\langle j_{||} \cdot B \rangle$ 91
5.3 Elm resolved E_ϕ and $\langle j_{||} \cdot B \rangle$ 93
5.4 Collisionality variation of bootstrap coefficients 95
5.5 Bootstrap drive comparison in fuelling study 96
5.6 ELM resolved $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ comparison for fuelling study 97
5.7 Inclusion of T_i measurements in current density calculation 99

6.1 T_e profiles and fitted ∇p comparison for database 103
6.2 Errors for all database fits against time 105
6.3 Peak edge j vs ∇P 107
6.4 Normalised j vs α 108
6.5 $\langle \boldsymbol{j}_{||} \cdot \boldsymbol{B} \rangle$ comparison from database 109
6.6 Scaling of LFS peak edge current density with pedestal top values 110
6.7 ELM frequency vs ∇n_e driven current fraction 112

6.8 Correlation of ELM loss and frequency with heating power 113
6.9 Scaled T_e values and interpreted collisionality values for a large database without advanced analysis 114
6.10 Projected collisionality against normalised pedestal beta poloidal and ion gyroradius 115
6.11 Scaling of separatrix contraction with ELM size 116
6.12 Regression of f_q and α with δ and j 118

7.1 Type-II ELMy discharge #25740 122
7.2 Electron temperature and density profiles used for Type-II ELM analysis 124
7.3 Edge current density profiles for Type-II ELM comparison 124
7.4 Electron temperature and density bootstrap components for Type-I and -II ELMs 125
7.5 Comparison of N_2 seeding with reference discharge 127
7.6 n_e and T_e profiles for N_2 seeding and reference discharges 128
7.7 ELM synchronised divertor current for N_2 seeding and reference discharges 129
7.8 N_2 seeded edge current density profile comparison 129
7.9 Pre- and post-ELM flux surface averaged current density profiles for N_2 seeding 131
7.10 Timetrace of key plasma parameters for discharge #28847 showing ELM mitigation 134
7.11 n_e (a) and T_e(b) profiles for discharge #28847 135
7.12 Current density, pressure gradient, and neoclassical current density for RMP discharge #28847 136
7.13 j-α diagram for RMP discharge #28847 137
7.14 Divertor currents showing ELM mitigation onset for MP comparison 138
7.15 Pedestal top evolution for ELM mitigation discharges 139
7.16 Evolution of ELM mitigation in j-α space for MP comparison . . 140

List of Tables

1.1 Engineering parameters for AUG 8

5.1 Fuelling study parameters (fuelling rate, ν^*, and f_{ELM}) 95

6.1 Engineering and pedestal top parameters for database 102
6.2 Correlation between $\nu*$, $p_{e,top}$, $T_{e,top}$, $n_{e,top}$, and max(j_{edge}) 111

7.1 Pedestal top and gradient evolution for discharge #28847 135

Abstract

The sudden decrease of plasma stored energy and subsequent power deposition on the first wall of a tokamak device due to edge localised modes (ELMs) is potentially detrimental to the success of a future fusion reactor. Understanding and control of ELMs is critical for the longevity of these devices and also to maximise their performance.

The commonly accepted picture of ELMs posits a critical pressure gradient and current density in the plasma edge, above which coupled magnetohydrodynamic (MHD) peeling-ballooning modes are driven unstable. Much analysis has been presented in recent years on the spatial and temporal evolution of the edge pressure gradient. However, the edge current density has typically been overlooked due to the difficulties in measuring this quantity. In this thesis, a novel method of current density recovery is presented, using the equilibrium solver CLISTE to reconstruct a high resolution equilibrium utilising both external magnetic and internal edge kinetic data measured on the ASDEX Upgrade (AUG) tokamak.

The evolution of the edge current density relative to an ELM crash is presented, showing that a resistive delay in the buildup of the current density is unlikely. An uncertainty analysis shows that the edge current density can be determined with an accuracy consistent with that of the kinetic data used. A comparison with neoclassical theory demonstrates excellent agreement between the current density determined by CLISTE and the calculated profiles.

Three ELM mitigation regimes are investigated: Type-II ELMs, ELMs suppressed by external magnetic perturbations (MPs), and Nitrogen seeded ELMs. In the first two cases, the current density is found to decrease as mitigation onsets, indicating a more ballooning-like plasma behaviour. In the latter case, the flux surface averaged current density can decrease while the local current density increases, thus providing a mechanism to suppress both the peeling and ballooning modes.

Chapter 1

Introduction

As we move through the early 21st century, climate change must be addressed. There is a general consensus among climate scientists that this climate change is man-made and due to long-lived greenhouse gases (LLGHG), with a 2007 inter-governmental panel[1] concluding that the combined contribution of man-made climate change is likely to be five times that of the solar irradiation changes. To combat this, a worldwide policy of capping global warming at a limit of 2 ° C has been adopted, with one estimate[2] concluding that less than half of the then proven economically recoverable oil, gas, and coal reserves may be burned before 2050 to achieve such a goal.

While great strides are being made to reduce power consumption, and hence reduce the demand on these fuels, extra sources of energy will certainly be required to replace the use of LLGHG emitting fossil fuels. This energy should be clean, cheap, and dependable. Attaining all three of these is a daunting challenge. Efforts to date have focussed on using solar cells and wind turbines to contribute the majority of the power.[3]. However, due to the inherent unpredictability in both of these means of power generation, the use of some sort of storage mechanism to balance the supply and demand of power is required. This can be achieved either through the use of batteries, or through water storage. The latter has been implemented in the Turlough Hill power station in Ireland and consists of two vertically separated reservoirs. The water is pumped from the lower to the higher when excess power is produced and flows back down through turbines when more power is needed.

1.1 The case for nuclear power

One obvious solution to the energy challenge is nuclear power. Nuclear fission has been in use across the world for several decades now and is by and large quite successful. An important issue for nuclear fission is the long term storage of nuclear waste. The radioactive output of these plants has a half life of tens of thousands of years, requiring the use of large facilities designed to contain this radiation. However, there is another option; nuclear fusion. This has been presented for the past 50 years as the solution to all our energy needs. At face value, it does seem to be the best solution. It promises to be clean, since there is no radioactive by-product from the reaction itself; there is concern regarding long lived radioactivity in that the components of such a device would be activated by high energy neutrons. However, research is currently underway which is investigating the use of low-activation steel as the material of choice for such a reactor[4]. Fusion power should be cheap, since the inputs are isotopes of Hydrogen. It would also be dependable, since once a fusion reaction begins to burn, it becomes self sustaining and is limited only by the device in which it is contained. It is also inherently safe; once the hydrogen-like fuel has been burned, the reaction stops.

No dream solution is without its drawbacks, however. Over the course of approximately 50 years of fusion plasma physics and technology research, several issues have arisen. The high temperatures and pressures required by the fusion reactions give rise to several families of instabilities. These instabilities can limit the plasma performance, since the ratio of the kinetic pressure and magnetic pressure, or beta, must be kept below a critical threshold, the Troyon limit[5]. While the magnetic field can be increased to allow operation at higher pressure, this increases the cost of a reactor. Other relaxation-type instabilities expel large amounts of the plasma stored energy, reducing the lifetime of the components used in such a device. Added to this is a requirement of a minimum size; small reactors are not a possibility for fusion due to the required large temperatures; a minimum volume to surface area ratio is required. Additionally, confinement scaling laws reveal a dependance on the square of the major radius of a toroidal device, meaning that only very large machines, such as the Iter reactor currently under construction in Cadarache, France, can attain the break-even point of generating enough power to sustain itself, a condition known as ignition. These challenges require collaborations between engineers and scientists around the world to solve.

The reaction with the highest cross section is the fusion of two hydrogen

isotopes, deuterium and tritium, which proceeds with the reaction:

$$ {}^{2}_{1}D + {}^{3}_{1}T \rightarrow {}^{4}_{2}He + {}^{1}n + 17.6MeV \tag{1.1} $$

The energy output in this case is distributed between the $\alpha-$particle (^4_2He), carrying 3.5 MeV, and the neutron, carrying 14.1 MeV. The difficulty involved in achieving a fusion reaction is that the isotopes need to breach the Coulomb repulsion barrier which exists between the two positively charged nuclei. The easiest way to achieve this is by heating the isotopes towards a highly energetic Maxwellian gas and allowing them to interact via random collisions. Since the temperatures required for this are far in excess of the ionisation energy of hydrogen, an ionised plasma is formed, allowing confinement using magnetic fields.

In order to increase the plasma to this high temperature, external heating must be initially applied to the plasma. When the plasma is in a equilibrium state, it contains an energy given by

$$ \mathrm{W} = 3\overline{\mathrm{n}}\overline{\mathrm{T}}\mathrm{V} \tag{1.2} $$

where $\overline{\mathrm{n}}$ and $\overline{\mathrm{T}}$ are the average values of density and temperature in the plasma, and V is the plasma volume. In order to relate the rate of power input to the stored energy of the plasma, a confinement time may be defined, given by

$$ \tau_E = \frac{\mathrm{W}}{\mathrm{P_{heat}}} \tag{1.3} $$

As the temperature increases, alpha particle heating also begins to play a role. This heating can be written as $\mathrm{p}_\alpha = \frac{1}{4}\mathrm{n}^2\langle\sigma v\rangle\mathcal{E}_\alpha$, where $\langle\sigma v\rangle$ is the reaction cross section times the particle energy, which is a function of temperature, and $\mathcal{E}_\alpha$ is the energy carried by the alpha particle after a fusion reaction. Once the alpha particles start to heat the plasma, the external heating power required to sustain a given plasma stored energy decreases. Since the aim is to have self sustained energy production, this external heating power should eventually go to zero and the stored energy should be provided entirely by alpha particles, yielding a minimum condition for ignition as

$$ \mathrm{n}\tau_E > \frac{12}{\langle\sigma v\rangle}\frac{T}{\mathcal{E}_\alpha} \tag{1.4} $$

Since the right-hand side of this equation is determined by T only, $\mathrm{n}\tau_E$ can be plotted as a function of temperature, which is shown in figure 1.1. This figure shows a minimum in $\mathrm{n}\tau_E$ near 30 keV. However, since τ_E is also a function of temperature, this minimum is slightly lower, in the range of 10-20 keV. Expressing

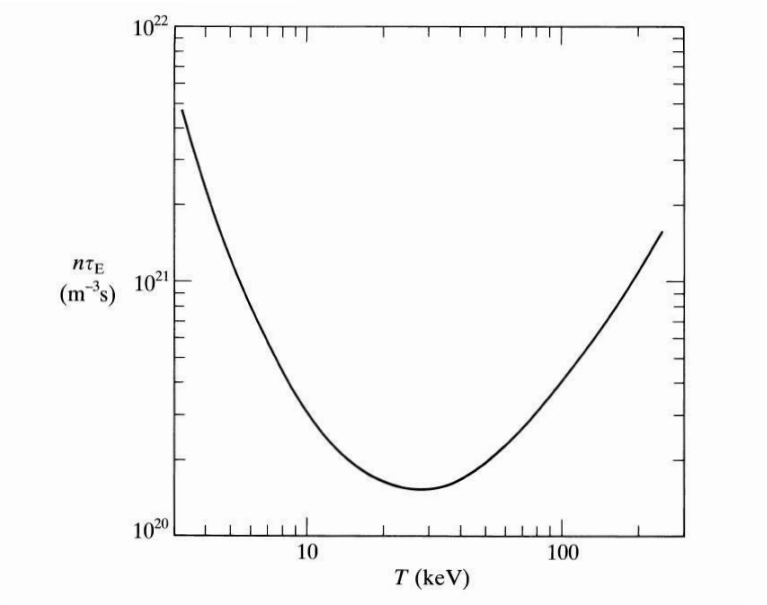


Figure 1.1: Product of density times confinement time as a function of temperature. This product shows a minimum near 30 keV, indicating an approximate range for fusion power plant operation. Figure reproduced from [6].

the three parameters, $nT\tau_E$, as a single unit, the criterion for ignition can be elegantly expressed in a Lawson criterion-like form:

$$nT\tau_E > 5 \times 10^{21} \frac{keVs}{m^3} \tag{1.5}$$

which is also known as the triple product. The temperature is fixed in the range of 10-20 eV, to give a minimum in the product $n\tau_E$, while there are two distinct approaches to attaining the $n\tau_E$ at a given temperature. One takes a very high density for a short confinement time and is called inertial confinement fusion. This uses many precision targeted lasers to compress a small fuel pellet for a fraction of a nanosecond. Another approach takes a density of $10^{20}m^{-3}$ and a confinement time of a few seconds. A temperature of 20 keV corresponds to approximately 200 million Kelvin; this is far in excess of what any present physical material can withstand. It is with regard to this challenge that the properties of a plasma are exploited to form magnetic confinement fusion.

1.2 Confinement

Since all the reaction particles will be ionised at these high temperatures, they can be trapped on applied magnetic field lines via the Lorentz force. This gives rise to a gyration around the magnetic field lines with a gyro-frequency given by

$$\omega_c = \frac{qB}{m} \tag{1.6}$$

where q is the charge of the particle, B is the magnetic field strength, and m is the particle mass. This then corresponds to a gyro-radius of

$$\rho = \frac{v_\perp}{\omega_c} \tag{1.7}$$

where $v_\perp$ is the velocity of the particle perpendicular to the magnetic field.

Early attempts at fusion reactors focussed on linear attempts, or so-called particle pinch machines. These machines, however, had very large losses at the end of the field lines, even with the use of magnetic mirror techniques. A different approach eliminated the requirement to have an end to the field lines, creating toroidal confinement systems. Several iterations of these systems have been made, starting from bending several well known linear systems, such as the Z- and θ-pinches, into toroidal shapes. The trade off of poor toroidal or magnetohydrodynamic (MHD) stability of either of these confinement types has led to the development of various combination devices, notably the screw pinch, the stellarator, and the tokamak. While these configurations feature good toroidal confinement, they are also subject to several other forces.

In a pure toroidal pinch device (a toroidal version of the linear θ-pinch) the magnetic field is generated solely by current flowing in external coils which wrap perpendicularly around the plasma. This gives rise to a $1/R$ dependence of the magnetic field, meaning that the gyrating particles have unequal gyro-radii in opposite halves of their orbits. As a result, electrons and ions drift to the top and bottom of the torus respectively. This sets up an electrostatic field and the associated $\mathrm{E} \times \mathrm{B}$ force acts in the radial direction, forcing both the electrons and ions to move as a single fluid towards the outer wall. By introducing a non-zero poloidal twist on the field lines, the orbits are averaged such that the vertical separation of the particles does not occur. This poloidal field can be created in one of two ways, both of which are being heavily researched at present. The first method is to impose this helical twist externally via special shaping of the external field coils, which is known as the stellarator concept. The other method is to drive a toroidal current in the plasma itself, adapted from the Z-pinch. Such a device, combining an applied toroidal magnetic field and a self generated poloidal magnetic field (as with the linear screw-pinch) is the tokamak[1] mentioned earlier.

The $\mathrm{E} \times \mathrm{B}$ drift force is not the only force which affects radial stability. The Hoop force and "tyre tube" forces also cause a radial movement of the plasma. The former is due to the poloidal field (which we have introduced to counter the radial drift) and the wrapping of the toroidal field. The poloidal flux, created

[1] A Russian acronym for "toroidal chamber with magnetic coils"

by the plasma current, defines a set of nested flux surfaces. These lines of equal flux have a larger surface area on the outboard side of the tokamak, meaning a lower poloidal magnetic field here. Since the magnetic force is proportional to B^2, there will be a net force in the outward direction. The tyre-tube force is related to the plasma pressure, which is also constant on flux surfaces; this will be shown later. Force due to pressure is directed radially outwards and also proportional to the surface area over which it acts. The larger surface area on the outboard side acts to produce a net outboard radial force. The 1/R dependence of the toroidal magnetic field is again an issue, since it is larger on the inboard side, creating a net outboard force adding to the hoop force.

These forces can be overcome by applying a uniform vertical field to the plasma. This field increases the z component of the poloidal field at the outboard side and decreases it at the low field side. The magnitude of this field can be varied such that the plasma has good radial stability. The vertical position of the plasma can be controlled in a similar manner, although this is prone to different instabilities. The principal components of a tokamak device, including the coils used to produce the toroidal and vertical fields are shown in figure 1.2.

This toroidal current of a tokamak is realised by applying a current through a solenoid in the hollow of the torus with the plasma acting as the secondary of a transformer. Driving a current in the plasma does come with drawbacks, however. For example, the duration of a single plasma discharge is limited by the flux in the central solenoid, also shown in figure 1.2, necessitating pulsed operation which is not ideal since it induces high loads on the components. Several quasi steady-state or long pulse operation scenarios are under development at the moment, including analysis of data from JT-60U in Japan, which can operate at Iter-relevant levels of plasma performance for up to 30 s[7]. This allows the impact of such long pulses on materials, both for the plasma facing components and for the magnetic coils to be determined.

1.3 Description of a typical tokamak device

The basic components of a tokamak were described in the previous section. In addition to the magnetic confinement components, there are also several other important pieces to a tokamak; the first wall, through which heat from the plasma is conducted, and which acts as shielding for diagnostics, the vacuum vessel, the heating systems, and, at a minimum, magnetic diagnostics to determine the plasma current, shape, and position. ASDEX Upgrade (AUG), the tokamak

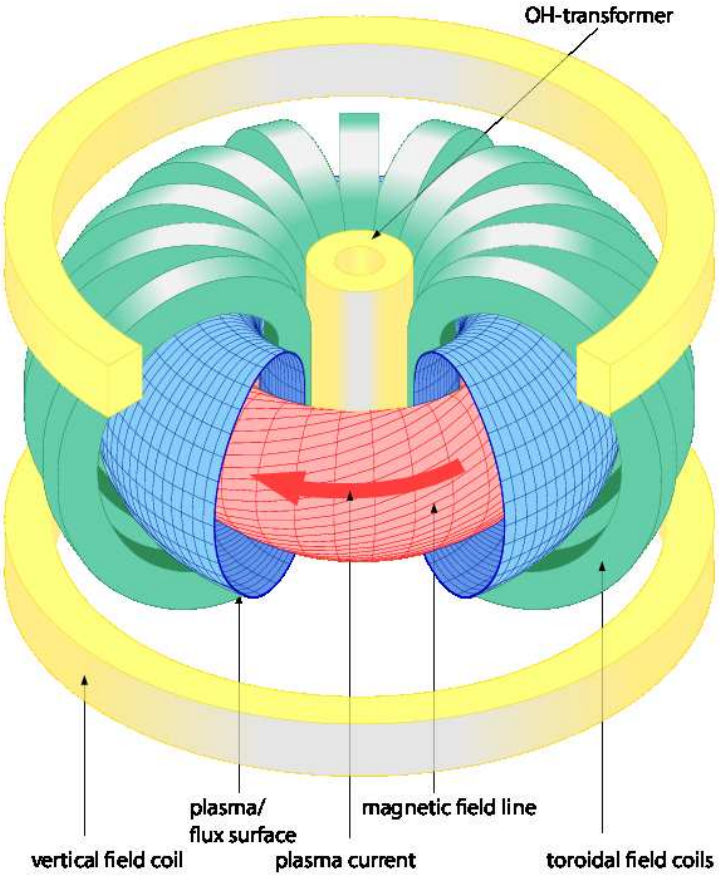


Figure 1.2: The principal components of a tokamak: the vertical field coils are shown in yellow; the toroidal field coils in green; a central solenoid is shown in the middle of the torus. Also shown are the plasma surface and the plasma current. Figure reproduced from AUG database.

which this work has been carried out on, is a mid-sized tokamak. The key engineering parameters for AUG are given in table 1.1. A description of the first wall and heating systems at AUG is given in this section, while the diagnostics, both for the magnetic and kinetic descriptions of the plasma will be described in chapter 3.

1.3.1 Poloidal cross section

To ensure the longevity of the plasma facing components, the power load from the fusion plasma on this "first wall" should be kept within tolerable limits. In smaller devices, it is possible to define a material limiter which sets a maximum allowable radius for the plasma; as the toroidal field and plasma current are in

Table 1.1: Engineering parameters for the ASDEX Upgrade tokamak.

Major radius (R_0)	1.65 m
Minor radius (a)	0.5 m
Plasma current (I_p)	0.4–1.2 MA
Toroidal magnetic field (B_T)	1.8–2.8 T
Heating power (P_{heat})	(up to) 30 MW
Plasma Volume (V)	14 m^{-3}

the toroidal direction, to first order, breaking the nested flux surface at a single toroidal location would then define the boundary of the plasma. This, however, comes with the drawback of high erosion rates on this limiter as well as a large impurity influx; such a limiter would have to be made from a material with a high melting point, such as tungsten, which radiates strongly when exposed to the high temperatures of a fusion plasma. This would be detrimental to eventual power production, as the plasma would most likely not be able to heat itself due to these losses.

A second option is the so-called divertor plasma. A large external current is applied parallel to the main plasma current, which causes the formation of a magnetic separatrix. This plasma configuration then consists of a set of closed nested internal field lines, the plasma boundary, or separatrix, and the open field lines situated outside the separatrix. An example of AUG geometry is shown in figure 1.3. The separatrix geometry has the obvious advantage that the majority of the plasma is contained inside the separatrix (also known as the last closed flux surface), with only a small portion touching the plasma facing components. The region of the poloidal cross section where the separatrix "legs" and the field lines connect to the plasma facing components is called the divertor region. This region experiences the highest heat loads in the device. The section of plasma outside the separatrix is called the scrape-off layer (SOL, purple shading) and the region between the two separatrix legs in the divertor is referred to as the private flux region.

1.3.2 Plasma facing materials

While the bulk of the plasma is, in theory, confined inside the separatrix region, radial particle and heat transport is unavoidable in a plasma device where temperature and density gradients exist. In addition, plasma instabilities can also cause the frequent, rapid, and large loss of particles and heat from the plasma. The key

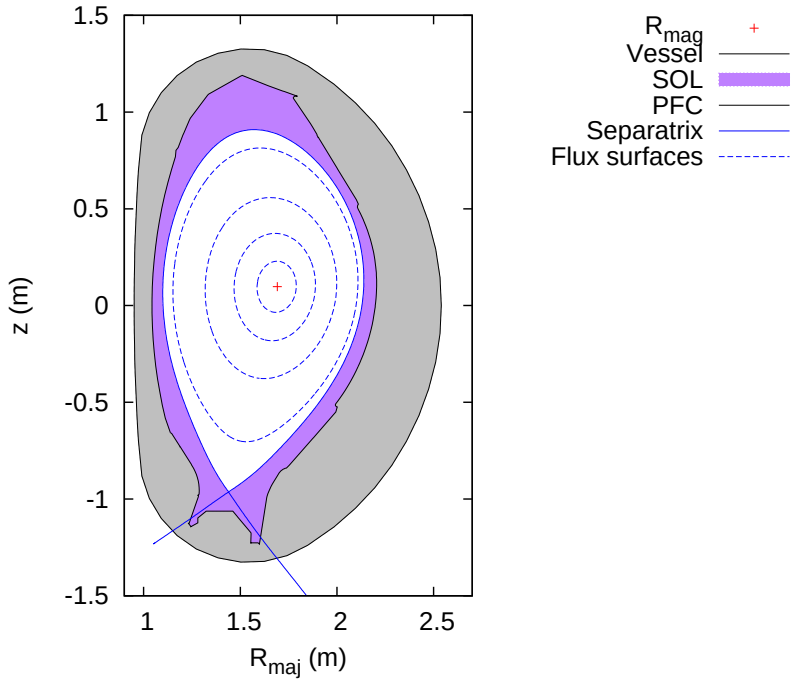


Figure 1.3: Separatrix geometry inside the vacuum vessel and plasma facing components (PFCs). The area between the vessel wall and the PFCs is shaded in grey, the scrape-off layer (SOL) is shown in purple, while the separatrix is shown as a solid blue line. Inside the separatrix, internal surfaces at 0.25, 0.50, 0.75, and 0.95 of the a poloidal flux radius (the square root of the normalised poloidal flux) are shown as dashed blue lines, while the magnetic axis is shown as a red cross.

components of the tokamak must be protected from this high energy flux, which is the job of the so-called "first wall". This requirement of the first wall suggests that it should be made from a material with a high melting point. However, it is not so straightforward to choose the material based on this. Many materials with high melting points are also high-Z materials, such as Tungsten. If these impurities were to be eroded from the first wall, they would radiate strongly in the hot core of a fusion device, lowering the available heating power and reducing the fusion yield.

In operation at present, there are two main choices of first wall material: Carbon and Tungsten. Carbon has the advantage of being a low-Z material, which is likely to be stripped of its electrons at or outside the plasma edge and thus does not radiate in the hotter core. Its disadvantage lies in its erosion rate, which is higher than that of higher-Z materials. In addition the tritium retention of Carbon is higher than Tungsten, making it a poor choice for use in a reactor device[8]. For the case of Tungsten, its low Tritium absorption and retention rates, as well as its high melting point, have made it the subject of intense study

over the past twenty years[9, 10, 11]. ASDEX Upgrade (AUG), the tokamak on which the present work has been carried out, has been converted from a full Carbon device to one with Tungsten coated carbon plasma facing components[12]. In the last few years, JET has also been upgraded to an all metal wall[13, 14, 15], though in this case only the divertor region has been coated in Tungsten while the rest of the device has been covered to Beryllium; this Beryllium/Tungsten mix is presently foreseen for eventual use in Iter, or potentially also as the initial material instead of the presently planned carbon plasma facing components.

1.3.3 Plasma heating techniques

While the Ohmic current in the plasma also applies a small amount of heating to the plasma, much more is required in order to reach fusion relevant temperatures. To this end, several methods of applying heat to the plasma have been developed. External power input comes from two main sources: microwave type heating, and injection of energetic particles. The former can be accomplished by heating at the cyclotron resonance frequency of either electrons or ions, while the latter relies on injecting highly energetic neutral particles which then interact with and scatter off the thermal particles in the plasma.

ASDEX Upgrade is equipped with several heating systems and all of these types of heating can be applied. Up to 5 MW of electron cyclotron resonance heating (ECRH), 7 MW ion cyclotron resonance heating (ICRH),and 20 MW of neutral beam injection heating (NBI) can be applied to the plasma. Since the installation of the Tungsten plasma facing components, ECRH is applied as standard to the core of AUG[10, 16] in order to reduce the accumulation of central impurities which leads to high radiation rates and discharge loss.

1.3.4 Plasma position control

Control of the position and shape of the plasma in the vacuum vessel is required for the successful management of the experiment. To this end, AUG is equipped with several coils, the current in which can be altered to produce the desired shape and location. The gross plasma position can be set via feed-forward programming of the magnitudes of the currents in the coils, while fine adjustments can be made while the discharge is running via a feedback control system based on the desired control parameters. These are determined in real time via a linear regression of signals calculated from an equilibrium database (see section 1.5.2), though other methods can also be used.

1.4 Magnetohydrodynamics

A magnetised plasma can be described in physical terms by magnetohydrodynamics (MHD). In this description, the plasma is treated as a fluid, either as two separate fluids to describe electrons and ions separately, or, more commonly, as a single fluid. Descriptions and derivations in the following sections have been taken from multiple sources, notably Wesson[6] and Freidberg[17].

In the simplest sense, the plasma can be considered as a single fluid, described by magneto-hydrodynamic (MHD) theory. The principal governing equations of MHD are:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho v) = 0, \tag{1.7a}$$

$$\rho \frac{\partial v}{\partial t} + \rho(v \cdot \nabla)v + \nabla p - j \times B = 0, \tag{1.7b}$$

$$\frac{\partial p}{\partial t} + v \cdot \nabla p + \frac{5}{3} p \nabla \cdot v = 0, \tag{1.7c}$$

$$J = \sigma(E + v \times B), \tag{1.7d}$$

$$\frac{\partial B}{\partial t} + \nabla \times E = 0, \tag{1.7e}$$

$$\mu_0 J = \nabla \times B, \tag{1.7f}$$

$$\nabla \cdot B = 0 \tag{1.7g}$$

where ρ is the mass density, v is the plasma velocity, p is the plasma pressure, j is the plasma current, B is the magnetic field, E is the electric field, μ_0 is the vacuum magnetic permeability, and σ is the conductivity. The first four equations are the continuity, momentum and energy equations, and Ohm's law, and the last three are three of Maxwell's equations (Faraday's law, Ampere's Law (low frequency), and Gauss' law for magnetism). If the plasma is considered to be in an equilibrium state (time derivatives = 0), several of these equations can be eliminated, reducing the number of equations further. In addition, the conductivity in Ohm's law can also be eliminated from an idealised plasma, simplifying the equations further. This leaves more limited applicability of the MHD description but can contain much of the physics, especially to leading order, when considering instabilities of the plasma.

1.4.1 Plasma equilibrium

The MHD description of a plasma requires knowledge about the magnetic structure. The toroidal field is well known and varies across the plasma with a 1/R dependence. However, the poloidal field, which is induced by the plasma current,

is not well known and depends on the distribution of the plasma current. Assuming that the plasma can be considered to be in an equilibrium state, the current density and magnetic field are related by Ampere's law: $\mu_0 \mathrm{J} = \nabla \times \mathrm{B}$.

The magnetic topology of the plasma in the poloidal plane is shown in figure 1.3 and can be considered as a set of nested flux surfaces; contours of equal magnetic flux. The existence of these contours can be deduced from $\nabla \cdot \mathrm{B} = 0$ and application of Gauss' theorem. We call any quantity which has the same value at all points on a flux surface a *surface quantity*; examples are pressure, which is taken as flux quantity in the ideal MHD framework considered here, and temperature and density. Temperature is the most rigorous flux quantity of these, due to the large parallel heat conductivity; any variations in temperature along the field lines are rapidly removed. Density can be considered a flux quantity in an isotropic plasma (or in the absence of rotation), though this is not strictly the case. These surface quantities allow us to make some simplifications in describing the plasma magnetic topology when constructing a magnetic equilibrium, that is, the map of contours of equal flux, examples of which are shown as dashed blue lines in figure 1.3.

If we start with the equation of force balance for an isotropic plasma in equilibrium, we have

$$\mathrm{j} \times \mathrm{B} = \nabla \mathrm{p} \tag{1.8}$$

From this we can deduce that pressure is a flux quantity (the dot product of B with either side is 0, implying no gradient of pressure along the magnetic field lines). Expanding on this equation to break the left-hand side into its component poloidal and toroidal parts we obtain

$$\mathrm{j}_\theta \mathrm{B}_\phi i_\phi + \mathrm{j}_\phi \mathrm{B}_\theta i_\theta = \nabla \mathrm{p} \tag{1.9}$$

where i_ϕ and i_θ are unit vectors in the toroidal and poloidal directions, respectively. The poloidal magnetic flux is defined as

$$\psi = \int_{\mathrm{S_{pol}}} \mathrm{B} \cdot \mathrm{d}S \tag{1.10}$$

The infinitesimally small poloidal area prescribed at a given major radius location is given by $\mathrm{Area_{pol}} = 2\pi R \mathrm{d}r$, giving

$$\mathrm{B_{pol}} = \frac{1}{2\pi R} \nabla \psi_{pol} \tag{1.11}$$

When a stream function $\psi \equiv \psi_{pol}/2\pi$ is defined, we are left with

$$\mathrm{B}_\theta = \frac{1}{R} \nabla \psi \times i_\phi \tag{1.12}$$

If we take $\mathrm{j}_\theta = \frac{1}{\mathrm{R}}(\nabla \mathrm{f} \times i_\phi)$, where f is the flux function describing the poloidal current, we arrive at

$$\frac{1}{\mathrm{R}}(\nabla \mathrm{f} \times i_\phi)\mathrm{B}_\phi i_\phi + \mathrm{j}_\phi \frac{1}{\mathrm{R}}(\nabla \psi \times i_\phi) i_\theta = \nabla \mathrm{p} \tag{1.13}$$

After some algebra and cancelling of parallel vectors we can derive

$$\mathrm{j}_\phi \nabla \psi = \mathrm{R}\mathrm{p}' \nabla \psi + \mathrm{B}_\phi \mathrm{f}' \nabla \psi \tag{1.14}$$

where $p' = \frac{\partial \mathrm{p}}{\partial \psi}$ and $\mathrm{f}' = \frac{\partial \mathrm{f}}{\partial \psi}$. Finally, setting $\mathrm{B}_\phi = \frac{\mu_0}{\mathrm{R}} f$ we obtain an expression for the current density as

$$\mathrm{j}_\phi = \mathrm{R}\mathrm{p}' + \frac{\mu_0}{\mathrm{R}} \mathrm{f}\mathrm{f}' \tag{1.15}$$

Since the current is related to the curl of the field, and the field is a derivative of the flux, we can also obtain an expression for the magnetic flux at a given point:

$$-\mu_0 \mathrm{j} = \mathrm{R} \frac{\partial}{\partial \mathrm{R}} \frac{1}{\mathrm{R}} \frac{\partial \psi}{\partial \mathrm{R}} + \frac{\partial^2 \psi}{\partial \mathrm{z}^2} \tag{1.16}$$

leaving us with

$$-\Delta^* \psi = \mu_0 R^2 \frac{dp}{d\psi} + f \frac{df}{d\psi} = \mu_0 R j_\phi \tag{1.17}$$

where $-\Delta^*$ is the Grad-Shafranov operator. This magnetic flux is then constructed as a 2D grid, describing the poloidal cross section of the plasma.

1.5 Equilibrium reconstruction

Determining the plasma equilibrium via the above Grad-Shafranov equation can be carried out in a number of ways. A wide variety of types of equilibrium codes exist; 2D or 3D, free- and fixed-boundary solvers, predictive and interpretive solvers. Each of these types of code has a specific function to perform and all, or at least many, are in regular use at most tokamaks. This section will give a broad overview of the principal reconstruction methods used at AUG and which will be used throughout this thesis. The descriptions given in this section are based on the work of Lackner[18] and McCarthy[19].

1.5.1 Predictive equilibrium calculation

A predictive equilibrium calculation is the calculation of a consistent equilibrium solution which satisfies given source profiles, p' and ff', prescribed poloidal field coil current, and the vessel geometry. From this knowledge, and knowing the plasma current, a magnetic equilibrium can be calculated. Due to the lack of

an explicit z dependence (unlike R) in the Grad-Shafranov equation, the vertical position of the plasma is numerically unstable. In order to enforce stability on the code, the magnetic axis location is also fixed by ensuring that $\nabla\psi = 0$ at a prescribed location. This form of calculation is used for generation of equilibrium databases and, due to the exact nature of the solution, for MHD stability analysis. The former application is used as a starting point for function parameterisation, a basic real-time plasma control and analysis tool.

1.5.2 Function parameterisation

The real-time control of the plasma position is vital for the success of the experiment. Typical parameters used are the current centre, the outer midplane separatrix location, or the location of the strikepoints. Using an equilibrium database, generated by taking basic current profile shape parameters and experimentally relevant plasma current and beta poloidal values combined with a Monte-Carlo sampling of possible control coil currents, the expected magnetic measurements (corresponding to magnetic field probes or flux loop difference values) are calculated[20, 21]. During the discharge, a linear regression is applied to a subset of the real time magnetic measurements to determine the desired controlled parameters (also called the feed-forward parameters). The currents in a subset of the plasma control coils (one control coil per degree of freedom) are altered to produce the desired values.

1.5.3 Interpretive equilibrium reconstruction

In addition to the predictive method of equilibrium calculation, where the form of the source profiles as well as the magnetic axis are prescribed, a very important tool is an interpretive equilibrium. With this formulation of the equilibrium problem, the control coil currents are known, but neither the magnetic axis location, or the shape of the source profiles are *a priori* known. Instead, a number of experimental measurements are used to reconstruct the equilibrium as a best fit to these measurements. A series of flux loop differences and poloidal magnetic field components are the basic diagnostics used in a magnetic reconstruction. In addition, several other constraints, such as knowledge of q at a given location, the measured pressure profile, currents in the divertor etc., can be used to more accurately constrain the solution. CLISTE[19, 22] and EFIT[23] are the two main codes used at various tokamaks at the moment. Both codes take similar approaches to solving the equilibrium problem, such as allowing both the mag-

netic topology and source profiles to vary at the same time. The same numerical vertical instability as described for the predictive calculation also exists for the predictive reconstruction. In the CLISTE code, this is solved by allowing a parasitic vertical shift of the plasma. This shift is of the order of 0.1 mm when the equilibrium is converged. Due to this small vertical shift force balance is not rigorously conserved; where strict force balance is required (for MHD stability analysis, etc.) this is corrected for with a predictive calculation, where the converged location of the magnetic axis is fixed.

Several other codes of this type also exist, with some of the most modern codes being designed to run in real-time in order to provide a more accurate picture of the internal magnetic structure in order to, for example, determine the location of particular resonant surfaces for stabilisation of neoclassical tearing modes. These solvers are optimised for speed at the expense of accuracy, using only a few basis functions and measurements to specify the equilibrium, but their power lies in their ability to provide an equilibrium on a timescale relevant to plasma control.

Throughout this work, the equilibrium code CLISTE has been used in interpretive mode to determine the axisymmetric equilibrium. In some cases, only external magnetic data have been used, though the majority of reconstructions presented also included the current in the SOL and an experimentally determined pressure profile as added constraints.

1.5.4 Fixed boundary solvers

If the boundary is held fixed and the source profiles known, the equilibrium solution can be obtained to a very high degree of accuracy, such as that required by transport or MHD stability codes. HELENA[24, 25] and SPIDER[26] are two examples of this class of codes. The former is used as standard at AUG for preparing equilibria described by field aligned coordinates for MHD stability analysis, while the latter is used as a solver for transport analysis. The VMEC solver[27, 28], originally designed for use in the non-axisymmetric system of stellarators, has also been applied to use in tokamaks where external magnetic perturbations are present. In this case, the Grad-Shafranov equation is not solved. Instead, the equation of force balance, or $j \times B - \nabla p = 0$ is enforced while the energy of the system is minimised. Instead of the p' and ff' source profiles being varied, Fourier functions of normalised radius, and poloidal and toroidal angles are varied to achieve these conditions.

1.5.5 Resonant surfaces

A useful description of the field line topology is given by the safety factor, q. The exact definition of q can be described as the gradient of the toroidal flux with the poloidal flux. More generally, q is defined as q = m/n, where m is the number of poloidal transits and n is the number of toroidal transits the field line undergoes before reconnecting with itself. As such, the q of a flux surface is a measure of the twist of the surface. Low order rational surfaces occur when m and n are small integers, making surfaces such as q = 1, q = 3/2 etc. The value of q increases towards the edge of the plasma, as the current enclosed by a flux surface increases.

In general, plasma instabilities can be described as having a toroidal and poloidal mode number; a mode will then resonate on a q-surface which has the same m and n values as it does. As an example, a peeling-ballooning mode might have a toroidal mode number of 10 and a poloidal mode number of 40 meaning that it resonates with a surface having a q of 4, which, for AUG, is close to the plasma edge in a typical plasma discharge.

1.6 H-mode

The H-mode (high confinement mode)[29, 30] is a favourable confinement mode first discovered on the ASDEX machine in 1982. It involves increasing the power input to the plasma by any of the means described previously. Once the input power reaches a density dependent threshold[31, 32], a transport barrier forms at the edge of the plasma, characterised by an increased pressure gradient in the last ~2 cm of the plasma radius. Confinement is also greatly improved in the plasma, shown by the increase in the electron density and electron temperature in figure 1.4. The large increases in both of these quantities across the entire plasma radius stem from the increased gradients in the outer region of the plasma. This transport barrier region is known as the pedestal, as the core plasma profiles sit atop it. The conditions for H-mode access are well documented[31]; the physics, however, is not.

Despite attempts to create a working predictive model of the transition to H-mode (L-H transition), no theory is yet widely accepted. Suppressed turbulence in H-mode compared to L-mode at the plasma edge has been well documented. As such, the theories describing the L-H transition focus on this as the principal source of improved confinement, with $E \times B$ shearing rates in the pedestal region considered to be the dominant driving mechanism. The formation of a strongly

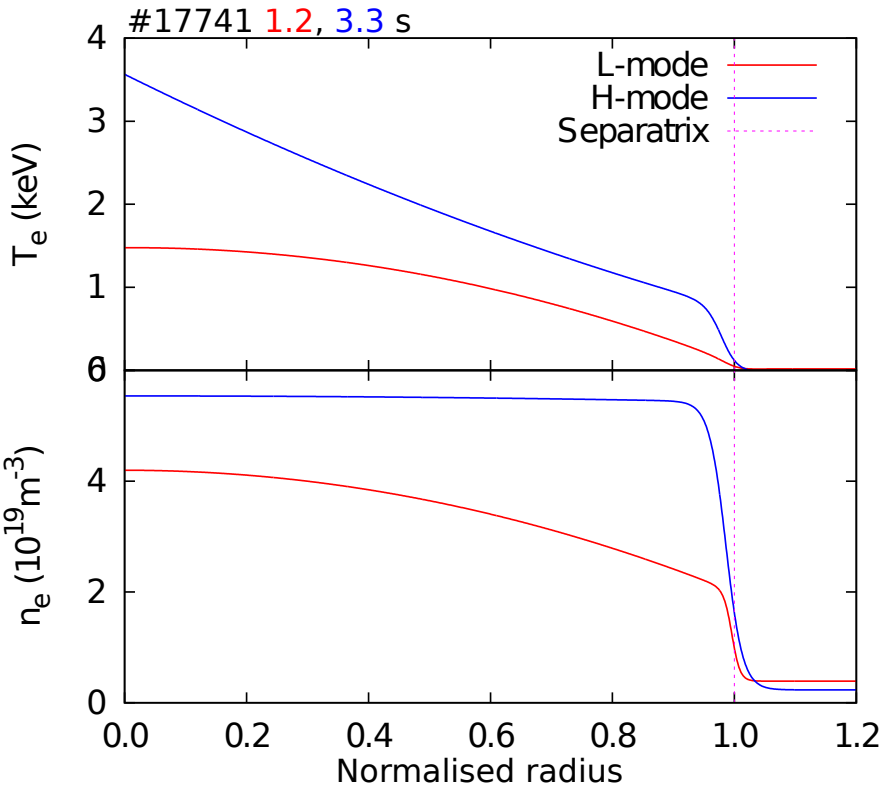


Figure 1.4: (a): Electron temperature profiles before (red) and after (blue) an L-H transition for AUG discharge #17741. The increased core temperatures and densities stem from the increased edge gradient region, known as the pedestal.

sheared radial electric field across an L-H transition has recently been reported at AUG by Viezzer et al.[33]. The properties of this edge transport barrier are fascinating. This barrier acts as a tremendous insulator, with both the temperature and density at its top and at the separatrix differing by approximately an order of magnitude. Due to the increased confinement, this is foreseen to be the eventual operating mode of the Iter experiment.

1.6.1 Plasma coordinate systems

Profiles of temperature and density (as well as several other plasma quantities) can be described in a number of coordinate systems. The most simple is a cylindrical (R,ϕ,z) system, where R is the major radius of the torus, ϕ is the toroidal angle, and z is the vertical coordinate. This can be easily converted into a Cartesian-cylindrical hybrid system of (ρ,ϕ,θ) with a normalised radius ρ and poloidal angle θ. The form of this normalised radius varies. One example is simply defined as $\rho = r/a$ where r is the distance along the major radius from the plasma center towards the separatrix and a is the location of the separatrix along this axis. A useful coordinate, given that the plasma consists of a set of

nested flux surfaces and many properties of interest can be considered constant on these surfaces, is a normalised flux radius. Using the poloidal magnetic flux ψ, for example yields a quantity $\rho_{\text{poloidal}} = \sqrt{\frac{\psi_{\text{axis}}-\psi}{\psi_{\text{axis}}-\psi_{\text{separatrix}}}}$. A similar coordinate ρ_{toroidal} used the toroidal magnetic flux ϕ and is very common in analysis of the core plasma. For the majority of this thesis, the poloidal flux will be used to define the plasma radius.

1.6.2 Edge Localised Modes - ELMs

Due to the existence of the edge pedestal and the associated high pressure gradient, several instabilities can potentially exist in the plasma edge. One example of such a plasma instability is an edge-localised mode (ELM), thought to be caused by a long wavelength MHD perturbation known as a peeling-ballooning mode. ELMs occur rapidly; a full crash can last mere hundreds of microseconds. During this crash, large amounts of particles and energy are expelled from the confined plasma which then travel along the open field lines in the SOL and impact the walls or divertor tiles. The frequency of these events varies from a few Hertz in larger machines, such as JET, to several hundred Hertz, in AUG for example. A broad overview of ELM phenomenology can be found in reference [34], and can be broken into three more or less distinct types:

- Type-I ELMs. These expel a large amount of stored energy and particles once the crash occurs. Since these ELMs occur in discharges with the best plasma performance characteristics, these ELMs are a concern for next generation devices, such as the Iter experiment currently under construction in Cadarache, France. It is predicted that the heat loads induced onto the surrounding wall of device by Type-I ELMs would be enough to cause substantial melting of any material[35, 36].

- Type-II ELMs/grassy ELMs. These ELMs are characterised by a much smaller size and higher frequency (several hundred Hertz) compared to Type-I ELMs. Each individual ELM has a much lower impact on the first wall of a device[37], reducing the damage to the plasma facing components.

- Type-III ELMs. Similar to Type-I ELMs, these have a smaller size and feature magnetic fluctuations just before the ELM crash occurs.

At present, ELMs are not fully understood. Several successful attempts to mitigate ELMs, via, for example, injecting small pellets into the plasma[38, 39], or by applying a toroidally varying magnetic perturbation[40, 41, 42, 43] to the

quasi-axisymmetric system of a tokamak have been made, but, as yet, these results are also not fully understood. The primary focus of much work in pedestal and edge physics at the moment is finding a description of ELMs, their causes, and modelling their expected frequencies and sizes[44, 45, 46, 36]. This is especially important when extrapolating results to larger machines where ELMs would be more detrimental to stable operation.

In order to understand and model ELM crashes and the recovery of the plasma from an ELM crash, tools are required. Many attempts have been made to predict the critical plasma parameters from different theories; the pedestal height and width are the dominant parameters. One of the more favoured theories, and certainly the most tested theory, is that of a combination of linear instabilities[47] in the pedestal, notably a kinetic ballooning mode (KBM) acting initially to clamp the pressure gradient, followed by the onset of a coupled peeling-ballooning mode once a critical height and width has been achieved. The former is linked to plasma turbulence while the latter is an ideal MHD instability. Peeling-ballooning theory posits a limiting pressure gradient, which varies as a function of the pedestal width, becoming lower at larger pedestal widths, and current density in the pedestal and is the principal motivation for this work. In order to test this theory, the magnitude and profile shape of both of these quantities is required. ELM resolved pressure gradients can be measured in a plasma device[48, 49] via temperature and density profiles as will be discussed in chapter 3, while this thesis will focus on the current density profile.

1.7 Focus of this thesis

There is a wide field of literature discussing the pressure build up to an ELM crash. Burckhart et al.[48] and Wolfrum et al.[49] reported extensively on the recovery of the pressure gradient and also on the separate temperature and density gradients from an ELM crash. One key observation in both these cases was the saturation of the pressure gradient before the ELM crash, which discounts a simple pressure gradient limit as being the cause of an ELM crash; since KBMs are pressure dependent, they are also linked to the pressure profile, which does not always exhibit substantial change leading up to an ELM crash. A possible explanation for this is that there is a resistive delay (~1 ms) in the current recovery after the crash and that the current density driven peeling mode is ultimately responsible for the ELM crash.

Potential methods to measure the current density are varied, and difficult;

probes cannot be used as the heat loads from the plasma are much too high. A rather elegant method exploits the motional stark effect (MSE) diagnostics[50, 51], which measure the relative strengths of the magnetic fields (also known as the pitch angle) deduced from the diagnostic line splitting and shift. It is a challenging technique and is made even more so at the edge of a tokamak where large changes in the local current density have only a small impact on the overall field line angle. All variations on this theme, such as Zeeman splitting of a diagnostic Lithium beam[52], field line imaging using electron Bernstein wave emission suffer from the same issue, as well as being even more technically complex. Success has been had on MAST[53], where the relatively weak toroidal magnetic field at the low-field side (LFS) has been exploited to make pitch angle measurements using an MSE diagnostic. This thesis is concerned with another method, directly related to the magnetic equilibrium where the plasma current is given by equation 1.17.

Here, the focus will be on the behaviour of the current density in relation to ELMs. Several aspects will be taken into account, including the critical values before an ELM crash, the recovery phase, and also the total current in the plasma edge. The driving mechanisms for the current will be discussed as well as engineering parameter (e.g. plasma current, heating power, etc.) dependencies. The next chapter will discuss the ELM phenomenon and some common observations at the ELM crash and the subsequent plasma recovery. Chapter 3 will introduce the method and diagnostics used to recover the current density and describe the current density evolution relative to the ELM crash. A sensitivity study is presented in chapter 4 and gives an estimate of the uncertainties expected from the analysis. Results from this analysis will be compared with theoretical calculations of the edge current density in chapter 5. A broader approach to understanding the current density magnitude is presented via the use of a database in chapter 6. Results from some special cases (Type-II ELMs, the smaller ELMs in nitrogen seeded discharges, and small ELMs induced by external magnetic perturbations), are shown in chapter 7.

Chapter 2

ELM phenomenology

The introduction described the basic effect of an ELM crash, and the three main types of ELMs. In this chapter specific aspects of a Type-I ELM crash, which is by far the most common ELM type, and the plasma recovery from this crash will be detailed.

There are several key markers of ELM crashes, though they all stem from the same source; a loss of confinement in the plasma edge. When an ELM crash occurs, the edge pressure gradient drops, lowering the plasma stored energy by an amount $\Delta W_{mhd} = \frac{3}{2} \int \Delta P.dV$. For a reactor plasma, this causes unwanted large releases of energy along the open field lines in the SOL which eventually impacts the first wall, in particular the divertor tiles. Since an ELM crash decreases the plasma stored energy by the order of 10%, this implies large heat loads which have the potential to melt the divertors of future devices[35].

Figure 2.1 shows several ELM markers including W_{mhd} (a) and the peak power load on the divertor (b) for ASDEX Upgrade discharge #27963 between 3 and 3.5 s. The ELMs in this discharge are regular with a frequency of approximately 40 Hz. The discharge had a 1 MA total plasma current, a -2.5 T toroidal magnetic field, 5 MW of NBI heating, and 1.5 MW of ECRH heating. What can be seen in this figure, apart from the very strong effect of the ELM crash on the plasma, is that the ELMs appear to be very regular and similar. This similarity of ELMs allows us to consider an "average" ELM cycle by taking all the data in a given time window and synchronising them into one composite ELM time trace. In doing so we can increase diagnostic sensitivity and reduce the effect of outlying data points in the overall analysis of ELM cycles. This technique will be referred to throughout this thesis as ELM synchronising. Historically, the D_α signal in the divertor was used as the ELM marker. However, since the installation of the full Tungsten divertor at AUG, the quality of this signal has degraded. A high quality

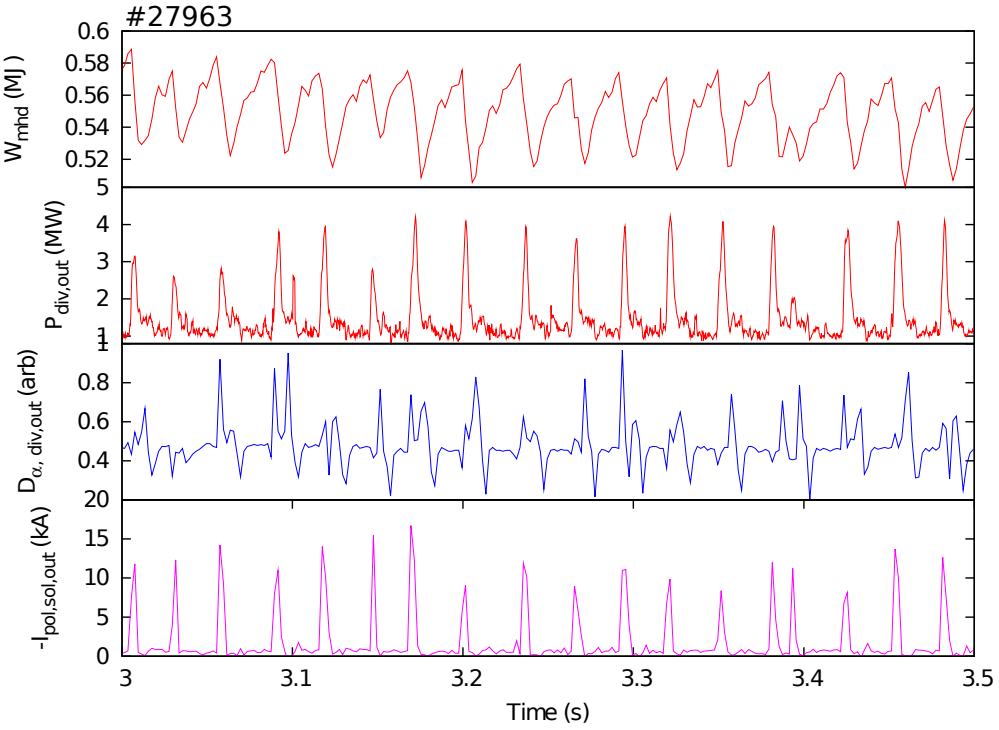


Figure 2.1: (a) timetrace of the plasma stored energy due to ELMs for a typical AUG discharge. (b) plots the peak power load on the outer divertor, (c) the corresponding D_α line emission, (d) SOL current measured at the outer lower divertor. The strong drops in stored energy and high peaks in divertor power are of particular concern for next generation devices

alternative is used in the form of SOL currents measured via shunt resistances measured via the metal divertor. The time trace of this measurement is shown in figure 2.1(d) where a large signal to noise ratio at the ELM crash is evident. In addition, it is sampled at a 100 kHz frequency, making it ideal to ascertain the exact start time of an ELM. In this thesis, the start time of the ELM crash is taken to be where the divertor currents start to rise, unless otherwise stated.

2.1 Pedestal evolution

Such an ELM synchronised time trace of (a): electron temperature, (b): density, and (c): pressure is shown in figure 2.2 for a radial location of $\rho_{pol} = 0.95$, which lies radially just inside the pedestal top (pre-ELM location = 0.954). These data have been determined via integrated temperature and density analysis and using electron cyclotron forward modelling (ECFM)[54] to determine accurate pedestal electron temperatures. All three time traces display a large drop at the ELM crash with the electron pressure pedestal top value decreasing by almost a factor

of 5 in a fraction of a millisecond. A typical feature of ELM recovery can be seen

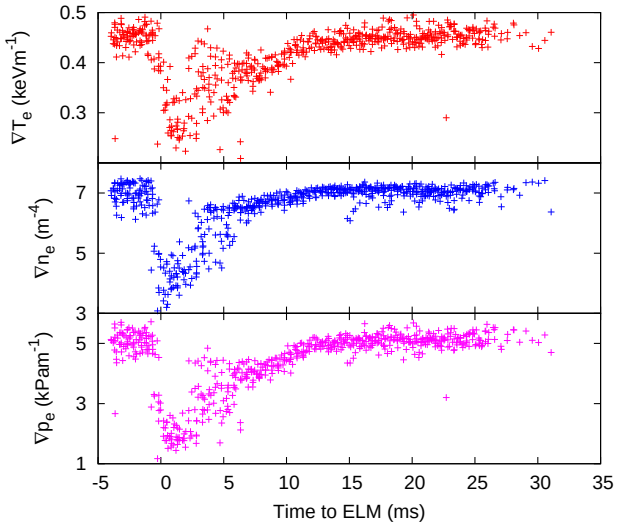


Figure 2.2: (a): electron temperature (b): pedestal top electron density (c): pedestal top electron pressure at the pedestal top relative to an ELM crash. The saturation of all three quantities significantly prior to the ELM crash can clearly be seen in this case.

in the behaviour of the pedestal top; the temperature and density recovery occurs at different rates. The density recovers rapidly in the 5 ms after the ELM crash, almost to its pre-ELM value, and then continues to recover at a gradual rate until it saturates before the onset of the next ELM. In the case of the temperature, however, there is a large scatter in the first 5 ms, and then a stagnation in the growth of the pedestal top temperature. At 10 ms, the temperature increases once more before saturating at its pre-ELM value at 15 ms.

The leading theory to explain ELMs is that of coupled peeling-ballooning MHD modes and a turbulence driven microinstability which limits the pedestal width. The ballooning MHD mode is sensitive to the local pressure gradient - in this case at the LFS of the plasma, where the curvature of the field lines is destabilising. Figure 2.3 shows the evolution of the electron temperature (a), density (b), and pressure (c) gradients over an ELM cycle. It is typical that that the electron temperature is the limiting factor in the recovery of the electron pressure[48], which is shown clearly in figure2.2 for this discharge. This indicates that the evolution of the pressure pedestal top and width is governed by the electron temperature dynamics.

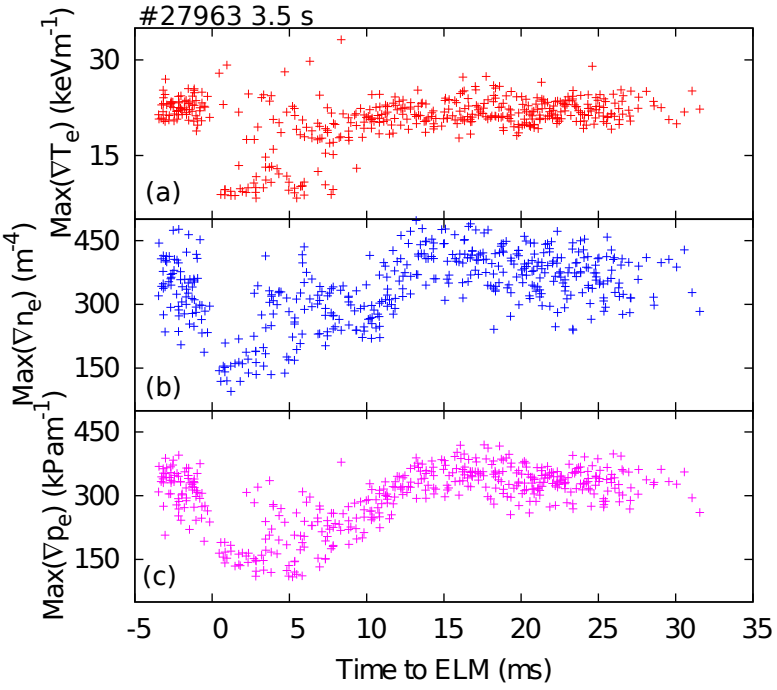


Figure 2.3: (a): maximal electron temperature gradient, (b): maximal electron density gradient, (c): maximal electron pressure gradient relative to an ELM crash.

2.1.1 Gradient length evolution

The gradient length ($= x/\frac{dx}{dr}$) is considered to be one of the main driving forces for microinstabilities in the plasma, such as electron/ion temperature gradient modes (ETGs/ITGs), or trapped electron modes (TEMs). Since one possible mechanism to describe the pedestal width is the limitation of radial transport, assumed to be turbulence dominated in a plasma, it is also useful to plot the gradient lengths as a function of time. Shown in figure 2.4 are the gradient lengths of electron temperature and density, as well as their ratio, $\eta_e = L_{\mathrm{n_e}}/L_{\mathrm{T_e}}$. The gradient lengths for the electron temperature ($L_{\mathrm{T_e}}$) and density ($L_{\mathrm{n_e}}$) show opposite reactions to the ELM crash, with $L_{\mathrm{T_e}}$ decreasing by approximately a factor of two and $L_{\mathrm{n_e}}$ increasing by almost a factor of 3. The distinct phases of ELM cycle recovery, which are clear in figures 2.2 and 2.3 cannot be clearly seen for either of these quantities, although they both reach their respective pre-ELM values at approximately 15 ms. A much more dominant structure can be seen in the evolution of η_e. A strong decrease at the ELM crash is followed by a phase featuring high scatter and a slightly increasing trend. This is then followed by a constant phase between 5 and 10 ms, reminiscent of the electron temperature gradient saturation reported by Burckhart et al.[48]. Between 10 and 15 ms η_e recovers fully and remains at its pre-ELM value for $15 - 20$ ms until the next ELM crash.

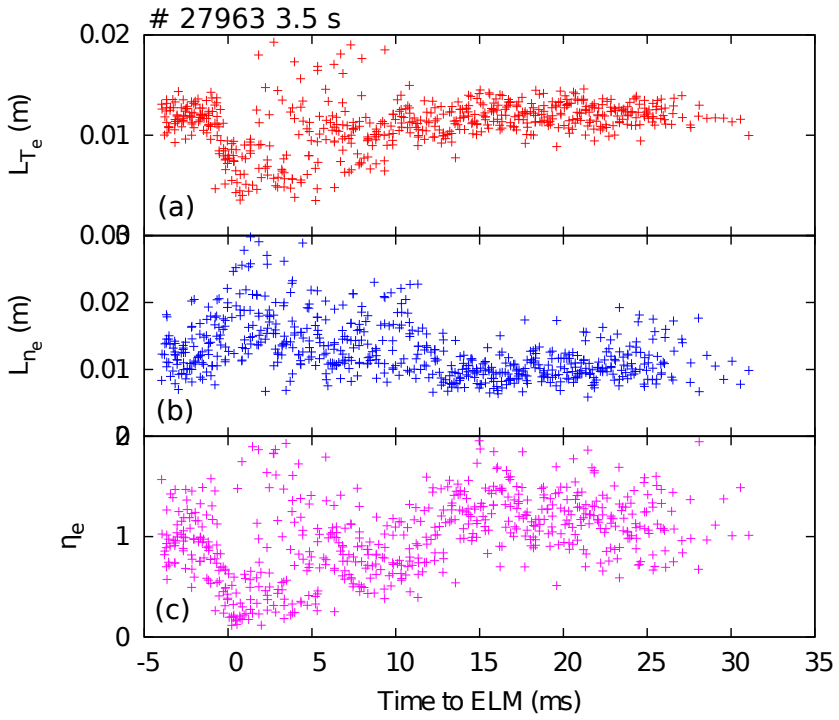


Figure 2.4: The gradient lengths of (a): electron temperature and (b):density. (c): evolution of η_e. A large scatter is evident in the gradient lengths, although there is a general trend matching the gradient evolution. The phases of η_e evolution closely resemble those reported for electron temperature and density in [48].

2.1.2 Comparison to ion measurements

This far we have focussed on the dynamics of electrons over the ELM cycle. The ions will also be affected by the ELM crash. Two new edge charge exchange recombination spectroscopy (CXRS) diagnostics have been recently installed at ASDEX Upgrade, one with a toroidal viewing angle and one with a poloidal viewing angle. Both diagnostics use the same NBI injection beam as a diagnostic beam and so measure at the same toroidal angle and have similar spatial recombination volumes. The combination of the high quality data from both diagnostics allows for new insights into the ion dynamics relative to an ELM crash. These diagnostics have a 2.1 ms integration time meaning that the post-ELM measurements are influenced by the temperature drop at the ELM crash. Thus, the actual recovery rate of the ion temperature could be faster than the current system can measure.

Shown in figure 2.5 is a profile of the data points taken from both systems between 4.5 and 1.5 ms prior to the ELM crash in the same time window as above. Overlayed in red is a hyperbolic tangent (tanh) fit to the data which gives a good fit to the data within the uncertainties and scatter. The blue line

shows the corresponding gradient, which is calculated analytically from the tanh fit, while the purple line denotes the location of the separatrix. Important points

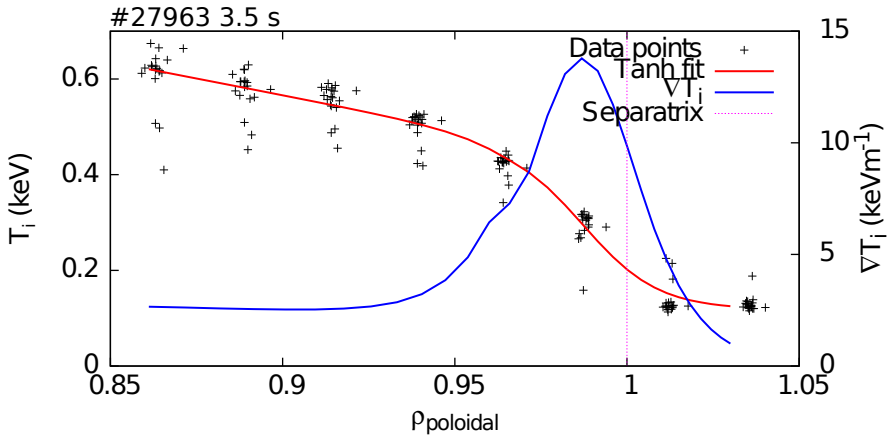


Figure 2.5: Pre-ELM ion temperature profile and ion temperature gradient, plotted with the data points used to generate the fit. The key features are a pedestal top temperature of approximately 450 eV at $\rho_{poloidal} = 0.95$, which is very similar to the electron temperature shown in figure 2.2, and a separatrix temperature of 200 eV.

to note about the form of the ion temperature profile are: the pedestal top temperature and location ($\approx$ 450eV located at $\rho_{poloidal} = 0.95$) agree well with those of the electron temperature, as shown in figures 2.2 and 2.6. The separatrix ion temperature is regularly observed to be 200 eV, a factor of two higher than observed for the electron temperature. CXRS data shown in the SOL should be ignored in all figures, as the impurity density (and hence the active signal, from which the ion temperature is derived) drops rapidly outside the separatrix[55]. Coupled with an increased passive spectrum, this makes the measurement of the active CXRS signal, and hence the ion temperatures, unreliable. These data are used in the fitting routine to provide a boundary condition for automated fitting routines, but are removed when fitting profiles "by-hand". The radial resolution of the ion temperature diagnostics can be increased by adding virtual lines of sight via a radial sweep of the plasma by 2 cm. This was done at a later time point in the discharge and no significant changes of either the pedestal top or separatrix temperatures, or the peak gradient were noticed during the sweep.

An interesting comparison is the evolution of the respective temperature profiles. Shown in figure 2.6 is the evolution of the temperature at $\rho_{poloidal} = 0.95$ (near the pedestal top) (a), peak gradient (b) and gradient length (c). What is clear from looking at (a) and (b) in figure 2.6 is that the pedestal top ion temperature decreases much less than that of the electron temperature, although the

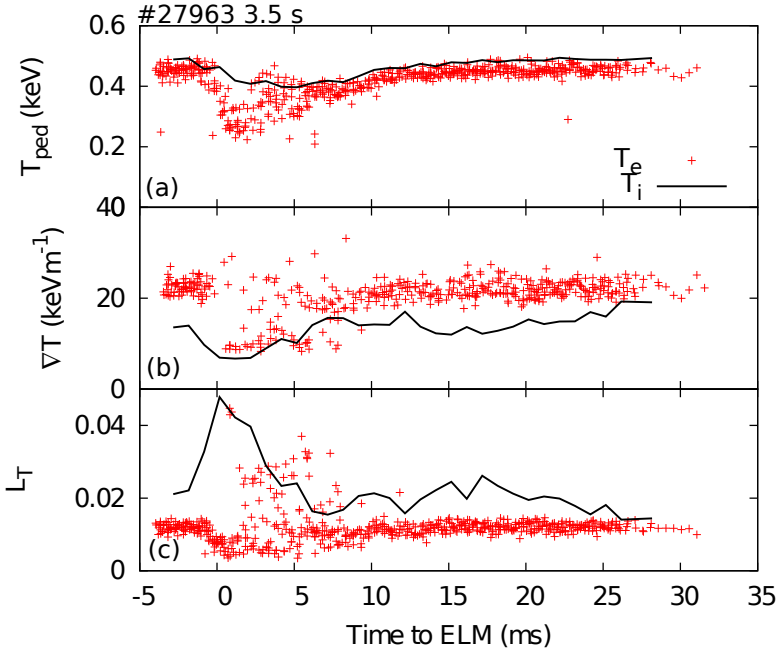


Figure 2.6: (a): evolution of the temperatures at $\rho_{poloidal} = 0.95$ which are noted to be the same throughout most of the ELM cycle. (b): peak edge gradient from both T_i and T_e.(c): gradient lengths corresponding to (a) and (b).

gradient appears to have a similar drop. While a first assumption may be that the lower T_i drop could be due to an artificial smoothing of the real situation caused by the 2.1 ms integraiton time of the edge CXRS systems, averaging the electron temperature data into the same bins still shows a drop of the pedestal top electron temperature to just above 300 eV. While not as large as the drop in the raw T_e data, this is still significantly higher, indicating that the ion and electron pedestals may have different loss mechanisms. A very pronounced increase in the gradient length is seen in (c), indicating that there are different loss mechanisms affecting the electrons and the ions. This broadening of the ion temperature profile echoes the broadening of the electron density profile which was shown in figure 2.4. Between 5 and 10 ms the gradients are almost equal. This is due partially to the recovery time of the electron temperature gradient and what seems to be an overshoot in the ion temperature gradient, which is followed by a slight decrease in the gradient length before its pre-ELM saturation.

2.1.3 Evolution of pedestal in real space

While it is desirable to analyse the plasma in terms of normalised internal coordinates, much useful information can also be gleaned by comparing time traces at different radial locations. Shown in figure 2.7 are the evolutions of the ELM

averaged electron density and temperature at fixed radii over the ELM cycle. The vertical dashed lines indicate different characteristic timepoints throughout the ELM cycle. The recovery from the ELM crash has already been well documented by Burckhart et al.[48] and the corresponding initial recovery, T_e saturation, final recovery and saturated gradient regions are marked on this plot.

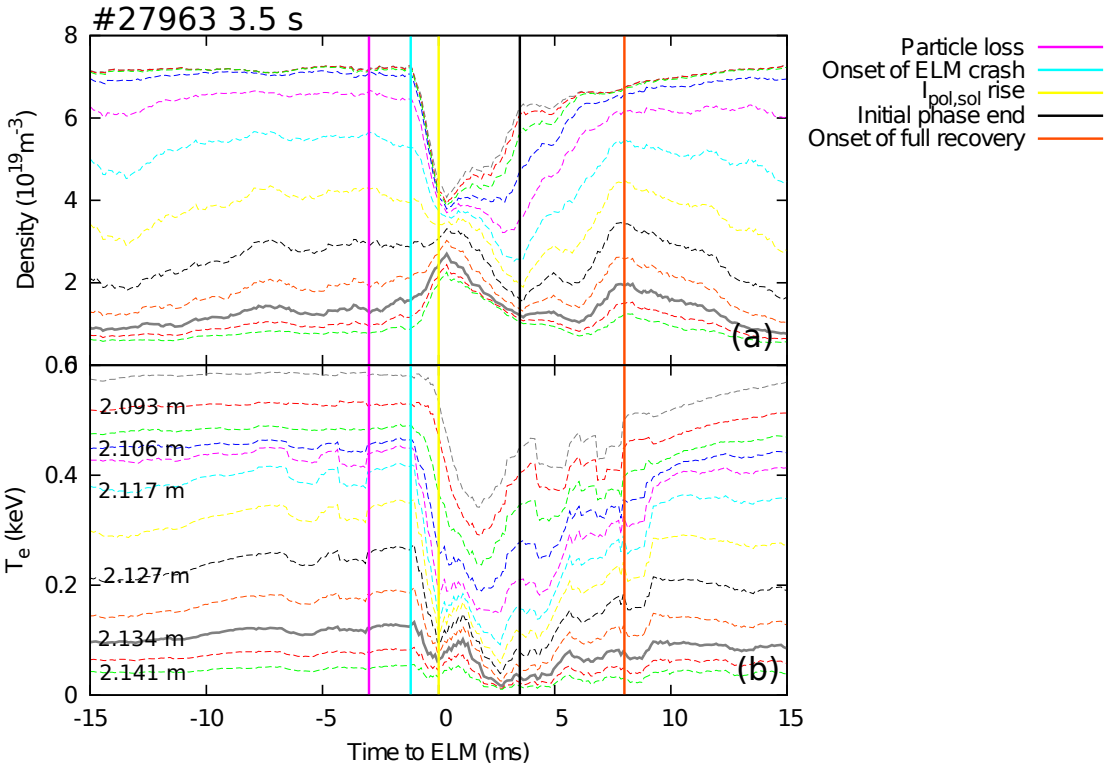


Figure 2.7: (a): horizontal dashed lines show the evolution of the electron density at fixed points in real space relative to the ELM crash. (b): the same information but for the electron temperature. The position of these lines in real space are also shown. Vertical solid lines indicate the timepoints corresponding to increased radial particle transport (purple), heat and particle confinement loss (cyan), rise of the divertor currents (yellow), onset of gradient recovery saturation (black), and onset of final recovery (orange).

It is interesting to note that while the edge gradients remain the same for long periods of time leading up to the ELM crash the same cannot be seen for the local temperatures and densities. The approximate separatrix location (assuming $T_{e_{sep}} = 100$ eV) is denoted by the horizontal solid grey line. There appears to be a general increase in the density both just inside and outside this point until a saturation is reached some 10 ms prior to the ELM crash. Similar is seen for the temperature gradient, although this continues for a few ms longer. A period of increased radial particle transport onsets at 3 ms before the ELM crash, where the inner density decreases and the outer density increases. Since we can expect

the same heat flux to the pedestal (assuming that there are no major changes inside the pedestal on this timescale, which is reasonable), this also explains the slightly increased electron temperatures near the pedestal top. In this case, the electron temperature pedestal appears to be extending further inside the plasma, but at a constant gradient. What follows is a sudden sharp loss of particle and heat confinement at 1.2 ms (cyan line) before the ELM crash and finally the ELM crash as defined by the rise of the divertor currents (yellow line). Note that this is not an inconsistent picture, since the divertor currents will not rise until there is a substantial difference in the inner and outer divertor temperatures. This requires a large heat flux to the outer divertor, which is delayed with respect to the ELM itself; this can in fact be seen in the sharp rise of the outer temperature points at 0 ms.

What follows next has already been reported by Burckhart et al.[48]; following the crash, there is an initial saturation phase, followed by a slight recovery up to, in this case, 3.5 ms (black line). After this, the gradients start to recover, first with both the density and temperature increasing, then with a saturation in the temperature until finally both enter the final recovery phase (orange line). After this point (approximately 15 ms after the ELM crash) both electron temperature and density slowly increase at a constant gradient (shown in figure 2.3) with another saturation phase beginning 7 ms before onset of the radial particle transport (again, purple line). The nature of each of these phases is currently undergoing intense investigation both from a peeling-ballooning and turbulence perspective with, respectively, the ILSA[56, 57] code and the GENE code[58]. The former is a linear MHD code, which will test the theory of peaked edge gradients and current densities while the latter is a turbulence code which can analyse the plasma from both global and local perspectives, as well as using linear and nonlinear growth and damping mechanisms.

2.2 ELM mitigation techniques

The high heat fluxes and reduced plasma confinement caused by ELMs must be mitigated if the divertor in a reactor scale device is to survive. There have been many attempts made to mitigate either the ELMs themselves, or decrease their size. The latter can be attempted by noting a logarithmic scaling for the relative ELM size against the inverse of the normalised ELM frequency ($= f_{ELM} \times \tau$, where τ is the confinement time), shown in figure 2.8. This figure was made using data from approximately 270 AUG discharges in the 2012 campaign and

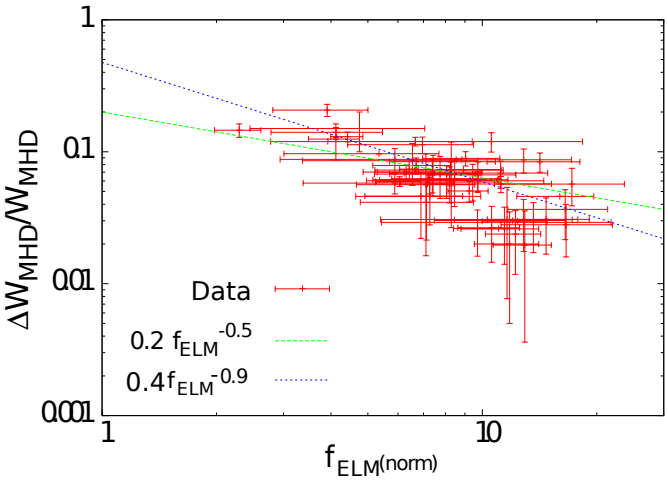


Figure 2.8: ELM size plotted against ELM frequency for AUG 2012 experimental campaign. Erorr bars indicate the standard deviation of the quantities during the time range analysed. The inverse trend, despite the standard deviations, is clear.

includes a variation of heating power and type, magnetic field, plasma current, and density. The uncertainties shown in this plot indicate the standard deviation of the data in the time window analysed. While there is a large scatter in the data points around the line shown (which is a best linear fit to the data), there is clearly a decrease of one order of magnitude of the ELM size with a corresponding increase in the normalised ELM frequency. The density has also been observed to be a significant factor in determining the ELM size (with a positive square root dependence), though the scatter remains large.

This particular scaling, and in particular the choice of normalisation of the ELM frequency, is derived from a model created by Fishpool et al.[59]. This model supposes that the energy lost due to an ELM crash, ΔW_{MHD}, would be restored in a time determined by the available reheating power. If all the ELMs in a given window occur at the same level of plasma stored energy, they would have a regular frequency, thus allowing an inter-ELM time, τ_{ELM}, to be determined as:

$$\tau_{ELM} = \frac{\Delta W_{MHD}}{P_{reheat}} \tag{2.1}$$

where P_{reheat} is the power available for reheating the plasma, i.e. the total input power less any power radiated from the plasma (this is of the order of 50%). In order to normalise the separate quantities on the right hand side of this equation, the plasma stored energy, W_{MHD}, and total heating power, P_{in} are used. Since the ratio P_{in}/W_{MHD} is the definition of the confinement time, this allows equation 2.1 to be rewritten as:

$$\tau_{ELM} = \frac{\Delta W_{MHD}}{W_{MHD}} \frac{P_{in}}{P_{reheat}} \tau \tag{2.2}$$

Since τ_{ELM} is simply the inverse of the ELM frequency, it is clear how the normalisation of the ELM frequency is defined. Additionally, this formulation reveals a key assumption in the above plot; the fraction of power available for reheating the plasma is assumed to be constant (or, more directly, the radiated power fraction is assumed to be constant), which gives rise to the factor of 0.4 in the inverse fit.

Several methods to increase the ELM frequency have been found experimentally. One is to inject fuelling pellets into the plasma edge. This causes a spontaneous triggering of an ELM crash via small modifications to the pressure profile[60]. In addition, it has been shown by a simple MHD model[61] that pellet fuelling cannot be used to decrease the ELM frequency, only increase it. There are also several operating scenarios which produce small frequent ELMs, such as Type-II ELMs, or so-called "nitrogen seeded" ELMs. Both of these ELM regimes will be analysed in chapter 7.

A further method has undergone intense investigation in recent years at DIII-D, ASDEX Upgrade, MAST, and JET - the use of non-axisymmetric magnetic perturbations. In AUG, two rows of coils (B-coils) above and below the midplane have been installed which can be used to generate an n=0,1,2,4 perturbation to break the toroidal symmetry of the system. These coils are shown as the smaller silver loops in figure 2.9. The red and purple loops are the A-coil conductors which are not yet installed. Once the plasma density reaches a critical threshold the large Type-I ELMs previously present vanish. Smaller ELM-like MHD events are still present in the plasma, however, although the exact nature of the change is unknown. This will also be investigated in more detail in chapter 7.

2.3 Pedestal relevant instabilities

Due to the high gradients in the plasma edge, one can expect several classes of instabilities to be observable. These can be split into two broad categories: MHD-driven instabilities, and turbulence-driven instabilities. The main instabilities relevant to ELMs are described in the following sections.

2.3.1 MHD modes

This section is principally based on a review of linear and non-linear MHD in relation to ELMs by Wilson et al.[62]. The most basic MHD mode that we can think of existing in the pedestal is the ballooning mode. This MHD mode is dependent on the field line curvature and the pressure gradient, and exists in ideal MHD. This mode features a short wavelength perpendicular to the field

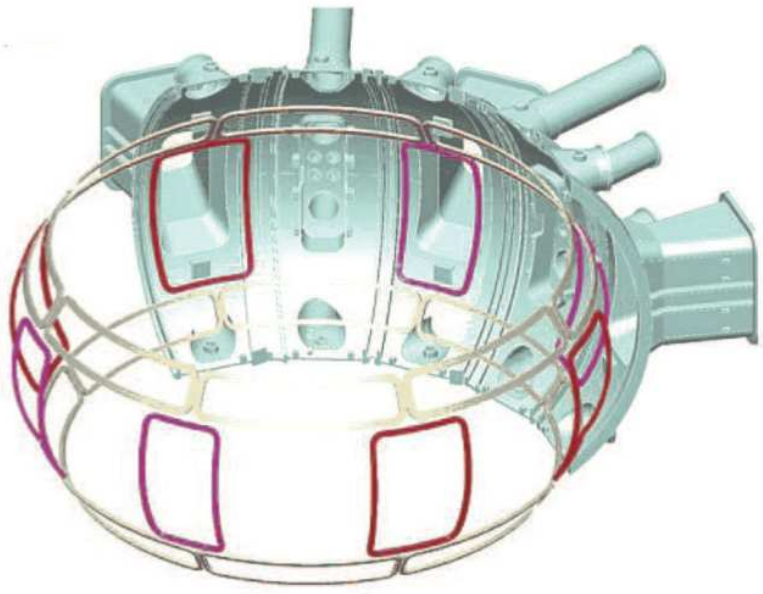


Figure 2.9: Layout of the B-coils at ASDEX Upgrade (silver) with respect to the vacuum vessel. Shown in red and purple are the planned A-coils, intended for use in resistive wall mode studies as well as further investigations on ELM suppression and mitigation.

line (high $k_{\perp}$) and a long wavelength parallel to the field line (low $k_{||}$) and is a highly localised mode; as such it is approximated as having a toroidal mode number of n=∞. Connor et al.[63] derived a minimised equation to describe the free energy of a system to such n=∞ modes which allows a simple ideal analysis of any plasma configuration. These modes are observed at the low field side of the plasma, where the magnetic field lines feature "bad curvature"; that is, the pressure gradient and the curvature vector point in the same direction (inwards). This acts to magnify an initial pressure-driven instability. In general, ideal ballooning modes exist at high toroidal mode numbers.

In addition to ballooning modes, kink or peeling modes can also exist in a plasma. These modes are essentially the same, with the kink mode depending on the gradient of the parallel current density and the peeling mode on the value of the current density. The peeling mode is essentially the limiting case of the kink mode in an infinitesimally small region in the vicinity of the separatrix. While the kink mode exists only at low toroidal mode numbers, the peeling mode can also exhibit higher mode numbers, which gives rise to the peeling-ballooning coupling[64]. This coupled mode is thought to be responsible for Type-I ELMs, since it can tap and release the free energy from both the pressure gradient and the current density[64, 65, 44].

Simple peeling-ballooning theory posits a linear limit cycle instability in the plasma edge. It is driven by the edge pressure gradient (ballooning component) and edge current density (peeling component). The simple description of an ELM cycle using this theory is sketched in figure 2.10. After the previous ELM crash

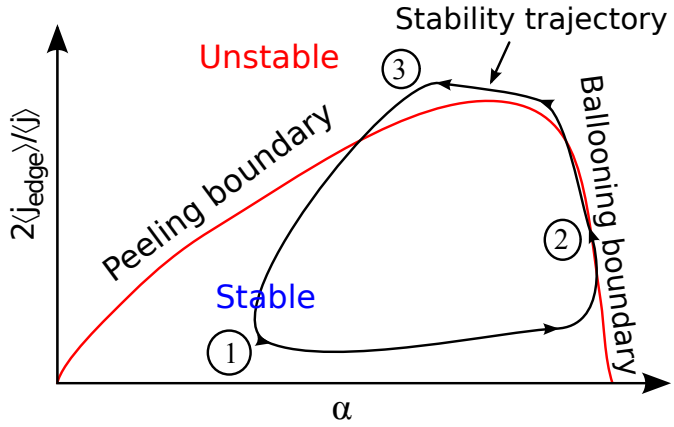


Figure 2.10: Simple cartoon showing the ideal MHD stability cycle of an ELM crash, starting at position one, reaching the ballooning boundary and finally triggering when it becomes unstable to the peeling mode.

(1), the plasma pressure gradient begins to increase. Since it is well observed that the pedestal pressure gradient does not change for a long time before the ELM crash, it is probable that the ballooning mode is not unstable during this time range, but rather the plasma sits at the ballooning stability boundary (2), which is therefore limiting the peak pressure gradient. The hypothesis for the delayed ELM crash is a delay of the current density building up due to resistive effects. This would therefore limit the drive for the peeling mode, which ultimately limits the plasma stability. Once the current density has increased to its critical value (3) an ELM crash occurs.

The x axis in figure 2.10 shows a value of increasing α. This is the normalised pressure gradient, which is taken to be a representative measure for the drive of the ballooning mode and is defined as[66]

$$\alpha = -2\mu_0 \frac{\partial V}{\partial \psi} \frac{1}{(2\pi)^2} \left(\frac{V}{2\pi^2 R_0} \right)^{1/2} \frac{\partial p}{\partial \psi} \tag{2.3}$$

where V is the plasma volume at a given ψ. The current density on the y axis is also a normalised value, although there is no solid consensus on the best normalisation to use; here, the definition used in the EPED model[47] is taken, given by

$$\mathrm{j}_{\mathrm{stability}} = \frac{\langle \mathrm{j} \rangle_{\mathrm{peak,edge}} + \langle \mathrm{j} \rangle_{\mathrm{sep}}}{\langle \mathrm{j} \rangle_{\mathrm{plasma}}} \tag{2.4}$$

This value indicates the average current density in the pedestal, though it is not clear if this particular quantity is more relevant than the peak edge current density prior to the ELM crash. The EPED model also goes one step further and posits that a separate instability, the kinetic ballooning mode (KBM), also acts as a constraint on the pedestal width/height relationship. This would also allow for a prolonged ELM cycle with a constant gradient and current density. In this case, the pedestal width increases at constant gradient, thus increasing the pedestal top values.

The peeling-ballooning theory is modelled via linear MHD stability analysis using codes such as CASTOR[67, 68], ELITE[69], and ILSA[56, 57]. This approach takes equilibrium plasma quantities and the plasma magnetic geometry and solves a set of equations for an exponential growth rate and unstable poloidal mode structure - where in the plasma the unstable toroidal modes are resonant. The linear analysis is limited in its application; it cannot resolve any non-linear effects such as growth rate saturation, the size of the ELM crash, or when exactly the ELM will occur. However, a linear analysis can still provide a useful picture of the basic physics involved and help, in part at least, to understand the ELM cycle evolution and the stability of different plasma configurations, such as the presence and location of the separatrix x-point[70, 71].

However, the hypotheses of this model must be tested, as well as the effect of experimental and numerical parameters on the outcome. The latter is the subject of a work being carried out in parallel with this[72] while we will focus in the coming chapters on the evolution of the pedestal, specifically the current density which has not yet been investigated in a systematic manner.

2.3.2 Kinetic ballooning modes

Snyder et al.[47, 73] have developed the EPED predictive model for critical pedestal parameters. Since both the pedestal height and pedestal width (or either of these and the pedestal gradient) must be resolved, two instabilities are analysed with this model. These instabilities are the MHD peeling-ballooning mode discussed above, and the turbulence driven kinetic ballooning mode (KBM). The latter has similar characteristics as the MHD ideal ballooning mode (long parallel wavelength, short perpendicular wavelength), but is driven by electromagnetic fluctuations which are caused by the interaction of the magnetic field with local density or temperature fluctuations[74] and has a similar high toroidal mode number. This instability was noted by Snyder et al. to appear at approximately the same critical α as the ideal ballooning mode, so it is thought to serve as a suitable

second parameter.

In addition, the KBM was also noted to appear independently of $E \times B$ shearing rates, which strongly suppress turbulence in the plasma edge. This particular feature makes the KBM a good candidate for a limiting pedestal instability. Since it is also dependent on the pedestal gradient, this gives, in effect, a second constraint for the two unknowns, being the pedestal top value and pedestal width. Saarelma et al.[45] tested the applicability of both the peeling-ballooning and KBM stability limits for MAST and JET plasmas and found that, while peeling-ballooning modes were found to be unstable, the KBM was stabilised at low edge collisionalities, as found in JET, and expected for Iter. It was found, however, that microtearing modes (MTM) were unstable at the pedestal top in all discharges. These modes have also been observed in gyrokinetic simulations of the outer core region of a Type-II ELMy scenario at AUG[75]. MTMs are gyrokinetic analogues of MHD tearing modes, which draw energy by relaxing the magnetic field geometry via island formation. MTM modes, unlike their MHD counterparts, draw energy from the local temperature gradients but also form a narrow current sheet about a resonant mode surface. More detailed analysis concerning the AUG plasma edge and the existence of MTM/ITG/ETG and KBMs is currently ongoing.

2.3.3 The ELM cycle

The question of how the pedestal recovers after an ELM crash and leading up to the next ELM still remains. A brief description of a purely MHD dominated ELM cycle was given above, though this is only used for basic illustrative purposes; typical resistive timescales in the pedestal are short, of the order of 1 ms, and certainly not long enough to allow the critical pedestal gradient to remain constant for 15 ms, as observed in discharge #27963. Instead, an interplay between a transport based gradient and the MHD based instabilities has been proposed by Snyder et al.[73] while Schneider et al.[76] have also proposed that the critical pedestal parameters are set by such a combination.

The EPED model, while a static model, has been used to chart the development of an ELM cycle by the onset of the gradient limiting KBM early in the ELM cycle[73]. This is then followed by the further increase of the pedestal width and height, lowering the peeling-ballooning boundary until a crossing point of the two instabilities is reached. This is illustrated in figures 2.11(a) and (b). Figure 2.11(a) shows a cartoon stability plot in j-α space, but this time indicating the movement of the peeling ballooning boundary as the pedestal widens. Figure

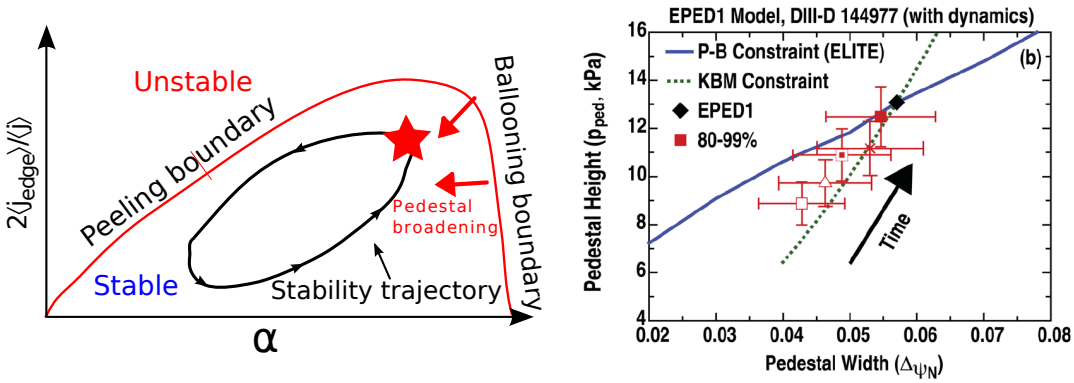


Figure 2.11: (a): Modified cartoon depicting the effect of a wider pedestal on the peeling ballooning stability boundary. The star indicates the operational point, which remains at the same values of pressure gradient and current density for a prolonged time. (b): EPED stability diagram in terms of the pedestal top as a function of pedestal width. The diagonal green line indicates the kinetic ballooning mode stability threshold, while the more horizontal blue line indicates the peeling-balloning threshold. Overlayed are data points from a DIII-D discharge showing that the model can predict the evolution of the pedestal over the full ELM cycle and also the final trigger of the crash. Reprinted from [73].

2.11(b) (reprinted from [73]) shows the evolution of the pedestal in terms of the pedestal width and pedestal top, the values which EPED predicts, over an ELM cycle. The diagnonal green line indicates a line of constant gradient, marking the stability limit for the kinetic ballooning mode, while the almost horizontal blue line indicates the peeling-ballooning stability limit. This theory of the ELM cycle predicts that at the point of intersection of the two lines an ELM will occur. Overplotted on this stability diagram are experimental data points from the DIII-D discharge which the diagram corresponds to. Throughout the ELM cycle, a constant increase of the pedestal top and width are evident, with the critical pedestal parameters matching with the theoretical prediction.

Since the model is static and linear, it makes no attempt to predict an ELM frequency for a given plasma shape or heating power, rather it focusses on a basic development of the pedestal. In addition, it also does not describe the separate growth of temperature and density gradients described by Burckhart et al.[48] and in this chapter. The separate growth of ion and electron temperatures could be due to different heat fluxes to the ions and electrons. The saturation of the electron temperature growth could be reconciled by a recovery of the density with a constant heat flux. Whether this difference then corresponds to longer ELM cycle lengths is difficult to tell with accuracy, though different fuelling rates will

be analysed in chapter 3.

2.4 Experimental evidence for MHD ELM description

If ELMs are MHD events, one would expect to see some magnetic perturbation signals leading up to and during the ELM event. Gas puff imaging of field aligned structures at the ELM crash have been reported at MAST[77], while divertor power deposition studies at AUG[78] have also shown evidence of such structures. In both cases, mode numbers were found to be approximately in the range 5-20. These filamentary structures are thought to be the long wavelength ballooning like filaments which, according to Wilson et al.[62], would slide out into the SOL between the field lines on the low field side of the plasma while remaining connected to the pedestal top at the high field side, thus transferring heat outside the confined region. If this is extended further and a reconnection event occurs in the SOL, energy could be removed on a faster timescale.

Toroidally localised rotating magnetic perturbations have also been reported in some discharges at ASDEX Upgrade[79]. This research reached the conclusion that, in the 100 μs preceding the ELM crash (as determined from the SOL current measurements), a strong n = 1 magnetic perturbation arises and that this is the dominant cause of energy release from the ELM. This has been supported by initial results from nonlinear-MHD calculations using the JOREK code[80]. However, from figure 2.7 also shows that much activity also takes place in the milliseconds prior to the ELM crash. It is this phase which is predicted to be dominated by linear instabilities, such as the peeling-ballooning mode, which then gives rise to the large release of energy.

Linear stability analysis of peeling-ballooning modes has been the main focus of comparison between theory and experiment when studying ELM crashes. This comparison takes the form of a stability diagram, which describes the stable and unstable regions in normalised pressure gradient/current density space, along with the mode numbers of the most unstable toroidal mode number. These diagrams are made by taking a reference equilibrium with an experimental pressure profile and then varying the pressure profile and corresponding current density profile while keeping the plasma current, toroidal magnetic field, and plasma boundary constant. Typically, the width of the edge pressure pedestal is scaled while keeping the pedestal top constant; this changes the pressure gradient in the pedestal while keeping its location the same. The edge current density profile is

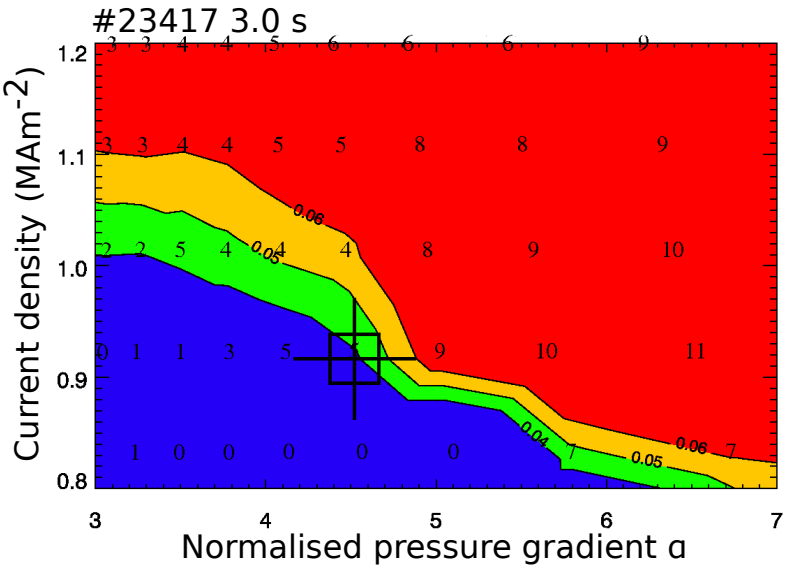


Figure 2.12: Stability diagram for AUG discharge #23417. The red area shows the region unstable to peeling-ballooning modes, the blue the stable region. The green and yellow areas correspond to marginally unstable points, while the black cross indicates the operational point. The numbers show the mode number of the most unstable toroidal mode.

simply scaled while keeping the total current the same. These modifications have an impact on the core profiles as well, though the change here is much smaller since it is spread out over a larger area.

The HELENA equilibrium solver is used to determine the new equilibria based on the modified profiles and a fixed plasma boundary. A 2D grid in pressure gradient-current density space is thus created and each point in this grid is then analysed separately by, for example, the ILSA stability code. Shown in figure 2.12 below is a stability diagram calculated using actual AUG data and magnetic equilibria. The black cross in the figure corresponds to the experimental profiles. The yellow boundary denotes the "stability boundary" beyond which unstable peeling-ballooning modes can be expected to appear in the plasma. The numbers printed correspond to the most unstable toroidal mode numbers. It is important to note that this analysis does not, and cannot, take nonlinear saturation of these modes into account. The rise of the strong n=1 perturbations is a purely nonlinear process. Nevertheless, the linear analysis is important if we are to understand the basic clamping mechanisms in the plasma, which appear approximately half way through the ELM cycle and limit the total performance of the discharge. The triggering of the ELM and the associated losses are a separate matter and we can only speculate on the nature of the relationship between linear theory and the eventual ELM size, if any.

Shown in figure 2.13 is a spectrogram of the Mirnov coil data taken at one toroidal location for the duration of the time range of interest in discharge #27963. Shown below the spectrogram is the inner divertor current signal, used to indicate the ELM onset time. While large fluctuations (indicated by a red colour) can be seen at low frequencies (10-20 kHz), a second set of modes also onsets approximately half way through the inter-ELM phase at frequencies of 100 and 200 kHz (shown in yellow); this is correlated with the pedestal reaching its critical gradient. This 200 kHz mode is possibly due to the onset of a KBM in the pedestal as the critical pressure gradient is reached[81]; its exact nature is currently under investigation.

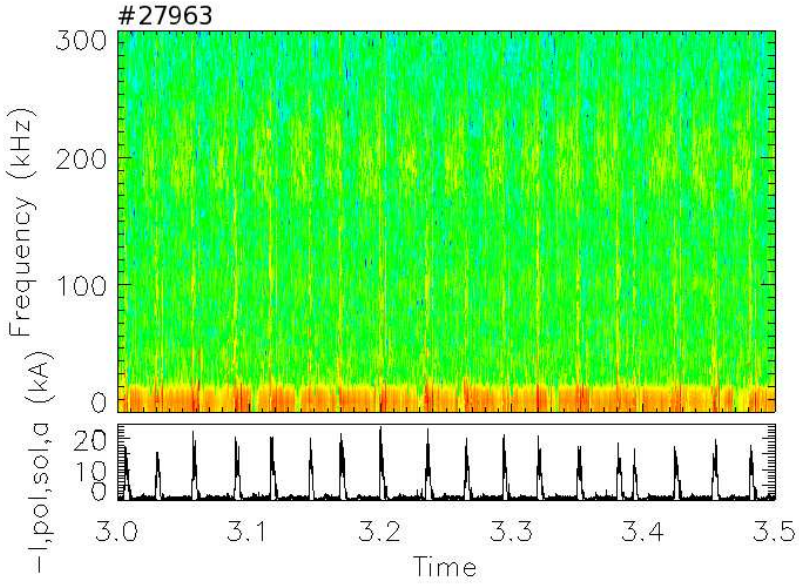


Figure 2.13: Spectrogram of Mirnov coil data showing magnetic fluctuations relative to the ELM crash (indicated via divertor current signal, shown below). There is a strong signal at approximately 200 kHz; this is believed to be indicative of a KBM in the pedestal[81]

This KBM could be the mode responsible for the initial restriction of the pedestal gradients, in agreement with the assumptions of the EPED model. In the case of this discharge with extremely long ELM cycles (for ASDEX Upgrade; larger devices have ELM frequencies of the order of Hertz), it is possible that we see the separation of the initial transport limitation and subsequent MHD perturbation which triggers the ELM. It is possible that this KBM does not appear in every discharge, although the peeling-ballooning mode is thought to be a limiting instability in all cases. The appearance of transport limitations could be due to a heat flux limit; when too large a heat source is present, especially near the pedestal, large bursty ELMs could be expected. This is in agreement with

the smaller more frequent ELMs seen when N_2 seeding is applied at JET[82] and AUG[83]; the change in the radiated power level could somehow suppress some modes while allowing others to surface. In particular, if the hypothesised initial transport limitation is suppressed then the MHD modes could grow unstable at a faster rate, leading to the shorter ELM cycles. Type-II ELMs cannot be reconciled with this picture, though it has also been speculated that the higher density and strong shaping (the conditions to enter this regime) suppress the peeling mode while denying the plasma access to the "second stability" region of shear-α space, which permits higher pressure gradients. As a consequence, an ELM is hypothesised to be triggered only by a "soft" ballooning limit, whereby only the pressure gradient is relaxed periodically[62, 84, 82].

Chapter 3

Current density analysis

Current drives in a tokamak plasma come from several sources. The first, and most obvious from the introduction, is the Ohmic current which is applied via the central transformer. In addition to this is any current due to the external heating systems; these can apply a torque to the plasma, causing a current. However, there are also two other sources of current. The first is driven purely by pressure gradients and is called the Pfirsch-Schlüter current. This current is derived from the diamagnetic drift in the plasma and is due to the return flows cancelling out charge accumulation at the top and bottom of the confined region. However, assuming axisymmetry, this current is negative at the HFS and positive at the LFS and carries a zero net current. Bird and Hegna[85] have investigated the effects of 3D asymmetries, such as those induced by the external magnetic perturbations discussed in the previous chapter, showing that the Pfirsch-Schlüter current can adopt a helical structure. This helical structure then alters the local shear profile significantly, even for small external perturbations, allowing ideal MHD ballooning modes to be triggered at a lower pressure gradient than in the axisymmetric case. This will be discussed further in chapter 7. The final current drive is called the bootstrap current, so called as it is self generated by density and temperature gradients in the plasma. It is a neoclassical effect (meaning that it arises due to the toroidal nature of the system and collisions) and has its basis in particles trapped at the LFS of the torus due to the toroidal magnetic field gradient. A rough sketch of this current will be given in chapter 5.

The bootstrap current is also a particularly useful current as it can provide a substantial fraction of the current required to create and sustain the poloidal field for confinement, lowering the input energy requirements of a future fusion reactor. Several theoretical steady state reactors have been planned which utilise a “bootstrap fraction” in excess of 80% which would all but eliminate the re-

liance on the central transformer. This would facilitate steady state (or at least quasi steady state) plasma operation with the additional current supplied by electron-cyclotron current drive (ECCD) or neutral beam current drive (NBCD), for example.

Since the Pfirsh-Schlüter and bootstrap currents are pressure gradient dependent, it should be apparent that they make a large contribution to the current density in the edge of the plasma when the pedestal is present. Shown in figure 3.1 are the pressure profiles and corresponding LFS current density profiles for the L- and H-mode phases of discharge #17741. The current density profiles have been calculated using the CLISTE equilibrium solver, using a combination of external magnetic data, pedestal pressure measurements, and SOL currents as constraints. The process involved in this fitting will be described in more detail throughout this chapter. What is notable is that at the plasma edge there is a

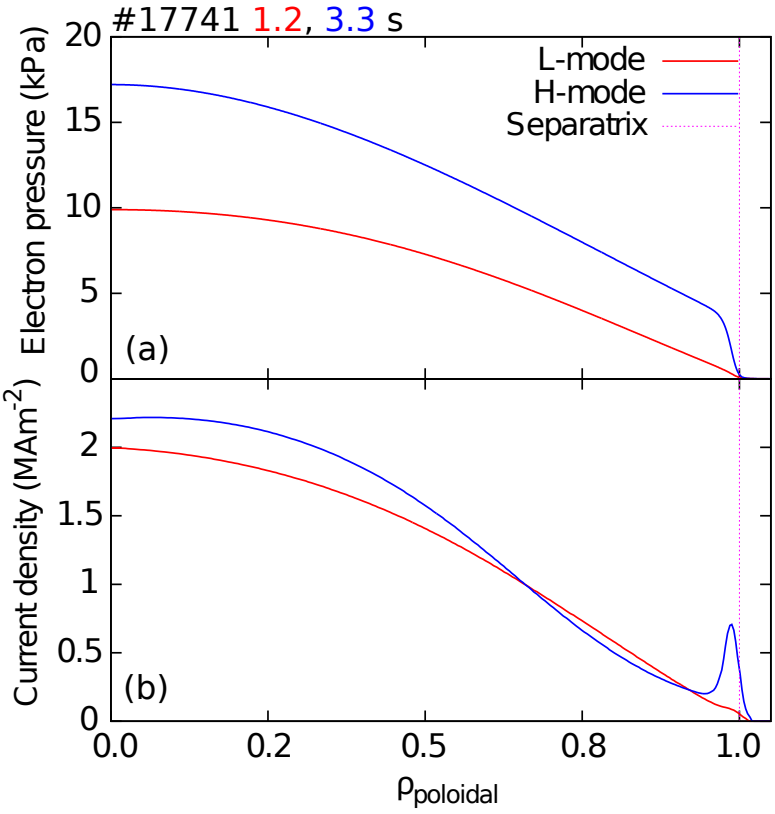


Figure 3.1: Comparison of electron pressure and LFS current density between low-confinement (L-) and high-confinement (H-) mode portions of AUG discharge #17741. The increased pressure gradient, seen in (a) creates the corresponding edge current density peak in (b).

large peak in the current density profile. While part of this peak is a symptom of the beneficial bootstrap current, it does not come without a price. The high current which is created in the edge of the plasma supplies a drive to the peeling

mode, which is one part of the ELM relevant peeling-ballooning mode. However, this current density also acts as a stabilising factor, as it changes the magnetic shear profile at the plasma edge, reducing the drive for the ballooning modes. Above a critical threshold (or rather, below a critical shear), access to a "second stability" regime of high pressure gradient operation is possible.

The current density in a tokamak cannot be directly measured as the temperatures reached in a fusion plasma are much too great to allow probes access to the main plasma. Instead, common methods to determine the internal magnetic structure (and hence the current density) measure the pitch angle of the magnetic field lines, yielding a constraint which can be used in an equilibrium solver. This technique is most often realised in the form of a motional Stark effect (MSE) diagnostic. Here, neutral particles experience an electric field in their rest frame as they move relative to a magnetic field which causes a splitting of, in the case of Hydrogen or deuterium plasmas, the Balmer-α line. Neutral beam injection is a common heating method on most tokamaks and the beam of injected particles can be used as the diagnostic beam. By measuring the ratio of the linear and circular polarisations of the emitted light, the pitch angle of the magnetic field lines can be determined. This is used to great effect in the plasma core, where the poloidal magnetic field changes strongly with radius. However, at the plasma edge, the field line angle changes only slightly with a large current peak, making such a determination difficult. Figure 3.2 shows the change of the field line angle as predicted by CLISTE at several radial locations over an ELM crash. The difficulty of the measurement at the plasma edge is illustrated by this figure. An ELM crash occurs at t=0 ms, which, as discussed in chapter 2, causes a large drop in the edge pressure gradient and hence current density. The pitch angle deflects by approximately 0.2 degrees, which requires a very sensitive measurement of the relative polarisations if one is to resolve changes in the pedestal current density relative to an ELM crash. More importantly for the determination of the current density, the field line angles of neighbouring radial locations do not show a large relative difference. This technique has, however, been used successfully at MAST[53] and first edge current density measurements have been obtained. A second option is to measure the Zeeman splitting of lithium 2s-2p transition lines, which also gives the magnetic field line angle. This has been done in the past at DIII-D[52], and is currently under re-investigation there.

In order to determine the edge current density at ASDEX Upgrade, the free boundary equilibrium solver CLISTE[19, 22] is used. In essence, this performs an iterative least squares fit to a set of internal and external measurements, with

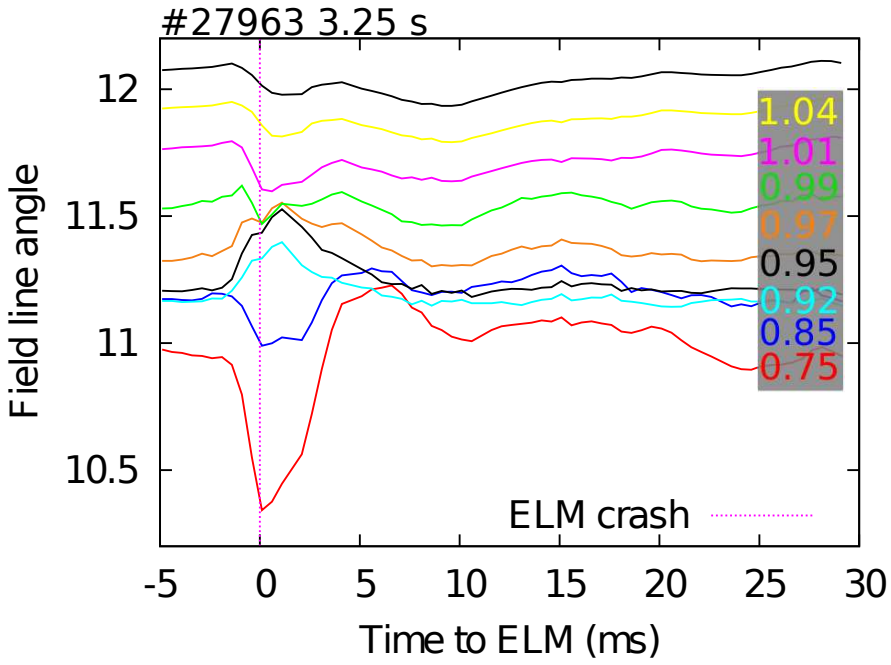


Figure 3.2: Timetrace of magnetic field pitch angle relative to an ELM crash. Different colours correspond to different radial locations (along the magnetic mid-plane) in the plasma (ρ_{poloidal} values are marked on the right hand side). The small change in pitch angle over the time trace indicates the difficulty of such measurements.

an extra condition of force balance on the plasma. The code fits the data by choosing free parameters for the p' and ff' magnetic flux source functions such that agreement with set of experimental data is reached. These source functions are defined by a set of basis functions; in CLISTE, this is realised in the form of spline functions. Via the Grad-Shafranov equation, derived in chapter 1 and given by

$$-\Delta^*\psi = \mu_0 R^2 \frac{dp}{d\psi} + f\frac{df}{d\psi} = \mu_0 R j_\phi \tag{3.1}$$

the current density is deduced.

3.1 Diagnostics

Measurements of the plasma are divided into two main subsections: internal, and external. The measurement geometry of the diagnostics used in this thesis are shown in figure 3.3. This figure is split into several key areas by colour. The outer grey area shows the region between the vacuum vessel and the plasma facing components. The blue section is the SOL, where currents are considered to flow, in contrast with many equilibrium codes. By allowing currents to flow in the SOL, the current density is allowed to decay realistically. Additionally,

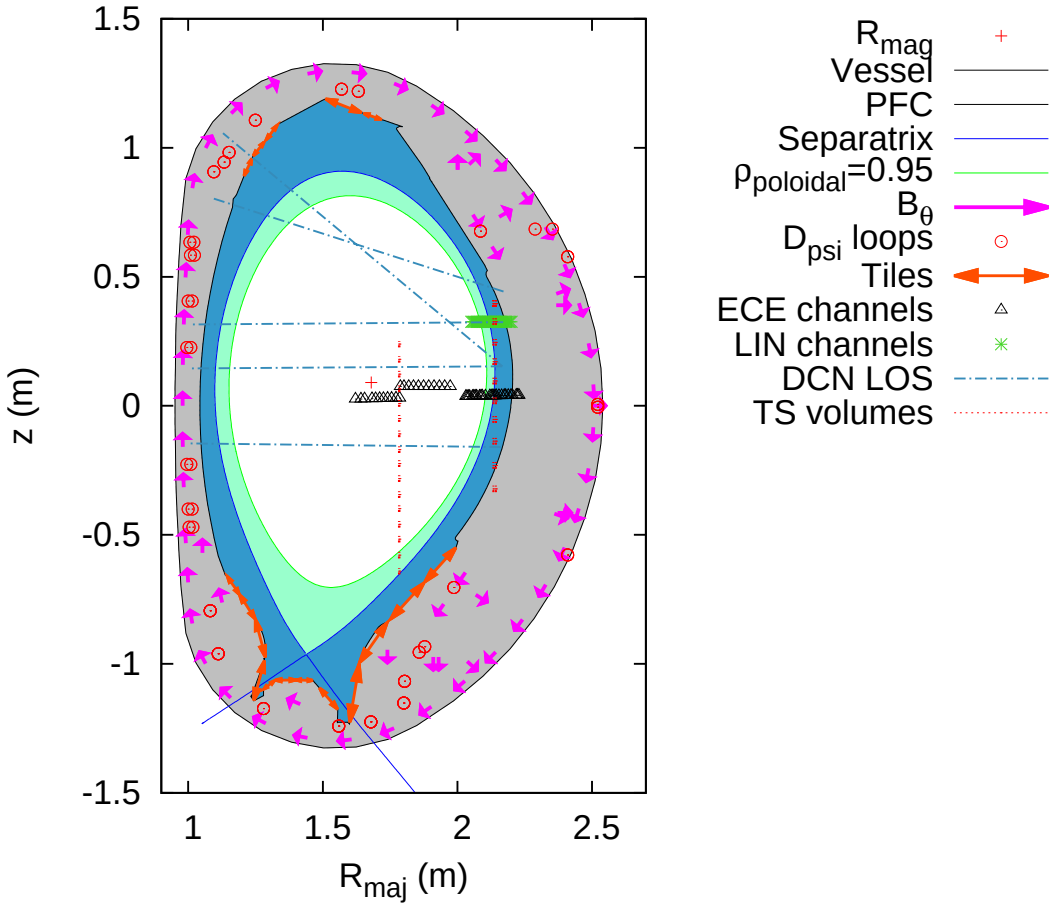


Figure 3.3: Several key diagnostics used in this thesis relative to the plasma itself, the plasma facing components and the vacuum vessel.

the large SOL currents which occur at an ELM crash are modelled correctly, allowing CLISTE to more accurately represent the current distribution between the confined plasma and the SOL at this time. The SOL region has a large suite of diagnostics to measure flows, Langmuir probes to measure the electron temperature, density and the ion saturation current at the target plates with a high spatial-temporal resolution as well as other probes which are used to determine changes in magnetic fields. The total current flowing in this region is also measured by shunts in the tiles, which have been used in this thesis. The orange arrows show the model of the divertor tiles used in CLISTE for comparison with these divertor current measurements. Inside the separatrix, the pedestal region, corresponding to $\rho_{poloidal} = 0.95$ is shown in green. Finally, the magnetic axis is shown as a red cross. The location of the separatrix and pedestal top were taken from AUG discharge #27963, the shape of which was optimised for edge kinetic measurements, which will be described in the next section.

3.2 Pressure measurements

The CLISTE equilibrium solution can be constrained by edge pressure data. There are two main methods by which this data was included in this work: by fitting a modified tanh (mtanh) function to a set of complementary diagnostics, or though the use of an integrated data analysis (IDA) routine[86]. IDA is a Bayesian approach to data analysis used to generate temperature and density profiles by combining complementary diagnostics. The pressure is made up of the temperature and density of both electron and ion species. In order to produce these fits, an initial equilibrium is required. It is for this purpose that a generic equilibrium reconstruction is made with a 1 ms time resolution on an inter-shot timescale. This generic reconstruction uses only external magnetic data and is run automatically, leaving the possibility for errors. However, the magnetic topology is well known at the plasma edge, meaning that the uncertainty due to this initial equilibrium reconstruction is small.

3.2.1 ECE

Shown as black triangles in figure 3.3 are the ECE measurement locations[87]. The precise location of these channels depends on the form of the magnetic field, but this is very well known, even for an equilibrium using only magnetic data. The ECE diagnostic measures the electron temperature via the intensity of emitted cyclotron radiaton. The frequency of the radiation is given by the well known $\omega = \mathrm{m}\frac{\mathrm{eB}}{\mathrm{m_{e,0}}}$, where m is the m*th* harmonic of the resonance, e is the charge on the electron, $\mathrm{m_{e,0}}$ is the electron rest mass, and B is the magnetic field. The 1/R dependence of the magnetic field allows a spatial localisation of the radiation. At AUG, the intensity of the radiation is measured by a heterodyne radiometer at a collection frequency of 31 kHz with a spatial resolution of 1 cm. In general, the relationship between the measured intensity and temperature is complicated. However, if the plasma is optically thick, the electron temperature can be approximated by the Rayleigh-Jeans approximation as

$$\mathrm{T_e} = \frac{8\pi^3 \mathrm{c}^2}{\mathrm{k_B}\omega^2}\mathrm{I} \tag{3.2}$$

where I is the intensity. This approach, however, is not valid at the edge of a H-mode plasma as the plasma is not optically thick [48, 54]. Rathgeber et al.[54] have developed an advanced analysis method which solves the full 1D radiation transport equations in the framework of Integrated Data Analysis (IDA). This solution has become available only recently and as such was not employed for

much of this work. For profile analysis, the ECE data was augmented by Thomson Scattering data, while temporal analysis was made by removing the edge electron temperature data and fitting a steep exponential profile for the final 1-2% of the plasma radius.

3.2.2 Thomson scattering

The Thomson scattering diagnostic measures both the electron temperature and density simultaneously. The diagnostic works by firing a high power laser pulse vertically through the plasma. The light is then scattered (predominantly) from the electrons and analysed by four-channel polychromaters outside the vacuum vessel. The light intensity gives the electron density while the Doppler broadening of the received light can be used to determine the electron temperature. The Thomson scattering diagnostic is a very valuable diagnostic in that it provides an intrinsic alignment between the electron temperature and density profiles.

The vertical red lines in figure 3.3 are the Thomson scattering volumes, with six separate lasers at the plasma edge and three in the core. These lasers are displaced by 3 mm radially and, in normal operation, fire sequentially. The laser pulses are 15 ns, meaning that a snapshot of the plasma at a given time is measured. Each laser has a repetition rate of 20 Hz, giving a maximum sampling rate of the diagnostic of 120 Hz. They can also be used in "profile mode", where all lasers are fired at 100 μs intervals. While this increases the effective radial resolution, temporal resolution is sacrificed. Although the core system has channels that reach the pedestal top and even inside, as in this case, there are issues with mapping the temperatures and densities from the bottom of the plasma to the midplane. That said, there is good agreement between the two TS systems and the ECE and LiBES diagnostics. The experimental setup at ASDEX Upgrade utilises a Nd:YAG laser system, operating at 1065 nm and a 1 J beam energy. Due to its relatively low sampling rate, this diagnostic is not typically used for time dependent analysis of the plasma. However, its high temporal localisation and radial resolution mean it is used extensively when analysing pre-ELM profiles.

3.2.3 Lithium beam emission spectroscopy

The green band in figure 3.3 show the extent of the sampling points along an injected Lithium beam[88, 89]. The emission intensity at these points is then used to calculate an electron density profile. The reliability of the electron den-

sity profiles drops further inside the plasma where the Lithium beam is strongly attentuated. The emission profile has a peak somewhere in the pedestal region; changes in the electron density profile inside this emission peak have little effect on the reconstructed electron density profile. These data are therefore combined with integrated measurements from interferometry in the IDA framework. The LiBES beam diagnostic at AUG has a sampling rate of 20 kHz and a spatial resolution of 5 mm.

3.2.4 Interferometry

Information on line integrated densities can be obtained by placing the plasma in one arm of an interferometer; the total density along this line is then proportional to the phase shift of the light. At AUG, 5 lines of sight from a Deuterium-Cyanide-Nitrogen laser pulse are used to determine densities along different chords in the plasma[90]. These lines of sight are shown as cyan lines in figure 3.3. The deconvolution of these data to provide densities in the core of the plasma is highly dependent on the equilibrium and appears to be especially sensitive to the vertical location of the magnetic axis. Although this work is concerned with the plasma edge, the integrated values provide a valuable boundary condition on the pedestal top density. In standard operation, the interferometers sample at 10 kHz, though this can be increased.

3.2.5 Charge-exchange recombination spectroscopy

While measurements of the electron properties are plentiful, the only reliable way to recover information on the thermal ions is via active charge-exchange recombination spectroscopy (CXRS). This requires the use of a diagnostic neutral beam, for which one of the heating beams at AUG is used. It can also be accomplished by puffing neutral gas into the plasma, though this technique requires more careful analysis. The injected neutral beam particles interact with impurities in the plasma, donating electrons to create hydrogenic excited impurity atoms which then decay with characteristic line emission. The Doppler broadening of this line gives information on the temperature of the impurity species, while the intensity of the emission is related to the impurity density. In addition, the rotation velocity of the impurity along the line of sight is given by the Doppler shift of the impurity emission. This information is important for determination of the radial electric field (for turbulence analysis) and also when inertia is taken into account in MHD.

AUG has recently been equipped with two edge CXRS systems with both poloidal and toroidal viewing angles, both of which feature 8 radial channels and a 2.1 ms integration time[91]. This long (for ELM resolved analysis) integration time presents challenges for interpretation of time series of data (meaning that only minimum recovery rates after an ELM crash can be obtained), but, similar to the Thomson scattering data, can be used for profile analysis. Most of the discharges presented in this thesis were made before these systems were installed, so the assumption of $T_i = T_e$ was used throughout most of the work.

3.2.6 Ion density

Information on the ion density is difficult to include with a high degree of accuracy. While the density of a single impurity species can be calculated from the CXRS data, this still does not take account of the other impurity species present. As such, several methods are used. The most direct is to make an assumption on the presence of other impurities based on the measured one (normally Boron) and then work from the measured and well known electron density to get an ion density. The other method uses the effective charge reconstruction to determine an "average" charge and again relate this to the measured electron density. This latter method has been used exclusively in this work. The effective charge (Z_{eff}) profile is estimated via forward modelling of Bremsstrahlung measurements[92], which, when combined with an assumed dominant impurity (Boron or Carbon are the dominant AUG relevant impurities) yields an ion density ratio.

3.3 Magnetic measurements

Outside the plasma facing components are the various magnetic diagnostics used as the minimal input for CLISTE. The purple arrows in figure 3.3 show the location and orientation of the B field probes which measure principally the tangential component of the poloidal magnetic field (at 38 poloidal locations), but also the radial component at several locations. The red circles plot the locations of the flux loops which are combined in pairs to make 18 flux difference measurements (the paired connections are not shown). A description of these diagnostics is given in Wesson[6]. Both of these sets of measurements can be used to diagnose several key plasma parameters[93]. In particular, integrated moments of the current density (total current, current centre) and information on the stored energy, beta poloidal β_{pol}, and internal inductance l_i can be well recovered from an elongated plasma. In addition, they form the minimal diagnostic inputs (also referred to as

"magnetics" for CLISTE input). The magnetic data for most of the discharges in this thesis had a time resolution of 1 ms, while any discharges analysed from 2012 onwards had a time resolution of 0.1 ms. A 1 s timeslice in all discharges prior to this also had a time resolution of 0.1 ms, though this was not exploited for the present work.

3.3.1 Current measurements

The current in the scrape off layer is measured via shunt resistance measurements of currents in the tungsten divertor. These provide a valuable constraint of the integral of the ff' profile in the SOL, allowing a better determination of the current density in the plasma edge. Each tile current gives the difference in the flux function $F(\psi) = \mu_0(I_{pol}(\psi)/2\pi)$ across the tile. The orange arrows in figure 3.3 show the tile model used in CLISTE for the comparison between prediction and measurement, though normally only the integrated currents from the inner and outer lower divertors are used. In cases where the upper x-point comes close to the main plasma (either in cases of upper single null, or double null where both x-points are inside the plasma facing components) the upper divertor currents can also be taken into account; this becomes important when analysing Type-II ELMs. The lower currents have a 0.01 ms time resolution. Only one toroidal location is used to determine the integrated current measurement, though some measurements are also routinely collected at other toroidal locations.

It should be pointed out that all of these measurements, both internal and external are taken at different toroidal locations around the plasma. This opens the possibility of systematic shifts between the equilibrium and the diagnostics since the assumption of zero toroidal gradient is not necessarily valid. At the very least, there exists an n=16 perturbation of the plasma corresponding to the 16 toroidal field coils. This so-called ripple effect, while small, could induce extra currents which modify the plasma edge, changing the mapping of the diagnostics slightly. This issue is amplified when using the magnetic perturbation coils (B-coils) recently installed at ASDEX Upgrade. These coils are located close to the plasma and can be run in configurations producing n=1,2,4 toroidal perturbations with even or odd up/down parity. Deflections of several millimeters have been reported, and toroidal asymmetries in the magnetic measurements used for equilibrium reconstruction have also been noted [94].

3.4 Current density determination

Measurement of the edge current density relies on the existence of the x-point. From Ampere's law, it can be shown for a straight cylinder with a current distribution depending only on the radius (that is, the equilibrium current distribution consists of concentric circles) that the only recoverable moment of the current distribution is the total current. Once the cylinder is bent into a tokamak, the concentric circles are shifted outwards by the Hoop force, otherwise known as the Shafranov shift. This shift is recoverable by the magnetics, yielding the coupled parameter $(\beta_{pol} + l_i/2)$[95], where β_{pol} is the ratio of pressure to the poloidal magnetic field and l_i is the internal inductance of the plasma. This coupled parameter may be separated into its component parts in an elongated plasma[93]. However, additionally, due to flux expansion at the x-point of an elongated plasma, seen in figure 3.3, it is possible to detect the current at the x-point as a separate wire current[22], whose value approximates that of the current flowing in the outermost 1% radial annulus. It has been shown that it is possible to detect this current using only the magnetic measurements, thus giving the current enclosed in the final 1% radial annulus.

While the power of this measurement lies in its interpretation as an integral quantity, it can be greatly enhanced through the use of the internal pressure profile constraints. In addition, the measured poloidal SOL current can be used to constrain the integral of ff' in the scrape off layer. While this is not enough to fully specify the current density profile in this location, it is a substantial help, especially since the current in the SOL is largely force-free. As the integral edge current is known, and a boundary condition at the separatrix is known, the internal current can be very well defined. Due to the curvature regularisation of the spline model used, this allows a consistent SOL and edge plasma current density profile to be deduced.

For discharges 2322x which are presented in this chapter, ion temperatures have been taken to be equal to the electron temperature, as shown for these cases by Wolfrum et al.[49], and the ion density to be 0.85 that of the electron density, as derived from a typical value for $Z_{eff} \approx 1.8$ and the main impurities being Boron and Carbon.

One of the assumptions made in this analysis is that the contribution of fast ions to equilibrium force balance in the plasma edge is negligible. The fast ion pressure profile was calculated with the FAFNER code[96] for discharge #23221 which verified that the contribution of the fast ions is of the order of a few percent at the pedestal top and effectively vanishes in the pedestal. This can

be seen in figure 3.4. In the core of the plasma, the fast ion pressure become

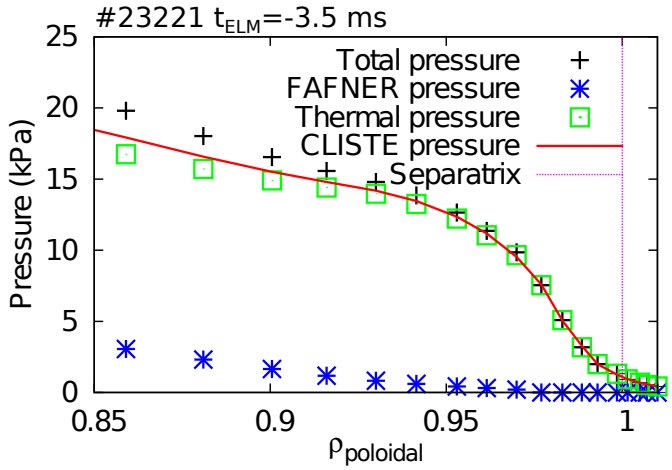


Figure 3.4: Radial profile of total pressure from CLISTE (red) with contributions from fast ions (blue), thermal pressure (green), and the total input pressure (black) for discharge #23221. In this case, fast ions are effectively negligible in the pedestal region and just inside the pedestal top.

more prominent; in this case $\approx 30\%$ of the total core pressure is due to fast ions. While it is possible to include the fast ion pressure profile in CLISTE and thus constrain the entire equilibrium pressure profile, this was not done in the present analysis since the focus was principally on the edge region. Accordingly, thermal pressure measurements, although available over the full plasma radius, were used to constrain the equilibrium pressure in the edge region only. It should be noted that in discharges with NBI power input off-axis a larger contribution of fast ions into the pedestal region via diffusion is possible.

3.5 Example current density profiles

Results presented in this section are from ASDEX Upgrade discharge #23221, which is a Type-I ELMy H-mode discharge, having 7.5 MW of Neutral Beam Injection heating, 1.3 MW of Electron Cyclotron Resonance Heating, a 1 MA plasma current, -2.5 T toroidal magnetic field and 5.8×10^{21} s^{-1} Deuterium gas injection. Data were analysed between 3.6 and 4.4 seconds of the discharge. In this case, ELM synchronised temperature and density profiles were taken from IDA analysis in a 0.1 ms timeslice, 3 ms before the ELM crash.

Figure 3.5 shows an example of the difference the pressure constraint makes to the edge current density as calculated in CLISTE. The two profiles presented were created with the same curvature constraints, the red with and blue without

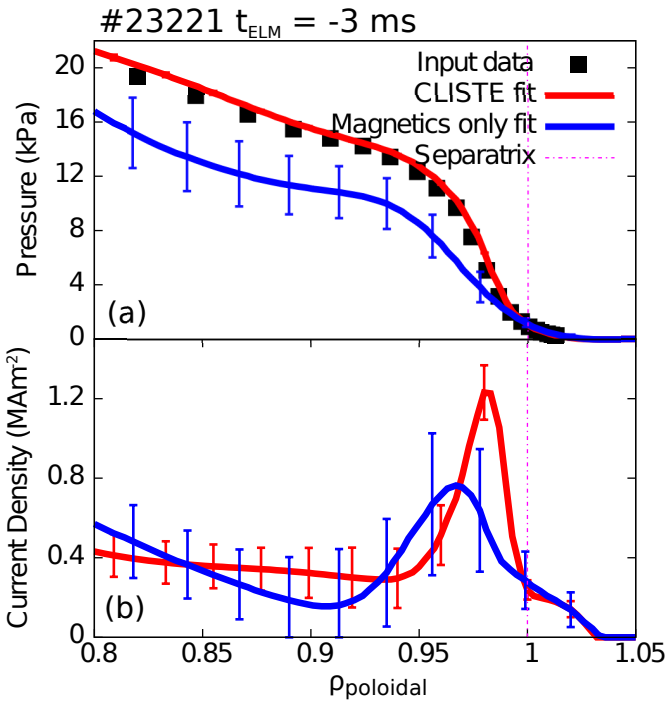


Figure 3.5: (a): Pressure profile. (b): current density profile. The blue line denotes a fit made only with magnetic and divertor current constraints, while the red line is the fit using an additionally prescribed pressure profile. Black boxes represent input data points. Error bars are 1 sigma confidence bands.

an input pressure profile. This smooth pressure profile was created by averaging IDA fits to the electron temperature and density data in 0.1 ms timeslices relative to the ELM crash. All kinetic data in this section, as well as in sections 3.7 and 3.8 were made with this method. The boxes show the input pressure data points. The effect of the divertor currents can be clearly seen in the SOL region of the current density fit where the error bars are much smaller than inside the separatrix and both profiles agree closely. Since these currents are measurements of the poloidal SOL current, this provides an integral constraint of the ff' source profile on and outside the separatrix.

Once we move radially inwards from the SOL, the two current density fits begin to diverge, and the error bars increase substantially in the magnetics-only case. The error bars shown here (and in all subsequent plots) are one sigma confidence bands, as calculated by CLISTE, and described in section 3.5.1. The flatter current density profile in the magnetics-only case is due to the internal curvature constraint in CLISTE, which aims to minimise the value and curvature of the knots in the profile spline fit.

However, despite these differences, the overall profile structure is similar in both cases. A strong peaking of the edge current density is observed with the

principal difference being the localisation of the peak. From the pressure gradient dependent drives of edge current in a tokamak it is not surprising that the form of the current density peak is similar to that of the pressure gradient.

3.5.1 Confidence band calculation

This section is reproduced from Dunne et al.[97]

The starting point for calculation of error bars for individual scalar parameters and confidence bands for plasma profiles is the NxN variance-covariance matrix V of the N free spline coefficients that parameterize the p' and ff' source profiles and hence the current density profile. This matrix is returned by the linear regression routine that optimizes the fit to the data at each iteration cycle. The diagonal of this matrix holds the variances, i.e. squared standard deviations, for the fitted parameters, while the off diagonal terms hold the covariances.

If p is any parameter of the interpreted equilibrium and $\nabla_\alpha p$ is the gradient vector of p with respect to the set of fitted parameters α_i, then $\sigma(p)$, the standard deviation for p, is given by:

$$\sigma^2(p) = (\nabla_\alpha p)^\top \cdot V \cdot \nabla_\alpha p \tag{3.3}$$

Confidence bands for spatial profiles (e.g. as a function of major radius R) can be constructed in a pointwise manner by treating each element of a regularly spaced array of profile values as a separate parameter and interpolating the calculated array of standard deviations to form a continuous function of R.

Equilibrium reconstruction is an ill-conditioned inverse problem and the error bars obtained by the above procedure are specific to the choice of current profile parameterisation and the choice of regularisation penalty. Tikhonov-type regularisation is used here, where the magnitude of the spline coefficients and the curvature at each knot location are penalised. The choice of the regularisation parameter is guided by methods such as the L-curve and Morozov's discrepancy principle [98].

In the case of the current density, the magnitude of the confidence bands is calculated from the covariances with respect to the ff' and p' source profiles, since j depends on both of these. In the magnetics-only case, both source profiles are internally free and constrained only by external information; in the kinetically constrained case the p' profile in internally constrained in the pedestal region, lowering the contribution of these coefficients to the confidence band calculation.

3.6 Power scan

As a first test of the dependence of the current density on external parameters, a discharge featuring a power scan was analysed. Discharge #23223 had broadly similar plasma parameters to #23221 described above, but with the NBI power varied between 5 and 10 MW. This gives a corresponding variation in the pedestal top pressure and hence the edge pressure gradient. The pressure profiles used as input were, in this case, fit using a modified tanh function to Thomson Scattering (T_e, n_e), Li Beam and DCN (n_e), and ECE (T_e)data. The ion temperature, as for #23221 was assumed to be equal to the electron temperature. An example of the input profiles is shown for the high power phase in figure 3.6. Since two of the

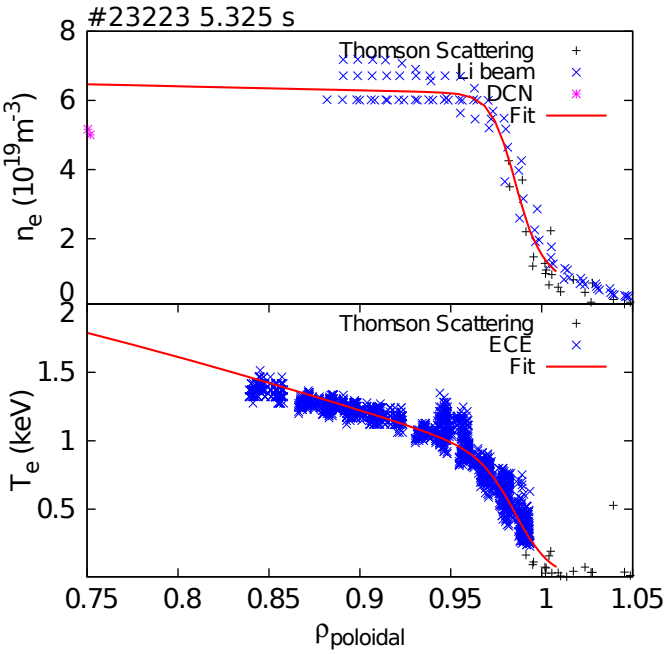


Figure 3.6: Input profiles fitted to raw data for discharge #23223. (a): electron density profile (red) fitted from Thomson scattering (black), LiBES (blue) and DCN interferometry (purple). (b): electron temperature fitted from Thomson scattering (black) and ECE data (blue).

principal drives of current density in the plasma edge are based on kinetic gradients, it is expected that an increase in heating power would lead to an increase in the current density. Shown in figure 3.7(a) are the input pressure profiles and corresponding CLISTE pressure while (b) shows the current density fits for all three heating levels in this discharge. The peak edge current density is seen, as expected, to increase with applied heating power. On examining the pressure profiles in closer detail, it can also be seen that for the larger increase in the pedestal top pressure (medium to high power) the current density increase is also

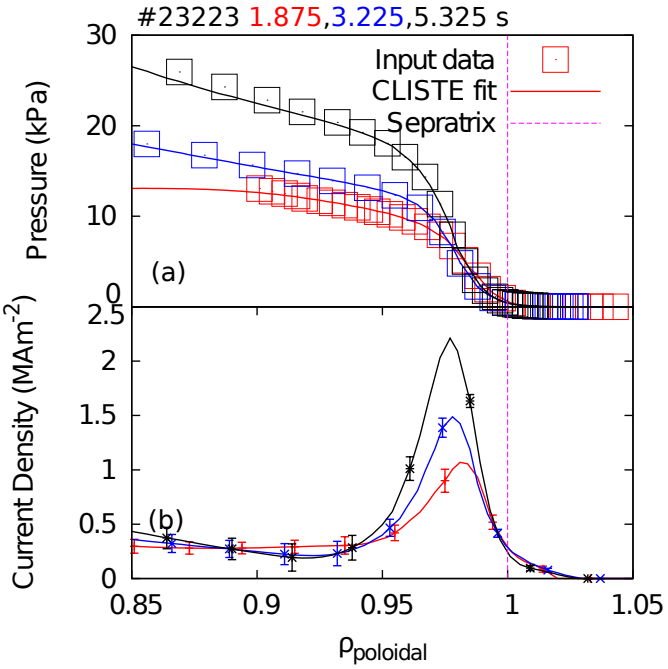


Figure 3.7: (a): input pressure data (squares) and fitted pressure profiles for low (red), medium (blue), and high (black) power. (b): corresponding current density profiles with one sigma confidence bands.

larger. Another point to note in this case is that the separatrix current density remains the same for all three timepoints. This could be simply coincidence, or it could also be due to an MHD effect, clamping the current value at the plasma surface. The SOL current varies substantially between the three different timepoints in the discharge, increasing when the gas puff is switched off during the second heating phase and subsequently decreasing for the highest heating power. That there is no corresponding trend in the separatrix current density could be indicative that this value is fixed either by current diffusion from the pedestal across the separatrix and into the SOL, since the resistivity is the same in all cases, or by a surface current constraint, such as the kink mode.

3.7 Recovery from ELM crash

Data and comments from this section, as well as section 3.8 are based on the work presented in Dunne et al.[97]. Results in the intervening sections are based on the same series of reconstructions presented in this reference.

One of the principal interests of this study is the current density evolution relative to the ELM crash. To examine this, a time series of CLISTE runs was made with ELM synchronised data. This section includes data from discharge #23221. As with the single profile detailed above, data were taken between 3.6

and 4.2 seconds in the discharge and synchronised into a single composite ELM cycle. Due to the high time resolution of the diagnostics used as well as the ELM frequency of ≈ 100 Hz, a time resolution of 0.1 ms was attained for the CLISTE reconstructions.

Time traces of the peak input pressure gradient (black squares), fitted peak pressure gradient (red line), the peak pressure gradient in the absence of kinetic constraints in CLISTE (blue) and the corresponding low field side peak edge current densities are shown in figure 3.8. The peak pressure gradient exhibits a sharp drop at the ELM crash, which can also be seen in the current density behaviour. In addition, the current density then recovers on a similar timescale to the pressure gradient. This finding, coupled with the constant pre-ELM current density and consistent with the model in [48], is therefore inconsistent with the theory of a resistive delay in the current density growth being responsible for the ELM crash. Also interesting in figure 3.8 is how the current density evolves when

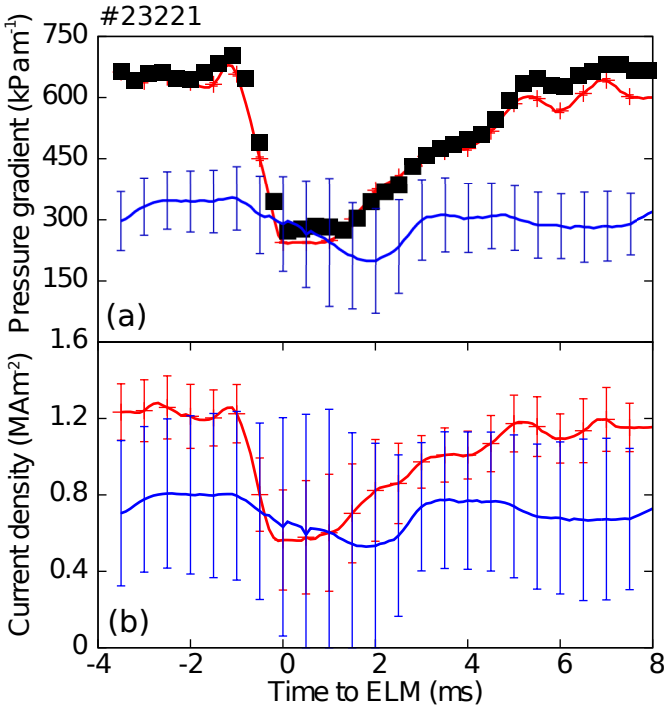


Figure 3.8: (a): maximum pressure gradient time trace. (b): peak edge current density time trace. The colours indicate the different constraints, as in figure 3.5. The current density closely follows the evolution of the pressure gradient in both cases.

the kinetic constraints are relaxed. Again, the trend of the peak edge current density appears to follow that of the peak edge pressure gradient closely. The largest contribution to the LFS edge current density comes from the Pfirsch-Schlüter current, which is proportional to the pressure gradient. From a physics point of view, this matches well with the similar trends in the pressure gradient

and current density, and also implies that the shape of the local edge current density peak is highly dependent on the form of the pressure gradient at the plasma edge.

The rapid drop of the edge current density at the ELM crash initially seems somewhat surprising; one would expect a decrease (and subsequent increase) in line with a resistive delay relative to the pressure gradient collapse. However, this becomes easier to understand when one sees that the plasma circumference shrinks rapidly at the ELM crash, as shown in figure 3.9. This corresponds to a

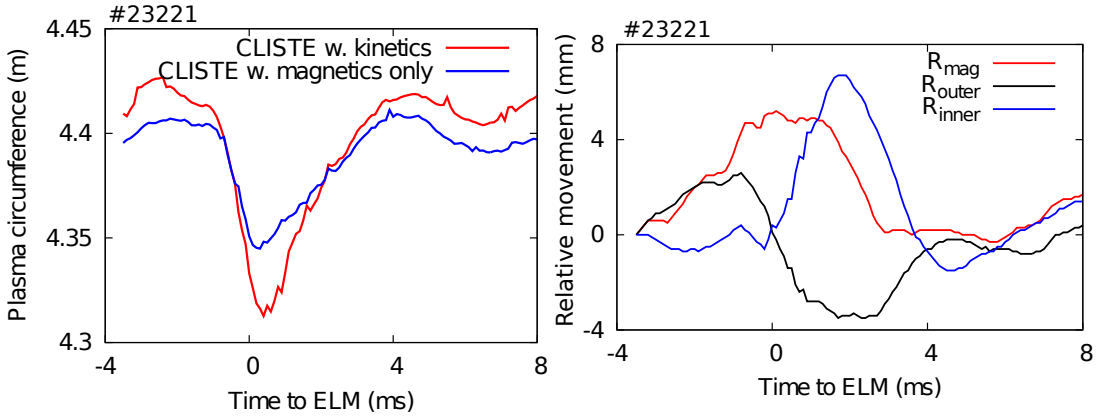


Figure 3.9: (a): Time trace of plasma circumference relative to ELM. Shown in red is the circumference corresponding to the kinetically constrained CLISTE, while the blue shows CLISTE with magnetics measurements only. (b): Time traces of movement of magnetic axis (red), outer midplane separatrix location (black) and inner midplane separatrix location (blue). All traces are taken relative to the first timepoint at 3.5 ms prior to the ELM crash.

movement of the inner and outer separatrix location of approximately 5 mm in each case, shown in figure 3.9. This reduction in the size of the plasma could lead to a large portion of the edge region (the pedestal being $\approx$ 20 mm wide) being removed to the SOL, allowing the plasma to lose current faster than expected. Support for this hypothesis is given by a larger movement of the inner separatrix location, which, due to the larger flux expansion, implies that a fixed amount of flux is separated from the entire plasma surface. Results in figure 3.9 (and all further plots) are from the CLISTE results constrained with magnetic, kinetic and divertor current data.

It can also been seen in figure 3.9 that the inward motion of the magnetic axis is delayed by $\approx 2-3$ ms with respect to the separatrix contraction. This inward movement after the ELM crash is expected due to the loss of pressure, and the delay corresponds to the pedestal resistive timescale of $\approx$ 1.5 ms; this was determined using an average pedestal resistivity of $3 \times 10^{-6}\,\Omega$m (calculated

from neoclassical resistivity) and a pedestal width of 2 cm.

The question naturally arises: what happes to the plasma current during the ELM crash? Is it lost, or redistributed? The fast drop in the edge current density might imply that there is some sort of fast redistribution effect. If this were the case, the lost current should be detected in the SOL. Figure 3.10(a) shows that, indeed, a significant current is detected in the SOL during the ELM crash. However, Pitts et al.[99] have shown that for ELMs on the TCV tokamak the currents detected in the SOL were of a predominantly thermoelectric nature, i.e that they were driven by the temperature difference between the inner and outer divertor.

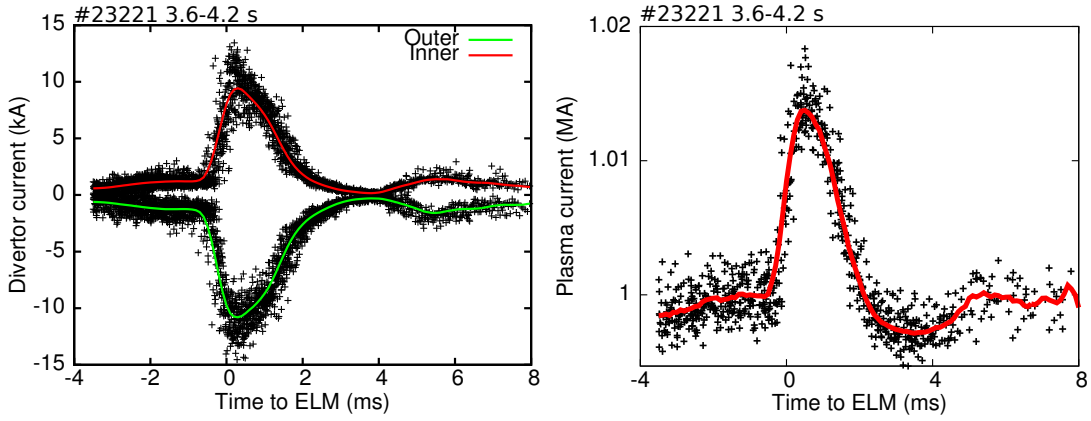


Figure 3.10: (a): Inner (red) and outer (green) plasma currents during the ELM cycle of discharge #23221, showing an increase in magnitude at the ELM crash. (b): Total plasma current relative to the ELM crash, showing an increase of ~1% at the ELM crash, corresponding to the inductive response of the plasma to the loss of edge current at the crash. The black point currounding each of the smoothed averages are the data points from the individual ELM cycles throughout the timerange analysed.

A further useful piece of information comes from the evolution of the total plasma current during the ELM cycle, plotted in figure 3.10(b), which shows an *increase* at the ELM crash and subsequent decay. This behaviour is reminiscent of the plasma current during the early phases of a disruption[100]. A disruption is the uncontrolled ending of a plasma in which confinement is suddenly lost due to, for example, a large MHD mode in the plasma core. In the case of a disruption, the core pressure profile, and hence current profile, collapses, which ejects a large amount of magnetic flux from the plasma core. By Faraday's law, this induces a voltage to balance the loss of magnetic flux, causing the plasma current to increase briefly before it ramps down completely in a short space of time. The ELM event is similar, in that the pressure gradient driven current

suddenly stops at the ELM crash, causing a loss of flux in the plasma, which is compensated for by an increased local voltage, as will be shown in chapter 5. In order to explain the increase in the total plasma current, a portion of the flux, and hence current, would have to be completely lost, as would be the case for the current density driven peeling mode.

3.7.1 Radial and temporal ELM resolved behaviour

The figures shown in the previous sections have focussed on the temporal behaviour of the edge current density, notably the peak value of the local LFS current density. Figure 3.11 shows a surface plot of this quantity relative to the ELM crash. The drop of the current density peak relative to the ELM crash can

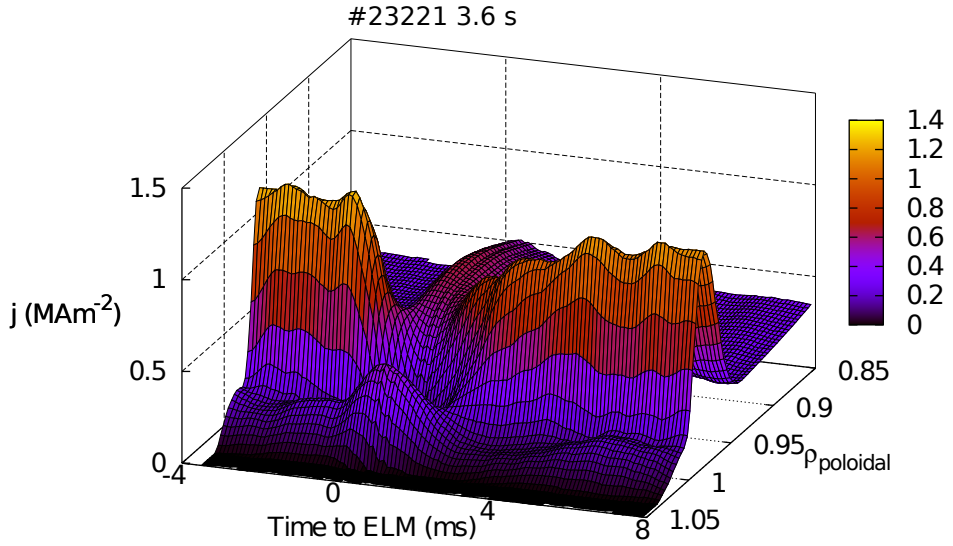


Figure 3.11: Surface plot of LFS edge current density from discharge #23221. This figure shows both the temporal and radial behaviour of the current density relative to the ELM crash and the pre-ELM peak location.

be very clearly seen in this figure. What can also be seen is that the separatrix current density does not change substantially, but rather acts as a fulcrum with the current density dropping inside and rising on the outside. At least part of this rise in the SOL is due to thermocurrents arising from a potential difference between the hotter outer divertor and cooler inner divertor[101].

The resistive spread of the current can also be seen in this case, with an increased current density propagating inwards. This is also due to an increased toroidal electric field (despite a reduced conductivity) and the slightly increased pressure gradient inside the pedestal top caused by the collapse of the pedestal; the core pressure and separatrix pressure are approximately the same as their

pre-ELM values, but the transport barrier is now reduced, meaning a smaller average gradient, but a steeper local gradient between $\rho_{\text{poloidal}} = 0.85$ and 0.95.

In the lead up to the ELM crash, figure 3.11 indicates that, in addition to a slight peaking of the edge current density, there is also a change in the outer gradient of this peak, signified by a spreading of the contour lines. The gradient of the edge current density is a strong driving force behind the peeling mode and indicates a compression of the edge current, so changes in this parameter could be indicative of an increase in the strength of this mode. Figure 3.12 shows the time trace of the peak value of this quantity relative to the ELM crash. This figure shows more clearly the sharp increase in the gradient of the

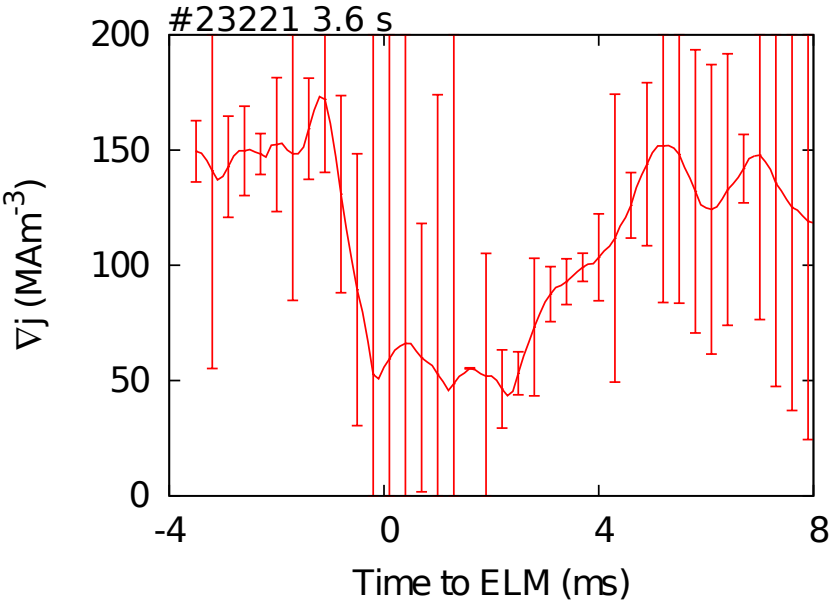


Figure 3.12: ELM synchronised peak current density gradient for discharge #23221. Prior to the ELM cycle, there is a steepening of this gradient, which then reduces and recovers on a timescale similar to that of the pressure gradient.

current density, always located outboard of the peak, just before the ELM crash. The uncertainties in this quantity (which have been derived analytically from the current density uncertainties and reduced by a factor of 1/(n-1), where n is the number of data points from the magnetic inputs used to create the averaged time series) are large, making a definite statement difficult. However, there is a possibility of a higher current density gradient just before the ELM crash which could be relevant to a final ELM trigger. The large uncertainties at and just after the ELM crash are to be expected due the disturbed equilibrium. The uncertainties later in the recovery cycle are due to few points being present; since there are many "fast" ELMs present in this discharge which are triggered just as the pressure gradient fully recovers at approximately 5-6 ms, the long ELM

cycles are not well populated.

3.7.2 Normalised pressure gradient and current density evolution

Up to this point we have mainly focussed on the behaviour of local quantities relative to the ELM crash. However, for comparison with theory, specifically linear MHD instability analysis, normalised values are required. As detailed in chapter 2, the normalised pressure gradient α and a normalised current density are used to parameterise the plasma edge. At the ELM crash, one expects a rapid drop in both the edge pressure gradient and current density as already shown in previous sections. However, figure 3.13(a) shows that this is not the case. Here, we see that the pressure gradient drops rapidly (red phase), but is followed by a slower initial decrease in the current density which then "catches up" with the pressure gradient before reaching its minimum. The pressure gradient then continues to fall and remains constant during the initial recovery phase (red leading to blue phases) while the current density fluctuates and then begins to increase. The two then evolve together as the density gradient recovers (blue phase). This is then followed by a drop and saturation of the current density as the temperature gradient recovers (green phase). Finally, in the purple phase, the current density can be seen to fluctuate despite a relatively constant pressure gradient.

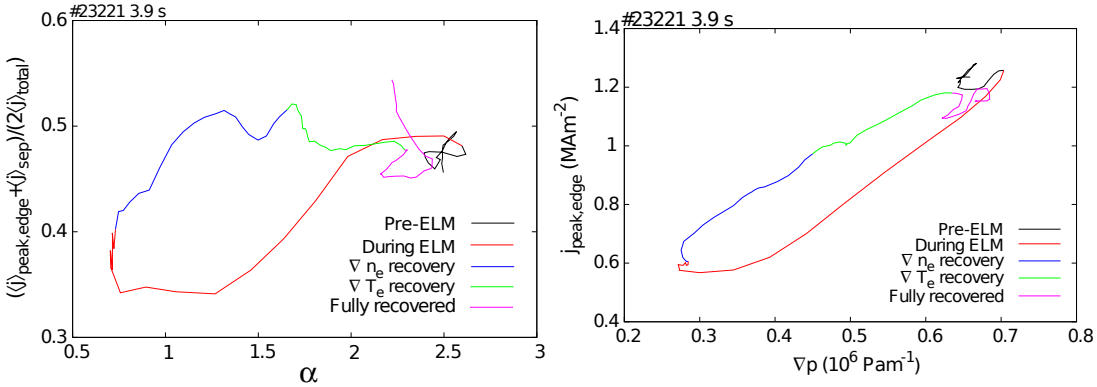


Figure 3.13: (a): evolution of the normalised parameters describing the peeling-ballooning stability (edge current density and α). Colours correspond to different phases in the ELM cycle: before the ELM crash (black), until the plasma current peaks (red phase), until electron density gradient recovery (blue), until electron temperature gradient recovery (green), and fully recovered (magenta). (b): evolution of the local pressure gradient and edge current density peak for the same five phases of the ELM cycle.

It is possible that this large variation after the recovery is due to too few points present in this phase of the ELM - many ELMs have already been triggered, and no data points within 3.5 ms of the next ELM were included here. Thus, for the pre-ELM values of α and j_{norm} the black phase represents the best data. A variation of ~10% in both of these quantities can be seen in this black phase with the maximum alpha being reached within 100 μs of the ELM onset, as determined by when the density gradient in the pedestal reaches its maximum. For comparison, the same phases are shown for the local peak edge pressure gradient and current density in figure 3.13(b). In this case, a more or less simple linear trend in both quantities can be seen throughout the majority of the evolution. The local current density also increases at constant pressure gradient as the plasma current reaches its peak and the two then increase steadily as both the density gradient and temperature gradient recover. Once the saturated conditions are reached, there is again a scatter in both parameters.

3.8 Fuelling study

Following on from the work of Burckhart et al.[48], the same discharges were analysed for a fuelling study. Using the methodology outlined above, additional series of CLISTE reconstructions were carried out for ASDEX Upgrade discharges #23225 and #23226. The same broad parameters were used in all discharges, (B_0 = 2.5 T, I_p = 1 MA, P_{NBI} = 7.5 MW), though the ECRH heating power was varied over the discharges. Feed forward fuelling for these two discharges was 9×10^{21} s^{-1} and 0 s^{-1} respectively, giving a wide range of collisionalities across the three discharges.

The low field side local current density profiles from all three cases 3.5 ms prior to the ELM crash are shown in figure 3.14(a) below. Significant differences can be seen between the three profiles. In the case of #23221, which featured the highest ELM frequency, the current density peak is both higher and positioned closer to the separatrix than in the other two cases, while the same separatrix current density remains. For the case of #23225, which had the lowest ELM frequency, the current density peak is the lowest of the set. These differences may be able to explain why the ELM frequency differs between the three discharges, although advanced MHD stability analysis is required to confirm this. The location and value of the current density peak could be obtained with higher accuracy if one were to move the plasma position radially in order to add artificial lines of sight; this has the disadvantage of changing the ELM behaviour. Since insight into

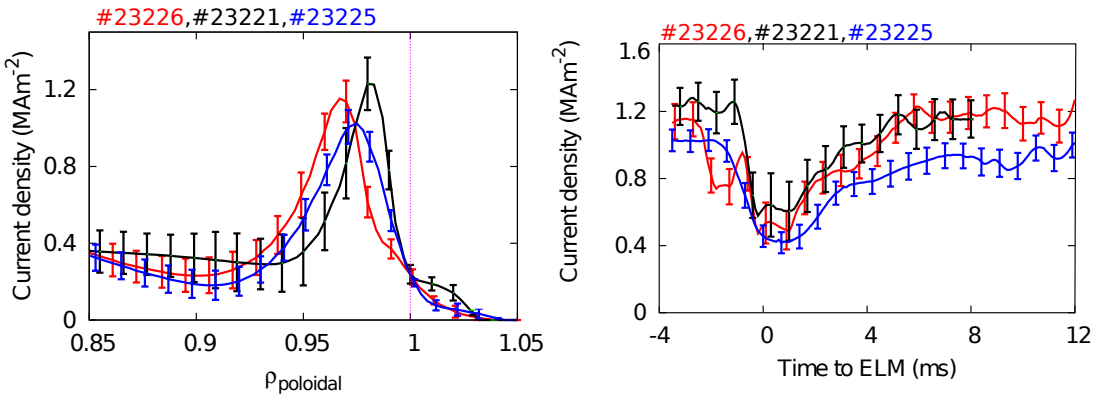


Figure 3.14: (a):low field side local current density profiles for discharges #23226 (red), #23221 (black), and #23225 (blue). Each discharge exhibits a significant localised edge current density peak, as was described in section3. (b): the evolution of the value of this peak for the same discharges. Pre ELM behaviour is similar for the discharges, in that no major fluctuations can be observed, but the recovery behaviour is different for the three cases.

the fine dynamics of ELM recovery was desired in this case, it was decided to sacrifice extra radial resolution in favour of the higher time resolution offered by ELM synchronisation. Data taken during radial plasma scans have been analysed and are shown in chapter 6.

If we now compare the time evolution of the peak edge current density in figure 3.14(b) of the three discharges, we again see significant differences. Prior to the ELM crash, some fluctuation can be seen in the current density values, although the general trend is a rapid decrease of the edge current density at the ELM onset. This is followed by a phase with no recovery of the current density, which lasts 1-2 ms and then the recovery begins in all three cases. It can also be seen that the onset of recovery is slightly slower for discharge #23225. This mirrors the finding by Burckhart et al. that the electron temperature gradient recovery onsets at 2 ms after the ELM crash in this case vs. 1.7 ms for discharge #23221 and 1.5 ms for discharge #23226. In terms of full recovery of the edge current density, discharge #23221 recovers fastest, in line with its higher ELM frequency, while 23225 is by far the slowest, which reaches its pre-ELM state only 10-12 ms after the ELM crash. This is again similar to the findings by Burckhart et al. where it was shown that the recovery of the pressure gradient was limited by the slow recovery of the temperature gradient. It is likely that this effect is due to the higher rate of feed-forward fuelling in this discharge which limits the temperature gradients in the pedestal and also, evidently, the rate of recovery to this lower temperature gradient.

3.9 Conclusions

This chapter has presented the first results of current density recovery using the CLISTE equilibrium code. It is robust for different heating powers and fuelling rates, and the form of the edge current density profiles has been shown to be strongly dependent on the edge pressure gradient profile. The various diagnostics used to determine this edge pressure gradient have been described and their strengths and shortcomings noted. It was shown that the edge current density decreases strongly at an ELM crash and recovers on a timescale similar to the pressure gradient. It was also speculated that this large drop in edge current density could be related to the similarly rapid contraction of the plasma radius at the ELM crash. In terms of spatial-temporal evolution, the edge current density gradient appears to increase just before the ELM crash, while the edge current density spreads out from the steep pressure gradient region following an ELM crash.

The evolution of the normalised pressure gradient and edge current density was also analysed and showed significant deviations from the evolution of the local profiles. In particular, a stalling of the normalised pressure gradient while the ELM was ongoing was observed, with a strongly fluctuating edge current density. Both j and α were seen to evolve separately while both the electron density and temperature gradients recovered. Fluctuations in the normalised pressure gradient and edge current density were also seen just prior to the ELM crash.

In a fuelling study, significant differences in the form of the local edge current density profile and in its recovery from the ELM crash were seen. It was noted that in the discharge with the highest ELM frequency (with medium fuelling), the current density peak was both higher and closer to the separatrix than the other two discharges analysed. The discharge with the lowest ELM frequency (highest fuelling) featured the lowest edge current density peak as well as a much slower recovery time. No firm conclusions on the role of a delayed edge current density rise leading to an ELM crash could be made, but the data presented here do not point to this being a factor in eventually triggering an ELM. The smaller fluctuations of the normalised parameters should however be checked against theoretical calculations, which is the subject of ongoing investigations.

Chapter 4

Sensitivity and accuracy of current density profiles and separatrix location

The previous chapter described the general form and time evolution of the edge curent density. It was also shown that the peak height of the edge current density depended strongly on the edge pressure gradient. This chapter will analyse the impact of uncertainties in the pressure gradient and how they affect the edge current density profile. Systematic effects, notably shifts of the profiles relative to the equilibrium, will also be analysed.

It was shown by McCarthy[22] that the integral of the edge current density (in the last 1% radial annulus) can be determined with a high degree of accuracy using only external magnetic measurements. This chapter will therefore also assess if other higher order moments of the edge current can be derived from these measurements. Profile location and the shape of the profile are two moments which could potentially be recovered and which would aid greatly the analysis of this edge current. In addition, this chapter will also compare the separatrix location as determined by CLISTE with other diagnostics to determine if this is a robust quantity.

4.1 Methodology

The overall goal of this study is to understand the sensitivity of the reconstruction to deviations in the measurements. Discharge #23223 was chosen for this study as it varied the pedestal gradient and current substantially (as seen in the previous chapter), which also provides the opportunity to check magnitude dependencies

on the reconstruction uncertainties.

The first test conducted was to determine the sensitivity of the reconstructions to shifts of the pressure profile. This has been observed as necessary for advanced discharge analysis for a number of reasons. The first of these is an error in the magnetic reconstruction, such that the location of the separatrix is not known. It has already been shown by Schneider et al.[76] that the poloidal magnetic flux gradient in the pedestal has a 5-15% uncertainty, depending on the input power. This could also have an effect on the location of the peak gradient, although it was noted that the uncertainties due to this effect are reduced when the profiles are mapped back to real space, as is done for CLISTE input. The second reason for profile shifts is due to uncertainty in the spatial location of the kinetic data. This is of the order of 3-5 mm at ASDEX Upgrade, but the combination of this uncertainty with the equilibrium can be much larger.

The second test concerned the form of the pressure profile. It was observed in chapter 3 that the peak current density and the pedestal top pressure/edge pressure gradient are related. Since the input pressure profiles for #23223 were made with a modified hyperbolic tan function (mtanh), the width of the pedestal could be varied systematically. By using the sigma values returned by the fitting routine, the pedestal width could be scaled in a fashion meaningful to the data. In the cases presented below, the pedestal gradient was scaled inversely to the width in order to retain the same pedestal top value.

4.2 Profile shifts

The kinetic profiles used here have been aligned manually via Thomson scattering data. However, this still leaves the possibility of errors in the location of the scattering volumes of Thomson scattering (shown in figure 3.3). As such, the profiles were shifted in 2 mm steps from the default location of the Thomson scattering volumes of the edge channel to give a total range of 14 mm, which is well outside the uncertainties in the kinetic profiles.

Figure 4.1 shows the effect of these shifts on the two principal types of magnetic diagnostics. In both subfigures the low power timepoint is shown in red, the medium power in blue, and the high power in black. Tic marks on both figures correspond to the same magnitude, although the axes show different ranges. Shown in figure 4.1(a) are the residuals from the B_θ probes, which measure the tangential component of the magnetic field. Very little effect can be seen throughout the majority of the range covered, apart from at the largest outward shifts

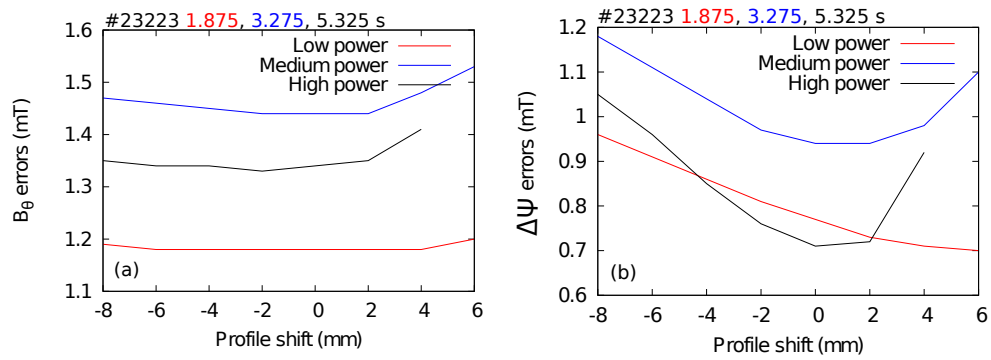


Figure 4.1: (a): residuals in B_θ probes plotted as a function of pressure profile shift for the low (red), medium (blue), and high (black) power phases of discharge #23223. (b): residuals in $\Delta\psi$ measurements corresponding to the same timepoints and showing a pronounced minimum in the case of medium and high power phases.

where a strong increase in the residuals is observed. Outside the range shown, converged equilibria could not be obtained. The $\Delta\psi$ flux difference residuals, corresponding to the normal component of the magnetic field, shown in 4.1(b), show a minimum for the medium and high power phases of the discharge. This minimum also appears to be more pronounced in the high power phase of the discharge, indicating a better localisation of the profiles when there is more current in the plasma edge. This is consistent with the theory behind the detection of the edge current as described by McCarthy[22]. This analysis would suggest that the profiles can be localised with an accuracy of up to ±2 mm when an advanced equilibrium is made.

Support for this statement is given by figure 4.2. According to the two-

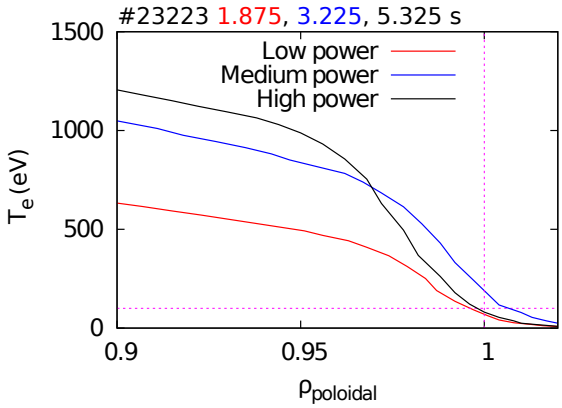


Figure 4.2: Electron temperature profiles at the positions of their respective minima in the position scan for the low (red), medium (blue) and high (black) power phases of discharge #23223. The purple lines show the location of the separatrix and T_e = 100 eV, indicating that the profiles can be well localised spatially.

point model, the separatrix temperature in an AUG H-mode is approximately 100 eV[102] with an uncertainty of 20-30 eV (generally above 100 eV). Figure 4.2 shows the electron temperature profiles from the best fits to the magnetic data. The separatrix is denoted by the vertical dashed purple line while $T_e = 100$ eV is shown by the horizontal dashed line. For the low and high power phases of the discharge the agreement between the fit and the expectation is excellent. In the medium power case, the agreement is poorer with a separatrix temperature of 175 eV, which is substantially outside the uncertainties. However, if this profile is shifted 2 mm further inboard the separatrix temperature is 130 eV, which is much more in line with the expectation. It should also be noted that the residuals from all positions in this phase are higher than in the low and high power phases.

4.3 Width scans

Once the profile locations were determined from the above scan (or taken to place $T_e = 100$ eV at the separatrix in the case of the low power timepoint), the width of the mtanh function was then varied in steps corresponding to a half sigma of the fitting error. The gradient was scaled such that the pedestal top value was conserved for all values of the width. Since the uncertainties in the electron temperature width are typically larger than the electron density, the pedestal width change is predominantly determined by T_e. Two types of width scan were conducted:

1. The pedestal width was varied around the centre point of the pedestal, meaning that the peak gradient location was the same (although its value varied) and the pedestal top and bottom values moved relative to this.

2. The location of the pedestal bottom was kept constant (i.e. a fixed separatrix boundary condition was imposed) and the maximal gradient and pedestal top were moved inwards from this point.

The first of these two tests assumed that the peak gradient could be well localised by the combination of the magnetic equilibrium and the kinetic measurements. The second makes the assumption that the separatrix is well known from the magnetic reconstruction and that the separatrix values of the kinetic data (temperature and density) do not change; this is a reasonable assumption. These two types of width scan will be dealt with in the following two subsections.

4.3.1 Width scan with peak gradient location fixed

Figure 4.3 shows the varied electron temperature (a), electron density (b), total pressure gradient (c), and resulting current density profiles (d) for the high power phase of the discharge in the first width scan. This first analysis makes the as-

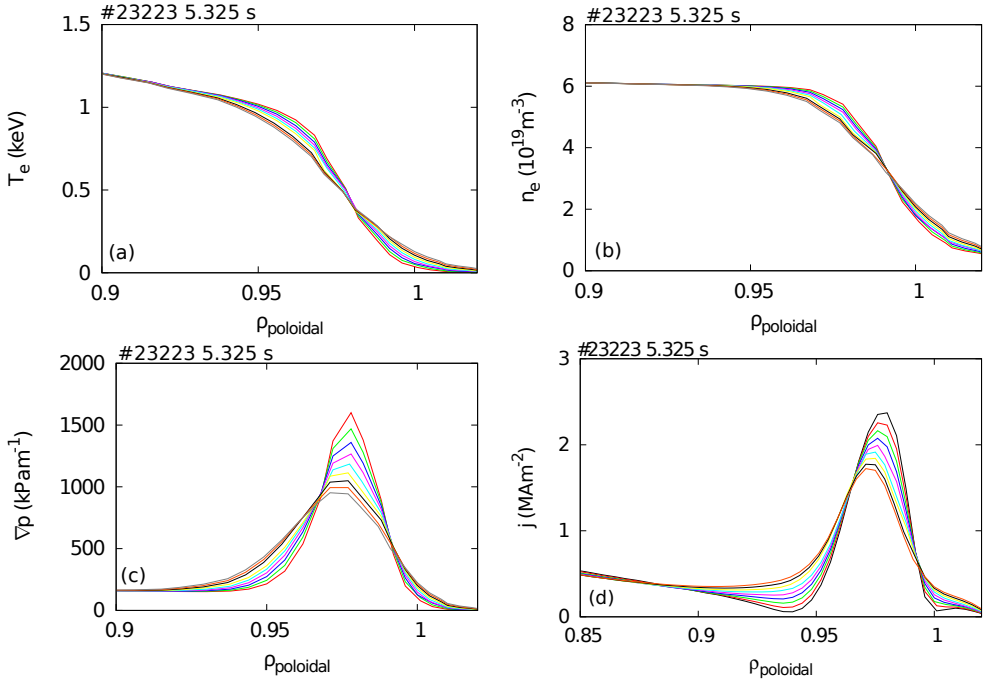


Figure 4.3: (a): variation of the electron temperature in the width scan, the extremes of which correspond to a 2σ uncertainties in the fit to the data from the high power phase of #23223. (b): electron density profiles, scaled in the same manner. (c): total pressure gradient constructed from these profiles. (d): output current density from CLISTE.

sumption that the location of the peak pressure gradient is well determined by the kinetic data and only the width of the pedestal around this point is varied, as can be seen from the input data. Correspondingly, the current density profiles steepen but the peak remains more or less localised (a small outward shift covering a range of 0.01 in ρ_{poloidal} is observed); if the current density is considered as a Gaussian shape, the test is essentially determining the sensitivity of the reconstructions the standard deviation of the current density profile. This can be interpreted as the deviation from a wire-like current distribution at the x-point to a flatter more smeared out current profile. The tangential and normal residuals can be seen in figure 4.4. The interpretation of this result is not straightforward. Almost no effect is seen in the tangential component of the field, even at the extremes of the variation. Two small steps can be seen in the tangential residuals for the high power phase, though the differences are small. However, it is striking

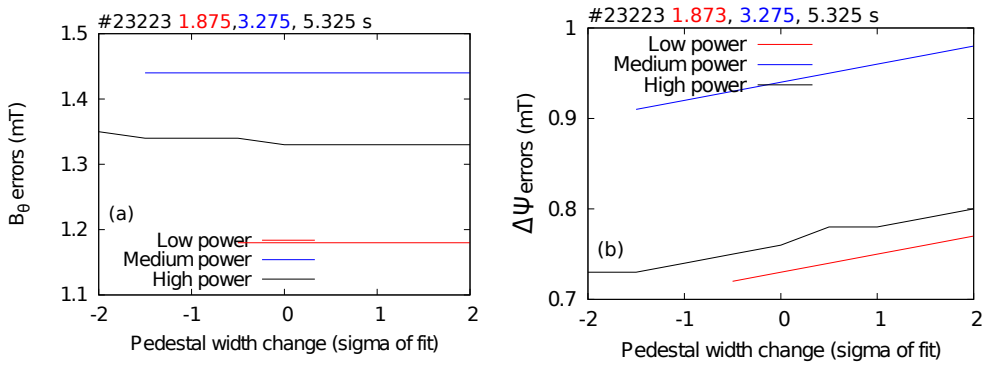


Figure 4.4: (a): residuals in B_θ against the change in pedestal width (as a function of σ_{width}. (b): $\Delta\psi$ residuals. The $\Delta\psi$ residuals are proportional to the pedestal width.

that the residuals in the normal errors increase almost linearly with increasing pedestal width. The difference between the two measurement sets is likely to be due to different responses of the two measurement types to a change in the local current profile. The large variation in the separatrix current density between the extremes of the pedestal width appears to have little effect on the B_θ measurements, which implies that they are more sensitive to the entire current profile. On the other hand, the $\Delta\psi$ measurements are sensitive to the localised changes in the flux, and hence are sensitive to the current distribution around the final 1% radial annulus. The residuals in this case indicate that the current density in this final 1% radial annulus should be lower than modelled with the measured kinetic profiles.

4.3.2 Alternative width scan

The previous section made the assumption that the location of the peak gradient is well localised and that the location of the pedestal top and bottom were unknown. In this section, we will examine the effect of a well defined pedestal bottom and uncertain pedestal width extending further into the plasma. The reasoning for this particular approach is that, for example, the electron temperature has been shown to be 100 eV at the separatrix in AUG, as has been determined from power balance calculations[102]. In addition, section 4.5 will demonstrate that the separatrix location can be well determined from the magnetic equilibrium. To examine this, not only has the pedestal width been changed, the centre point has also been shifted to take this change into account. Again, the pedestal top values have been kept constant and the gradients scaled accordingly. Figure 4.5 shows the input electron temperature profiles as an example, and the output

current density profiles for the high power phase of the discharge. Note that the

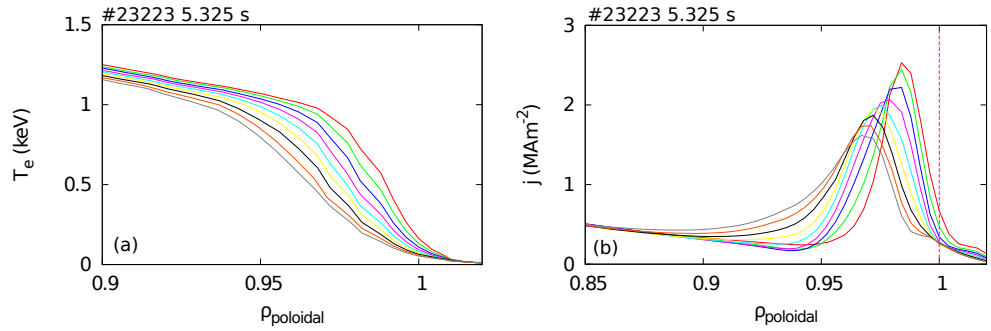


Figure 4.5: (a): electron temperature profile change when keeping the pedestal bottom at the same value and location and increasing only the extent into the plasma. (b): corresponding current density profiles, showing the lower peak which is located further inside the separatrix with increasing pedestal width.

separatrix temperature is not kept exactly constant as the bottom of the pedestal (for the mtanh fit) is slightly outside the separatrix. The same is true for the current density profile, particularly in the cases of the narrowest two pedestal widths which show a large current density in the SOL.

The effect of this scan on the magnetic residuals is shown in figure 4.6. In this case, where the boundary is fixed, a minimum in the residuals can be located in both the tangential and normal residuals which is in strong contrast with the "smeared" profiles shown in the previous section. Despite the strongly localised

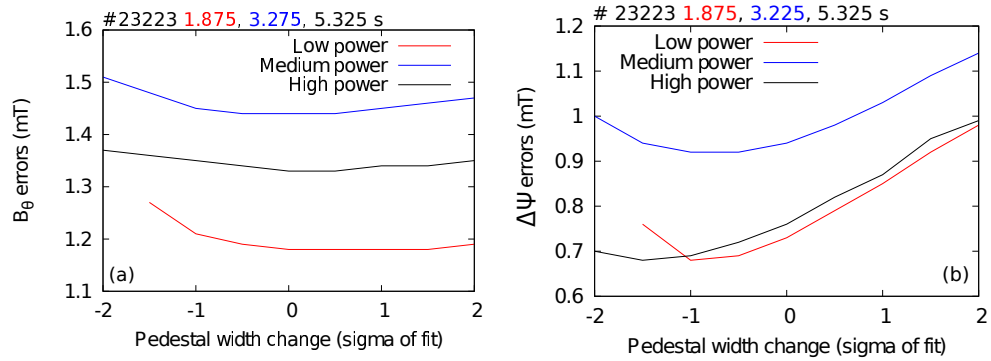


Figure 4.6: (a): residuals in B_θ probes for modified width scan. (b): corresponding residuals in $\Delta\psi$. A trade off between the two measurements can be seen, the minimum of which occurs at approximately the value of the nominal fit. The $\Delta\psi$ residuals alone indicate a minimum at narrower pedestal widths.

minimum in the flux difference measurements at narrower pedestal widths, it seems that there is a better minimum for the combined set closer to the nominal profile fit. This is due to the large number of tangential field measurements (56 vs. 18 normal measurements), meaning that the location of the minimum in this

measurement set, if it exists, carries much more weight. As can be seen in figure

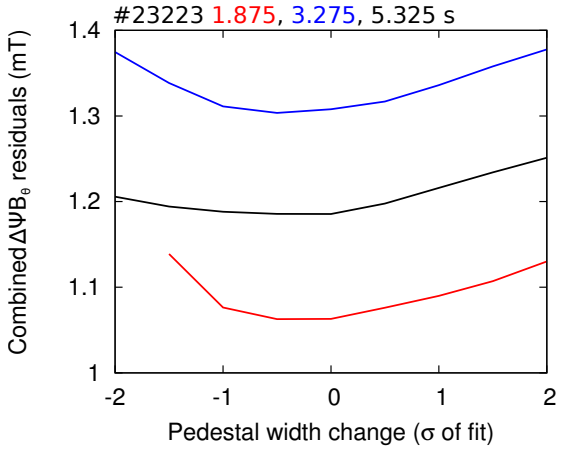


Figure 4.7: Combined residuals from both magnetic diagnostic groups plotted as a function of the pedestal width change. The uncertainties from both groups have been weighted with the number of measurements they contain. The minimum in the width for all three timepoints moves towards the nominal pedestal width.

4.6(a) the minimum, although slight, is shifted towards wider pedestal widths. This is confirmed in figure 4.7 which shows the combined residuals from the $\Delta\Psi$ and B_θ measurements as a function of pedestal width. The minimum in the respective curves is indeed shifted more towards the "nominal" fit than in only the $\Delta\Psi$ residuals.

Narrower pedestal widths can be ruled out with certainty due to the strong increase in the residuals of the tangential measurements. Additionally, due to experimental smearing, the measured pedestal width will always be wider than the real pedestal width. This means that the true pedestal width should be well determined by the combination of the kinetic measurements and the analysis of the magnetic residuals from the equilibrium solution. On comparing the minimum residuals found in both cases, it can be seen that the minimum normal residual corresponds appromximately to the smallest residual found in the "smeared profiles" scan. This indicates two things: first, the intial scan did not extend far enough to provide a minimum and; second, the inital scan convolved the effect of changing the current densty in the final 1% annulus with changing the current density at and outside the separatrix. When keeping the separatrix and SOL current density almost constant, a more clear effect on the residuals is seen.

4.4 Interpretation of effects

It is necessary to look at which signals in the B_θ and $\Delta\psi$ measurement sets provoke this strong variation in the residuals. In all three power phases the residuals of the flux difference measurements which cross both of the strikelines have been observed to have the same trend as the total residuals while the other measurements remain constant. This could imply that there is some property of the strikelines which is changed by the shape of the local current density profile (or the total current profile) inside the plasma.

On investigating the second width scan, it was observed that the strikelines follow different paths depending on the current profile inside the separatrix. This variation can be traced back to the location of the x-point, which varies with a similar form in both the R and z directions as the residuals for the flux differences; the basic movement is that narrower pedestal widths give an x point which is higher and further outboard than wider widths in all three cases, although the movement inside the 1σ width range is within ± 5 mm. This corresponds to a movement along the target plate of approximately 10 mm. This leaves only the question of how realistic each of these strikepoint locations and corresponding x-point locations are.

If the fitting formula given by Eich et al.[103] is used to fit the the power to the outboard divertor we can determine this piece of information. This power is measured by infra-red thermography using an IR camera focussed on the lower divertor tiles. The fit of the power profile along the divertor target plates is given by

$$q(\bar{s}) = \frac{q_0}{2}\exp\left[\left(\frac{S}{2\lambda_q f_x}\right)^2 - \frac{\bar{s}}{\lambda_q f_x}\right]\mathrm{erfc}\left(\frac{S}{2\lambda_q f_x} - \frac{\bar{s}}{S}\right) + q_{BG} \tag{4.1}$$

where q is the target heat flux, q_0 is the initial height for the exponential fall off, $\bar{s} = s - s_0$ is the coordinate along the target plates with s_0 being the strikepoint (used as a free parameter in the fit), S is the width of the Gaussian (representing power leakage over the strikeline and into the private flux region), λ_q is the power fall off length, f_x is the mapping function from the divertor to the midplane, and q_{BG} is the background heat flux. In order for the IR camera to attain a good signal to noise ratio, power attached conditions are required; these conditions are fulfilled in the medium and high power phases of this discharge. Figure 4.8(a) shows a poloidal view of the divertor region with the tiles (solid black line), the strikelines (dashed black lines), strikepoint locations from the equilibrium (red points) and strikepoint locations in the pre-ELM phase from the Eich et al. fitting function (blue points). Figure 4.8(b) shows the development of the strikepoint

location in divertor tile coordinates (along the tile from a fixed starting point) relative to the ELM crash with the equilibrium point superimposed; the point corresponds the nominal profiles while the error bars denote the extremes of the pedestal width change.

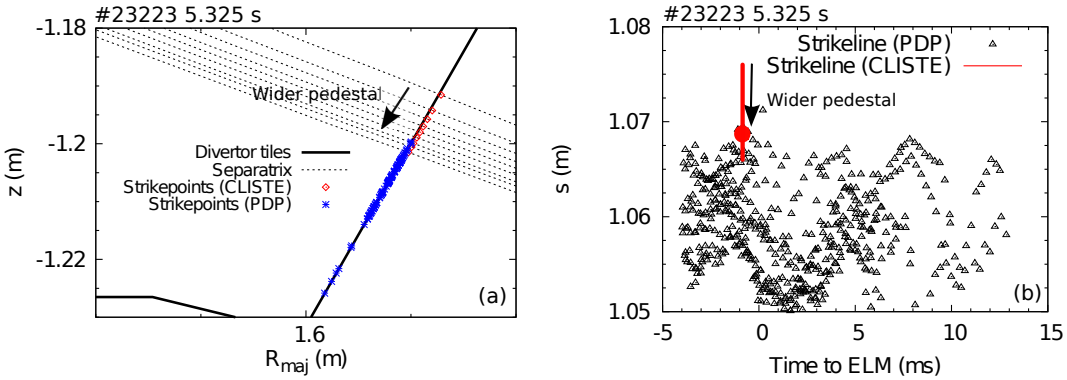


Figure 4.8: (a): strikepoint locations from the equilibrium reconstruction (red points) with different pedestal width constraints and the fitted strikepoint locations for all pre-ELM timepoints using data from the IR camera. (b): ELM resolved strikepoint location in divertor coordinates from the IR camera with the red point indicating the location determined from the equilibrium. The error bars on this point correspond to the upper and lower extremes of the pedestal width.

The large scatter of the IR data points is due to the radial shift of the plasma during the time range of interest. This shift was conducted, as previously mentioned, to increase the radial resolution of the edge diagnostics, such as Thomson scattering. The two methods do not agree perfectly, with the equilibrium giving values which are, on average, approximately 1 cm too far along the tile. The reasons for this are not certain, although the generic equilibrium reconstruction gives values which are closer to those determined from the power deposition. One possible reason could be the effect of $\mathbf{E} \times \mathbf{B}$ drifts in the SOL. The electric field has two components perpendicular to the field line, radial and poloidal. Since the potential always drops between the main plasma and the divertor tiles, the poloidal electric field creates a drift which always points radially inwards in the usual case of "negative" (anticlockwise) toroidal field[104]. The sign of the radial electric field can vary bewtween discharges, but is generally positive, creating a drift downwards, towards the outer divertor[105].

To make an estimate of the size of this effect, appromximate values will be taken for the relevant parameters. Typical values for both of these electric fields are between 0.1-1 V/mm, yielding an average E × B velocity of ~200 ms^{-1} radially inwards and downwards, or a net velocity of 280 ms^{-1} at 45 °. To calculate the distance this corresponds to, the time taken to travel from midplane is required,

which is define as $\tau_{||} = L_c/c_s$, where L_c is the connection length and c_s is the sound seed. An average connection length, or distance along the field line, between the plasma midplane and the divertor of 15 m[106] of 15 m is taken, though this varies with plasma current and toroidal field. The ion sound speed is given by the equation: calculated as[99]:

$$c_s = \sqrt{e\left(\mathrm{T_{e,sep}} + \mathrm{T_{i,sep}}\right)/m_i} \tag{4.2}$$

where values for $\mathrm{T_{e,sep}}$ and $\mathrm{T_{i,sep}}$ are 100 and 200 eV, respectively. This yields a sound speed in the SOL of ~1.2×10^5 $\mathrm{ms^{-1}}$. These approximate values yield a drift distance of 3.5 cm inwards and downwards. This is qualitatively in the correct direction to correct for the discrepancy observed here, though the magnitude is much too high; quantitative agreement would require more in-depth modelling or particle transport in the SOL, such as with the SOLPS package[107]. A simple test would be to perform matched experiments with opposite directions of magnetic field, though such experiments are typically difficult at AUG.

The nominal profile does reach inside the upper end of the fitted points with wider pedestal widths extending further inside this range, indicating that the discrepancy is relatively small. It is assumed that the current distribution in the plasma edge influences the x-point location and causes this systematic shift of the strikepoint location, which are mainly derived from a radially outward shift of the x-point. The evidence of figure 4.8, coupled with the increased errors seen at lower pedestal widths in figures 4.4(a) and 4.6(a) and (b) would suggest that the equilibrium is capable of resolving a minimum pedestal width, below which magnetic residuals are often observed to rise to unacceptable levels. In some extreme cases a lack of convergence of these equilibria has also been noted.

4.4.1 Comparison of errors

A useful comparison is between the one σ current density confidence bands returned by CLISTE and the one σ pedestal width from the pressure data, which is shown in figure 4.9. The profiles shown are $\langle \boldsymbol{j}\cdot\boldsymbol{B}\rangle$ since the uncertainties returned by CLISTE for this profile are more balanced between both source profiles (p' and ff'), rather than being overly weighted by the low uncertainty in p' for the local current density on the LFS midplane. For each timepoint, three equilibria were compared. The initial one, shown in red in all cases, was calculated using the best-fit pressure profile from the mtanh fitting routine. The 1σ $\langle \boldsymbol{j}\cdot\boldsymbol{B}\rangle$ profiles, shown in blue, were calculated by CLISTE using the pressure profiles scaled according to the $\pm1\sigma$ pedestal width. Interestingly, it appears that the CLISTE

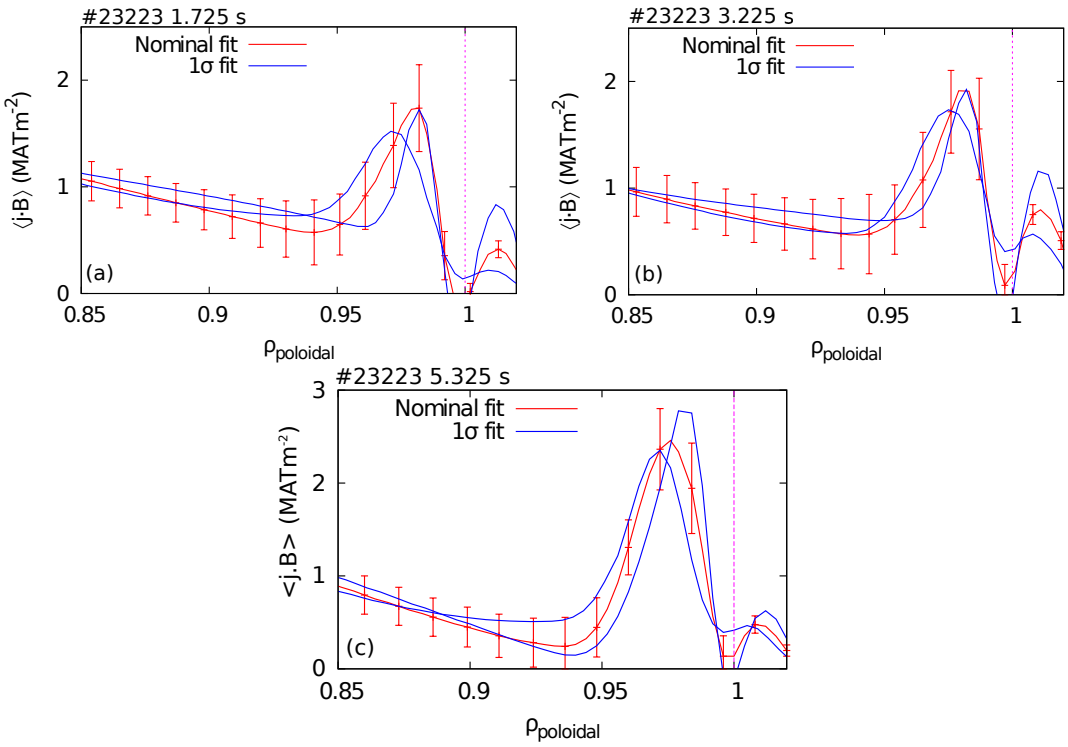


Figure 4.9: Comparison between the one sigma confidence intervals of the pressure data and the CLISTE fit for the three power phases (low (a), medium (b), and high (c)) of discharge #23223.

confidence interval corresponds well to the 1σ confidence interval in the pressure data for all three cases. It could be coincidental that this is the case, but the inclusion of realistic pressure measurement uncertainties as the weights for these data in the covariance fitting matrix is likely to be an important factor. A second observation from these figues is that the uncertainties inside the pedestal top appear to be systematic in nature and do not necessarily correspond to changes in the shape of the current density profile outside of them. The size of these error bars also reflects the lack of knowledge of the shape of the $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ profile in this location; its curvature is dictated by the strength of the curvature penalisation on each of the spline basis functions, which make up the p' and ff' source profiles, which set prior to the CLISTE run.

4.5 Separatrix location

One of the main findings in the above sections was that the equilibrium is intolerant to shifts of the pressure profile of more than a few millimeters in either direction. Some further analysis is required to determine whether this implies that the equilibrium reconstruction can accurately recover the separatrix loca-

tion, or if this is some other effect coming into play, such as higher order effects on the magnetic signals. Since direct imaging of the plasma separatrix layer is difficult (though it can be accomplished) and time consuming, this section will focus on theory based comparisons. As already mentioned, the separatrix electron temperature at ASDEX Upgrade in Type-I ELMy H-mode plasmas has been determined to be approximately 100 eV[102] for a typical Type-I ELMy H-mode plasma. The location of this point can be found by making a fit to data from Electron cyclotron emission (ECE) data, Thomson scattering (TS) data, or a combination of both. It has also been hypothesised that a sharp kink in the density profile could also mark the location of the separatrix due to increased parallel transport on the open field lines[108]. Since resolving such a kink is also difficult, the location of the maximal curvature in the density profile suffices as an indicator of this change. This was investigated by P.A. Schneider in his PhD thesis[76] and will be discussed below.

The TS diagnostic has the advantage of directly comparing both of these hypotheses. However, at ASDEX Upgrade, the edge TS diagnostic exhibits a large scatter, making it less ideal for this use. As well as this, a standard analysis of the ECE temperatures in the vicinity of the separatrix suffer from the so-called "shine through" effect; due to the low densities here, the plasma is no longer optically thick making the classical analysis of the radiation temperatures invalid. While a method which solves the full radiation transport equation has been deployed at ASDEX Upgrade[54], this is not yet available as a routine analysis and is also time consuming. Ideally, a large number of data points would be analysed and a statistical comparison made.

The principal comparison will therefore be between the (mapped) point of maximal curvature of the density profile and the outer midplane location as determined by the inter-shot generated magnetics-only CLISTE run. Modelled electron density (and temperature, though with the aforementioned quality issues) profiles have been determined by the IDA routine which was introduced in chapter 3. These profiles have been made available for a large number of ASDEX Upgrade discharges from 2012 onwards constructing the electron temperature from the ECE and the electron density from a combination of the LiBES data and line integrated measurements from the DCN interferometer. While it is possible to include prior information on separatrix conditions within IDA, this has not been included as standard in this dataset.

Shown in figure 4.10 is a comparison of the separatrix location produced by this method, by the magnetics-only CLISTE runs, and the function parameter-

isation (FP)[21] result. Each input point is taken from averaged data from 150

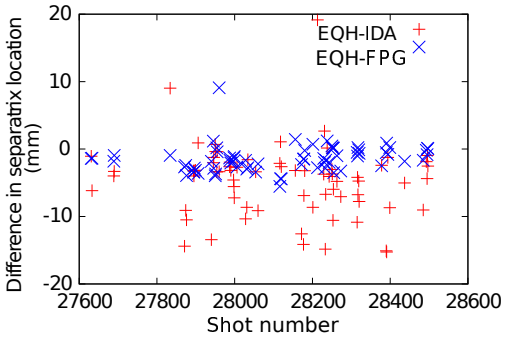


Figure 4.10: Difference in separatrix locations determined by CLISTE and maximal n_e curvature (red), and function parameterisation (blue) with time. No significant trend can be seen in either time or separatrix location

ms either side of the same timepoint (2.3 s) in approximately 200 discharges from the 2012 ASDEX Upgrade experimental campaign. These discharges all featured H-mode behaviour with a lower single null (LSN) plasma configuration. Where ELMs were present, the values were averaged in the time period of -4 ms to -1 ms relative to an ELM crash. In the cases without ELMs, the data were averaged over the whole time range. No discharges featuring non-axisymmetric magnetic perturbations were analysed.

The plot shows a random distribution overlaid on a systematic offset in the location of the separatrix determined via CLISTE and the density profile. This systematic offset, calculated as the mean of the differences, is -5 mm (that is, the separatrix determined from the density profile lies 5 mm outside the CLISTE one) and the distribution has a standard deviation of 6 mm. The mean difference between the CLISTE separatrix location and the FP location is -1.7 mm, with a standard deviation of 2 mm. This systematic error of approximately 5 mm can perhaps be accounted for, however. An investigation into the critical pedestal parameters at ASDEX Upgrade[76] demonstrated that there is a systematic shift between the location of the separatrix determined by $R_{sep,T_e(ECE)} = 100\text{eV}$ and $R_{sep,\nabla^2 n_e}$ of approximately 5 mm. When the two profiles are aligned via the TS diagnostic, the average relative misalignment of the LiBES profile with the ECE profile was determined to be approximately +3 mm. This would then correspond to a systematic shift of the separatrix location (assuming $R_{sep,T_e(ECE)} = 100\text{eV}$ and that this is a simple relative shift valid in all discharges) of -2 mm with an uncertainty of ±6 mm. Since the resolution of the LiBES diagnostic is of the order of 5 mm, this points to a low uncertainty in the separatrix location determined by CLISTE.

Upon examining some of the outlying discharges, it was found that the distribution of the value $R_{sep,\nabla^2 n_e} - R_{sep,CLISTE}$ was somewhat skewed to more positive values (i.e. the distribution featured a heavy tail). The reasons for this heavy tail are unclear, although it may be related to increased radial transport preceding the ELM crash (see chapter 2 for further details of this). It is notable that the general shape of this distribution is very similar to the inter-ELM power deposition profile on the divertor target plates[103].

By taking this heavy tail into account and fitting the distribution of the separatrix differences with a Gaussian and a polynomial, the heavy tail could be taken into account. The peak of the Gaussian was then taken as the expected value of $R_{sep,\nabla^2 n_e} - R_{sep,CLISTE}$ rather than the simple mean. This method was applied to the same database of discharges shown in figure 4.10 and produced the results shown in figure 4.11(a). The mean of this distribution now indicates that

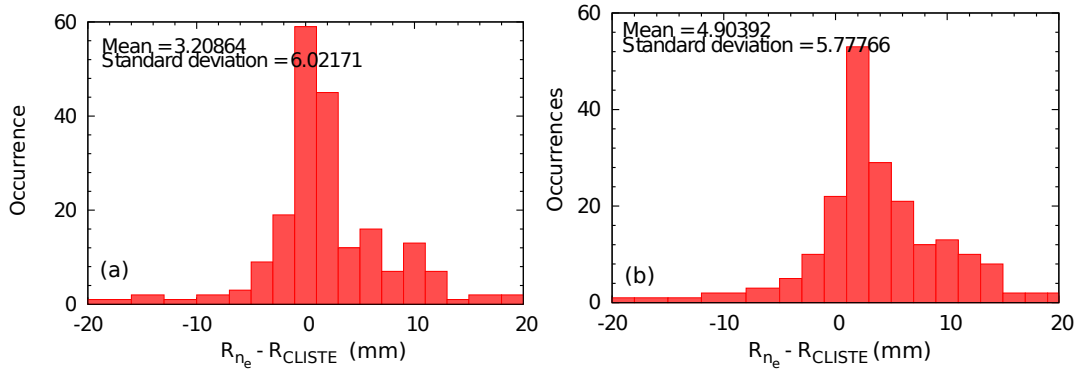


Figure 4.11: (a): improved agreement between $R_{sep,CLISTE}$ and $R_{sep,\nabla^2 n_e}-$, but with the same standard deviation as figure 4.10. (b): the same binning applied to the original data, illustrating the heavy tail of the data. It may be of significance that the mean of the distribution in (a) is at approximately 2-3 mm, given that the power fall-off length is of this order.

the separatrix denoted by $\max(\nabla(n_e)$ lies 2-3 mm outside that determined from CLISTE. The standard deviation is the same in this case, but the total spread is less. In fact, when the same binning method is applied to the (EQH-IDA) values from figure 4.10, shown in in figure 4.11(b), a strong peak in the occurrences is seen at 2 mm, with the Gaussian peak likely to lie at approximately 2-3 mm. This difference between the mean and most-likely locations could also be the reason for the mean outward shift of the LiBES diagnostic against the ECE of 3 mm reported by P.A. Schneider[76].

That the separatrix as determined from the LiBES diagnostic lies 2-3 mm outside that from CLISTE is also not necessarily indicative of a systematic error in either the diagnostic or the equilibrium reconstruction. The SOL power fall-

off length has been measured at both ASDEX Upgrade and JET and has been found to be in the range of 1-3 mm[103]. A realisation of this fall-off length would be that radial transport is enhanced outside which leads to the kink in the density profile (and also in the electron temperature profile). So, while there is not perfect agreement between the separatrix location determined via max($\nabla^2 n_e$) and CLISTE, there are several potential reasons for this. It is difficult to deconvolve the 6 mm standard deviation in these fits into the components due to CLISTE and the LiBES diagnostic.

A more detailed analysis utilising radial sweeps of the plasma to create virtual lines of sight for the LiBES diagnostic should reduce this scatter more, giving a better idea of the uncertainty in the CLISTE determination of the separatrix location. In addition, comparison with the location of T_e = 100 eV could also give more confidence, although this is more difficult due to the large errors in determining the edge electron temperature measurements at ASDEX Upgrade[48] and suffers from the fact that this separatrix temperature can vary from discharge to discharge by 20-30 eV.

4.5.1 Separatrix location using electron cyclotron forward modelling

Accurate high-resolution edge electron temperature profiles can be obtained through use of a Bayesian radiation transport analysis code[54] recently made available at ASDEX Upgrade. Figure 4.12 shows a comparison between electron temperatures determined via the classical analysis of the ECE diagnostic and the electron cyclotron forward modelling (ECFM) results. In the case of discharge #27963,

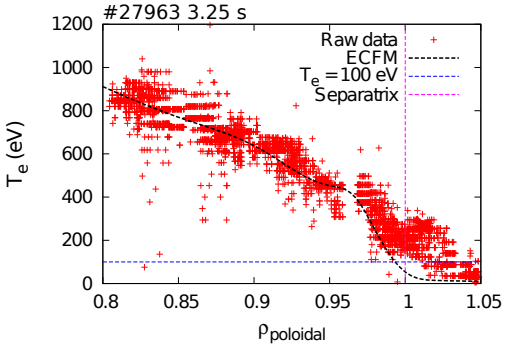


Figure 4.12: ELM synchronised temperature data from the ECE diagnostic (red points) and the ECFM modelled electron temperature (black dashed line) from AUG discharge #27963. Also marked are the separatrix and the $T_e = 100eV$ points.

the LiBES data were shifted by 10 mm inwards in order to obtain the lowest residual fit error in the Bayesian routine. The pre-ELM difference between the separatrix location predicted by the ECFM temperature profile and the generic CLISTE equilibrium is -4.5 ± 0.1 mm. When the edge kinetic pressure is used as a constraint in the CLISTE code, the equilibrium separatrix location for this timepoint moves inwards from 2.144 m to 2.139 m. The density profile, once aligned with the temperature profile, gives a separatrix location of 2.134 ± 0.003 m. The original profile gives a separatrix location of 2.145 ± 0.003 m, which is, incidentally, consistent with the standard magnetic equilibrium. The best fit to the magnetic data is achieved when the pressure profile as a whole (both temperature and density) is shifted inwards by 2 mm. This implies that while good agreement has been found between the LiBES separatrix and that from EQH, the true separatrix could lie radially inside this point. This is difficult to accurately diagnose, as several possible uncertainties are conflated here, notably the physical reasoning, the uncertainty in the equilibrium, and the resolution of the LiBES diagnostic.

The separatrix location determined via the electron temperature should be a more robust quantity compared to that from the density profile. Since the ECE diagnostic measures at the magnetic midplane, when the edge electron temperatures have been accurately reconstructed using the ECFM method there are fewer mapping uncertainties to corrupt the data. The LiBES diagnostic, however, measures 32 cm above the midplane and as such is susceptible to errors in the flux surface shape reconstruction. This could explain some of the scatter in the data, although it is unlikely to account for the systematic offset. In addition, the uncertainty in the separatrix temperature is approximately 20-30 eV, which corresponds to a very small shift due to the high gradients in the pedestal region; an electron temperature gradient of the order of 10^4 keVm^{-1} is typical, so a shift of some tens of eV in one millimeter would not be unusual.

Separatrix movement relative to the ELM crash

Section 3.7 and Dunne et al.[97] showed that the separatrix moves rapidly inwards at an ELM crash and slowly recovers to its original location on a timescale comparable with the profile recovery time. Figure 4.13 plots the movement of the predicted separatrix location from the magnetic equilibrium, $T_e = 100$ eV, and $\max(\nabla^2 n_e)$ relative to their respective pre-ELM locations. What can clearly be seen in the case presented here is that despite slight inconsistencies in the determination of the separatrix location from the various methods, the general trend

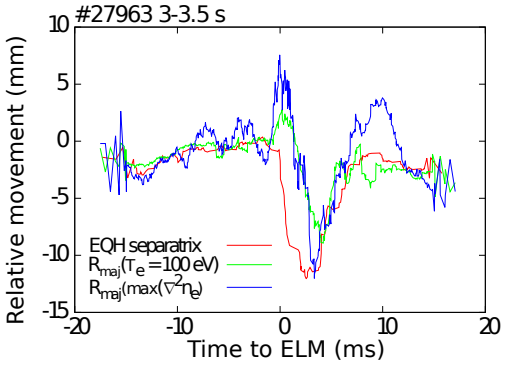


Figure 4.13: Separatrix movement from the three methods presented above plotted relative to the pre-ELM value. All time traces have been smoothed in time to allow easier interpretation of the general trend. Some outlying points have also been removed from the $T_e = 100$ eV points as erroneous and lying far outside the standard distribution.

is the same for all three. The pre-ELM behaviour of the locations is similar, with very little deviation for some time before the ELM crash occurs. The post-ELM behaviour is, however, quite different. The magnetic separatrix moves inwards rapidly, as reported before, but the kinetic profiles are slower to react. The most probable reason for this is a change of profile shape due to the ELM crash; the profiles become more L-mode like. In addition, the ELM expels large amounts of heat and particles, which changes the location of these critical points relative to the separatrix.

Despite the multiple possibilities for uncertainty, the magnitude and direction of the shift is consistent for all three cases, as is the recovery time. All three return to their approximate pre-ELM values in ≈ 7 ms. Interestingly, the $T_e = 100$ eV point is steady throughout the entire ELM cycle (once recovered), while the density location varies substantially. This could, however, simply be due to an increased uncertainty, since this determination depends on the second derivative of the density profile. This would explain the pre-ELM variability well, but the post-recovery overshoot points to a more physics based interpretation, possibly changes in the location of the diffusion kink which is the physical basis for using $\max(\nabla^2 n_e)$ as a separatrix marker.

4.6 Conclusions

In conclusion, we have seen in this chapter that the magnetic measurements are sufficient to localise the pressure profile (via the current density) with an accuracy of ± 2 mm in certain cases. In addition, a minimum pedestal width can

be determined, and, in the case of a fixed separatrix condition, also a maximum pedestal width. The one sigma uncertainties in the CLISTE current density profiles agree well with the one sigma uncertainties in the kinetic data in all three timepoints analysed in this chapter. This indicates that the confidence bands from CLISTE are experimentally relevant when the experimental uncertainties are included as weights for the minimisation of the difference between the CLISTE prediction and the measured values at each iteration of the code.

In order to validate the finding that the pressure profile can be localised to ± 2 mm, it is important to also validate the location of the separatrix as determined by CLISTE. It was found that the generic equilibrium produced for each AUG discharge agrees well with the determination of the separatrix from a kink, or peaked second derivative, in the density profile. However, it was also acknowledged that this density profile is often subject to a shift towards lower values of R_{maj} and that this should be investigated further. On comparing the separatrix location from an advanced CLISTE equilibrium and the location of T_e from an advanced electron temperature profile reconstruction good agreement was found. In addition, the movement of all three ways of determining the separatrix location relative to an ELM crash was also found to be consistent.

Further detailed analysis with high quality edge temperature measurements is necessary to fully investigate this topic and resolve this remaining uncertainty. However, this first investigation would suggest that an advanced equilibrium combining both external magnetic measurements and internal kinetic profiles is required for an accurate separatrix location. Given this, it is then possible to accurately determine the location of the profiles. As an additional step for high quality analysis of the plasma edge, it may be necessary to perform a second iteration by fitting the edge profiles as a consistency check, and to determine if there is a substantial effect on the mapping of, say, the LiBES diagnostic data to the midplane.

Chapter 5

Theoretical edge current density

As described in chapter 3, contributions to the local current density in the confined edge region of a tokamak plasma were considered from three main sources: Ohmic current, Pfirsch-Schlüter current, and bootstrap current[109]. Fast ions, while present in the core of the plasma, were assumed to have a negligible impact on both the pressure and current drive in the pedestal region. Accurately calculating the Ohmic, Pfirsch-Schlüter, and bootstrap contributions requires knowledge of the local plasma magnetic geometry as well as temperature and density conditions. The Ohmic drive also requires a model for the electric field, which was calculated from current diffusion. The constraints on this model and its limitations will be described in section 5.2.1.

5.1 Model description

This section, as well as the data, results, and comments in sections 5.2 and 5.3 are drawn from Dunne et al.[97].

The bootstrap current, so called as it provides a self generated current drive, effectively pulling the plasma to higher levels of efficiency, is a consequence of the magnetic mirror effect in a tokamak, and a density gradient. To first order, the gyro-centres of particles simply follow the field lines from the low field side of the tokamak to the high field side and back again. Now, since there is a gradient in the magnetic field parallel to the field lines, the particles experience a force on the magnetic moment of their orbit:

$$F = -\mu \nabla_{||} B \tag{5.1}$$

with

$$\mu = \frac{\frac{1}{2} m v_{\perp}^2}{B} \tag{5.2}$$

Since μ is more or less constant throughout the particle's orbit, with increasing B, $v_\perp$ must also increase. Through conservation of energy, this implies that the parallel velocity of the particle decreases as it moves along this magnetic field gradient until it reaches zero, at which point it is bounced back along the field line. This happens both above and below the magnetic midplane, creating, in combination with a vertical drift (also caused by the toroidal field gradient and curvature), a poloidal orbit in the shape of a banana, as shown in figure 5.1.

If we then consider two neighbouring banana orbits which share a common flux surface in the middle (figure 5.1), and a density gradient, we can then see that there is an asymmetry on this flux surface in terms of the poloidally flowing current; the banana orbit on the outside has a lower density, but the same velocity distribution as the inner banana orbit. This asymmetry, which is opposite for

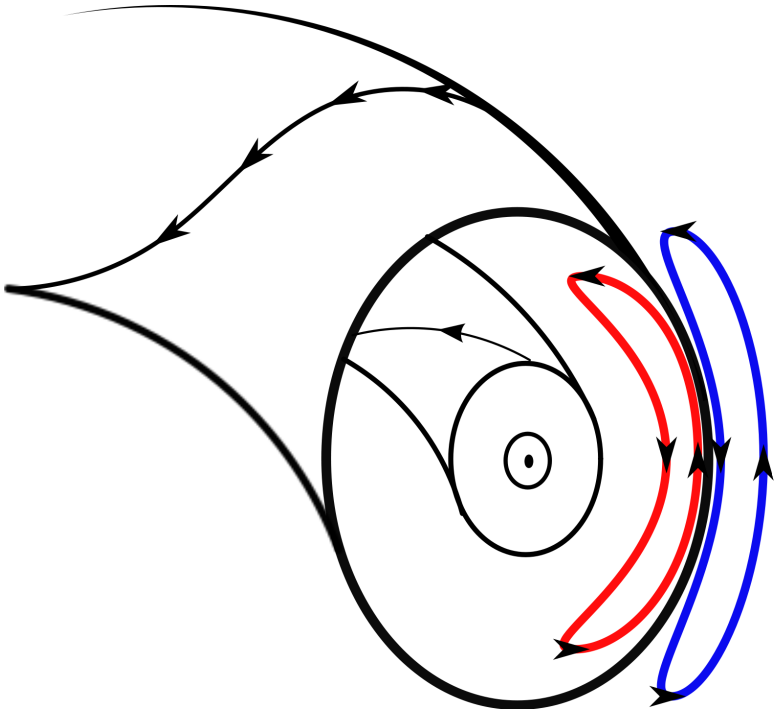


Figure 5.1: Rough sketch of the bootstrap current mechanism. The blue line shows the banana orbit in the region of lower density, while the red shows that of the higher density.

electrons and ions, gives rise to a poloidal banana current, which is then transfered via collisions to a parallel current around the entire torus. The bootstrap current is then the current resulting from the difference in the parallel currents of the electrons and ions.

A general formula to describe the bootstrap current density, referred to here as j_{boot}, for arbitrary collisionality has been derived by Sauter et al.[110, 111]. The equations and definitions used in the calculation of the bootstrap current

and conductivity in the following are taken from Sauter et al. where care has been taken to include the erratum (notably the difference between Z, the main ion charge, and Z_{eff}, the effective charge). Collisionality is a key ingredient to the bootstrap current since the duration of time a particle spends in a trapped orbit is dependent on whether it receives energy from other particles via collisions. In addition, the distribution of the banana current into the bootstrap current also relies on collisions. This formula is dependent on temperature and density gradients in the plasma and is given by

$$\begin{aligned} \langle \mathrm{j_{boot}} \cdot \mathrm{B} \rangle &= f(\psi)p(\psi)\Bigg[\mathcal{L}_{31}\frac{\partial \ln \mathrm{n_e}}{\partial \psi} \\ &+ \mathrm{R_{pe}}(\mathcal{L}_{31}+\mathcal{L}_{32})\frac{\partial \ln \mathrm{T_e}}{\partial \psi} \\ &+ (1-\mathrm{R_{pe}}) \times \left(1+\frac{\mathcal{L}_{34}}{\mathcal{L}_{31}}\alpha\right)\mathcal{L}_{31}\frac{\partial \ln T_i}{\partial \psi}\Bigg] \end{aligned} \tag{5.3}$$

where $\mathcal{L}_{31}, \mathcal{L}_{32}$, and $\mathcal{L}_{34}$ are the collisionality dependent coefficients given by Sauter et al., R_{pe} is the ratio of electron pressure to total pressure (approximately 1/2), and α is also related to the collisionality. $\mathcal{L}_{31}, \mathcal{L}_{32}$, and $\mathcal{L}_{34}$ are also dependent on the trapped particle fraction and the effective charge. The full forms of these coefficients are reproduced in appendix A. This neoclassical approach has previously been validated by Kelliher et al.[112] and Wade et al.[113] in core and edge plasmas respectively. These analyses focussed on the evolution of the toroidal electric field determined by the TRANSP code and a series of equilibrium reconstructions (using the EFIT code, constrained by MSE measurements) in a steady state plasma, respectively. The neoclassically predicted current diffusion was then calculated and found to agree well with the experimental values. The analysis by Wade et al. found poor agreement with the distribution of the edge current density from the EFIT code, which is a quantity of interest for edge stability analysis. In addition, only long timescales were investigated (the flux grid assumed to vary in a linear fashion over 200 ms), which is unsuitable for ELM resolved analysis. Since we have validated the edge current density profile shape and magnitude, it would be of interest to directly compare the current density from CLISTE and neoclassical theory. Therefore, we use a current diffusion model to predict the toroidal electric field while comparing the calculated and measured $\langle j_{||} \cdot B \rangle$.

The Ohmic current is calculated from the toroidal electric field and the parallel conductivity. In a normal plasma, the conductivity can be described by the

Spitzer conductivity, given by:

$$\sigma_{\text{Spitzer}} = 1.9012 \times 10^4 \frac{\mathrm{T_e[eV]}^{3/2}}{\mathrm{Z_{eff} N(Z_{eff})} ln\Lambda_\mathrm{e}} \tag{5.4}$$

where $ln\Lambda_\mathrm{e}$ is the Coulomb logarithm for electrons, defined as

$$ln\Lambda_\mathrm{e} = 31.3 - ln\left(\frac{\sqrt{\mathrm{n_e}}}{\mathrm{T_e}}\right) \tag{5.5}$$

and $\mathrm{N(Z_{eff})}$ is a function of the effective charge, given by:

$$\mathrm{N(Z)} = 0.58 + \frac{0.74}{0.76 + \mathrm{Z_{eff}}} \tag{5.6}$$

Since the resistivity is caused by collisional friction and there is a large trapped particle fraction in the plasma edge, this should be taken into account. Sauter et al. also derived a scaling factor to obtain a neoclassical resistivity which is given by:

$$\frac{\sigma_{\text{neo}}}{\sigma_{\text{Spitzer}}} = 1 - \left(1 + \frac{0.36}{\mathrm{Z_{eff}}}\right) X + \frac{0.59}{\mathrm{Z_{eff}}} X^2 - \frac{0.23}{\mathrm{Z_{eff}}} X^3 \tag{5.7}$$

where X is proportional to the trapped particle fraction and $1/\nu_\mathrm{e}^*$. Due to the dependence on temperature, the conductivity tends to decrease towards the separatrix. At the pedestal top, the conductivity scales with the electron temperature.

To calculate the toroidal electric field, the model from [48] was used. This time dependent model, which is analagous to the well known heat diffusion equation, is described by

$$\frac{\partial \mathrm{E}}{\partial \mathrm{t}} = \frac{1}{\sigma}\left[\frac{\nabla^2 \mathrm{E}}{\mu_0} - \mathrm{E}\frac{\partial \sigma}{\partial \mathrm{t}} - \frac{\partial \mathrm{j_{aux}}}{\partial \mathrm{t}}\right] \tag{5.8}$$

and required a starting assumption for the electric field, E; the loop voltage measured at the vessel wall was taken to be equal at all places during the steady state time before an ELM. It should be noted that a flat voltage profile was not assumed throughout the duration of the simulation, only as a starting approximation. The only auxiliary current, $\mathrm{j_{aux}}$, was taken to be the bootstrap current; as this analysis was concerned only with the pedestal region, this is a reasonable assumption. Using the plasma geometry and experimental temperature and density profiles, a neoclassical conductivity profile, σ, was obtained for each time point. This allowed an Ohmic current density profile, $\langle \mathrm{j_{Ohmic}} \cdot \mathrm{B} \rangle$ to be calculated via[97]

$$\langle \mathrm{j_{Ohmic}} \cdot \mathrm{B} \rangle = \sigma_{\text{neo}} \frac{\mathrm{E}}{\mathrm{B}} \langle \mathrm{B} \rangle^2 \tag{5.9}$$

The collisionality is an important factor in many of these equations, due to its impact on the likelihood that a trapped particle will escape its orbit. The electron collisionality as used by Sauter et al. is defined as

$$\nu_\mathrm{e}^* = 6.921 \times 10^{-18} \frac{\mathrm{qRn_e Z_{eff}} ln\Lambda}{\mathrm{T_e}^2 \epsilon^{3/2}} \tag{5.10}$$

where $\epsilon^{3/2}$ is used as an assumption for the trapped particle fraction. The collisionality increases sharply towards the separatrix as the electron temperature lowers and, more importantly, as q increases towards infinity. A typical pedestal top collisionality range at AUG is between approximately 0.5 and 10 with the main variation being due to the pedestal top electron temperature. Effects on the edge current density due to the collisionality will be described in chapter 6.

In order to make a comparison between the CLISTE output and j_{neo}, the flux surface averaged current density was analysed. The calculation of j_{neo} results in a value of $\langle \mathrm{j}_{||} \cdot \mathrm{B} \rangle$, which is also a standard output of the CLISTE code (given by $\langle j_{||} \cdot \mathrm{B} \rangle = \mathrm{fp}' + \mathrm{f}' \frac{\langle \mathrm{B} \rangle}{\mu_0}$ where p' and ff' are the source profiles from the Grad-Shafranov equation). By using these values, the Pfirsch Schlüter current is also rigorously eliminated in neoclassical transport calculations[109], making the analysis easier as well as allowing a natural extension to peeling-ballooning theory.

5.2 Comparison to CLISTE

A comparison between $\langle j_{boot} \cdot B \rangle$, CLISTE, and $\langle j_{neo} \cdot B \rangle$, given by $\langle j_{boot} \cdot B \rangle + \langle j_{Ohmic} \cdot B \rangle$, is shown in figure 5.2 for discharge #23221. The same reconstructions used in chapter 3 were used for this analysis. The red line shown is the CLISTE

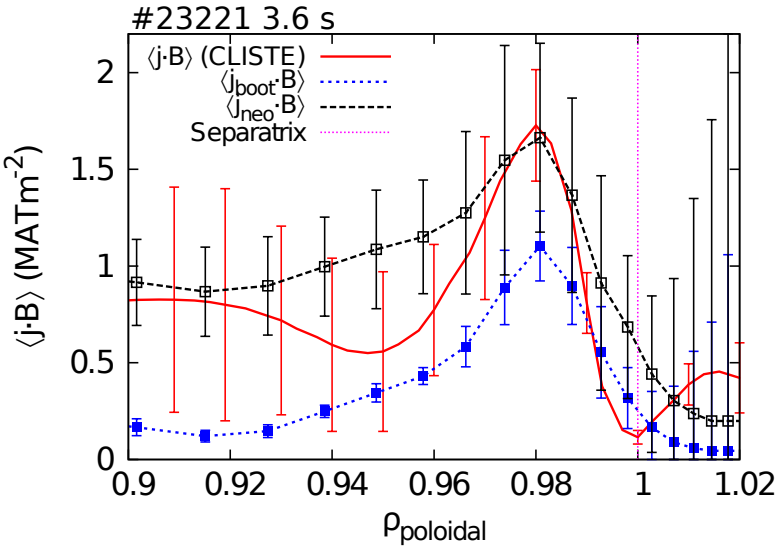


Figure 5.2: Profile of $\langle j_{||} \cdot B \rangle$ taken at $t_{ELM} = -3$ ms. The red line marks the CLISTE output and 1σ confidence bands, the blue is $\langle j_{boot} \cdot B \rangle$, and the black shows $\langle j_{neo} \cdot B \rangle$. The error bars on the latter two were calculated via partial differential error propagation.

output $\langle j_{||} \cdot B \rangle$ with 1σ confidence bands. The blue line shows the calculated $\langle j_{boot} \cdot B \rangle$ and the black line is $\langle j_{neo} \cdot B \rangle$. These profiles were taken 3 ms prior to

the ELM crash. The excellent agreement between $\langle j_{neo} \cdot B \rangle$ and the CLISTE result demonstrates the accuracy of the neoclassical bootstrap prediction, to within the confidence bands from the CLISTE calculation. Also shown are uncertainties for the bootstrap and neoclassical currents, which were propagated via partial differentiation. The total magnitude and location of the edge current peak is well matched between both methods, as is the spatial distribution. The broadening of the calculated current peak that can be seen here has also been noticed, though not systematically analysed, for many other discharges. However, this broadening is within the uncertainties from both methods.

Figure 5.3(a) shows the time evolution of the calculated toroidal electric field at the position of the maximum edge current density. The increase at the ELM is predicted from Faraday's Law in order to conserve the total flux in the plasma, and decays away according to the plasma L/R time calculated for a 2 cm slab at the plasma edge. The red line shown in figure 5.3(b) is the time evolution of the peak $\langle j_{||} \cdot B \rangle$ as output from CLISTE, with associated 1σ confidence bands. The blue line represents the calculated $\langle j_{boot} \cdot B \rangle$ value, and the black shows $\langle j_{neo} \cdot B \rangle$, as in figure 5.2. Error bars have not been included for the time trace of the calculated currents, though since the uncertainties in the kinetic data are similar for all times during the ELM cycle, the error bars can be considered as being similar to the pre-ELM profile. The overall trend agrees with that seen in chapter 3, in that the current density drops sharply at the onset time of the ELM and slowly recovers. There is a delay seen in the recovery of j_{neo} compared to the CLISTE output, despite the increase of the toroidal electric field at the ELM crash; it is possible that this is due to ∇T_i being taken as equal to ∇T_e. The T_i data at the time of this discharge were not of a sufficient temporal resolution to be used in an ELM synchronised fashion. It is also possible that the discrepancy is due to a neglected orbit squeezing modification to the neoclassical resistivity model used here. An analysis using ELM resolved T_i data at a lower time resolution will be shown in section 5.4.

Two things are important to note in this analysis: (i) the current density found by CLISTE has not been constrained by j_{neo} and is therefore independent of it and (ii) while all the input profiles have been smoothed in time, each equilibrium reconstruction is independent; CLISTE is in no way constrained by Faraday's law, meaning the reconstructions are temporally localised. The smoothness of the CLISTE output indicates that there are few systematic uncertainties in the equilibrium reconstruction and that, despite the fact that the equilibrium reconstruction is derived from a nonlinear process, the regularisation employed in

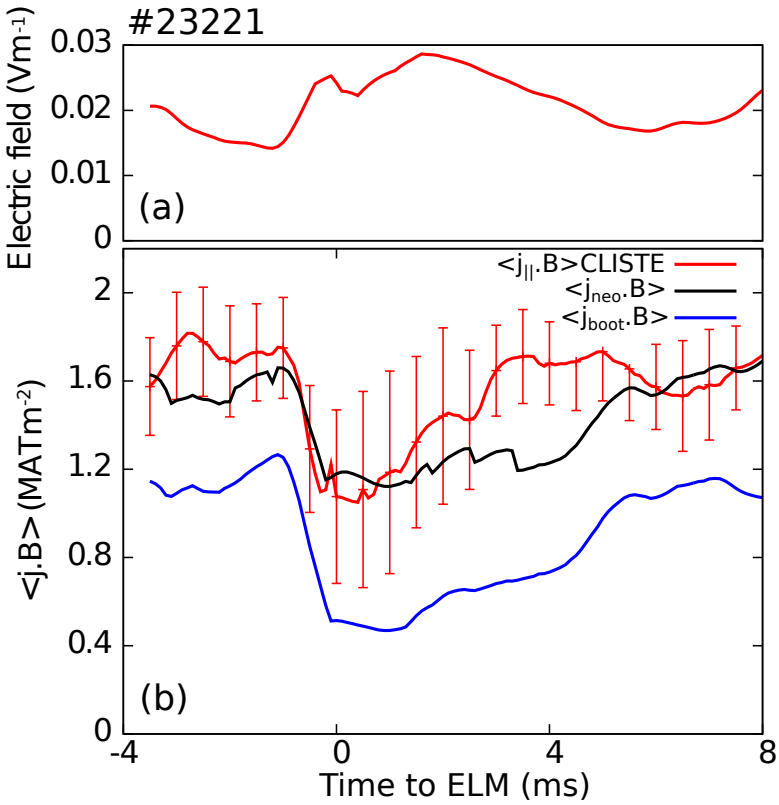


Figure 5.3: (a): electric field calculated at the peak bootstrap position using a resistive current propagation model. (b): $\langle j_{||} \cdot B \rangle$ from CLISTE (red), $\langle j_{boot} \cdot B \rangle$ (blue) and $\langle j_{neo} \cdot B \rangle$ (black).

CLISTE are suitable to produce reasonable results.

5.2.1 Current diffusion boundary conditions

Although the starting assumption for the model was explained, there is still the question of how to deal with the current as it reaches the wall (or, more accurately in this case, the last current carrying surface). It should be constrained in such a way as to give an accurate reproduction of the plasma flux and how the total current changes. Several possibilities were investigated in [48] and it was decided that using the formula $E_{sep} = E_0 - \frac{1}{2\pi R_0} L_{ext} \frac{dI}{dt}$, where E_{sep} is the electric field at the separatrix, E_0 is the initial electric field, and I is the plasma current, to determine the separatrix electric field should be used. It was noted, however, that this gave similar results to simply setting the gradient of the electric field at the boundary to be zero(since the calculation involves a radial derivative, some form of constraint is necessary). The latter was the constraint used in this thesis, which performed rather well. However, an explanation as to why this works was lacking and is given below.

Starting by assuming a circular cross section as the limiting case and taking

the second derivative of the above definition for E_{sep} gives

$$\nabla^2 E_{sep} = \nabla^2 E_0 - \frac{1}{2\pi R_0}\nabla^2(L_{ext}\frac{\partial I}{\partial t}) \tag{5.11}$$

where I is the plasma current. Next, assuming flux surface averaged profiles for the electric field and the current density, ∇^2 I is given by $2\pi(\sigma E + j_{aux})$. Continuing, an expression for $\nabla^2 E$ is given by:

$$\nabla^2 E = \frac{1}{R_0} L_{ext}\left[E\frac{\partial \sigma}{\partial t} + \sigma\frac{\partial E}{\partial t} + \frac{\partial j_{aux}}{\partial t}\right] \tag{5.12}$$

Rearranging this gives an equation resembling equation 5.8:

$$\frac{\partial E}{\partial t} = \frac{1}{\sigma}\left[R_0\frac{\nabla^2 E}{L} - E\frac{\partial \sigma}{\partial t} - \frac{\partial j_{aux}}{\partial t}\right] \tag{5.13}$$

The external inductance for a circular loop is given analytically by

$$L_{ext} = \mu_0 R_0 (ln(\frac{8R_0}{a} - 2) \tag{5.14}$$

which for a typical AUG equilibrium gives an area equivalent radius of 0.7 m with $R_0 = 1.65$ m and hence a value for L_{ext} of $\approx 1.07/R_0\mu_0$. With this in mind, an exact (approximate) expression for the surface electric field is given by:

$$\frac{\partial E}{\partial t} = \frac{1}{\sigma}\left[1.07\frac{\nabla^2 E}{\mu_0} - E\frac{\partial \sigma}{\partial t} - \frac{\partial j_{aux}}{\partial t}\right] \tag{5.15}$$

The only difference to equation 5.8 is a factor of ~1.07 modifying the Laplacian of the electric field. Effectively, the medium aspect ratio of ASDEX Upgrade allows the analysis to be undertaken with the standard equation and includes a built in boundary condition. It would however be important to take this factor into account in, say, spherical tokamaks, or very low or high aspect ratio devices. This factor also becomes significant at higher major radius. For example, in JET it is ≈ 2.8 and in Iter it will be ≈ 3.

5.3 Fuelling study

Because of the dependence of the j_{boot} forefactors on collisionality, it is expected that increasing the density should have a large impact on the total bootstrap current. This can be checked by revisiting the discharges first discussed in section 3.8. Table 5.1 lists the fuelling rate, pedestal collisionality, and ELM frequency of these discharges. Several differences were seen between discharges #23226, #23221 and #23225 in terms of the local current density. These differences in both radial profiles and temporal evolution can now be assessed from the point of view of the bootstrap current, which is the main drive of the parallel edge current density.

Table 5.1: Parameters used in the fuelling study; fuelling rate, collisionality variance (calculated using the formula provided by Sauter et al.[110, 111]), and resulting ELM frequency.

Discharge #	Fuelling (s^{-1})	ν^* range	f_{ELM} (Hz)
23226	0	$0.8 - 2.8$	104 ± 36
23221	5×10^{21}	$1.5 - 4.2$	125 ± 24
23225	9×10^{21}	$3.5 - 6.0$	80 ± 17

The value of the density and temperature bootstrap coefficients at the point of peak bootstrap current as a function of collisionality is shown for each timepoint in the three discharges in figure 5.4. These coefficients are calculated for the density and temperature drives from equation 5.3, i.e. the plot shows the values of $\mathcal{L}_{31}(\mathrm{n_e}$, red), $\mathrm{R_{pe}}(\mathcal{L}_{31} + \mathcal{L}_{32})(\mathrm{T_e}$, blue), and (1-$\mathrm{R_{pe}}$) $\times (1 + \frac{\mathcal{L}_{34}}{\mathcal{L}_{31}}\alpha)\mathcal{L}_{31}(\mathrm{T_i}$, black). The collisionality range for each of the discharges is indicated by arrows, and the separate forefactors are colour coded. The density forefactor decreases strongly

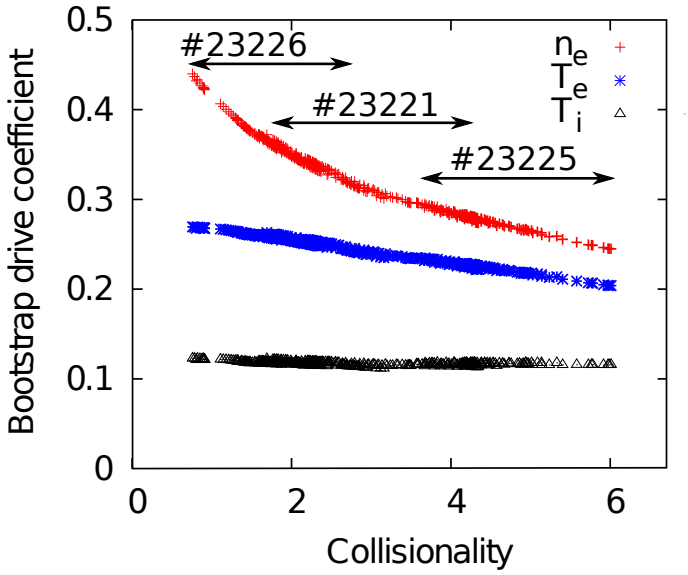


Figure 5.4: Coefficients for n_e, T_e and T_i drives, calculated over the large range of collisionalities presented in these three discharges. The approximate range of collisionalities over the ELM cycle for each of the three discharges is indicated. A clear drop in the density gradient term can be seen at higher collisionalities. This is also seen, though to a lesser extent, for the electron temperature gradient term.

with increasing collisionality (~45% between $\nu^* = 1$ and $\nu^* = 6$), meaning a drop in the efficiency of the electron density drive. In order to maintain the same overall current drive, a larger density gradient must be present. Since the collisionality was increased through fuelling, this could be the case here. Although the electron

temperature gradient drive is also affected by this drop, it is less significant (~25% between $\nu^* = 1$ and $\nu^* = 6$). The ion temperature contribution is not more than half that of the electron temperature contribution, indicating that, for the same gradient lengths in all three components, the ion temperature would form 15-20% of the total bootstrap current.

The difference between in the recovery of the components of the bootstrap drive (electron density (a), electron temperature (b), and ion temperature (c)) are plotted in figure 5.5 Corresponding to the slow recovery of the electron tem-

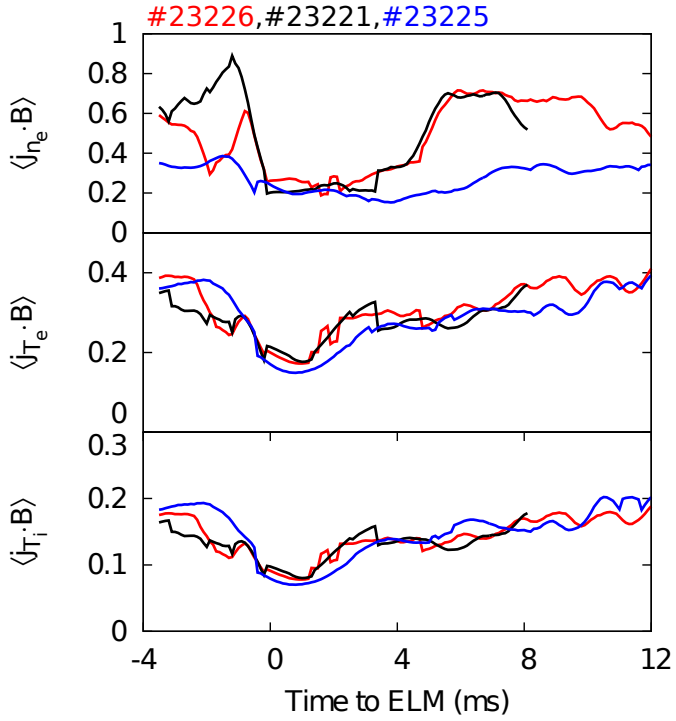


Figure 5.5: Comparison of the three bootstrap component drives for discharges #23226, #23221, #23225. (a): n_e drive, showing a clear drop for the highly fuelled case of #23225. (b): T_e drive, which remains more or less constant, though recovery rates differ. (c): T_i drive, which contributes only 15-20% of the total bootstrap current.

perature gradient reported for discharge #23225 by Burckhart et al. [48], the electron temperature drive of the bootstrap current also recovers slowly, as can be seen in figure 5.5(b). As expected from the forefactors calculated above, the density gradient term is typically dominant, though becomes less so in the highly fuelled case when it is of the order of the electron temperature gradient term. From these discharges, it appears that there is an operational point where the total bootstrap current can be peaked, above which increased fuelling simply damps the bootstrap current, possibly by lowering the trapped particle fraction to such a point that the saturated gradient can no longer drive the same current.

The ion temperature gradient term contributes ~$15-20\%$ to the total bootstrap current, though may play a more significant role in overall recovery if its recovery rate were different to that of the electron temperature gradient.

The accuracy of the neoclassical prediction was also checked for the additional two cases presented here. Figure 5.6 shows the evolution of the edge peak $\langle \boldsymbol{j}\cdot\boldsymbol{B}\rangle$ from CLISTE (red), the bootstrap current (blue) and the total neoclassical current density (black). Again, the toroidal electric field was calculated via the

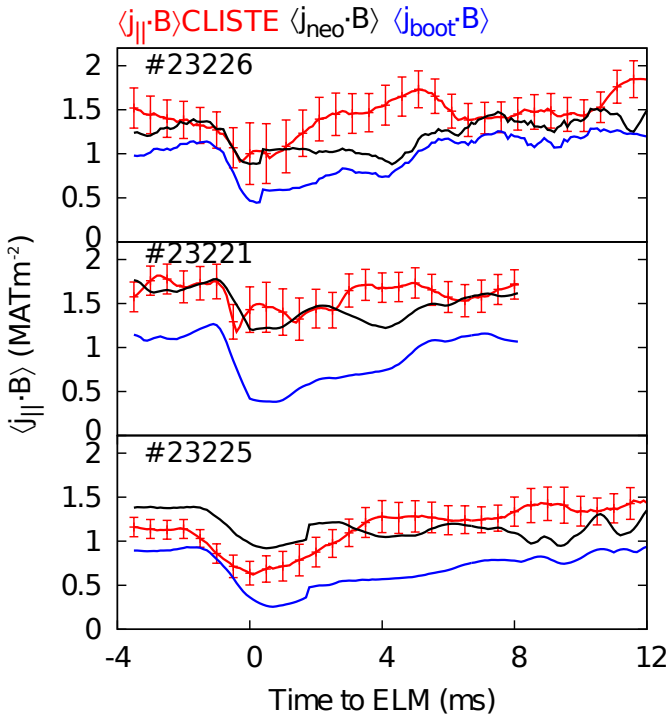


Figure 5.6: Time evolution of the CLISTE $\langle \boldsymbol{j}\cdot\boldsymbol{B}\rangle$ (red), with associated confidence bands, the bootstrap current (blue) and the total neoclassical current (black). A generally good agreement between the two approaches can be observed, with a disagreement between the prediction and measurement shortly following the ELM crash evident, as in the case of #23221.

current diffusion model. The agreement in all three cases is excellent throughout most of the ELM cycle. A poorer agreement between the CLISTE prediction and the neoclassical calculation can again be seen for the case of discharge #23226 between 2-5 ms after the ELM crash. This discharge has the lowest collisionality range, implying a possibly larger separation of the electron and ion temperatures in the non steady-state conditions at and just after an ELM crash. This disagreement is not so evident in the case of discharge #23225, which features both a high collisionality and the slowest recovery of the electron temperature gradient. However, there is possibly a disagreement at approximately 8 ms after the ELM crash. These effects must be analysed with newer discharges which have high

resolution edge ion temperature measurements.

In the cases of these discharges, the ELM frequency follows the inverse trend of the density gradient contribution to the bootstrap current drive. It is acknowledged that this is, however, a very small sample size and much further study is needed in this area before any firm conclusions can be drawn. This provided the motivation behind the creation of a large database of ASDEX Upgrade discharges, which is presented in chapter 6. With this database, systematic effects, such as changes in the divertor geometry, in addition to basic physics can be checked.

5.4 Importance of T_i in current density recovery

Sections 5.2 and 5.3 showed the comparison between the evolution of the parallel current density from CLISTE and as calculated from neoclassical theory. The conclusions showed that there was excellent agreement over most of the ELM cycle, apart from timepoints a few milliseconds after the crash. Given that this effect was reduced in the case with highest collisionality, it was hypothesised that this was due to a difference in the electron and ion temperature recovery rates after an ELM crash. This difference was shown in chapter 2 for the well diagnosed discharge #27963, and was taken into account for a second analysis of the current density evolution.

The edge electron temperature and density profiles were created using the ECFM analysis technique[54]. Similar to the analysis presented already in this chapter, the ion temperature was assumed to be equal to the electron temperature as a first assumption. Figure 5.7 shows the evolution of the CLISTE $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ (red, with associated confidence bands). As has been shown several times already, the general trend of rapidly decreasing peak current density at the ELM crash is reproduced. This is then followed by a period of, in this case, rapid recovery, saturation, and final recovery to pre-ELM values approximately 15 ms after the ELM crash. The calculated bootstrap $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ (blue) and neoclassical $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ (black) assuming $T_i = T_e$ are also shown as solid lines. Again, the toroidal electric field was calculated from a current diffusion model. As shown in the fuelling study, the calculated peak edge current density matches the CLISTE prediction well over the majority of the ELM cycle. Deviations are again seen in the rapid recovery phase, and, in this case, also in the slow recovery phase. This latter difference may be specific to discharges featuring such long ELM cycles. ELMs usually occur at AUG at a rate of 50-150 Hz, while this discharge features regular 40 Hz ELMs. However, such deviations between the CLISTE prediction and the calculation can

also be observed at the end of the ELM cycles shown in figure 5.6 for discharges #23225 and #23226 in the fuelling study.

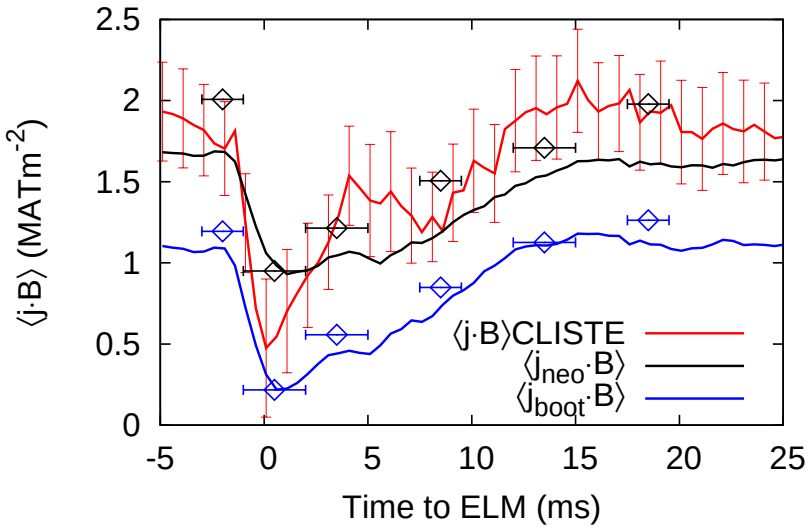


Figure 5.7: $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ from CLISTE (red) over an ELM cycle in discharge #27963. The solid black line shows the neoclassical current density, and the blue shows the bootstrap current density calculated with the assumption that $T_i = T_e$. The diamonds (same colour code) show the respective current densities when the real ion temperature is used in the calculation. The horizontal lines on the boxes indicate the time range the data were taken over. An excellent agreement is observed at the timepoints just after the ELM crash when these points are taken, without any reduction in agreement at other timepoints.

For the next step in this analysis, the edge ion temperature was fitted with a modified tanh function using data from the edge charge exchange systems at several points during the ELM cycle. A CLISTE reconstruction was then made at each of these timepoints using the real ion temperatures. The calculated bootstrap and neoclassical current density from these timepoints are plotted in figure 5.7. The horizontal lines indicate the timerange that the kinetic data were averaged over; since the edge CXRS systems have a typical integration time of 2.1 ms, a relatively wide window was taken. This sacrifices time resolution, but affords a better comparison with the neoclassical current density. The toroidal electric field was assumed to be the same as predicted by the current diffusion analysis from first step. The initial discrepancies seen when T_i is assumed equal to T_e are no longer seen, though the predicted current density is shifted towards higher values at all timepoints. This is very obvious for the first timepoint after the ELM crash, where the new calculated current density agrees much better with the fast initial recovery predicted by CLISTE. The subsequent saturation

and slow recovery to the pre-ELM value are now more accurately predicted by the calculations. The fast recovery of edge ion profiles was also documented by Viezzer et al.[33] in an analysis of the radial electric field. In this work, it was found that the radial electric field is almost fully recovered within 2 ms of the ELM crash and is completely recovered by 4-6 ms after the ELM crash. Further analysis of this effect is required to fully verify the ion temperature recovery, which will be undertaken at AUG in the coming campaigns with faster edge ion temperature measurements.

Additionally, using the experimental ion temperature did not impact the CLISTE result significantly. The only exceptions to this were the timepoints just at the ELM crash (which are challenging to reconstruct due to the strong dynamics in the plasma) and the first post-ELM timepoint, where the ion-temperature constrained case returned a slightly higher peak current density. Both points where, however, still within the uncertainties of the original CLISTE fit. While the pre-ELM timepoints from both sets of results agreed almost perfectly, if a full ELM-cycle evolution analysis with stability codes is desired, it would be preferable to use the experimental ion temperatures. In particular, while the pressure gradient and current density may not be impacted significantly by the addition of the actual ion temperatures, these real data are very important for microinstability analysis[75]. Both the type and growth rate of microinstabilities are strongly affected by the gradient lengths of the ion and electron temperatures.

Chapter 6

Analysis of Type-I ELMy H-mode database

In the previous chapters a method for analysis of the edge current density has been established and its efficacy and applicability over a wide range of cases verified. In addition, neoclassical theory has been shown to describe the edge current density for a number of discharges. It is now possible to undertake a different analysis based on a database of kinetic fits made for the ASDEX Upgrade tokamak. These temperature and density profiles cover a wide range of plasma shapes, currents, and heating powers. Data from the critical phase just before the ELM crash were taken and averaged in a ~2 ms time window. This approach then delivers the plasma parameters at the peak values of pressure gradient and current density. This chapter will focus on the analysis of results from the reconstruction of approximately 40 ASDEX Upgrade equilibria and the comparison of these reconstructions with neoclassical theory.

6.1 Input data

The kinetic data presented in this chapter were taken from a database created to analyse the edge pedestal at AUG. As such, it covers a wide range of plasma parameters and makes use of the high spatially and temporally resolved diagnostics at the plasma edge. The range of engineering parameters and pedestal top values is given in table 6.1. Results based on an in-depth analysis of this database are presented by Schneider et al.[114]. One of the key conclusions in this paper is that, at AUG, the pedestal width is reasonably constant over the parameter range studied, meaning that the peak pedestal gradient is linearly correlated with the pedestal top height. Another finding was that the density pedestal has a smaller

Table 6.1: Engineering and pedestal top parameters for database. Indicated are the range of values, the mean value, and the standard deviation used in this database.

	Range	mean	σ
I_p (MA)	0.6–1.1	1.0	0.1
B_T (T)	-2.2– -2.8	-2.49	0.07
q_{95}	3.8–7.1	4.9	0.6
δ_{upper}	0.28–0.51	0.41	0.06
δ_{lower}	0.28–0.40	0.33	0.03
$T_{e,ped}$ (keV)	0.2–0.94	0.45	0.15
$n_{e,ped}$ (10^{19} m^{-3})	3.4–7.9	6.0	1.0
$p_{e,ped}$ (kPa)	1.9–7.7	4.5	1.2

width than the temperature pedestal, which has implications for the bootstrap current. This will be analysed in section 6.3.2 and was already mentioned in passing in chapter 5.

6.1.1 Kinetic profiles

The kinetic profiles were created using an mtanh function to fit electron temperature and density data, and, where available, edge ion temperature data. The diagnostics used are the same as described in section 3.1, with all fits being made by hand thus ensuring the highest possible accuracy of the output profiles. Care was also taken to align each of the profiles using the Thomson scattering diagnostic, which, despite its lower time resolution could be used to good effect for this approach. While there is an inherent symmetry bias in the mtanh fit function around the mid-point of the pedestal, the fit is still acceptable within the tolerances of the data as well as the equilibrium solution, as shown in chapter 4. These inherent errors were analysed in detail by P.A. Schneider in his thesis[76] and could be important for characterisation of the critical parameters. Nevertheless, a smoothly varying pressure profile is required for the reconstruction and the mtanh represents the best results from the presently available methods.

6.1.2 Magnetic data

Magnetic data were taken subject to the same averaging conditions as the kinetic data. All time points in the final $\approx$ 2 ms of the ELM cycle were averaged to create a single timepoint for analysis. The same procedure was performed

on the poloidal scrape off layer current measurements. By taking data from several campaigns, a divertor condition based analysis of differences between the measurements and the CLISTE prediction is possible. Unfortunately, no data from the pre-Tungsten stages of ASDEX Upgrade, i.e. during operation with a carbon wall, were included in this analysis.

6.1.3 Fitting

As in previous chapters, the CLISTE fits were performed via error minimisation. In most cases, a separatrix temperature of 100 eV was used to determine the shift of the pressure profile used by CLISTE. In some cases, this method produced unacceptably high residuals in the magnetic data; the profile location was scanned over a large range in these cases and errors were minimised, resulting in a separatrix temperature between 60 and 100 eV. The fitted temperature profiles mapped to the advanced equilibrium are shown in figure 6.1(a). The vertical line

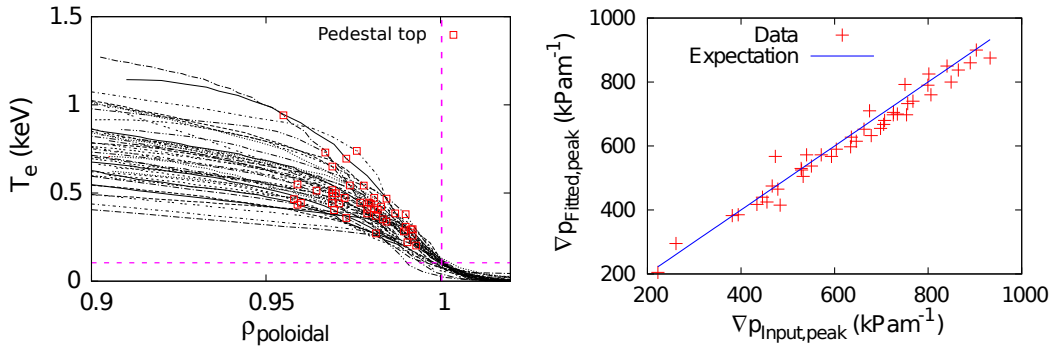


Figure 6.1: (a): T_e profiles used as input for the database. The separatrix location (vertical line) and T_e = 100 eV are shown in purple, demonstrating the generally good CLISTE fit achieved with this value. The red boxes indicate the temperature pedestal top. (b): maximal fitted pressure gradient against the maximal input pressure gradient, with the expectation line shown in blue.

marks the separatrix location, while the horizontal line indicates T_e = 100 eV. In addition, the temperature pedestal top values are also indicated. As can be seen, there is generally good agreement, with some outlying cases. The mean separatrix temperature of these profiles is 100 eV with a standard deviation of 25 eV. It is hypothesised that divertor cooling from unusually high SOL radiation or a large cool density blob, as is observed to form when the plasma nears power detachment conditions, causes the outlying low separatrix temperatures due to the limitations of the tanh fitting method. In some of these cases, the $I_{pol,sol}$ measurements were noted to be large, which is indicative of a large temperature

difference between inner and outer divertors, a large conductivity in the SOL, or a combination of both. This is reflected in a large fitted SOL current, leading to a "bump" in the typically smooth decay of the flux surface averaged current density into the SOL.

Shown in figure 6.1(b) is the plot of fitted maximal edge pressure gradient against maximal input pressure gradient. The blue line indicates the value of the maximal input pressure gradient. In most cases the fitted CLISTE value is slightly lower than the input value; this is due mainly to the curvature constraints in CLISTE which seek to minimise the total curvature. However, the rms value of this difference is 33 kPam^{-1}, which corresponds to a 5.9% error in the fitted pressure gradients. For detailed stability analysis this uncertainty is likely to be unacceptable. However, it can also be seen in this plot that there are several gradients which have a larger deviation. When the three largest residuals are removed, the rms error reduces to 4.5%. These values correspond to an average pressure profiles fit error of 5.4%; this is for all fitted input pressure data points.

6.2 Time resolved equilibrium uncertainty

Since the data were taken over the course of several years, this analysis allows any changes of the fitting errors to be charted over time. Since each CLISTE fit returns an r.m.s. error of the pressure fit as well as the tangential and normal poloidal field measurements, in addition to an overall fit and the convergence error, we can accurately chart any changes that occur. Figure 6.2 charts the progression of errors by shot number where the 22xxx campaign corresponds to 2006 and the 27/28xxx campaign corresponds to 2012. There is a sudden sharp jump in the fitting error of the $\Delta\psi$ coils in the 2011 campaign (shot number > 26000) with no corresponding decrease in the other parameters. This is reflected in a relatively modest increase in the total fitting error for the same discharges. There are a number of possible reasons for this, including a shift in the phenomenology of the ELMs; the ELMs in this campaign lasted for a longer time. This means that instead of a 1 ms rise and subsequent exponential fall off of heat deposited in the divertor the ELMs exhibited a quasi-flattop behaviour which lasted approximately 5 ms. While this may not seem like a potential source for error, given that the data were taken prior to the ELM crash and not during it, the issue arises from the fact that the ELMs were now long enough for the plasma control system to react.

As shown in chapters 2, 3, and 4 the plasma experiences several changes

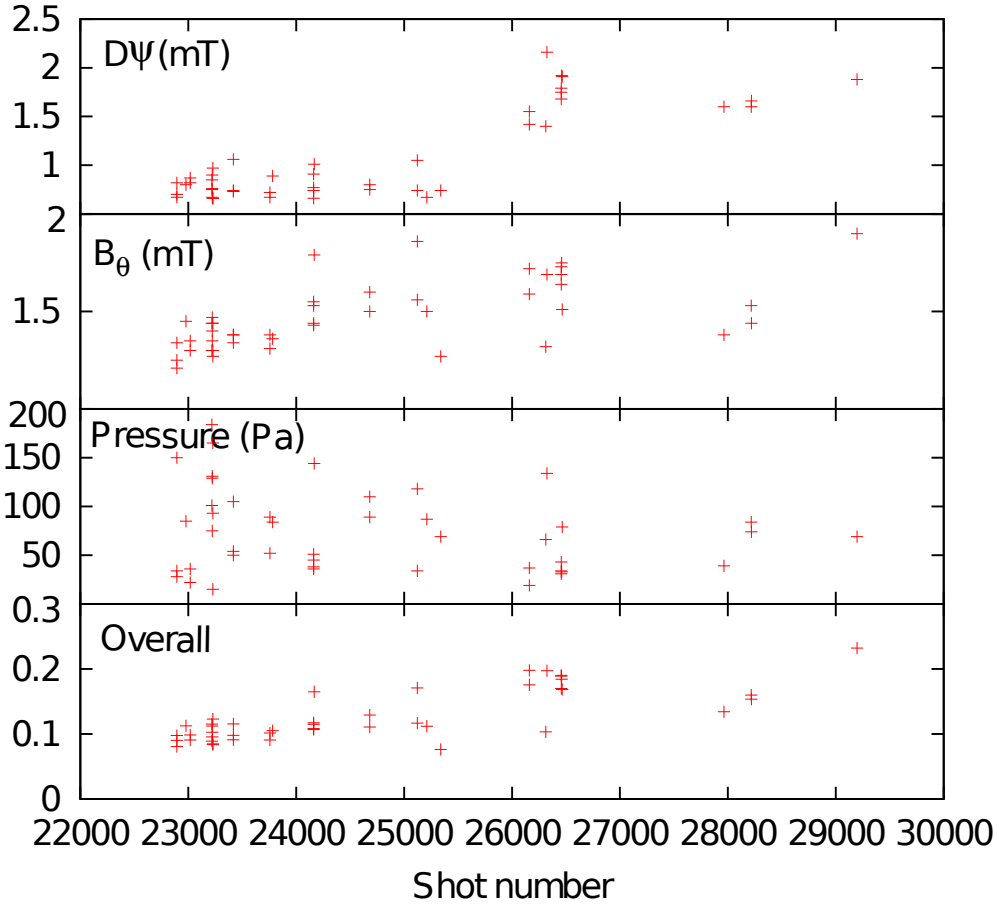


Figure 6.2: Variation of the fitting errors of the $\Delta\psi$ and B_θ coils, the pressure, and the overall statistical fit error with time. In general, there is a scatter about a mean error, although the $\Delta\psi$ error exhibits a sharp jump upwards from the 2011 campaign (shot=26000) onwards.

due to an ELM crash, including an increase of plasma current (which is always controlled) and a general movement of the magnetic axis as well as the separatrix location. The plasma control system rapidly changed the currents in the poloidal field coils, located outside the vacuum vessel, to correct these changes in the plasma. This rapid change induced a skin current in the vessel, which then diffuses with a $\sim$ 5 ms time constant. This skin current affects the magnetic measurements inside the vessel via a current which is not detectable but can only be inferred from detailed analysis. In particular, large deviations of the order of 10-20 mT in the B_θ coils were observed. While several of these large deviations can be filtered out automatically by CLISTE, the quality of the fits to the magnetic data was poor. Changes were made to allow CLISTE extra freedom to handle these currents, and errors were in general reduced. However, the $\Delta\psi$ errors have persisted. It could also be the case that there was a systematic error in the measurements throughout this campaign, although this does not appear to be

the case after the signals have been re-calibrated. Another possibility is that the installation of the external magnetic perturbation coils, or B-coils, as introduced in chapter 2, were installed just prior to this jump in the $\Delta\psi$ residuals. While the coils themselves were not used during any of the discharges analysed in this database, their installation could have perturbed some of the flux loops used for these measurements, changing their geometry and thus causing the observed jump. Analysis of this in-vessel will hopefully be possible in the near future.

Apart from this jump in the $\Delta\psi$ measurements, however, the errors are well within the acceptable range determined from sensitivity studies. It would be desirable to expand the database somewhat to fill in the gaps in time in order to asses whether there was a well defined jump in the $\Delta\psi$ errors and whether they reduced on average with time or not. In particular, the timing of the jump in the $\Delta\psi$ residuals appears to coincide with the installation of the external magnetic perturbation coils (B-coils) at AUG. While the effect of the B-coils on the measurements is small, errors in the diagnostic geometry created by the installation cannot be ruled out.

6.3 Output current density

6.3.1 Dependence on pressure gradient

From the description of the edge current density given in section 3.5, we know that the current density depends on the pressure gradient via the Pfirsch-Schlüter current and, less directly, the bootstrap current. It was seen in section 3.7 that the height of the edge current peak was strongly dependent on the pressure gradient and appeared to have a similar variation in time. Figure 6.3(a) shows that this is also the case over the discharges sampled for this analysis. Indeed, the linear dependence is very strong. This could, however, be expected as the dominant contribution to the local current density at the low-field side is from the Pfirsch-Schlüter currents, which depend linearly on the pressure gradient. To remove the effect of the Pfirsch-Schlüter currents, the flux surface averaged current (defined as $\langle j \cdot B \rangle$) is plotted against the pressure gradient in figure 6.3(b).

As expected, there is still a generally increasing trend with the pressure gradient, although now with a much increased relative scatter (though the absolute scatter is comparable to that in figure 6.3(a)). This relative scatter is due to the effect of collisionality on the bootstrap current, and, to a lesser extent, the effect of temperature on resistivity and hence the Ohmic current drive. It was noted in

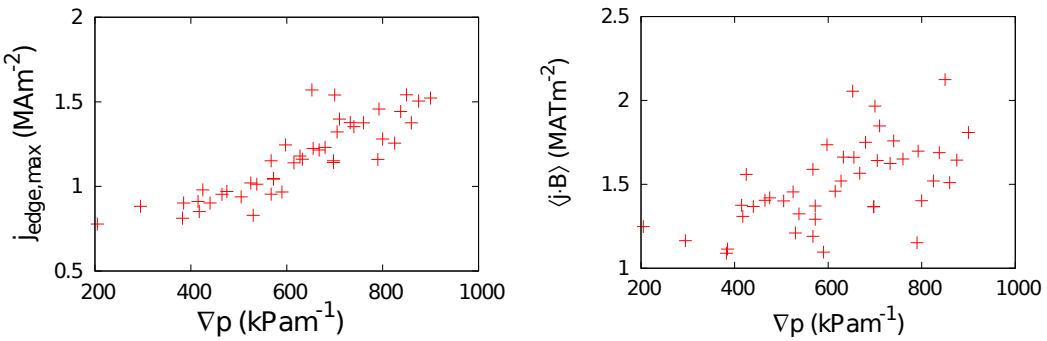


Figure 6.3: (a): peak edge current density as a function of pressure gradient. (b): $\langle j \cdot B \rangle$ against pressure gradient. A very clear linear dependence can be seen in (a) due to the contribution of bootstrap and Pfirsch-Schlüter currents. The offset from zero is due to the Ohmic contribution to the edge current density. This trend disappears in (b) as the Pfirsch-Schlüter current is averaged out.[1]

[97] that it is not clear which current plays a larger role in the overall stability of the plasma; the LFS local current density, or the flux surface averaged value. One of the aims of this database is to determine if either are more dominant in the pre-ELM phase. It can be seen in this figure that the flux surface averaged current density can vary by up to a factor of two for the same pressure gradient, which could be indicative of different ELM behaviour, such as more peeling-like or ballooning-like behaviour of the mode.

In order to asses the plasma stability from an MHD viewpoint, it is also interesting to plot the "operational space" of the plasma in terms of normalised pressure gradient and current density. This is shown in figure 6.4. The edge current density is parameterised by

$$\mathrm{j}_{\mathrm{stability}} = (\langle j \rangle_{edge,max} + \langle j \rangle_{sep})/(2\langle j \rangle_{total}) \tag{6.1}$$

where $\langle j \rangle_{total}$ is the total plasma current divided by the poloidal surface area. The normalised pressure gradient is the well-known[66]

$$\alpha = -2\mu_0 \frac{\partial V}{\partial \psi} \frac{1}{(2\pi)^2} \left(\frac{V}{2\pi^2 R_0} \right)^{1/2} \frac{\partial p}{\partial \psi} \tag{6.2}$$

Interpretation of this figure is not straightforward; it cannot be considered as a quasi time-trace of the progression of an ELM cycle, but rather as a map of the maximum normalised pressure gradients and corresponding current densities allowed by different plasma configurations at AUG. Due to the wide range of plasma parameters covered by this database, it can be considered a good approximation of the operation space at AUG. Investigations into the MHD stability of

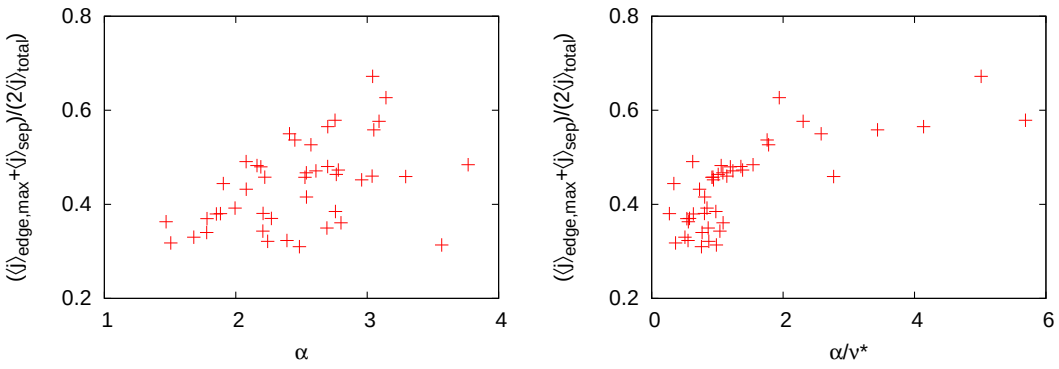


Figure 6.4: (a): normalised edge current density against normalised pressure gradient, α. In this case, the linear dependence of j on ∇p is no longer seen; instead, there is a wide range of pre-ELM current densities and pressure gradients. (b): collisionality included as a controlling parameter, showing reduced scatter, though there appears to be a saturation of j_{norm} with a scatter at lower values.

some of the bounding points on this diagram are currently underway to determine the nature of these points; i.e. whether the stability is determined from a pressure gradient or current density limit, or some combination of both.

The scatter between values of $2 < \alpha < 3$ is of particular interest since this difference in current densities could be due to a collisionality effect; it was shown in chapter 5 that the efficiency of the edge bootstrap current was highly dependent on collisionality. A strong inverse dependence on collisionality is seen, although when both α and $\nu*$ are combined, there appears to be no significant trend. This is shown in figure 6.4(b), where a large spread can be seen at lower values of j_{norm} which then appears to reach a saturated value. If a power law fit is made to the data, the best fit is found to be

$$j_{stability} = 0.449\alpha^{0.2}\nu*^{-0.25} \tag{6.3}$$

with an rmse of 13%. However, this fit also does not substantially reduce the scatter at lower values of j_{norm} and instead averages the saturated values with the scatter, meaning the physics is lost. However, despite the nonlinear trend shown here, it should be noted that the edge bootstrap current peak increases linearly with the real space pressure gradient, meaning that high power discharges are more likely to have a larger fraction of their total current driven by the bootstrap effect.

6.3.2 Theoretical current density

Chapter 5 showed the excellent agreement between CLISTE and the current density derived from neoclassical theory. This database allows us to examine the applicability of this description to a large number of different discharges in the pre-ELM phase. In figure 6.5 the CLISTE output $\langle \boldsymbol{j}_{||} \cdot \boldsymbol{B} \rangle$ is plotted as a function of the calculated value, $\langle \boldsymbol{j}_{\boldsymbol{neo}} \cdot \boldsymbol{B} \rangle$, described in section 5.1. The toroidal electric field for these discharges was determined from the loop voltage measured outside the plasma boundary; since the plasma is in a quasi steady state before an ELM crash, and since figure 5.3(a) showed that the toroidal electric field more or less returns to its pre-ELM value within 10 ms of an ELM crash this is a reasonable assumption. The uncertainties in this figure correspond to the one sigma uncertainties calculated in CLISTE. There is generally good

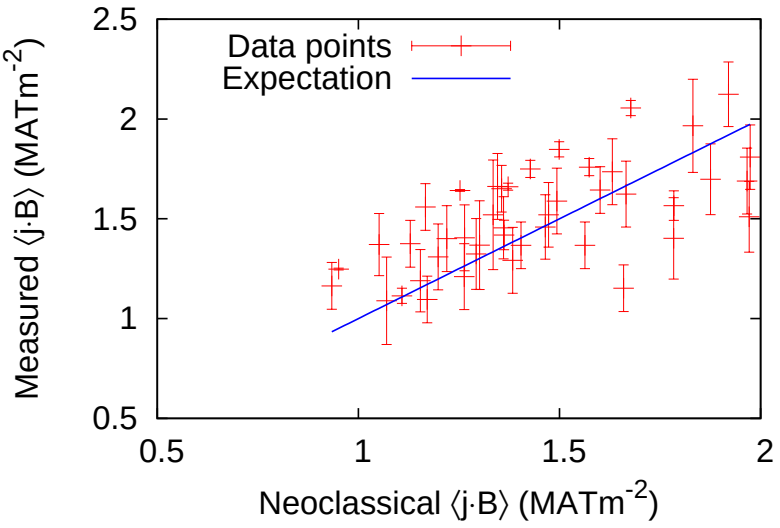


Figure 6.5: $\langle \boldsymbol{j}_{||} \cdot \boldsymbol{B} \rangle$ from CLISTE plotted against $\langle \boldsymbol{j}_{\boldsymbol{neo}} \cdot \boldsymbol{B} \rangle$ with associated confidence intervals. The equality line is drawn for comparison.

agreement, certainly within the uncertainties of the CLISTE reconstructions. In some cases, CLISTE returns a current density which lies substantially away from the calculated current density; it has been noted that a small (~4 mm) shift in the pressure profile can often reduce this difference significantly and that it occurs in discharges made after the start of the 2011 campaign, i.e. when there are greater uncertainties in the $\Delta\psi$ measurements used to determine the profile location. However, the rmse is 12%, indicating that, on avgerage, there is excellent agreement. Ion temperature data were taken into account in the cases featuring low collisionality and high pressure gradient in order to reduce uncertainty.

In addition to testing the specific validity of the theoretical predictions of the edge current density, this database also allows some implications of the theory to be tested. It was shown in chapter 5 that the bootstrap current coefficients depend strongly on the collisionality. As such, one would expect some sort of trend of the edge current density with this parameter. However, it may not be a particularly strong variation as this can be compensated by an increasing the density gradient, which was shown for discharge #23221 in chapter 5. Due to the linear dependence of the collisionality on density, a higher pedestal top density (and hence electron density gradient) is likely to be linked with increased collisionality.

Shown in figure 6.6(a) is an inverse trend of the peak LFS current density with increasing pedestal top collisionality. Figures (b) through (d) show the

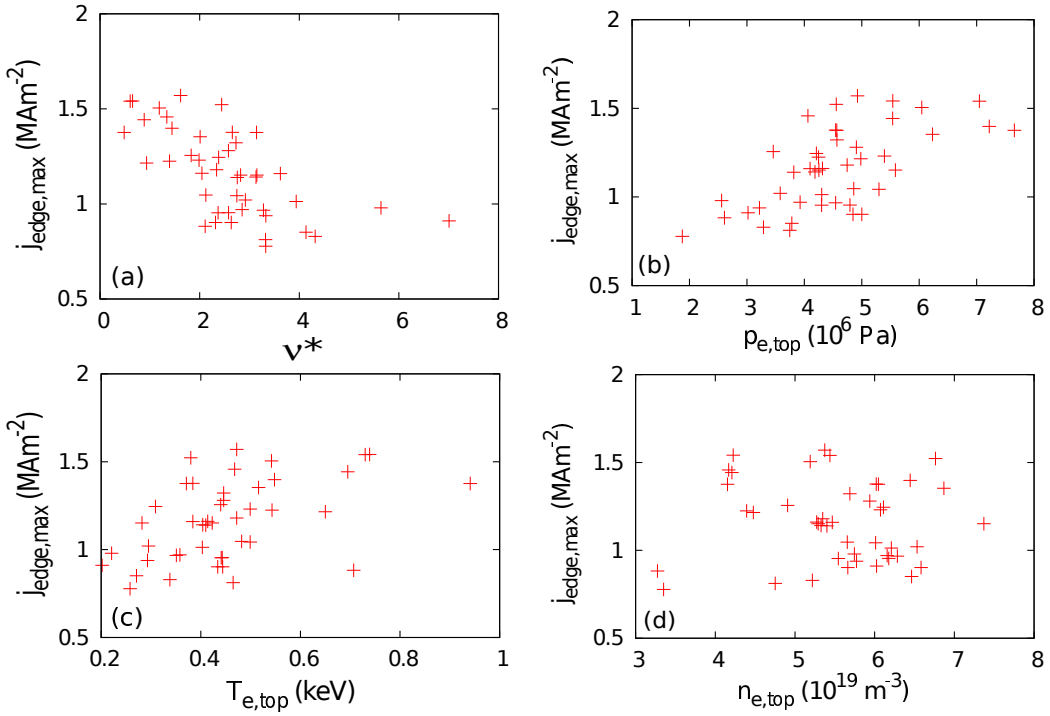


Figure 6.6: Variation of the peak LFS edge current density with pedestal top (a): collisionality, (b): electron pressure, (c): electron temperature, and (d): electron density.

trend with pedestal top electron pressure, temperature, and density, respectively. What is interesting is that the variation in the current density appears to come from a variation in the electron temperature, though table 6.2 shows a stronger dependence of the peak edge current density on the pedestal top electron pressure. That the current density depends on the temperature is is not entirely surprising, as an increased temperature lowers the collisionality. In addition, an increased temperature pedestal top has been correlated with an increased tem-

Table 6.2: Correlation between $\nu*$, $p_{e,top}$, $T_{e,top}$, $n_{e,top}$, and $\max(j_{edge})$. The data shown illustrates the high correlation of the edge current density on the pedestal top temperature, acting via the pressure gradient and collisionality, which are also highly correlated to this value.

	$\nu*$	$p_{e,top}$	$T_{e,top}$	$n_{e,top}$	$\max(j_{edge})$
$\nu*$	1.00	-0.64	-0.80	0.36	-0.64
$p_{e,top}$	-0.64	1.00	0.66	0.21	0.63
$T_{e,top}$	-0.80	0.66	1.00	-0.51	0.49
$n_{e,top}$	0.36	0.21	-0.51	1.00	-0.04
$\max(j_{edge})$	-0.64	0.63	0.49	-0.04	1.00

perature gradient[76] which further increases the electron temperature drive for the bootstrap current.

Table 6.2 shows the correlation matrix between these parameters, demonstrating the high dependence of $\nu*$ and $p_{e,top}$ on $T_{e,top}$. This is then further reflected in the correlation strength between the peak LFS edge j and $T_{e,top}$. This table shows very clearly that there is no trend of the current density with the pedestal top density. However, the electron pressure pedestal top, the combination of both temperature and density, has the strongest correlation with the edge current density, demonstrating that the density does have an important role to play in determining the peak edge current density.

6.4 ELM frequency dependence

It was reported in [97] that the ELM frequency for a limited number of discharges increased with increasing fraction of density gradient driven current. With this large database, it is possible to test the extent of this observation and its applicability to different ELM regimes. The results of this extended correlation are shown in figure 6.7.

This figure shows the normalised ELM frequency, defined as $f_{ELM,norm} = f_{ELM} \times \tau$, where τ is the confinement time, plotted against the fraction of bootstrap current driven by the density gradient relative to that driven by the electron temperature gradient; current density driven by the ion temperature gradient is usually negligible and its efficiency does not vary with collisionality. The plot is divided into three separate areas: the lower denotes discharges from the 2011 experimental campaign onwards; the middle contains discharges after the installation of the full tungsten wall at ASDEX Upgrade and before 2011; and the

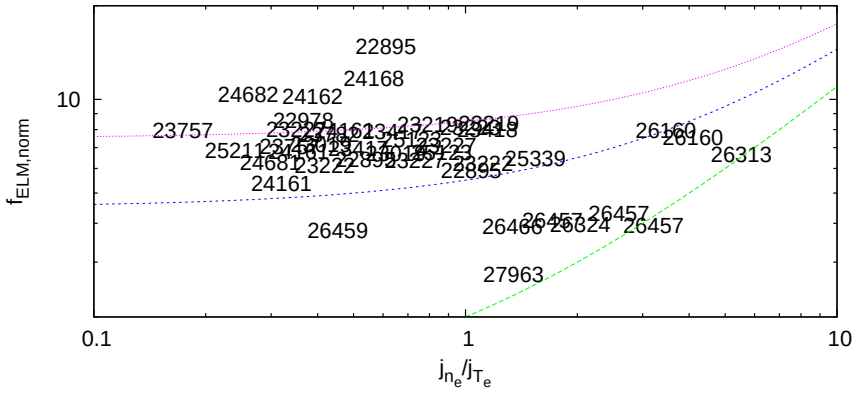


Figure 6.7: Normalised ELM frequency plotted as a function of the fraction of bootstrap current driven by the electron density gradient. The dashed lines distinguish different segments of the plot within which the trend is more or less linear. The lower part of the plot shows discharges from the 2011 experimental campaign onwards, the middle discharges before this, and the upper several "special" types of discharges, such as improved H-modes, N_2 seeded discharges, and W erosion studies.

upper shows slightly esoteric discharges from this time frame, e.g. N_2 seeded discharges, improved H-modes, W erosion studies. No clear trend can be seen in any of the bands as there is a large scatter in the ELM frequency. The finding in [97] is thus likely to be coincidental and that a scaling for the ELM frequency, if indeed one exists, depends on other parameters.

Regarding the separation of the experimental campaigns, it was noted throughout the 2011 experimental campaign that the ELM crash itself was much prolonged for a large number of discharges. A classical ELM crash lasts of the order of a millisecond or so, with the fall off time of the divertor current a few milliseconds longer. However, the 2011 campaign saw ELM lengths (not the ELM cycle length, but quasi flattop behaviour of the SOL current and heat deposition) of 5-6 ms on a routine basis. These had also been noticed towards the end of the 2009 campaign but were not common. Similar observations have been made at JET since the installation of the all-metal wall[115]. It was also reported at JET that the time over which the electron temperature pedestal changed increased during these long ELMs, but this does not appear to be the case for these diacharges at AUG. Instead, the profiles remain at a constant low value for a longer period of time after the ELM crash, which then corresponds to a longer heat flux into the divertor, causing these observed prolonged SOL currents. More detailed analysis into these longer ELMs is required before a definitive statement about the losses

during them can be made.

This shift away from classical Type-I ELMs is not easily explained. It is hypothesised here that different divertor conditions are responsible; for example, a different rate of clearing of neutrals or different wall recycling conditions. This could have the effect of lowering the divertor temperatures, thus lowering the amount of power radiated which would also explain the upper level of figure 6.7; a much larger amount of power is radiated in the divertor via N_2 or Tungsten sputtering in these cases. While it has not been analysed here, it is also hypothesised that the higher impurity content associated with a carbon wall should shift the behaviour towards the top of the figures, alongside N_2 seeding. A further description of N_2 seeded discharges will be given in the following chapter. More detailed analysis of the long ELMs is required to fully justify this observation, and also to explain this phenomenology.

6.5 Further observations of ELM frequency variation

A second way to observe this phenomenology is to plot the ELM frequency times the energy released by an ELM against the heating power applied to the plasma. This is shown in figure 6.8(a) for a database of 2012 AUG H-mode plasmas and 6.8(b) for the database presented in this chapter. While largely scattered, there is a very clear trend of increasing $f_{ELM} \times \Delta W_{MHD}$ against the applied heating power. The energy released by an individual ELM (relative to the stored energy) has

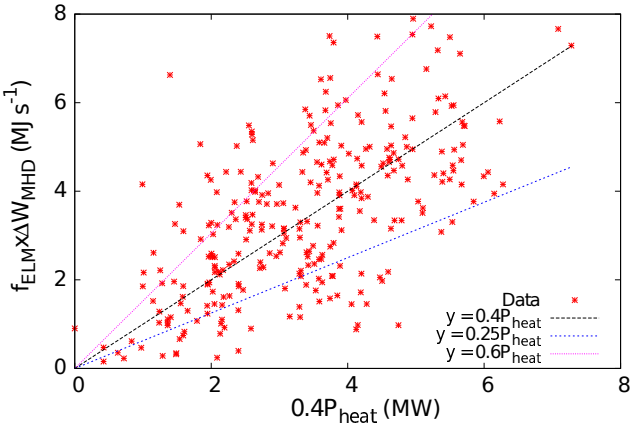


Figure 6.8: Energy lost per ELM times the ELM frequency plotted against the heating power, showing a generally increasing trend, albeit with a large scatter. This scatter could be due to collisionality, density, or q_{95} effects, for example.

been correlated to the Greenwald fraction and, more strongly, to the collisionality,

decreasing as these quantities increase[116, 36]. It is difficult to test these theories with a single machine as the variations of density and collisionality are too small. For AUG, the variation in each is approximately one order of magnitude, which becomes lost in the scatter of such an analysis. As such, in order to make any attempt at this analysis using AUG data, a very large database is required. Since measuring information on the temperature pedestal and collisionality for, in this case, hundreds of discharges, is rather difficult and time consuming, the pedestal top temperature scaling given by Schneider et al. is employed. This is given by

$$\mathrm{T_{e,ped}} = 1.3\mathrm{n_e}^{-1.3}\mathrm{P_{heat}}^{0.7}\mathrm{I_p}^{0.9} \tag{6.4}$$

The density used for the large database was taken from the edge LOS of the DCN interferometer. The collisionality is given by

$$\nu* = 0.0012\frac{\mathrm{q_{95}}R_0^{3/2}\mathrm{Z_{eff}}\mathrm{n_e}}{\mathrm{a}^{1/2}\epsilon\mathrm{T_e}^2} \tag{6.5}$$

The scaling for $\mathrm{T_{e,ped}}$ is plotted against the measured data for the well analysed discharges of the database in figure 6.9(a). The $\approx 13\%$ scatter, corresponding to the rmse of the scaling given by Schneider et al., is also observed here. This will translate into a 25% uncertainty in the calculation of the collisionality, but this trade off is acceptable for an initial analysis. Figure 6.9(b) shows the calculated collisionality for the large database against heating power. Values of collisionality

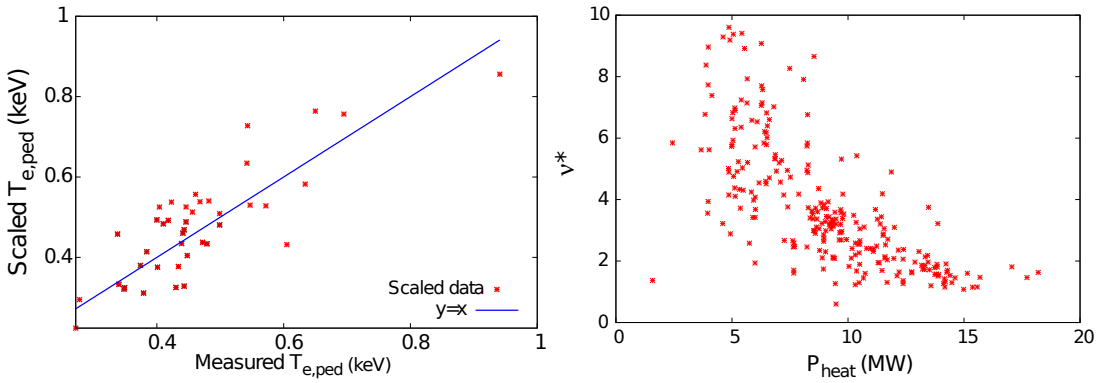


Figure 6.9: (a): comparison of pedestal top temperature values used in this database with the machine parameter scaling from Schneider et al.[76]. (b): projected pedestal top collisionality using the temperature from Schneider et al.[76] and an edge line integrated density measurement.

larger than 10 have been neglected for further usage; this removed only 5% of the total number of points. As a semi-independent test, the ion Larmor radius, $\rho*$, and the pedestal poloidal beta, $\beta_{\mathrm{pol,ped}}$, were also calculated as

$$\rho* = 0.00646\frac{\mathrm{T_{e,ped}}^{0.5}}{\mathrm{aB_{t,ped}}} \tag{6.6}$$

$$\beta_{\mathrm{pol,ped}} \quad = \quad 0.00251 \frac{\mathrm{p_{ped}}}{\langle \mathrm{B_p} \rangle^2} \tag{6.7}$$

where $\langle \mathrm{B_p} \rangle^2 \equiv \mu_0 \mathrm{I_p}/\mathrm{l_{lcfs}}$ and $\mathrm{l_{lcfs}}$ is the plasma circumference, is evaluated at the last closed flux surface instead of locally; it is assumed that the poloidal field does not change much in this region[2] at the midplane. These quantities were then compared to the calculated collisionality, which is shown in figure 6.10. While a

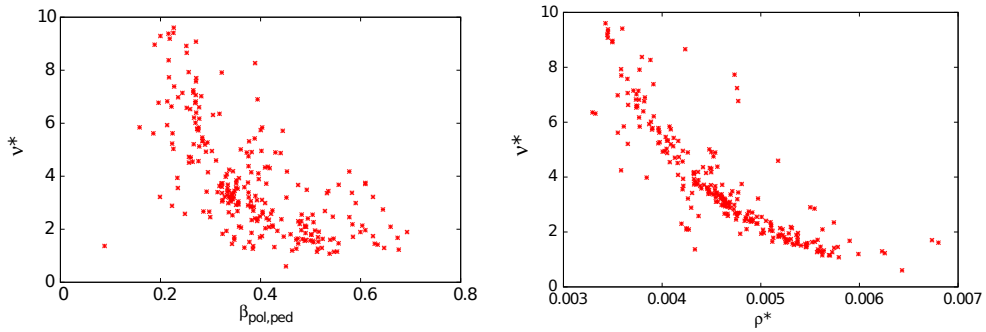


Figure 6.10: Check of the calculation of the pedestal top collisionality by plotting it against the normalised pedestal poloidal beta (which depends on a separate pressure scaling) and the normalised ion gyroradius. The general trends shown in Schneider et al. are reproduced here, although no further specific analysis was undertaken.

detailed analysis has not been made of the trends, they appear to reproduce the findings and general correlations given by Schneider et al.

Including the inverse of collisionality as a controlling parameter in figure 6.8 does not appear to reduce the scatter of the data. In addition, including the inverse Greenwald fraction (a much more robust parameter from this broad approach) also does not reduces this scatter. It would however be extremely beneficial to undertake this analysis using data from several machines in order to correctly determine the effect of very high or low collisionality. This is especially the case since Loarte et al.[36] distinguish two separate types of ELM: conductive (driven by a temperature pedestal collapse) and convective (density driven), which have very different heat and particle loss mechanisms. The differences between these mechanisms have implications for the "per-ELM" energy loss and sustainability of a divertor in larger machines, such as Iter. AUG alone does not provide the required parameter variations to see any changes in the power released per ELM. In addition, it does not appear to be possible to reduce the scatter of the parameter $\mathrm{f_{ELM}} \times \Delta \mathrm{W_{MHD}}$ using either the collisionality or the Greenwald fraction as moderators.

[2]The small change in field line angle over the pedestal was already shown in chapter 3.

6.6 Separatrix movement at ELM crash

It was shown in chapters 3 and 4 that the separatrix contracts at the mid-plane during an ELM crash and subsequently returns to its pre-ELM state on a timescale similar to the recovery of the pressure gradient. This is also backed up by determining the location of the separatrix from electron temperature and density measurements. Use of this database, with its variation in both absolute and relative ELM size and frequency allows analysis of this separatrix movement. In this case, a post-ELM advanced reconstruction was not made as figure 4.13 demonstrated that the generic reconstruction is sufficient to determine the ELM resolved separatrix movement, if not the actual location of the separatrix.

Figure 6.11 shows the normalised decrease in the plasma diameter due to the ELM plotted against the normalised ELM loss (red points). The blue line is a

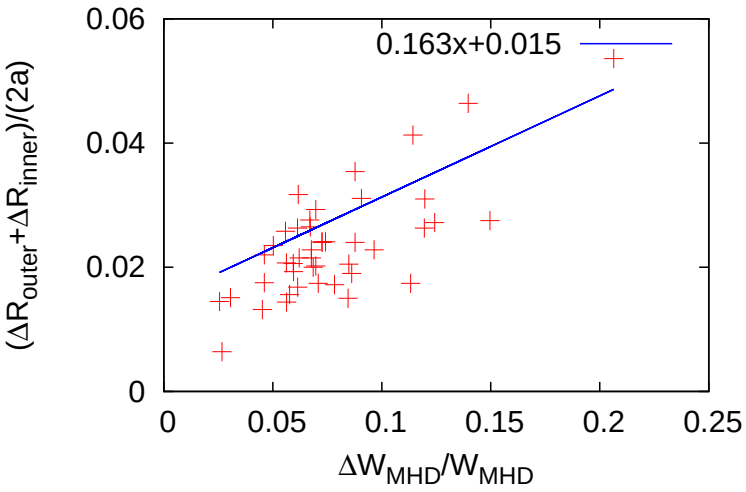


Figure 6.11: Separatrix deflection at the ELM crash, defined as (ΔR_{outer} + ΔR_{inner})/(2*a), against relative ELM losses. A generally increasing trend is found, though there are only few points at higher values.

linear regression to the data, showing a posisitve correlation between the diameter decrease and the size of the ELM. The only remaining question is the driving mechanism behind this effect. Hobirk[117] has suggested that it is caused by an adiabatic shrinking of the plasma at the ELM crash. Dunne et al.[97] posited that the ELM could cause a fixed amount of flux to be separated from the main plasma. This latter explanation was based on a larger movement of the inner separatrix at an ELM crash. This phenomenology could not be reproduced for all discharges in the database, based on the generic equilibrium reconstructions. It was, however, found that the contraction of the outer separatrix location rel-

ative to the inner location increased in magnitude as the pedestal top pressure increased. Based on the findings of Schneider et al.[114], it is known that the edge flux surface compression can be ~15% lower in a generic equilibrium when compared to an advanced reconstruction. As such, neither explanation given above can be discounted based on the evidence obtained to date. Further analysis involving both pre- and post-ELM advanced reconstructions is necessary to reduce the uncertainties involved.

6.7 Magnetic shaping

Schneider et al.[76, 114], among others, have often found that it is necessary to include a "magnetic shaping" parameter, called f_q when fitting the normalised edge pressure gradient. This parameter is defined as

$$f_q = \frac{q_{95}}{q_{\mathrm{cyl}}} \tag{6.8}$$

where q_{cyl} is given by

$$\frac{2\pi a^2}{RI_p}\frac{B_{t_0}}{\mu_0}\frac{1+k^2}{2} \tag{6.9}$$

and q_{95} is the value of q at 95% of the normalised flux. The exact meaning of f_q is not intuitive, but, due to its dependence on q_{95} it can be speculated that it represents a measure for the edge current. In addition to this, it will also have a dependence on the overall shape of the plasma; specifically, its departure from a true ellipse. Working from these assumptions, f_q should depend on the current in the pedestal and the plasma triangularity. Additionally, the current in the pedestal can be characterised by the average current density in the pedestal and the poloidal area which the pedestal defines. This average current density in the pedestal, when normalised to the plasma current, is nothing more than the current density parameterisation described and plotted in figure 6.4. The poloidal area will depend on the plasma radius and triangularity. However, it is not possible to vary the plasma radius significantly at AUG alone, so we can only analyse the triangularity dependence of this parameter.

Performing a power law regression of the upper and lower triangularities and the edge current density parameter against f_q we find an excellent fit, as shown in figure 6.12(a). The rmse error of this fit is 1.5%, and is given by

$$f_q = 1.4 \times ({\delta_{upp}}^{0.09} \times {\delta_{low}}^{0.06} \times {j_{stability}}^{0.04}) \tag{6.10}$$

where $\delta_{o,u}$ are the upper and lower triangularities, respectively. Schneider also

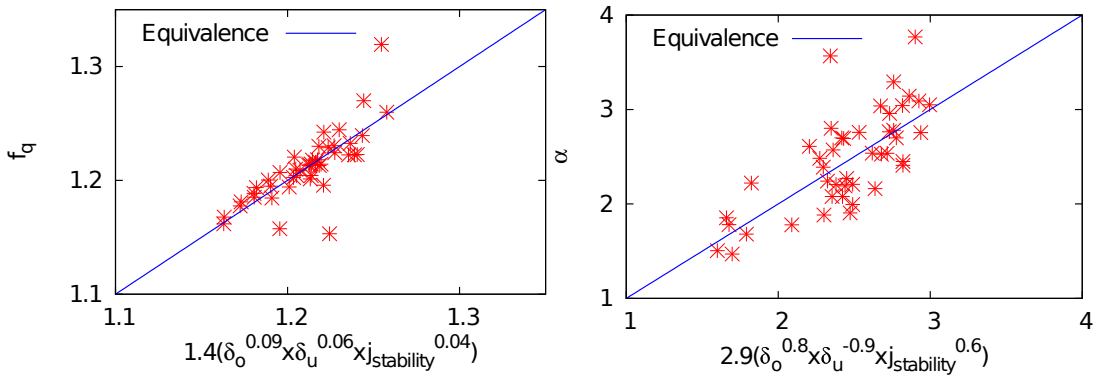


Figure 6.12: (a): f_q plotted against a power law regression to $\delta_{\text{upper,lower}}$ and $j_{\text{stability}}$. (b): α against a separate regression to the same parameters. In the former case, excellent agreement is reached with the linear regression. In the latter, the rmse is 14% and different trends for upper and lower triangularities are observed. Both f_q and α are correlated with increasing edge current density.

found a strong trend of α with f_q, although α was not seen to increase substantially at the AUG values of f_q. However, no trend of α with total triangularity ($(\delta_{upp} + \delta_{low})/2$) was found. The reason for this finding becomes clear when the same regression parameters are used to predict α. In this case, the same positive trend is seen with the upper triangularity, but an inverse relationship with lower triangularity is now apparent. The full fit is given by

$$\alpha = 2.9 \times ({\delta_{upp}}^{0.8} \times {\delta_{low}}^{-0.9} \times {j_{stability}}^{0.6}) \tag{6.11}$$

where alpha is given as the Miller formulation[66] described in equation 6.2. A plot of the data against the fit is shown in figure 6.12(b) and has an rmse of 14%. From this we can infer that a higher normalised pressure gradient can be supported at higher edge current density, higher upper triangularity, and reduced lower triangularity. This fit explains the scatter of $j_{\text{stability}}$ with α in figure 6.4 which could not be explained by a collisionality variation; rather, it is due to the upper and lower plasma triangularity. This is not to say that collisionality does not play a role, simply that no such role can be determined from the present database. It could be that the collisionality sets the edge current density, and that the edge pressure gradient α is determined by the edge current, given by the edge current density and the upper and lower triangularities. Again, data from other machines are required to determine the importance of plasma radius and elongation effects. This effect of increased edge current density on α is also supported by theory as the magnetic shear is reduced by the increased edge current, which stabilises high-n ballooning modes and increases the MHD pressure gradient limit[44]. Synder et al.[44] reported the same increasing trend of normalised

pressure gradient with upper triangularity which is shown in figure 6.12(b).

What cannot be tested using this database of pre-ELM discharges is the causal relationship between the normalised edge current density and pressure gradient. This was briefly examined in section 3.7.2 where a scatter in $j_{stability}$ and α was observed in the pre-ELM phase. For this single discharge, $j_{stability}$ remains constant just before and after the ELM crash with the predominant change being in α, though this is most likely not a finding which can be generalised to all discharges. A similar analysis should be carried out for different values of α but with similar triangularity values to attempt to disentangle these observations. A triangularity scan separating the upper and lower triangularities should also be performed to test if this effects the edge current density parameter or just the pressure gradient. This could then identify if a particular ELM regime has more peeling-like or ballooning-like characteristics and any effects on the ELM losses or frequency resulting from this can also be analysed.

Chapter 7

ELM mitigation scenarios

The bulk of this thesis has, up to this point, dealt only with Type-I ELMs. However, as discussed in the introduction and the previous chapter, there are several other ELM types. Type-II and Type-III ELMs are well classified and have been studied for many years[37, 42, 83]. However, potentially new types of ELMs also exist, notably so-called nitrogen ELMs, seen when N_2 is used to cool the divertor region[16], and the small Type-II-like MHD events seen when external magnetic perturbations (RMPs) are applied to the plasma[42, 43].

All of the above mentioned ELMs occur with a higher frequency, and hence have a lower per-ELM loss compared to Type-I ELMs. This makes them better suited to reactor plasmas, which will not be able to withstand the heavy heat loads from large ELMs. Type-II ELMy plasmas can be operated with high confinement (although less than the corresponding Type-I plasmas), though Type-III ELMy plasmas feature poor confinement. The N_2 seeded ELMs, however, are accompanied by an increase in overall confinement[118], as well as lower steady state temperatures on the divertor plates. Understanding this phenomenon and verifying if we can expect the same phenomenology in reactors is of great importance. The RMP ELMs are again a different phenomenon. When RMPs are applied to the plasma and a density threshold (an edge line integrated density above 0.63 $n_{Greenwald}$[1] for 800 kA and 1 MA plasmas[43]) is reached the ELM size reduces and resembles that of Type-II or high frequency Type-III ELMs. None of these effects are fully understood, though there is much speculation on the nature of the transition to small ELM regimes and many experimental observations linked with the regimes. This chapter will add observations on the edge current density

[1] The Greenwald density, $n_{Greenwald}$ is an experimentally derived scaling law for the maximum density achievable in Ohmic plasmas before the discharge collapses. It is given by the equation $n_{GW} = \frac{I_p}{\pi a^2}$, with n_{GW} having units of $10^{20} m^{-3}$.

to the already observed phenomena and speculate on possible associated changes to the MHD behaviour.

7.1 Type-II ELMs

Type-II ELMs have long been studied in several machines, such as JET, MAST, JT-60U, and AUG[37, 82, 42, 83, 119]. This regime is accessed by beginning with a normal Type-I ELMy plasma and increasing the plasma triangularity, moving it vertically upwards towards a double null plasma (two x-points on the separatrix), and increasing the density. Figure 7.1 shows the transition from Type-I to Type-II ELMs in the edge line integrated density measurements, plasma current, plasma stored energy, and divertor current.

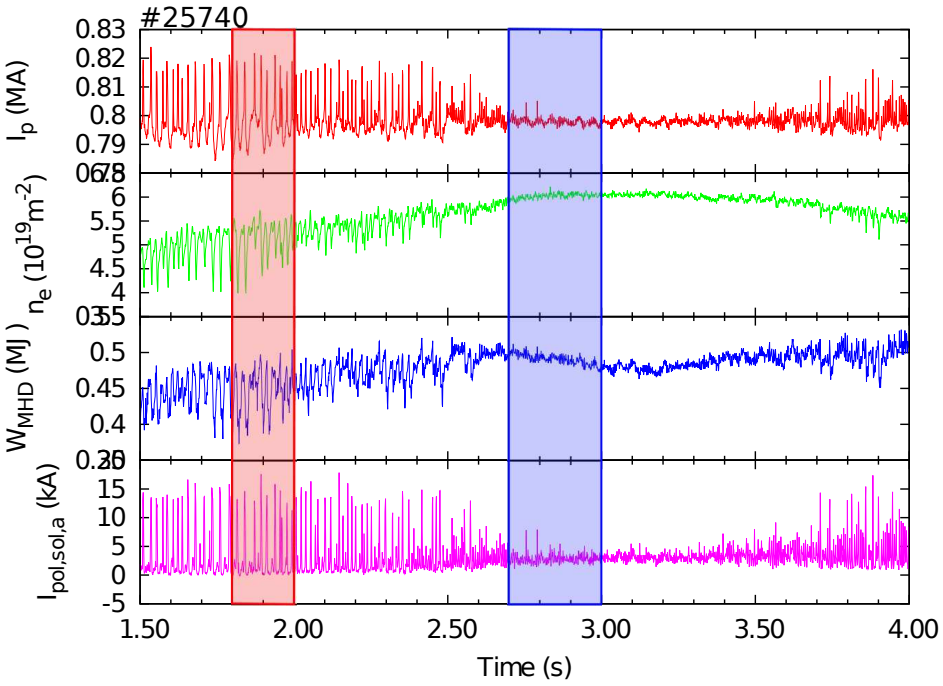


Figure 7.1: Timetraces showing (respectively) the total plasma current, the edge line integrated density, plasma stored energy, and divertor current measured at the outer lower target plate during discharge #25740. In this discharge, a Type-I ELMy phase was initiated and slowly morphed into a Type-II ELMy phase (between approximately 2.5 and 3.5 s).

Type-II ELMs have a much increased frequency compared to Type-I ELMs, and a correspondingly smaller size meaning that the heat loads on the divertor tiles are effectively steady state during this phase. Confinement is good in these discharge types, with H factors of approximately one. The higher density means that H-factors above one are normally not accessible. This is very clear from the

timetrace of the plasma stored energy. W_{MHD} begins to recover when the edge density decreases from around 3 s onwards. An interesting effect of the approach to double null shaping is that current flows through both the upper and lower divertor target plates. Due to the different heat loads at the upper and lower divertors, the baseline $I_{pol,sol}$ measurements rise during the double null phase.

Attempts have been made to classify the change in ELM behaviour via ideal-MHD analysis[84, 82]. In [82] it was postulated that, for JET, the density increase corresponds to a decrease in the edge current density, such that the low magnetic shear induced second stability access is closed off. The plasma is then unstable to n=∞ ballooning modes which have a much lower pressure gradient threshold. To test this hypothesis, current density profiles for the two phases marked in red (Type-I ELMs) and blue (Type-II ELMs) in figure 7.1 were constructed. As input, the electron density (left) and temperature profiles (right) in figure 7.2 were used. The Type-I profile (red) was made with data corresponding to the highest 20^{th} percentile of the edge electron pressure gradient in the time window between 5 and 3 ms prior to an ELM crash, while the Type-II profile (blue) was made using the highest 10^{th} percentile (or decile) in the time range between 2.7 and 3 s in the discharge; the different percentiles are used in order to keep a similar number of points in both cases. From these figures the increase in the pedestal top density can be seen, though it should also be noted that the gradient has not increased. The temperature pedestal top is significantly lowered in the Type-II phase, and the temperature pedestal width has also increased. Both of these effects combine to produce a much lower electron temperature pedestal gradient. Magnetic data were also taken during the same timeranges, though they were not specifically selected to correspond to the highest gradient timepoints; of interest here are the highest gradients an average plasma shape can support before becoming unstable.

LFS edge current density profiles are shown in figure 7.3(a). The colours here correspond to those in figures 7.1 and 7.2. As expected, a significant decrease in the edge current density is clearly visible. In addition, the profile is also broadened slightly, but retains the same separatrix current density, despite increased thermal currents in the divertor. Shown in figure 7.3(b) are the flux surface averaged parallel current densities for the two phases, along with the neoclassical predictions (boxed lines).

Radially inward of the pedestal, the neoclassical prediction differs significantly from the CLISTE result in the Type-I phase. In the pedestal itself ($\rho_{poloidal} > 0.95$), there are still differences between the CLISTE result and the neoclassical

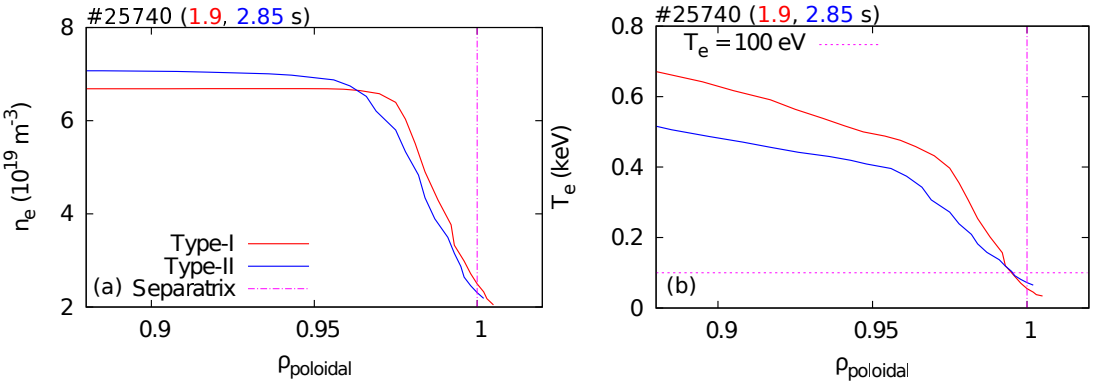


Figure 7.2: (a): electron density profiles for the Type-I ELMy phase (red, 1.8-2 s, ELM synchronised) and the Type-II ELMy phase (blue, 2.7-3 s, highest decile). (b): corresponding electron temperature profiles. When the plasma reaches the Type-II phase, the density gradient remains constant with an increased pedestal top value, while the electron temperature pedestal top becomes lower at a larger width, yielding a much lower gradient.

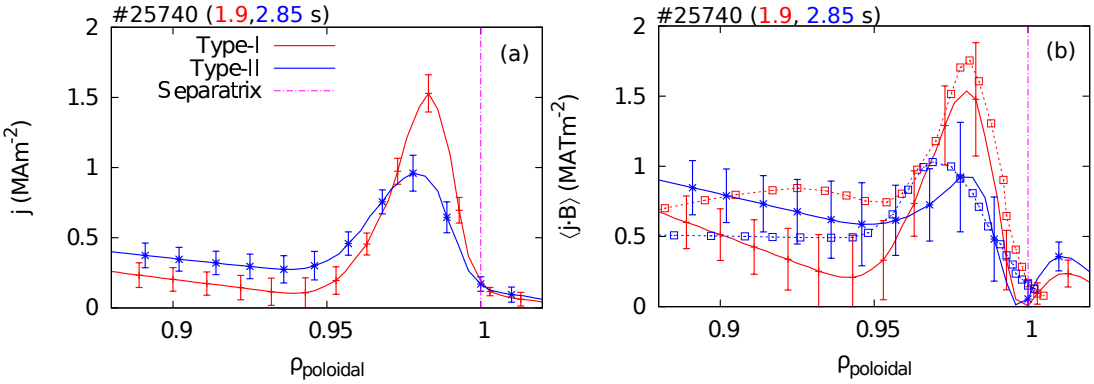


Figure 7.3: (a): LFS edge current density profiles for the Type-I (red) and Type-II (blue) ELMy phases. (b): corresponding $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ profiles with the neoclassical predictions (boxed lines).

calculation, though the calculation is still within the CLISTE uncertainties. The same trend is seen in both; once the plasma is moved to a Type-II ELMy phase, the peak edge current density decreases. The CLISTE result shows no broadening of the profile while the neoclassical prediction shows the peak moving away from the separatrix and the profile broadening. In both cases, the separatrix current density remains constant. Figure 7.4 shows the change in the bootstrap current driven by the electron density (left) and the electron temperature (right) with the same colour coding as the previous plots. Both components are strongly affected by the transition to a Type-II ELMy regime. The electron density term is most likely decreased due to the higher pedestal top collisionality, while the lower electron temperature term corresponds to the wider, lower temperature

pedestal.

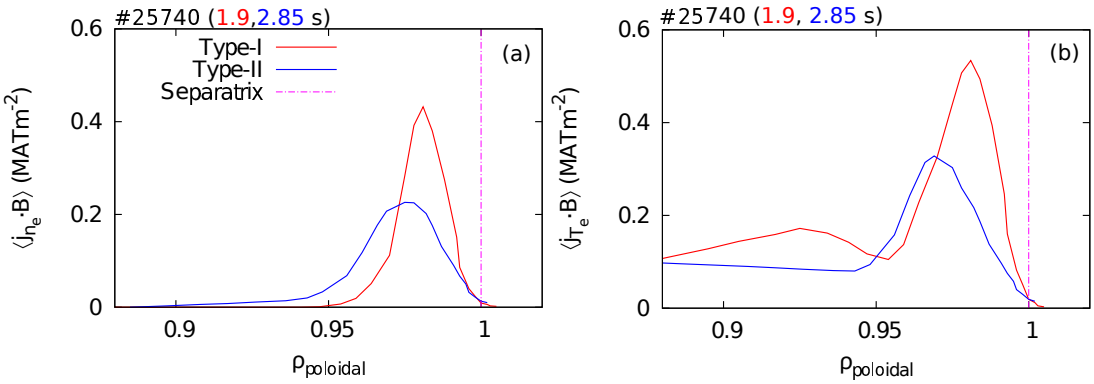


Figure 7.4: (a): electron density and (b): temperature contributions to the pedestal bootstrap current. The decrease in the total edge current density can be attributed to both components, with the density component affected by the higher pedestal top collisionality and the temperature component corresponding to the lower wider electron temperature pedestal in this phase.

The loss of both local edge current density (corresponding to a lowered pressure gradient) and the flux surface averaged current density means that the plasma is likely to be stabilised against ideal peeling-ballooning coupled modes. It is possible that the plasma enters an n=∞ unstable regime, as speculated by Saarelma[82], but this has not been determined conclusively for this AUG discharge. However, modelling with the GENE code by Doerk et al. revealed that a similar Type-II plasma is unstable to microtearing modes[120] in the outer core plasma. Unfortunately, a comparison of the MTM behaviour in the Type-I and Type-II regimes has not been conducted so it is not known if this mode appears only in Type-II phases, or if its characteristics become more pronounced during Type-II phases. Doerk et al.[120] showed that the growth rate of the MTM depends strongly on collisionality, and was sensitive to the local electron temperature gradient. Additionally, it was reported by Boom et al.[121] that broadband high-n (~20) fluctuations of the electron temperature exist just inside the pedestal top in off-midplane locations during the Type-II ELMy phase, which could be indicative of the ballooning instability hypothesised by Saarelema et al.[82].

Magnetic perturbations preceding Type-II ELMs with mode numbers of 3-4 have also been reported[37], which suggests that a large scale mode resembling a coupled peeling-ballooning mode may also be present. Linear stability analysis by Saarelma et al.[84] for an AUG Type-II ELMy discharge also showed that the lower edge current density prohibits access to the second stability regime, meaning

that the ELM is more likely to be triggered by an intermediate-n ballooning mode (n=8 was the maximum studied in that particular analysis), and that modes with a more peeling-like characteristic were eliminated in the Type-II regime. The higher-n mode was also shown to have a narrower radial extent, which could correspond with less plasma loss once the ELM onsets.

The phenomenology of the edge current density presented in this section is consistent with the assumptions made by Saarelma (higher density leading to a lowered bootstrap current). The decrease of the edge pressure gradient and both local and flux surface averaged current densities could indicate a stabilisation of the ideal peeling-ballooning modes thought to eventually trigger a Type-I ELM. Conclusive evidence on the existence of higher-n instabilities cannot be obtained from the current density analysis alone, though the combination of MTMs and pedestal top electron temperature fluctuations would point to a combination of the small scale turbulence driven instability either coupling to, or acting as a second limiting instability with, an intermediate- to high-n ballooning mode.

7.2 N_2 seeded ELMs

The improved H-mode (IPHM) is a scenario that has been under intense investigation [122, 123, 118] as it offers very high confinement properties which are ideal for next generation reactors and devices. However, after the installation of the full tungsten wall at ASDEX Upgrade operation without a Deuterium gas puff was no longer possible, meaning that the same levels of high confinement were no longer possible. A solution was found in the form of N_2 seeding. When N_2 gas is puffed into the divertor, confinement was found to recover to levels seen with carbon plasma facing components[118]. In addition, the ELMs in these discharges were found to decrease in size and increase in frequency. Both of these phenomena are shown in figure 7.5 where a reference discharge (#24681, red) and the N_2 seeded discharge (#24682, blue) are shown. The inset shows a close up of the ELM behaviour during a short time, where the smaller higher-frequency nature of the ELMs is shown.

The improvement in confinement was recently analysed by Tardini et al. [124], which focussed on a gyrokinetic heat transport analysis. It was found that the increased W_{MHD} came from an increase in electron and ion temperatures while the electron density remained the same. These increases stem from a phenomenon known as “stiff profiles”, whereby the pedestal top essentially sets a strong boundary on the rest of the plasma; increasing (decreasing) the pedestal top will force

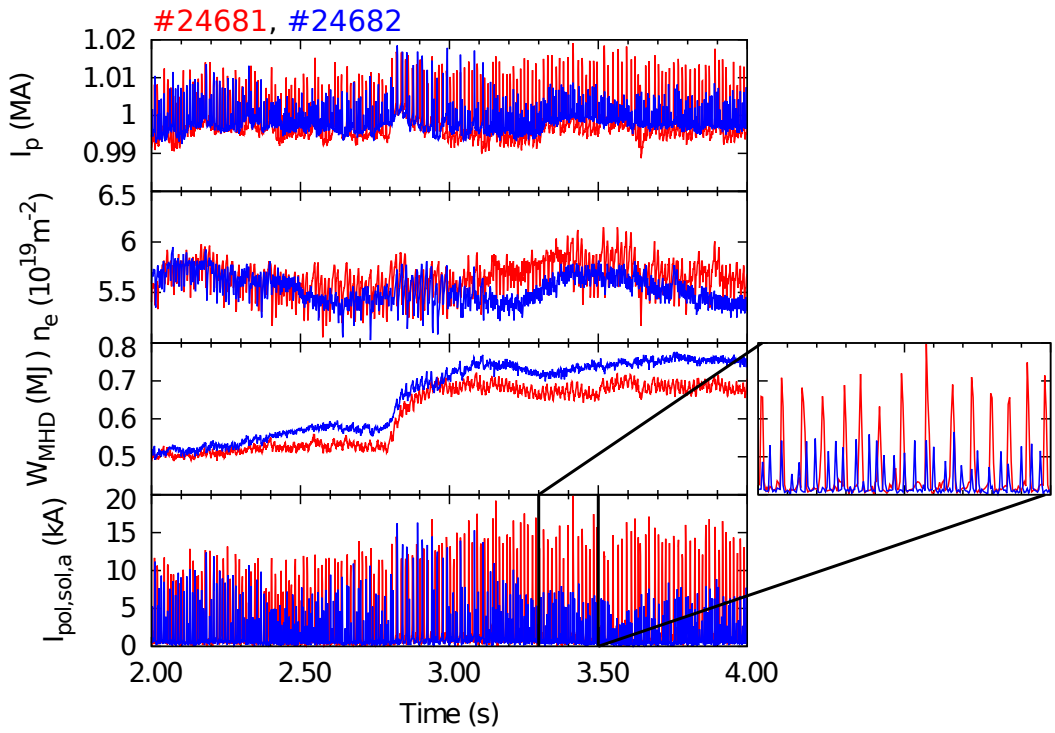


Figure 7.5: Effect of N_2 seeding (blue) on a discharge as compared to a reference (red), which was completed just prior to the discharge of interest. The increase in W_{MHD} is evident, as is the smaller higher-frequency nature of the ELMs, which can be seen in the inset.

a corresponding increase (decrease) in the plasma core. It was found that the gradient lengths of the electron and ion temperature profiles remained at some critical value throughout the plasma core, meaning all confinement improvement is based on changes to the edge pedestal. Figure 7.6 shows the pre- and post-ELM profiles for discharges #24681 (reference, red lines) and #24682 (N_2 seeded comparison, blue lines). Here, pre-ELM means between 3.5 and 1.5 ms prior to the next ELM, while post-ELM varies between the two discharges. In the case of the reference discharge, the post-ELM data were taken between 0 and 2 ms, while in the seeded case between 0 and 1 ms. The first subfigure shows the electron density profiles, which are seen to decrease at the pedestal top when N_2 seeding is applied. The second subfigure shows the increase of the electron temperature pedestal top. These inverse trends lead to the electron pressure being the same in the two discharges.

The smaller ELM size in the N_2 seeded discharge is the principal interest for this work. This smaller size is clearly indicated in figure 7.6, where the dashed lines show the post-ELM profiles in comparison to the pre-ELM profiles as solid lines. Again, the red lines pertain to the reference case, the blue to

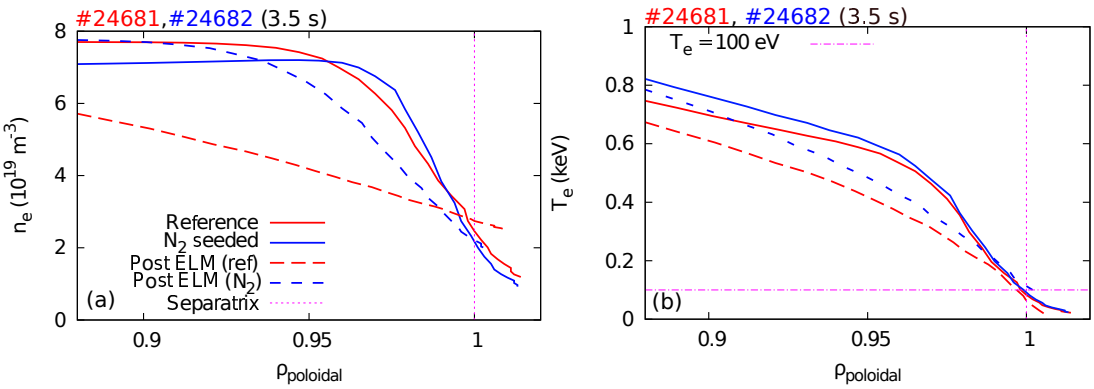


Figure 7.6: Profiles of (a): electron density and (b): electron temperature for the reference (red) and N_2 seeded (blue) discharges. Shown are the pre-ELM (solid) and post-ELM (dashed) profiles for each discharge, indicating in particular the difference in the behaviour of the electron density at the ELM crash.

the seeded case. The temperature pedestal drops by a similar amount in both cases at the ELM crash, but the density pedestal shows only a small reaction in the seeded case. The density has already been shown by Schneider et al.[76] to decrease as far inside as $\rho_{poloidal} = 0.8$ in the reference case, compared to 0.9 in the seeded case. The lower confinement loss can also be seen in the divertor current measurements (used as ELM markers), where the reference case features a larger absolute magnitude of the current as well as a longer duration at its peak value. This is shown in figure 7.7 where the divertor currents in both cases rise to their respective peak values on the same timescale. In the seeded case, however, the current then decreases immediately, indicating improved post-ELM confinement. The reference case features a quasi flattop phase lasting for approximately 2 ms before it decreases, indicating the onset of recovery.

We can now compare the pre- and post-ELM edge current densities to see if there are any significant differences in the profiles or evolution. These profiles were made using the data shown in figure 7.6 as input and assuming $T_i = T_e$. The results from the CLISTE fits are shown in figure 7.8. The left figure shows the pre-ELM profiles of the LFS edge current density for the reference (red) and seeded (blue) cases, while the right hand figure shows the post-ELM profiles. The pre-ELM profiles are similar, with the edge current density peak being higher and closer to the separatrix in the seeded case. The edge current density profile is also narrower in the seeded case. These differences could account for the change in ELM behaviour, though a detailed stability analysis is required to determine the extent of this effect; it is possible that the steeper edge current density gradient could provide a stronger drive to the peeling mode, rather than to a coupled

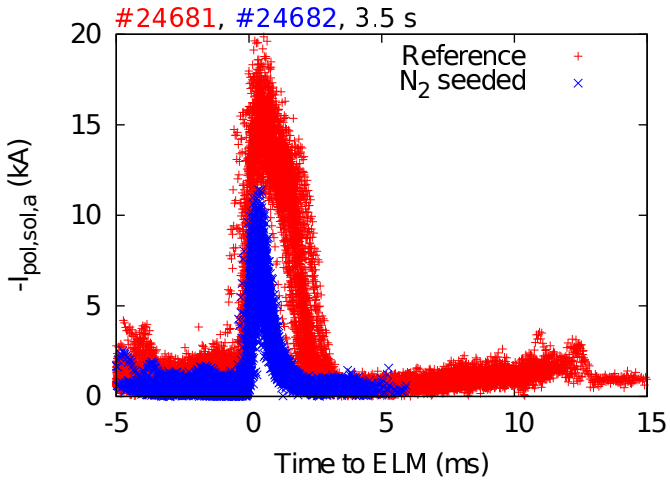


Figure 7.7: ELM synchronised outer divertor currents for the discharge pair showing the striking difference in both the peak magnitude and duration of this peak between the reference discharge (red) and the seeded case (blue). Onset of recovery from the crash, as well as the total ELM cycle, is significantly faster in the seeded case.

peeling-ballooning mode.

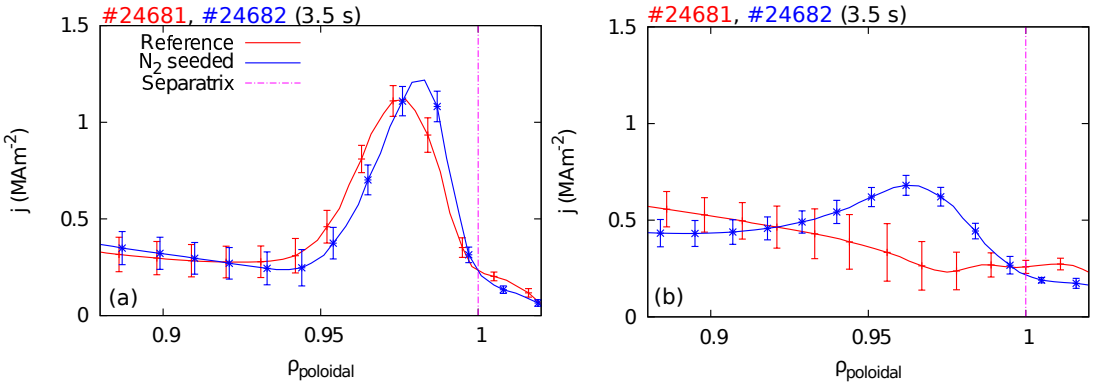


Figure 7.8: Edge current density for (a): pre- and (b): post-ELM timepoints in the reference (red) and seeded discharge (blue). The pre-ELM profiles have a similar basic shape, with the seeded case having a higher peak which is closer to the separatrix while the post-ELM profiles show large differences, reflecting the differences in the crash sizes.

Much more striking is the difference in the current density profiles just after the ELM crash. In the reference case, the edge current density profile has completely flattened, as is typical of a Type-I ELM. However, in the seeded case, there is still a strong edge peak indicating that the plasma has not completely relaxed the free energy available to it for some reason, or that some hypothetical second phase of the crash was somehow prohibited in the nitrogen seeded case. This gives rise to several questions, notably, what could this second phase of the

crash be? The flattop-like divertor current behaviour in the reference discharge corresponds to the phase of fluctuating edge current density at constant low α shown in figure 3.13. Here, it was seen that once the current density and pressure gradient reduce after the ELM crash, the pressure gradient remains fixed while the current density oscillates before both begin to recover. There are some possibilities:

1. There is a density loss channel which does not activate when N_2 seeding is applied; the remaining gradient is still enough to retain a current density peak

2. The finite edge curent density peak acts to stabilise a secondary mode which usually occurs after the initial crash

3. A current density driven mode is not activated when seeding is applied

To evaluate these, it is also necessary to examine the flux surface averaged current density. This is shown in figure 7.9, where the CLISTE $\langle \boldsymbol{j} \cdot \boldsymbol{B} \rangle$ (solid lines) is plotted with the neoclassical predictions (square points). The left hand figure shows the striking observation that the flux surface averaged parallel current density is smaller in the seeded case compared to the reference case. Typically, the LFS peak size is a good indicator for this quantity. However, in this case, the separation could have important implications for the stabilisation of different modes. A strongly flattened (or reversed) local magnetic shear, caused by the strong local edge current density peak could stabilise ballooning modes, while the lower average edge current density is also stabilising for peeling modes.

The question of what causes this change now arises. It was found that the increased Z_{eff} caused by the nitrogen seeding decreased the neoclassical prediction for the edge current density; the excellent agreement seen for the seeded case (both pre- and post-ELM) is only possible with this increased Z_{eff} of 2.5. The decreased pre-ELM peak height for the seeded case and its proximity to the separatrix are reproduced by the neoclassical predictions giving further weight to these findings. Since the pedestal profiles are similar (inside the pedestal top there are differences), it is hypothesised that the change in the flux surface averaged current profiles comes mostly from the increased Z_{eff}, although there is also a change in the density gradient just inside the pedestal top (see figure 7.6).

The post-ELM profiles also show some interesting findings. The reference case shows a complete loss of the edge current density peak, while peaking strongly in the SOL; this corresponds to the larger divertor currents in the reference

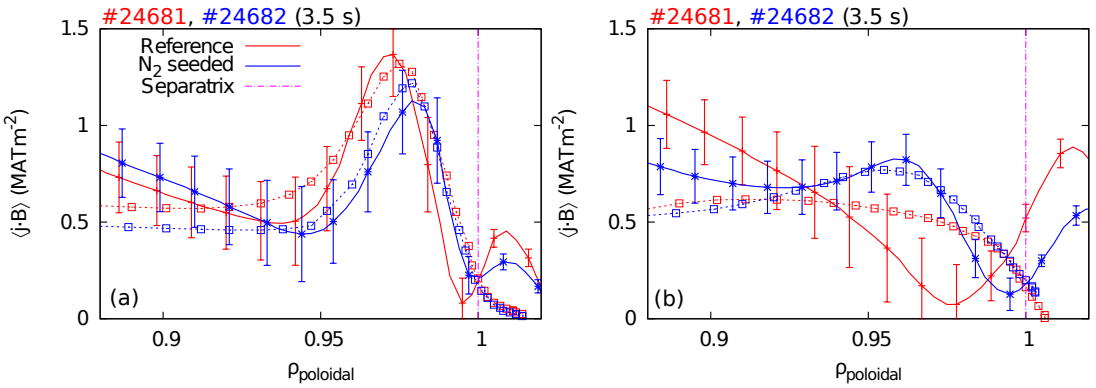


Figure 7.9: Profiles of the flux surface averaged current density for (a):pre- and (b)-post ELM profiles in the reference (red) and seeded (blue) discharges. The solid lines show the CLISTE results while the dashed lines with boxes show the neoclassical prediction. This figure shows the striking finding that the flux surface averaged current density is lower in the nitrogen seeded discharge than in the reference case, contrasting with the larger peak found for the local LFS current density.

case shown in figure 7.7. There is still a clear edge current density peak in the seeded case, similar to the findings for the local current density, which also compares well to the neoclassical prediction. Again, the higher Z_{eff} given by the Bayesian Bremsstrahlung analysis was necessary to obtain good agreement between CLISTE and neoclassics, implying that the impurity content in the pedestal is not decreased substantially by the ELM crash. For the reference case, however, the same quality of agreement cannot be found. The neoclassical prediction also demonstrates the flattening of the edge current density profile. In the two cases presented here, the ELM induced toroidal electric field was not included since only single timepoints relative to the ELM were taken. Instead, the toroidal loop voltage measured outside the plasma was used, which may not be an accurate indication of a complete profile in the plasma. However, it could be the case that neoclassical transport is not valid at the large Type-I ELM crash seen in the reference case. This is perhaps not surprising, given that ELMs expel large amounts of particles, heat, and, importantly, impurities, but implies that the N_2 ELMs are not "complete" in some way. This, in addition with the fact that the density pedestal is relatively unaffected by the ELM could mean that the seeding mechanism bypasses the usual impurity expulsion in some way. A detailed and careful ELM resolved transport simulation of the pedestal could help resolve this issue. In addition, measurements of the impurity density in the pedestal could also reveal the behaviour of confined nitrogen at an ELM crash, though the high

time resolutions required are difficult to achieve with spectroscopic methods.

To summarise, the application of nitrogen seeding in AUG has been found to both improve total confinement and to reduce the size of Type-I ELMs. It was shown that the smaller ELM size is mostly due to a reduction in the collapse of the electron density pedestal and the lack of a "second phase" of the ELM crash. A higher and narrower LFS edge current density peak is also observed when N_2 seeding is applied, which could act to stabilise ballooning modes. This current density also does not completely disappear following an ELM crash, indicating that Type-I ELMs normally consist of an initial loss of confinement which may then destabilise a further mode, or open a second loss channel. The flux surface averaged current density profile was also observed to be narrower but with a *lower* peak in the seeded case, indicating that peeling modes may also be stabilised in this regime. These features are in agreement with neoclassical predictions and partially due to the higher Z_{eff} in the seeded case. The post-ELM profiles of the flux surface averaged current density echoed those of the local current density, with the seeded case showing excellent agreement with the neoclassical prediction, while the reference case deviated significantly. This could imply that the losses in the seeded case proceed according to neoclassical transport while typical Type-I ELMs experience a more rapid loss mechanism, or that the ions have a different temporal behaviour in both cases.

7.3 External perturbations

A more recent attempt at mitigating ELMs has been presented at several tokamaks[40, 41, 125, 43] in the form of external non axis-symmetric magnetic perturbations (MPs). At ASDEX Upgrade, these take the form of 2 rows of 8 so-called B-coils mounted inside the vacuum vessel above and below the midplane. These coils can excite toroidal perturbations with mode numbers 1,2,4, all of which have been used to mitigate ELMs[126]. However, simply applying the perturbations is not enough, at least at ASDEX Upgrade. A density threshold has also been observed[43, 127], below which ELM mitigation has not been achieved. This is in contrast to results from DIII-D where complete ELM suppression has been observed at low density[128] (referred to as low collisionality ELM suppression), although ELM mitigation at higher density has also been observed at DIII-D[41].

This section will describe the edge current in three phases relative to the ELM mitigation: an initial phase where normal Type-I ELMs are present; phases in which Type-I ELMs and smaller ELMs with higher frequency are present

and; phases which feature mitigated ELMs. External perturbations were applied for the duration of the discharge, though ELM mitigation is not achieved until the documented critical density threshold is reached. Where Type-I ELMs are present, data are shown as ELM synchronised profiles, otherwise all data were averaged over the full time window. The kinetic data in the pre-ELM phases were selected such that only the highest decile of the pressure gradient between 4.5 and 1.5 ms prior to the ELM crash was taken as CLISTE input. For the mitigated phases, the highest decile of the entire time slice was taken. In both cases this is justified as being the largest pressure gradient the plasma can be expected to achieve given the average shaping, fuelling, etc., in these time ranges.

The calculation of an axis-symmetric equilibrium cannot be considered as absolutely valid when these perturbations are applied; since the diagnostics used to measure the kinetic profiles and the magnetic coils are not located in the same sector of the torus, there is an inherent mismatch in the location of the magnetic separatrix and the mapping of the profiles. The separatrix deflection due to MPs has been shown by both vacuum field calculations and profile measurements to vary with an amplitude of ~10 mm around the torus. To correct for this, the kinetic profiles have been shifted throughout the duration of the discharge by ~2 mm inwards (this MP induced shift is also convolved with the generic equilibrium separatrix uncertainty previously discussed). The profile location was chosen via the minimisation of the residuals of both the kinetic data and the magnetic measurements. Where possible within the 2 mm uncertainty of the equilibrium reconstruction, the density profiles from each timeslice outside the peak gradient have been aligned.

The discharge analysed here is #28847, which featured an 800 kA plasma current, -2.5 T toroidal magnetic field, and 5 MW of external heating. A Deuterium gas puff was applied to the plasma, increasing from 1×10^{22} s^{-1} to 2×10^{22} s^{-1} between 1 and 5 s, with a further ramp up to 4×10^{22} s^{-1} by 7 s. An n=2 resonant magnetic perturbation was applied between 1.5 and 7.5 s during the discharge. The timetraces of plasma current, edge line integrated electron density, W_{MHD}, and the SOL currents are shown in figure 7.10. As the density rises, a thinning out of the large ELMs is evident, which eventually gives way to a patchy grassy ELM regime from approximately 4.1 s onwards with the final Type-I ELM occurring at 5 s, which is concurrent with the application of the large divertor gas puff. It can also be seen from this figure that the plasma stored energy decreases as the gas puff is further increased from this point. However, the peak plasma stored energy does appear to increase before this large puff rate is applied, implying that

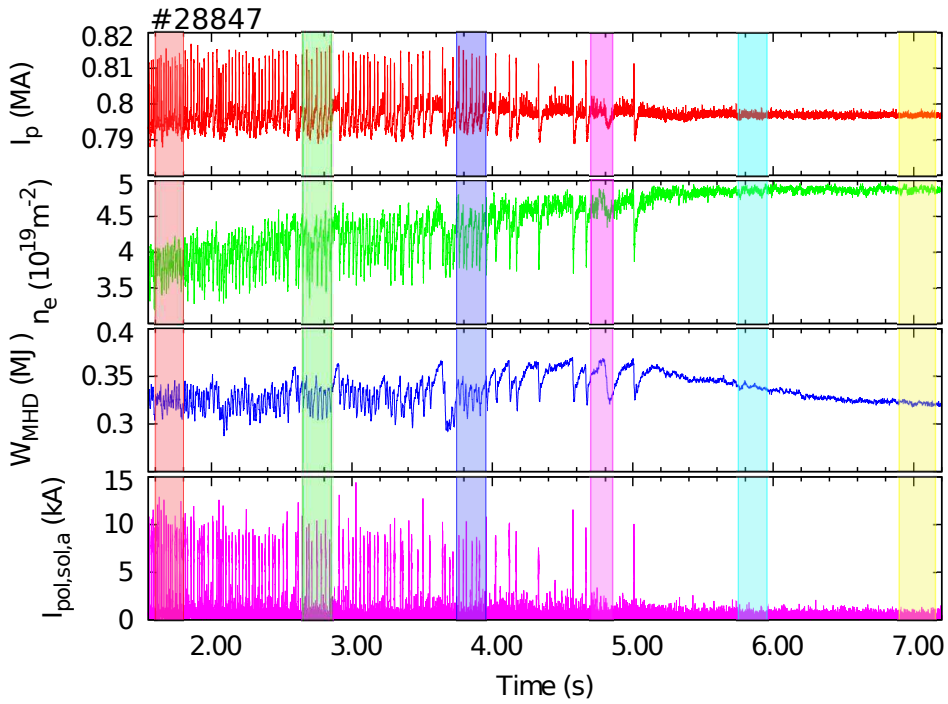


Figure 7.10: Timetraces showing (a): plasma current, (b): edge line integrated density, (c): plasma stored energy, (d): divertor current measurements. These timetraces indicate the onset of ELM mitigation with the application of external magnetic perturbations. The plasma stored energy can also be seen to decrease throughout the ELM mitigated phase though this behaviour is not typical and occurs here only because of the large gas puff applied after 5 s.

confinement is not unduly affected by the reduction of the ELM size.

The development of the electron temperature and density profiles is shown in figure 7.11 and also detailed in terms of the pedestal top and peak pedestal gradient values in table 7.1. The colour coding for the profiles corresponds to the colouring in the segments in figure 7.10. From the initial (red) profile we see a steadily increasing pedestal top density. In the second two phases (green and blue), this corresponds to a wider density pedestal at a similar gradient. Since the heating power applied throughout the discharge is the same, the temperature pedestal top decreases with the increasing density. As the temperature pedestal width initially increases, the temperature gradient decreases. This phenomenology is also correlated here with the appearance of smaller ELMs interspersed between the large Type-I ELMs. After 4 s in the discharge, a distinct thinning out of the large ELMs is evident, with the smaller ELMs appearing to take over. In the 4^{th} timeslice (purple), where full mitigation has almost been achieved, the density pedestal width shrinks again, with the temperature gradient recovering. Note that the pedestal top density continues to increase in this phase, correspond-

Table 7.1: Electron temperature and density pedestal top and gradient evolution for discharge #28847.

Time	Te, ped (keV)	∇Te (keVm^{-1})	ne, ped (10^{21}m^{-3})	∇ne (10^{22}m^{-4})
1.74	0.41	2.58	5.24	3.00
2.75	0.38	1.45	5.52	2.27
3.85	0.37	1.57	5.28	2.50
4.78	0.37	1.78	6.27	3.31
5.85	0.33	2.01	6.29	3.58
7.03	0.30	1.70	6.36	3.93

ing to the higher applied gas fuelling. In the final two phases (cyan and yellow),

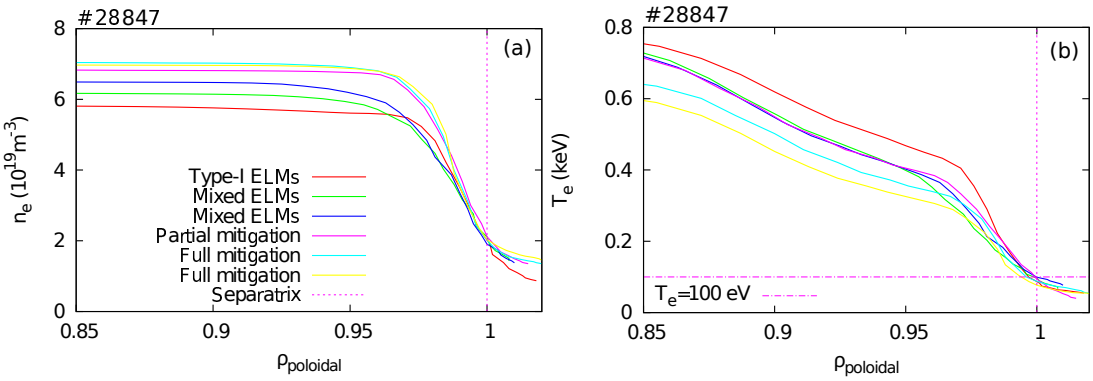


Figure 7.11: (a): electron density and (b): electron temperature profiles during discharge #28847. The colour coding corresponds to the timeslices in figure 7.10. The increasing pedestal top density can clearly be seen in (a), as well as a wider pedestal width. The temperature pedestal top value also decreases over time as the total plasma density continues to increase at a constant heating power. Also shown in the electron temperature plot are the separatrix temperatures and the T_e line in purple.

where Type-I ELMs are not present, the temperature profile shows some continuing evolution, with both the pedestal top and width decreasing. The density profile does not vary substantially in these phases.

Due to the collisionality increase leading up to the ELM suppression phase, we can speculate that there could be a decrease in the edge current density. In addition, the pressure gradient also decreases leading up to the ELM suppression phase, which alone is sufficient to lower the LFS peak edge current density. This is very clearly shown in figure 7.12(a), where the maximal pressure gradient and edge current density, as determined by CLISTE, first decrease and, once ELM mitigation has been achieved, subsequently increase. In the cases presented here,

the ion temperature has been taken to be the same as the electron temperature due to the high collisionality. A comparison with the neoclassically predicted current density is shown in figure 7.12(b) where generally good agreement is observed. The external loop voltage was used to determine the toroidal electric field for these timepoints; this may not be valid in the fluctuating phase in the mitigated phase of the discharge where rapid changes in the edge current density could be responsible for changes in the local electric field.

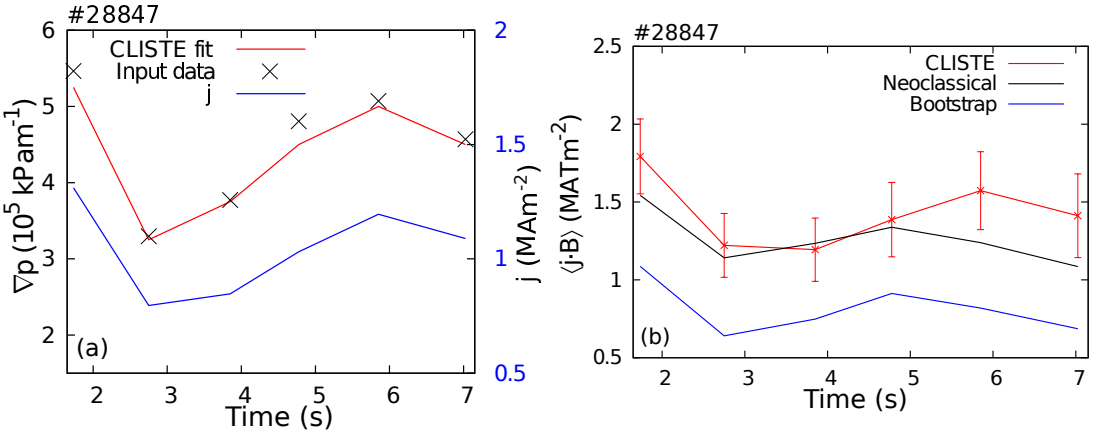


Figure 7.12: (a): maximal input pressure gradient (black crosses) compared with the CLISTE fit (red line) and the corresponding edge current density peak (blue line, right hand axis). The evolution of the latter closely follows the pressure gradient, consistent with previous findings. (b): $\langle j \cdot B \rangle$ from CLISTE (red line) with uncertainties, neoclassical calculations (black line), and the bootstrap current (blue line), showing good agreement between the CLISTE prediction and the calculation.

From results presented in section 6.7 and knowing that the edge current density, via the magnetic shear, can stabilise the pressure gradient driven ballooning mode, we can speculate on a possible causal relationship in this case. The chain of events would be as follows:

1. The high fuelling rate increases the density pedestal top and width, effecting a higher pedestal top collisionality

2. This higher collisionality, coupled with a lower temperature gradient, lowers the pedestal bootstrap current

3. Since the edge current is now lower, and the plasma shape has not been affected, a smaller normalised pressure is attainable (see figure 6.12(b))

4. A smaller edge current and pressure gradient leads to the introduction of a new ELM type in the pedestal. The high pressure gradients which cause

the large Type-I ELMs are eventually no longer reached, and the smaller ELMs take over as the dominant ELM type.

The Type-I ELMs in the first phase of the discharge exhibit a very small desnity pedestal top drop, with no change in the gradient, while both the temperature pedestal top and gradient collapse. Eventually, the ELMs develop a hybrid characteristic where both the temperature and density pedestal are affected by the ELM crash. The development of the normalised edge pressure gradient and current density are shown as a j-α diagram in figure 7.13. The timepoints in the discharge are marked on the plot. From this figure, we see that the maximum edge pressure gradient and current density is achieved in the first phase, with large Type-I ELMs only. This is followed by a rapid decrease in the edge current density and edge pressure gradient, leading to the phase with mixed Type I and small ELMs. Finally, both quantities increase again as the mitigated phase is entered and saturated for the final two timepoints. The plasma stored energy decreases in this phase, which is reflected in the lower pressure gradient. This behaviour is also seen for the local pressure gradient in figure 7.12(a).

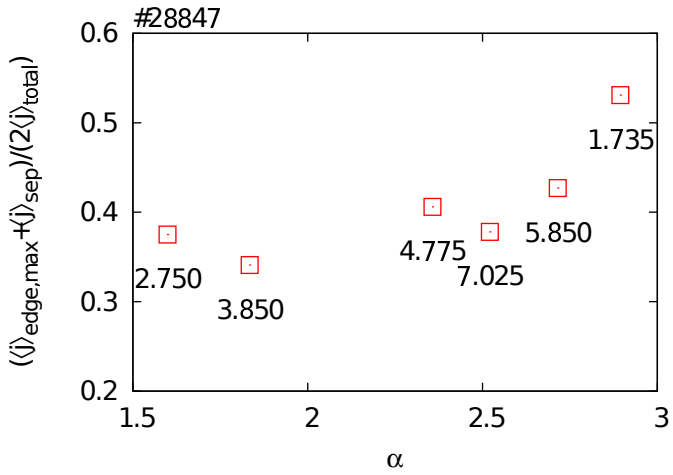


Figure 7.13: Development of the normalised pressure gradient and edge current density from a Type I ELMy phase (1.735 s), to a mixed phase (2.75-3.85 s) towards the mitigated ELMy phase (4.775 s onwards). The recovery of both parameters after the onset of full mitigation is clear and indicates that the discharge performance is not adversely impacted by limitations on the pressure gradient due to the B-coils.

7.3.1 Comparison with other regimes

In addition to discharge #28847, two other discharges were performed with non-resonant MPs (#28850) and no perturbations (#28848). The same heat-

ing power, plasma current, magnetic field, and gas puff rate were applied as in #28847, allowing an accurate comparison of all three regimes. The outer divertor current for all three discharges is shown in figure 7.14, with the coloured phases corresponding to the same timepoints analysed in the previous section. In this figure, we see that ELM mitigation is achieved in all three discharges, but at different times. In the non-resonant MP case, mitigation is achieved early in the discharge, while in the case without MPs, the ELMs are mitigated only when a very large gas puff is applied.

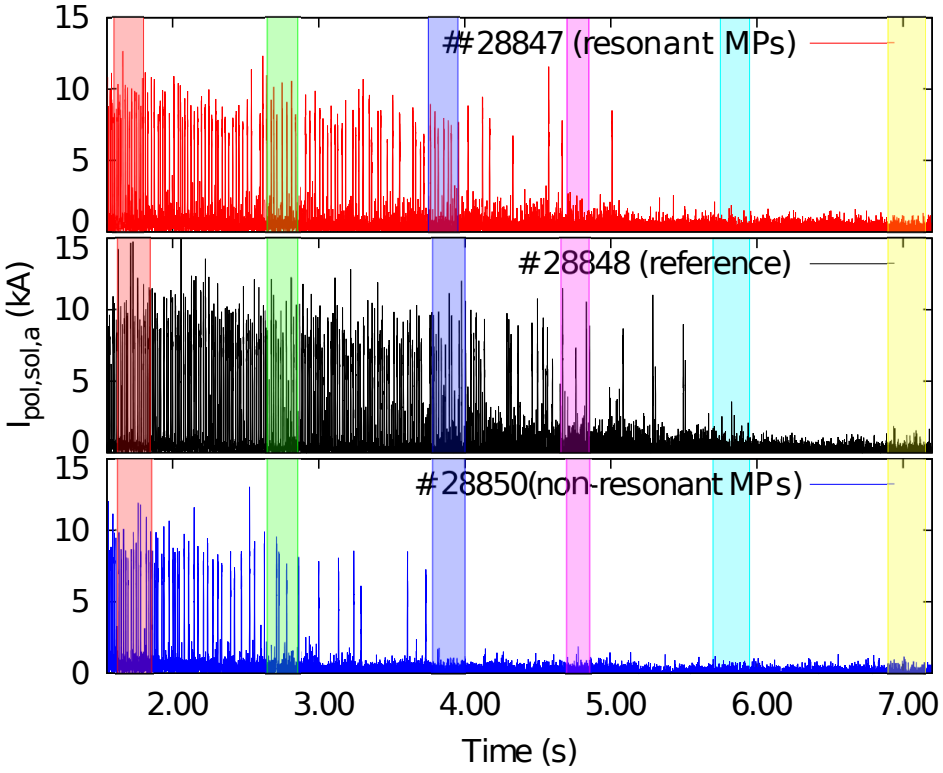


Figure 7.14: Comparison of ELM mitigation onset for resonant (top), no (middle), and non-resonant (bottom) magnetic perturbations, indicated by the divertor currents. The gas puff rate and general discharge parameters were the same in all three cases. Also shown are the same timepoints from the previous section where data will be analysed.

For the case with non-resonant MPs, there is only a short period in the second (green) phase where mixed ELMs are present. After this, mitigation is observed and full onset has occurred by 4 s, in the blue timeslice. A much slower evolution is evident in the reference discharge which proceeds in a similar fashion to the case with resonant MPs, showing first Type-I ELMs, then several time points in the mixed phase where the large ELMs are gradually thinned out, and finally two timepoints in the mitigated phase. The evolution of the temperature and density profiles is similar in all three discharges. The variation of the temperature and

density pedestal heights for the three discharges is shown in figure 7.15.

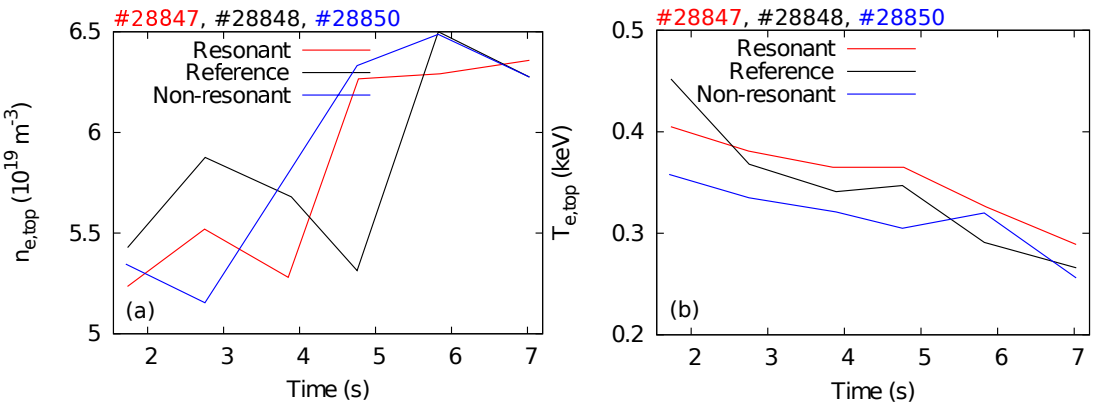


Figure 7.15: (a): pedestal top electron density as a function of time for resonant (#28847, red) and non-resonant (#28850, blue) MPs, and the reference case (#28848, black). (b): evolution of the pedestal top electron temperature for the same three discharges.

Here, we see very clearly that the density pedestal height increases substantially over the course of each of the discharges. Corresponding to the empirical finding that a critical density threshold must be reached for ELM mitigation to occur (in AUG), the density is seen to rise fastest in the non-resonant case, slightly slower in the resonant case, and slowest for the case without MPs. In addition, a saturated pedestal top height, which is the same in all discharges, is also observed when full mitigation has been achieved. The pedestal width also increases when mixed ELMs are present and subsequently decreases when mitigation has been achieved, meaning that the peak edge density gradient first decreases slightly and then increases to values above that attained in the Type-I ELMy phase.

The behaviour of the electron temperature pedestal top is in strong contrast to that of the electron density. The decreasing trend seen in figure 7.15(b) is not surprising, due to an increased plasma density. However, the almost monotonic decrease of the temperature pedestal tops at the same rate in all three discharges is not expected, given that there are different rates of increase of the pedestal top density. In all three cases, the pedestal top temperature decreases linearly and at the same rate. There is no corresponding saturation of the temperature as ELM mitigation is achieved. Instead, the edge temperature pedestal appears to become more and more eroded throughout the discharge. This is most likely due to the increasing edge gas fuelling rate as the line integrated densities continue to increase towards the end of the discharge, while the total heating power remains the same. The electron temperature pedestal width also increases as the Type-I ELMs thin out, decreasing again at the onset of full mitigation. This corresponds

to a strong drop in the edge electron temperature gradient and subsequent recovery to the Type-I ELMy levels.

Since the phenomenology of the pedestal top evolution is similar in all discharges, it would be interesting to plot the change of α and j_{norm} for each of the three cases. This is shown in figure 7.16 where the lines depict the time evolution of the discharges in j-α space. These lines are a smooth interpolation between the data points from each discharge. The arrows indicate the Type-I ELMy phase in all discharges and the direction of the trajectory. The stars show the final timepoint before full mitigation is achieved and the squares indicate the final timepoint for each discharge. The other timepoints are shown as crosses on the lines. What is very obvious is that all discharges drop sharply in α leading towards ELM mitigation, and that the mixed ELM phases all occur at low edge current density (i.e. below 0.4). ELM mitigation takes place when α drops below a value of 2; in the reference case, where mitigation proceeds at a slow rate, it is much easier to distinguish a critical value for α. Only in the resonant case is there more than one timepoint at these low values of α.

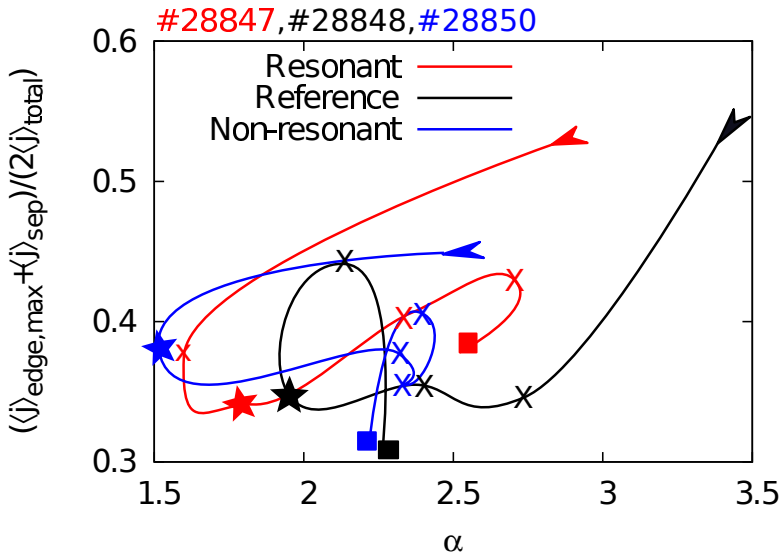


Figure 7.16: j-α evolution for each of the cases analysed; resonant MPs (#28847, red), no MPs (#28848, black), and non-resonant MPs (#28850, blue). Each discharge starts at high current density and pressure gradient (triangles), moves along its respective path to lower values of both of these parameters and a final mixed ELM phase (stars), and finally fully saturated ELM mitigation. The squares denote the final timepoint analysed in each discharge.

Once ELM mitigation has been achieved, all three discharges converge at a similar normalised pressure gradient between 2.2 and 2.5 and seem to remain here for the duration, but with fluctuations in both j_{norm} and α. The low values of

j_{norm} seem to indicate that the low current density could be responsible for the occurrence of the small ELMs. This is very similar to the finding by Saarelma et al.[82], where it was proposed that a high pedestal density, low pedestal temperature combined to lower the edge current density, closing access to the second stability regime and allowing access to the high-n or n=∞ ballooning mode regime. In this paper, Saarelma speculates that the small Type-II ELMs are due to the fact that only the pressure gradient must be relaxed to allow the plasma to return to a stable state, while for peeling-ballooning modes, both the pressure gradient and current density must be relaxed, allowing a much larger release of energy.

Another possibility is that the external magnetic perturbations cause local shear modulations linked to a helical form of the Pfirsch-Schlüter currents, as modelled by Bird and Hegna[85]. The shear modulation then lowers the threshold for ideal ballooning modes, halting the growth of the pedestal. This modelling assumed that the perturbation was resonant, but found that the effect is still oeprative when the resonant components were shielded by the plasma response. Additionally, the effects modelled corresponded to the analysis at a nearly resonant surface, since ideal MHD does not allow the resonant effects of such island overlap to be described. This may explain why the coil configuration does not largely impact the phenomenology of ELM mitigation at AUG. However, other tokamak devices typically find a dependence on whether the perturbation is resonant or non-resonant, and does not give an explanation of ELM mitigation in the high density regime. However, it is possible that the overall effect is the same; the plasma enters a regime which is unstable to ideal ballooning modes throughout the entire pedestal. It should be noted that the intent of Bird and Hegna was to use the ideal MHD ballooning mode stability criterion as an indicator for the onset of a kinetic ballooning mode, which could then allow either a steady rate of transport, or create the small bursty ELMs seen in these discharges. Nonlinear modelling of KBMs remains to be performed, meaning that it is not yet known what form the transport should take and if the KBM would be a candidate to explain these small ELMs.

To summarise this section; ELM mitigation has already been reported to depend on a critical plasma density. The results on the current density presented here show that this increased density (and correspondingly lowered temperature) acts to lower the normalised edge current density to a level where high-n ballooning modes could become dominant, resulting in smaller ELMs in a mixed phase which finally gives way to only small ELMs in a fully mitigated ELMy regime.

This transition to small ELMs appears to occur at normalised edge current densities below 0.4 and values of α below 2. After these criteria have been achieved, both normalised quantities can recover to higher values, though the Type-I ELMs remain mitigated.

7.4 Conclusions

In conclusion, this chapter has presented the general phenomenology associated with ELM mitigated regimes in AUG. The edge current density has been evaluated using the techniques developed for analysis of standard Type-I ELMs. Good agreement has also been found when comparing the current density from CLISTE with neoclassical predictions based on Sauter et al.[110, 111], adding much support to the findings discussed here. In the case of Type-II ELMs, both the local and flux surface averaged current densities were observed to decrease as the mitigated regimes were accessed. This agrees well with findings from linear stability analyses for JET[82] and AUG[84] where it was postulated that the second stability access for high pressure gradients was closed off due to this lower current density. In addition, the higher shaping required to achieve Type-II ELMs was also shown to narrow the unstable mode envelope, which may be linked to smaller ELM sizes.

ELM mitigation achieved using external magnetic perturbations was seen to have a similar phenomenology in terms of the edge current density as Type-II ELMs, with a decrease in both the normalised pressure gradient and edge current density evident in the three analysed discharges (including a reference case with no external perturbations). Some recovery of the normalised pressure gradient was evident, while the edge current density remained below a seemingly critical threshold; this could indicate that, again, second stability access for the pressure gradient has been closed off by the lower current density.

ELMs occurring in nitrogen seeded discharges show a very different phenomenology, with the local current density peak becoming higher and narrower, thus making it likely that higher edge pressure gradients can be achieved due to a local shear suppression of ballooning modes. However, the flux surface averaged current density decreases due to the higher Z_{eff}, thus also lowering the drive for a peeling mode. This means that there is likely a completely different instability present which must be analysed using non-ideal MHD and turbulence codes. No ideal MHD instability was identified in an analysis of highly radiating JET plasmas[82], which is consistent with the findings and hypotheses presented here.

In terms of projecting these results to Iter, it is possible that the two regimes which depend on density increases to lower the bootstrap current may be difficult to attain in the low collisionality device. Increasing a low pedestal top collisionality via density fuelling may not be as effective, meaning lowering the temperature (by, for example, reducing external heating) may be a more favourable method of achieving ELM mitigation from external MPs or access to a Type-II ELM regime. This method, however, reduces the attainable plasma pressures and the likelihood of achieving net fusion power. However, the nitrogen seeding regime should be robust; since the local edge current density increases slightly and the flux surface averaged current density decreases due to the higher Z_{eff}, this regime is hypothesised to be accessible also at low collisionality. The exact instability which eventually triggers the ELMs in this regime is presently unknown, meaning no firm predictions for future devices can be made. Unfortunately, in addition to the uncertainty in the trigger mechanism, there is also a significant problem in that ammonia is formed not only in the discharges when nitrogen is seeded, but also during subsequent discharges[129]. Ammonia may prove to be a serious issue for gas plants in Iter, which are designed to remove tritium from exhaust gases, and also for cryo pumps, which control the amount of neutral particles in the main chamber. Thus, an alternative edge radiator will need to be found which offers the same beneficial properties of nitrogen, but without this significant drawback.

Chapter 8

Summary and outlook

The initial goal of this work was to make observations on the reconstructed edge current density recovery after an ELM crash. In addition to this, other aspects of the edge current density were also analysed, notably its comparison to the predictions from neoclassical theory. A final goal of this work was to analyse the behaviour of the edge current density in small-ELM regimes to determine if its behaviour in these cases could explain the small rapid ELM behaviour with which they are characterised.

In chapter 3, the post-ELM LFS edge current density was seen to evolve on a similar timescale to the local pressure gradient. A rapid loss of current density at an ELM crash was attributed to the shearing off of outer flux surfaces at the ELM crash, allowing a rapid loss of particles and energy into the SOL. This effect was also analysed using a large database of discharges and was found to vary linearly with the relative ELM size. The onset of the ELM is determined from the increased thermal currents in the SOL, caused by the temperature difference between the inner and outer divertor target plates. The drop of internal current density and increase of SOL current density was well resolved by the CLISTE code through the use of internal kinetic profiles and external integrated current measurements. It was also shown that the separatrix current density remained constant throughout an ELM crash (see figures 3.11 and 7.9), despite the large changes on either side of the separatrix.

One interesting finding from the analysis of normalised edge current densities and pressure gradients was that in the short "flat-top" phase just after a typical Type-I ELM crash, the edge current density fluctuates at constant pressure gradient. This could be indicative of a second current driven mode which becomes destabilised at low pressure gradients; it is possible this is a peeling-like mode since this mode is stabilised by pressure gradients and x-point geometries. The

current density fluctuation could however also be due to an Ohmic redistribution of current following the ELM crash.

A sensitivity study of the current density profile shape and location was also carried out. This study determined that there is a ±2 mm uncertainty in the location of the profile as determined from a minimum in the magnetic residuals. In addition, for a discharge featuring a power scan, it was also shown that the profile location corresponding to this minimum placed the separatrix electron temperature at approximately 100 eV. In addition, it was shown that if the separatrix profile values are known, or can be fixed, that the pedestal width can also be determined. If this is not known, there is some evidence that a minimum pedestal width can still be recovered using just the magnetic measurements. Finally, once realistic measurement uncertainties for the pressure profile are included in the regression matrix for CLISTE the one sigma confidence bands returned by CLISTE correspond well to the one sigma uncertainties in the kinetic measurements; this bodes well for stability analysis and the interpretation of uncertainties on the equilibrium reconstruction.

A comparison with a theoretical model from Sauter et al.[110, 111] was also made. An initial investigation using three discharges from a fuelling study found that the neoclassically predicted bootstrap current and resistivity could be used to predict both the peak height and the shape of the pedestal flux surface averaged current density. This is a somewhat surprising result, given that it has long been held that neoclassical theory is not valid in the low gradient length region of the pedestal. In addition, it was also found that the edge current density evolution throughout the entire ELM cycle can be well described by this neoclassical model, apart from the phase where the electron temperature gradient recovery stalls between 3 and 6 ms after the ELM crash. Excellent agreement in this phase can, however, be recovered when real ion temperature data is used, rather than assuming that the electron and ion temperature profiles are the same after the ELM crash.

To test the wider applicability of this model, a database of some 40 AUG discharges was created which varied heating power, plasma current, heating type, magnetic field, fuelling rate, and plasma shape. It was shown that, once correct ion temperature are taken into account at low collisionality, there is excellent agreement throughout the parameter space between the CLISTE results and the neoclassical predictions. This database also showed that the peak of the LFS edge current density scales linearly with the local pressure gradient and inversely with the pedestal top collisionality, which is also in agreement with neoclassical theory.

A different trend was found, however, for the normalised edge current density and pressure gradient, where a square-root like dependence of the normalised pressure gradient on the edge current density was found. In addition, the plasma shaping also played a significant role in obtaining a good fit.

The final part of this work focussed on the other types of ELM which are observed regularly at AUG; Type-II ELMs, the smaller ELMs seen when nitrogen seeding is applied, and the mitigated ELMs observed when external magnetic perturbations are applied. In the case of Type-II ELMs, the findings echoed those hypothesised earlier by Saarelma et al.[84, 82] in that the current density, both local and flux surface averaged, decreased significantly once the Type-II regime was accessed. This most likely closes the second stability access for higher pressure gradients, limiting plasma performance (albeit still at high confinement) and giving way to smaller more rapid ELMs. A similar phenomenology was observed for the magnetic perturbation mitigation regime, where the normalised edge current density and pressure gradient dropped throughout the three analysed discharges (including a reference case) until a critical value appeared to be met. After this, saturated ELM mitigated phases with increased pressure gradients but low current densities were observed.

The nitrogen seeded ELMs presented interesting results. The local current density profile became more strongly peaked and narrower, which is likely to alter the local shear and stabilise the ballooning modes which grow on the LFS midplane. However, in contrast to the other discharges analysed throughout this work, the flux surface averaged current density was observed to drop during nitrogen seeding. This effect is due to the higher effective charge in the pedestal from the nitrogen. This lower edge current density would then correspond to a lack of a drive for peeling modes, meaning both the ideal modes thought to be responsible for ELMs (peeling and ballooning modes) are suppressed. This again echoes findings by Saarelma et al.[82] that ideal MHD modes are not responsible for ELMs in highly radiating plasmas. Instead, some other mode, likely a resistive ballooning mode, must be responsible for the ELM crash.

8.0.1 Future work

This work has laid the ground for much work to be done in relation to pedestal stability. Robust and reliable edge current density profiles and associated uncertainties have been provided and an effective workflow for their creation has been put in place at AUG. At present, there is a large effort underway in the international community to model the ELM cycle, based predominantly on the

E-PED model of two instabilities constraining the pedestal height and width. Work at AUG is focussing on ideal MHD stability, non-linear MHD evolution of the ELM cycle, and non-linear turbulence analysis. This complements the work by Saarelma et al.[45] at JET, Dickinson et al.[46] at MAST, and Snyder et al.[73] at DIII-D in understanding the edge pedestal.

To profit from this work, detailed analysis and reconstruction of post-ELM current density profiles for the discharges in the presented database should be made. This would also allow an analysis of the separatrix and shape changes following an ELM crash and, in parallel with electron temperature and density loss analysis, aid in the determination of ELM loss mechanisms and allow extrapolation to future devices. In addition, highly time resolved analysis of the ion temperature relative to the ELM cycle is of the utmost importance, as highlighted by the fast regrowth of the ion temperature gradient relative to the electron temperature gradient.

Insights into the dynamics of an ELM crash are also possible. The difference between the behaviour of the current density following a nitrogen mitigated ELM crash highlights that there could be a second current density driven mode which appears once the pressure gradient has been sufficiently lowered. The small drop of the electron density at a nitrogen mitigated ELM is also an interesting phenomenon, the understanding of which would be greatly enhanced by nitrogen density analysis from the edge CXRS diagnostic suite at AUG. Also, it has been speculated that the size of an ELM crash could be linked to the size of the unstable mode envelope which is determined from linear stability analysis. Again, the database which has been created here has made a good starting point to continue along this line of work and also determine the effect of shaping on ELM losses.

The analysis of pedestal stability is of great importance to future devices, both to predict the performance of such devices, and to understand the possible threats to machine integrity that stem from such instabilities as ELMs. The work presented here has contributed to the knowledge of the behaviour of the edge current density, in place of the heretofore general uncertainty of the reliability of equilibria for linear stability analysis. The application of this analysis to other devices would allow a size and wider shape dependence of the edge current density as well as other stability properties to be made. In addition, the access conditions to the ELM mitigated regimes are now well documented from several points of view and it is hoped that these results can be applied in a more general fashion to design experiments to maximise plasma performance while mitigating the effects of ELMs on the machine.

Bibliography

[1] David C D.C. Lowe, New Zealand, P. Forster, et al. Changes in Atmospheric Constituents and in Radiative Forcing. In S. Solomon, D. Qin, M. Manning, et al., editors, *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, volume AR4, chapter Changes in, pages 129–234. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2007.

[2] Malte Meinshausen, Nicolai Meinshausen, William Hare, et al. Greenhouse-gas emission targets for limiting global warming to 2 degrees C. *Nature*, 458(7242):1158–62, April 2009.

[3] Dominik Heide, Lueder von Bremen, Martin Greiner, et al. Seasonal optimal mix of wind and solar power in a future, highly renewable Europe. *Renewable Energy*, 35(11):2483–2489, November 2010.

[4] H. Tanigawa, K. Shiba, a. Möslang, et al. Status and key issues of reduced activation ferritic/martensitic steels as the structural material for a DEMO blanket. *Journal of Nuclear Materials*, 417(1-3):9–15, October 2011.

[5] F. Troyon, A Roy, W. A. Cooper, et al. Beta limit in tokamaks . Experimental and computational status. *Plasma Physics and Controlled Fusion*, 30(11):1597–1609, 1988.

[6] J.A. Wesson. *Tokamaks*. 3 edition, 2003.

[7] N. Oyama. Overview of JT-60U results towards the establishment of advanced tokamak operation. *Nuclear Fusion*, 49(10):104007, October 2009.

[8] G Federici, C.H. Skinner, J.N. Brooks, et al. Plasma material interactions in current tokamaks and their implications for next step fusion reactors. *Nuclear Fusion*, 41(12R), 2001.

[9] R Neu, K Asmussen, K Krieger, et al. The tungsten divertor experiment at ASDEX Upgrade The tungsten divertor experiment at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 165, 1996.

[10] R Neu, R Dux, a Kallenbach, et al. Tungsten: an option for divertor and main chamber plasma facing components in future fusion devices. *Nuclear Fusion*, 45(3):209–218, March 2005.

[11] R Neu, V Bobkov, R Dux, et al. Ten years of W programme in ASDEX Upgradechallenges and conclusions. *Physica Scripta*, T138:014038, December 2009.

[12] R Neu, M Balden, V Bobkov, et al. Plasma wall interaction and its implication in an all tungsten divertor tokamak. *Plasma Physics and Controlled Fusion*, 49(12B):B59–B70, December 2007.

[13] G.F. F Matthews, V. Philipps, R. Kamendje, et al. An ITER-like wall for JET. *Journal of Nuclear Materials*, 363-365:1–11, June 2007.

[14] V. Philipps, Ph. Mertens, G.F. F Matthews, et al. Overview of the JET ITER-like Wall Project. *Fusion Engineering and Design*, 85(7-9):1581–1586, December 2010.

[15] Lorne Horton. The JET ITER-like wall experiment: First results and lessons for ITER. *Fusion Engineering and Design*, 88(6-8):434–439, October 2013.

[16] A. Kallenbach, R. Dux, M. Mayer, et al. Non-boronized compared with boronized operation of ASDEX Upgrade with full-tungsten plasma facing components. *Nuclear Fusion*, 49(4):045007, April 2009.

[17] J. Freidberg. *Plasma physics and fusion energy.* Cambridge University Press, 2007.

[18] K Lackner. Computation of ideal mhd equilibria. *Computer Physics Communications*, 12(1):33–44, 1976.

[19] P.J. McCarthy. Analytical solutions to the GradShafranov equation for tokamak equilibrium with dissimilar source functions. *Physics of Plasmas*, 6(9):3554, 1999.

[20] B J Braams, W Jilge, and K Lackner. Fast determination of plasma parameters through function parametrization. *Nuclear Fusion*, 26(6):699–708, 1986.

[21] P J McCarthy. *An Integrated Data Interpretation System for Tokamak Discharges*. PhD thesis, University College Cork, 1992.

[22] P J Mc Carthy. Identification of edge-localized moments of the current density profile in a tokamak equilibrium from external magnetic measurements. *Plasma Physics and Controlled Fusion*, 54(1):015010, January 2012.

[23] LL Lao, H. St. John, R.D. Stambaugh, et al. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11), 1985.

[24] GTA Huysmans, J.P Goedbloed, and W Kerner. Helena. In *Conference on computational plasma physics (CP90)*, page 371, 1991.

[25] C Konz and R Zille. HELENA Fixed Boundary Equilibrium Solver. Technical Report October, Max-Planck Insititut fuer Plasmaphysik, Garching bei Muenchen, 2007.

[26] A A Ivanov, R R Khayrutdinov, S Yu Medvedev, et al. New Adaptive Grid Plasma Evolution Code SPIDER. 29(July):2–5, 2005.

[27] S. P. Hirshman and J. C. Whitson. Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Physics of Fluids*, 26(12):3553, 1983.

[28] S P Hirshman and D K Lee. MOMCON: A spectral code for obtaining three-dimensional magnetohydrodynamic equilibria. *Computer Physics Communications*, 39(2):161–172, 1986.

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

[30] F Wagner. A quarter-century of H-mode studies. *Plasma Physics and Controlled Fusion*, 49(12B):B1–B33, December 2007.

[31] F. Ryter, J.A. Snipes, R.S. Granetz, et al. H mode power threshold database for ITER. *Nuclear Fusion*, 36(9):1217, 1996.

[32] P. Sauter, T. Pütterich, F. Ryter, et al. L- to H-mode transitions at low density in ASDEX Upgrade. *Nuclear Fusion*, 52(1):012001, January 2012.

[33] E. Viezzer, T. Pütterich, G.D. Conway, et al. High-accuracy characterization of the edge radial electric field at ASDEX Upgrade. *Nuclear Fusion*, 53(5):053005, May 2013.

[34] H Zohm. Edge localized modes (ELMs). *Plasma Physics and Controlled Fusion*, 38:105–128, March 1996.

[35] G Federici, A Loarte, and G Strohmayer. Assessment of erosion of the ITER divertor targets during type I ELMs. *Plasma Physics and Controlled Fusion*, 45(9):1523–1547, September 2003.

[36] A Loarte, G Saibene, R Sartori, et al. Characteristics of type I ELM energy and particle losses in existing devices and their Characteristics of type I ELM energy and particle. *Plasma Physics and Controlled Fusion*, 45:1549–1569, 2003.

[37] J Stober, M Maraschek, G D Conway, et al. Type II ELMy H modes on ASDEX Upgrade with good confinement at high density. *Nuclear Fusion*, 1123, 2001.

[38] P.T Lang, J Neuhauser, L.D Horton, et al. ELM frequency control by continuous small pellet injection in ASDEX Upgrade. *Nuclear Fusion*, 43(10):1110–1120, October 2003.

[39] P.T. Lang, a. Alonso, B. Alper, et al. ELM pacing investigations at JET with the new pellet launcher. *Nuclear Fusion*, 51(3):033010, March 2011.

[40] Y. Liang, H. Koslowski, P. Thomas, et al. Active Control of Type-I Edge-Localized Modes with n=1 Perturbation Fields in the JET Tokamak. *Physical Review Letters*, 98(26):265004, June 2007.

[41] T. Evans, R. Moyer, P. Thomas, et al. Suppression of Large Edge-Localized Modes in High-Confinement DIII-D Plasmas with a Stochastic Magnetic Boundary. *Physical Review Letters*, 92(23):235003, June 2004.

[42] A Kirk, H W Muller, E Wolfrum, et al. Comparison of small edge-localized modes on MAST and ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 53(9):095008, September 2011.

[43] W. Suttrop, A. Herrmann, A. Kallenbach, et al. First Observation of Edge Localized Modes Mitigation with Resonant and Nonresonant Magnetic Perturbations in ASDEX Upgrade. *Physical Review Letters*, 106(22):1–4, June 2011.

[44] P B Snyder, H R Wilson, J R Ferron, et al. Edge localized modes and the pedestal: A model based on coupled peeling-ballooning modes. *Physics of Plasmas*, 9(5):2037, 2002.

[45] S Saarelma, M N A Beurskens, D Dickinson, et al. Pedestal Modelling Based MHD Analyses on MAST and JET Plasmas. In *EPS, Stockholm*, pages 2–5, 2012.

[46] D Dickinson, S Saarelma, R Scannell, et al. Towards the construction of a model to describe the inter-ELM evolution of the pedestal on MAST. *Plasma Physics and Controlled Fusion*, 53(11):115010, November 2011.

[47] P.B. Snyder, N. Aiba, M. Beurskens, et al. Pedestal stability comparison and ITER pedestal prediction. *Nuclear Fusion*, 49(8):085035, August 2009.

[48] A Burckhart, E Wolfrum, R Fischer, et al. Inter-ELM behaviour of the electron density and temperature pedestal in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 52(10):105010, October 2010.

[49] E Wolfrum, A Burckhart, R Fischer, et al. Investigation of inter-ELM pedestal profiles in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 51(12):124057, December 2009.

[50] D. Wroblewski, K. H. Burrell, L. L. Lao, et al. Motional Stark effect polarimetry for a current profile diagnostic in DIII-D. *Review of Scientific Instruments*, 61(11):3552, 1990.

[51] F. M. Levinton, R.J. Fonck, G.M. Gammel, et al. Magnetic field pitch-angle measurements in the PBX-M tokamak using the motional stark effect. *Physical Review Letters*, 63(19):2060–2063, 1989.

[52] D. M. Thomas, a. W. Leonard, R. J. Groebner, et al. Measurement of edge currents in DIII-D and their implication for pedestal stability. *Physics of Plasmas*, 12(5):056123, 2005.

[53] M F M De Bock, J Citrin, S Saarelma, et al. Measurements of the edge current evolution and comparison with neoclassical calculations during MAST

H-modes using motional Stark effect. *Plasma Physics and Controlled Fusion*, 54(2):025001, February 2012.

[54] S.K. Rathgeber, L. Barrera, T. Eich, et al. Estimation of edge electron temperature pro les via forward modelling of the electron cyclotron radiation transport at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 2012.

[55] Eleonora Viezzer. *Radial electric field studies in the plasma edge of ASDEX Upgrade*. PhD thesis, LMU, 2013.

[56] E Strumberger, S Günter, P Merkel, et al. Numerical MHD stability studies: toroidal rotation, viscosity, resistive walls and current holes. *Nuclear Fusion*, 45(9):1156–1167, September 2005.

[57] G. T. a. Huysmans, S. E. Sharapov, a. B. Mikhailovskii, et al. Modeling of diamagnetic stabilization of ideal magnetohydrodynamic instabilities associated with the transport barrier. *Physics of Plasmas*, 8(10):4292, 2001.

[58] F. Jenko, W. Dorland, M. Kotschenreuther, et al. Electron temperature gradient driven turbulence. *Physics of Plasmas*, 7(5):1904, 2000.

[59] G.M. Fishpool. Loss of confinement due to reduction of the edge pedestal in JET. *Nuclear Fusion*, 38(9):1373, 1998.

[60] S Futatani, G Huijsmans, A Loarte, et al. Nonlinear MHD Simulations of ELM Triggering by Pellets in DIII-D and Implications for ITER. In *39th EPS conference on Plasma Physics*, pages 3–6, Stockholm, 2012.

[61] V Igochine, D Constantinscu, O. Dumbrajs, et al. Minimal physical model for interaction of MHD instability with plasma. In *EPS, Stockholm*, pages 1–4, 2012.

[62] H R Wilson, S C Cowley, a Kirk, et al. Magneto-hydrodynamic stability of the H-mode transport barrier as a model for edge localized modes: an overview. *Plasma Physics and Controlled Fusion*, 48(5A):A71–A84, May 2006.

[63] J.W Connor, R. J. Hastie, and J.B. Taylor. Shear, periodicity, and plasma ballooning modes. *Physical Review Letters*, 40(6):396–399, 1978.

[64] J. W. Connor, R. J. Hastie, H. R. Wilson, et al. Magnetohydrodynamic stability of tokamak edge plasmas. *Physics of Plasmas*, 5(7):2687, 1998.

[65] H. R. Wilson and R. L. Miller. Access to second stability region for coupled peeling-ballooning modes in tokamaks. *Physics of Plasmas*, 6(3):873, 1999.

[66] R. L. Miller, M. S. Chu, J. M. Greene, et al. Noncircular, finite aspect ratio, local equilibrium model. *Physics of Plasmas*, 5(4):973, 1998.

[67] W Kerner, S Poedts, J.P Goedbloed, et al. No Title. In P. Bachman and D.C. Robinson, editors, *EPS, Berlin*, page 89, Berlin, 1991.

[68] G. T. a. Huysmans, J. P. Goedbloed, and W. Kerner. Free boundary resistive modes in tokamaks. *Physics of Fluids B: Plasma Physics*, 5(5):1545, 1993.

[69] H. R. Wilson, P. B. Snyder, G. T. a. Huysmans, et al. Numerical studies of edge localized instabilities in tokamaks. *Physics of Plasmas*, 9(4):1277, 2002.

[70] G T a Huysmans, Home Search, Collections Journals, et al. External kink (peeling) modes in x-point geometry. *Plasma Physics and Controlled Fusion*, 47(12):2107–2121, December 2005.

[71] S Saarelma, O J Kwon, and a Kirk. X-point effect on edge stability. *Plasma Physics and Controlled Fusion*, 53(2):025011, February 2011.

[72] Andreas O Burckhart. *Different ELM regimes at ASDEX Upgrade and their linear stability analysis*. PhD thesis, LMU, 2013.

[73] P. B. Snyder, T. H. Osborne, K. H. Burrell, et al. The EPED pedestal model and edge localized mode-suppressed regimes: Studies of quiescent H-mode and development of a model for edge localized mode suppression via resonant magnetic perturbations. *Physics of Plasmas*, 19(5):056115, 2012.

[74] P. B. Snyder and G. W. Hammett. Electromagnetic effects on plasma microturbulence and transport. *Physics of Plasmas*, 8(3):744, 2001.

[75] H. Doerk, F. Jenko, T. Gorler, et al. Gyrokinetic prediction of microtearing turbulence in standard tokamaks. *Physics of Plasmas*, 19(5):055907, 2012.

[76] Philip Adrian Schneider. *Characterization and scaling of the tokamak edge transport barrier*. PhD thesis, LMU, 2012.

[77] A. Kirk, H. Wilson, G. Counsell, et al. Spatial and Temporal Structure of Edge-Localized Modes. *Physical Review Letters*, 92(24):245002, June 2004.

[78] T. Eich, A. Herrmann, and J. Neuhauser. Nonaxisymmetric Energy Deposition Pattern on ASDEX Upgrade Divertor Target Plates during Type-I Edge-Localized Modes. *Physical Review Letters*, 91(19):1–4, November 2003.

[79] R.P. Wenninger, H. Zohm, J.E. Boom, et al. Solitary magnetic perturbations at the ELM onset. *Nuclear Fusion*, 52(11):114025, November 2012.

[80] Isabel Krebs. Non-Linear reduced MHD Simulations of Edge-Localized Modes in Realistic ASDEX Upgrade Geometry. 2012.

[81] D. R. Hatch. Private communication. *Private communication*, 2013.

[82] S Saarelma, a Alfier, M N a Beurskens, et al. MHD stability analysis of small ELM regimes in JET. *Plasma Physics and Controlled Fusion*, 51(3):035001, March 2009.

[83] E Wolfrum, M Bernert, J E Boom, et al. Characterization of edge profiles and fluctuations in discharges with type-II and nitrogen-mitigated edge localized modes in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 53(8):085026, August 2011.

[84] S Saarelma, S G nter, L.D Horton, et al. MHD stability analysis of type II ELMs in ASDEX Upgrade. *Nuclear Fusion*, 43(4):262–267, April 2003.

[85] T M Bird and C C Hegna. A model for microinstability destabilization and enhanced transport in the presence of shielded 3D magnetic perturbations. *Nuclear Fusion*, 53(1):013004, January 2013.

[86] R Fischer, C J Fuchs, B Kurzan, et al. Integrated data analysis of profile diagnostics at ASDEX Upgrade. *Fusion Science and Technology*, 58(2):675–684, 2010.

[87] N. A. Salmon. First electron temperature edge measurements on the ASDEX Upgrade tokamak using a heterodyne radiometer. *International Journal of Infrared and Millimeter Waves*, 15(1):53–60, 1994.

[88] E. Wolfrum, F. Aumayr, D. Wutte, et al. Fast lithium-beam spectroscopy of tokamak edge plasmas. *Review of Scientific Instruments*, 64(8):2285, 1993.

[89] R Fischer, E Wolfrum, J Schweinzer, et al. Probabilistic lithium beam data analysis. *Plasma Physics and Controlled Fusion*, 50(8):085009, August 2008.

[90] O. Gehre. The HCN-laser-interferometer of the divertor tokamak ASDEX. *International Journal of Infrared and Millimeter Waves*, 5(3):369–379, March 1984.

[91] E Viezzer, T Pütterich, R Dux, et al. High-resolution charge exchange measurements at ASDEX Upgrade. *The Review of scientific instruments*, 83(10):103501, October 2012.

[92] S K Rathgeber, R Fischer, S Fietz, et al. Estimation of profiles of the effective ion charge at ASDEX Upgrade with Integrated Data Analysis. *Plasma Physics and Controlled Fusion*, 52(9):095008, September 2010.

[93] B J Braams. The interpretation of tokamak magnetic diagnostics. *Plasma Physics and Controlled Fusion*, 33(7):715–748, 1991.

[94] J C Fuchs, M Dunne, R Fischer, et al. Equilibrium reconstruction in the presence of 3D external magnetic perturbations on ASDEX Upgrade. In *EPS, Stockholm*, number figure 1, 2012.

[95] V.S. Mukhovatov and V.D. Shafranov. Plasma equilibrium in a Tokamak. *Nuclear Fusion*, 11, 1971.

[96] G.G. Lister. FAFNER. Technical report, IPP Garching, 1985.

[97] M.G. Dunne, P.J. McCarthy, E. Wolfrum, et al. Measurement of neoclassically predicted edge current density at ASDEX Upgrade. *Nuclear Fusion*, 52(12):123014, December 2012.

[98] P.J. McCarthy. Direct analytic model of the L-curve for Tikhonov regularization parameter selection. *Inverse Problems*, 19:643–663, 2003.

[99] R.a Pitts, S Alberti, P Blanchard, et al. ELM driven divertor target currents on TCV. *Nuclear Fusion*, 43(10):1145–1166, October 2003.

[100] H. Zohm, K. Lackner, and C. Ludescher. Statistical analysis of disruptions in ASDEX. *Nuclear Fusion*, 33(4):655–662, April 1993.

[101] K. Lackner. Emissive double probes - a model for thermocurrents during welding, 2006.

[102] J. Neuhauser, D Coster, HU Fahrbach, et al. Transport into and across the scrape-off layer in the ASDEX Upgrade divertor tokamak. *Plasma Physics and Controlled Fusion*, 44(6):855–869, 2002.

[103] T. Eich, B. Sieglin, a. Scarabosio, et al. Inter-ELM Power Decay Length for JET and ASDEX Upgrade: Measurement and Comparison with Heuristic Drift-Based Model. *Physical Review Letters*, 107(21):1–4, November 2011.

[104] L. Aho-Mantila, M. Wischmeier, K. Krieger, et al. Effect of EB driven transport on the deposition of carbon in the outer divertor of ASDEX Upgrade. *Journal of Nuclear Materials*, 415(1):S231–S234, August 2011.

[105] L. Aho-Mantila, M. Wischmeier, M.I. Airila, et al. Modelling of Carbon Transport in the Outer Divertor Plasma of ASDEX Upgrade. *Contributions to Plasma Physics*, 50(3-5):439–444, May 2010.

[106] A. Herrmann. Overview on stationary and transient divertor heat. *Plasma Physics and Controlled Fusion*, 44:883–903, 2002.

[107] a V Chankin, D P Coster, R Dux, et al. SOLPS modelling of ASDEX upgrade H-mode plasma. *Plasma Physics and Controlled Fusion*, 48(6):839–868, June 2006.

[108] M. a. Mahdavi, R. Maingi, R. J. Groebner, et al. Physics of pedestal density profile formation and its impact on H-mode density limit in burning plasmas. *Physics of Plasmas*, 10(10):3984, 2003.

[109] M. Coronado and H. Wobig. On the definition of PfirschSchluter and bootstrap currents in toroidal systems. *Physics of Fluids B: Plasma Physics*, 4(5):1294, 1992.

[110] O. Sauter, C. Angioni, and Y. R. Lin-Liu. Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime. *Physics of Plasmas*, 6(7):2834, 1999.

[111] O. Sauter, C. Angioni, and Y. R. Lin-Liu. Erratum: Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime [Phys. Plasmas 6, 2834 (1999)]. *Physics of Plasmas*, 9(12):5140, 2002.

[112] D J Kelliher, N C Hawkes, P J Mc Carthy, et al. Comparison of TRANSP-evolved q-profiles with MSE constrained equilibrium fits on JET. *Plasma Physics and Controlled Fusion*, 47(9):1459–1473, September 2005.

[113] M. Wade, M. Murakami, and P. Politzer. Validation of Neoclassical Bootstrap Current Models in the Edge of an H-Mode Plasma. *Physical Review Letters*, 92(23):1–4, June 2004.

[114] Philip A Schneider, E Wolfrum, S Günter, et al. Alternative method for characterization of inter ELM edge profiles of type-I ELMy H-modes in ASDEX Upgrade. In *EPS, Dublin*, pages 2–5, 2010.

[115] G.F. Matthews. Plasma operation with an all metal first-wall: Comparison of an ITER-like wall with a carbon wall in JET. *Journal of Nuclear Materials*, 438:S2–S10, July 2013.

[116] A Loarte, M Becoulet, G Saibene, et al. Characteristics and scaling of energy and particle losses during Type I ELMs in JET H-modes. *Plasma Physics and Controlled Fusion*, 44:1815–1844, 2002.

[117] J Hobirk. Private communication. *Private communication*, 2013.

[118] J. Schweinzer, a.C.C. Sips, G. Tardini, et al. Confinement of improved H-modes in the all-tungsten ASDEX Upgrade with nitrogen seeding. *Nuclear Fusion*, 51(11):113003, November 2011.

[119] Y. Kamada, R. Yoshino, Y. Neyatani, et al. Onset condition for ELMs in JT-60U. *Plasma Physics and Controlled Fusion*, 38:1387–1391, 1996.

[120] H. Doerk, F. Jenko, M. Pueschel, et al. Gyrokinetic Microtearing Turbulence. *Physical Review Letters*, 106(15):1–4, April 2011.

[121] J.E. Boom, E. Wolfrum, I.G.J. Classen, et al. Characterization of broadband MHD fluctuations during type-II edge localized modes as measured in 2D with ECE-imaging at ASDEX Upgrade. *Nuclear Fusion*, 52(11):114004, November 2012.

[122] J Hobirk, C D Challis, R Fischer, et al. Comparison of different improved H-mode scenarios on ASDEX Upgrade and JET. In *EPS, Strasbourg*, pages 3–6, 2011.

[123] J Hobirk, J Schweinzer, L Barrera Orte, et al. Overview of ASDEX Upgrade Improved H-mode Scenario Developments. In *IAEA*, pages 1–8, 2012.

[124] G Tardini, R Fischer, F Jenko, et al. Core transport analysis of nitrogen seeded H-mode discharges in the ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 55(1):015010, January 2013.

[125] A. Kirk, E. Nardon, P. Tamain, et al. The effect of resonant magnetic perturbations on L and H-mode plasmas on MAST. *Journal of Nuclear Materials*, 415(1):S910–S913, August 2011.

[126] W Suttrop, L Barrera Orte, T Eich, et al. Mitigation of Edge Localised Modes with small non-axisymmetric magnetic perturbations in ASDEX Upgrade. In *IAEA*, pages 1–8, 2012.

[127] W Suttrop, L Barrera, a Herrmann, et al. Studies of edge localized mode mitigation with new active in-vessel saddle coils in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 53(12):124014, December 2011.

[128] T.E Evans, R.a Moyer, J.G Watkins, et al. Suppression of large edge localized modes with edge resonant magnetic fields in high confinement DIII-D plasmas. *Nuclear Fusion*, 45(7):595–607, July 2005.

[129] D Neuwirth, V Rohde, and T Schwarz-Selinger. Formation of ammonia during nitrogen-seeded discharges at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 54(8):085008, August 2012.

Appendix A

Formulae for neoclassical current density calculation

Listed below are the detailed equations contributing to the calculation of the bootstrap current, as referenced in chapter 5. These equations are reproduced from the papers by Sauter et al.[110, 111].

The Coulomb logarithms describing the electrogmatic interactions for electrons and ions are given by:

$$\ln\Lambda_e = 31.3 - \ln(\frac{\sqrt{n_e}}{T_e}) \tag{A.1}$$

$$\ln\Lambda_{ii} = 30 - \ln(Z_i^3\frac{\sqrt{n_e}}{T_e}) \tag{A.2}$$

where Z_i is the main ion charge.

The electron and ion collisionalities are calculated as:

$$\nu_e^* = 6.921\times10^{-18}\frac{qR_0n_eZ_{eff}\ln\Lambda_e}{{T_e}^2\epsilon^{3/2}} \tag{A.3}$$

$$\nu_i^* = 4.9\times10^{-18}\frac{qR_0n_eZ_i^4\ln\Lambda_{ii}}{{T_i}^2\epsilon^{3/2}} \tag{A.4}$$

where q is the safety factor, R is the major radius of the machine, and ϵ is the inverse aspect ratio.

The modifications of the trapped particle fraction for each of the bootstrap terms are given by the following equations:

$$f_{teff}^{31} = \frac{f_t}{1+(1-0.1ft)\sqrt{\nu_e^*}+0.5(1-f_t)\nu_e^*/Z_{eff}} \tag{A.5}$$

$$f_{teff}^{32_ee} = \frac{f_t}{1+0.26(1-f_t)\sqrt{\nu_e^*}+0.18(1-0.37f_t)\nu_e^*/\sqrt{Z_{eff}}} \tag{A.5a}$$

$$f_{teff}^{32_ei} = \frac{f_t}{1+(1+0.6f_t)\sqrt{\nu_e^*}+0.85(1-0.37f_t)\nu_e^*(1+Z_{eff})} \tag{A.5b}$$

$$f_{teff}^{34} = \frac{f_t}{1+(1-0.1f_t)\sqrt{\nu_e^*}+0.5(1-0.5f_t)\nu_e^*/Z_{eff}} \tag{A.6}$$

where f_t is the calculated geometric trapped particle fraction, and f_{teff} indicates the effecticve trapped particle fraction. From this set of equations it is clear where the inverse dependence of the bootstrap current on collisionality arises.

Finally, the bootstrap coefficients are calculated as:

$$\mathcal{L}_{31} = f_{\text{teff}}^{31}\left(1+\frac{1.4}{\text{Z}_{\text{eff}}+1}-f_{\text{teff}}^{31}\left(\frac{1.9}{\text{Z}_{\text{eff}}+1}-f_{\text{teff}}^{31}\left(\frac{0.3}{\text{Z}_{\text{eff}}+1}+\frac{0.2}{\text{Z}_{\text{eff}}+1}f_{\text{teff}}^{31}\right)\right)\right) \quad \text{(A.7)}$$

$$\begin{aligned}\mathcal{L}_{32} = {} & \frac{0.05+0.62\text{Z}_{\text{eff}}}{\text{Z}_{\text{eff}}(1+0.44\text{Z}_{\text{eff}})}\left(f_{\text{teff}}^{32_ee}-\left(f_{\text{teff}}^{32_ee}\right)^4\right) \\ & + \frac{\left(f_{\text{teff}}^{32_ee}\right)^2}{1+0.22\text{Z}_{\text{eff}}}\left(1-1.2f_{\text{teff}}^{32_ee}+0.2\left(f_{\text{teff}}^{32_ee}\right)^2\right) \\ & - \frac{0.56+1.93\text{Z}_{\text{eff}}}{\text{Z}_{\text{eff}}\left(1+0.44\text{Z}_{\text{eff}}\right)}\left(f_{\text{teff}}^{32_ei}-\left(f_{\text{teff}}^{32_ei}\right)^4\right) \\ & + \left(f_{\text{teff}}^{32_ei}\right)^2\left(1-0.55f_{\text{teff}}^{32_ei}-0.45\left(f_{\text{teff}}^{32_ei}\right)^2\right)\frac{4.95}{1+2.48\text{Z}_{\text{eff}}}\end{aligned} \quad \text{(A.8)}$$

$$\mathcal{L}_{34} = f_{\text{teff}}^{34}\left(1+\frac{1.4}{\text{Z}_{\text{eff}}+1}-f_{\text{teff}}^{34}\left(\frac{1.9}{\text{Z}_{\text{eff}}+1}-f_{\text{teff}}^{34}\left(\frac{0.3}{\text{Z}_{\text{eff}}+1}+\frac{0.2}{\text{Z}_{\text{eff}}+1}f_{\text{teff}}^{34}\right)\right)\right) \quad \text{(A.9)}$$

with the coefficient α given by:

$$\alpha_0 = -\frac{1.17(1-f_t)}{1-0.22f_t-0.19f_t^2} \quad \text{(A.10)}$$

$$\alpha = \left(\frac{\alpha_0+0.25\left(1-f_t^2\right)\sqrt{\nu_i^*}}{1+0.5\sqrt{\nu_i^*}}+0.315\left(\nu_i^*\right)^2 f_t^6\right)\frac{1}{1+0.15\left(\nu_i^*\right)^2 f_t^6} \quad \text{(A.11)}$$

AUG	ASDEX Upgrade (AXially Symmetric Divertor EXperiment). Medium sized tokamak situated in Garching bei München, Germany.
B-coils	Non-axisymmetric perturbation coils situation on the LFS of AUG. Used to induce small perturbations (~1 mT) and hence effect ELM mitigation.
DCN	Deuterium-Cyanide-Nitrogen laser. Used for interferometric measurements of the line-integrated plasma density.
DIII-D	Doublet III-D. Medium sized tokamak situated in San Diego, US.
ECE	Electron Cyclotron Emission.
ECRH	Electron Cyclotron Resonance Heating.
ELM	Edge Localised Mode.
ETG	Electron temperature gradient. Turbulenct instability driven by the electron temperature gradient.
HFS	High field side. Region of high toroidal magnetic field (i.e. inboard side of tokamak).
H-mode	High confinement mode. Mode of operation offering higher heat and particle confinement, characterised by steep edge gradients.
ICRH	Ion cyclotron resonance heating.
IPHM	Improved H-mode. Avanced operation scenario offering better confinement than usual H-mode.
IR	Infrared camera diagnostic
ITG	Ion temperature. Turbulent instability driven by the ion temperature gradient.
JET	Joint European Torus. Largest tokamak currently in operation. Situated in Culham, England.
KBM	Kinetic Ballooning Mode. Interchange-type instability driven by electromagnetic fluctuations.
LFS	Low field side. Region of low toroidal magnetic field (i.e. outboard side of tokamak).
LiBES	Lithium Beam Emission Spectroscopy.
LLGHG	Long-lived greenhouse gases.
MAST	Mega Ampere Spherical Tokamak. Medium sized spherical tokamak. Situated in Culham, England.
MHD	Magnetohydrodynamics. Fluid description of magnetised plasma.
MSE	Motional Stark Effect. Stark effect experienced in the rest frame of a particle moving through a magnetic field.
MTM	Microtearing mode.

NBI	Neutral Beam Injection heating. Fast neutral particles are injected into the plasma and ionised via charge exchange or electron impact excitation. Excited fast ions then lose energy via collisions with thermal plasma particles.
P-B mode	Peeling-ballooning mode. Mode thought to be responsible for ELMs. Driven by combination of current density and pressure gradient in plasma edge.
PFC	Plasma facing components.
RMP	Resonant magnetic perturbation. Frequently also used to mean any non-axisymmetric external magnetic perturbation. See also B-coils.
SOL	Scrape-off Layer. Region between confined plasma and wall.
TEM	Trapped electron mode. Turbulent instability linked with electrons in trapped orbits on the plasma LFS.
TS	Thomson Scattering. Diagnostic used to determine the electron temperature and density simultaneously via scattering of laser light by the electrons in the plasma.

a	Plasma minor radius
α	Pressure gradient normalised to field line tension
B	Magnetic field, has both toroidal ($B_{T,\phi}$) and poloidal ($B_{p,\theta}$) components
β	Ratio of kinetic to magnetic pressure
δ	Plasma triangularity
f_q	Ratio of q_{95} to q_{cyl}
E	Electric field
f	Flux function of poloidal current ($f = RB_\phi$)
I_p	Toroidal plasma current
l_i	Plasma internal inductance
L_{n_e}	Electron density gradient length = $n_e/\frac{\delta n}{\delta R}$
μ_0	Vaccuum permeability
n_e	Electron density (m^{-3})
n_{GW}	Greenwald density; an empirical scaling for the maximum density achievable in a tokamak, dependent on the plasma current and inversely proportional to the minor radius
η	Plasma resistivity
η_e	Ratio of electron density to electron temperature gradient length
p_e	Electron pressure (Pa)
ψ	Poloidal magnetic flux, used to determine "flux surfaces" for diagnostic mapping
$\rho_{poloidal}$	Normalised radius based on ψ, 0 at magnetic axis and 1 at last closed flux surface
q_{95}	Value of safety factor at 95% of normalised flux (close to plasma edge)
q_{cyl}	Value of q assuming cylindrical plasma
R_{maj}	Major radius, measured from the centre of the torus
R_0	Plasma major radius
σ	Plasma conductivity
τ	Confinement time, or ratio of stored energy to input power
$T_{e,i}$	Electron/ion temperature (eV)
ν^*	Normalised collision frequency, also referred to as "collisionality"
V	Plasma volume
W_{MHD}	Plasma stored energy (kinetic)
Z_{eff}	effective charge on ions in the plasma, usually calculated assuming one dominant impurity

Acknowledgements

This work would not have been possible without the support of many people over the last 3.5 years.

Firstly, I would like to thank my academic supervisor Dr. P.J. McCarthy for endless support, ideas, patience, and encouragement throughout my PhD. I was never left waiting for a reply, even when this reply meant making changes to CLISTE on short notice, and your insights were always valuable.

I also thank my supervisor in Garching Dr. Elisabeth Wolfrum for advice, and explanations, and always backing me up. Without your guidance, this work would be a much poorer effort.

I would also like to thank Prof. Karl Lackner, Prof. Harmut Zohm, and Prof. Sibylle Günter for making time to discuss results and make suggestions for improvements throughout the work.

I am very grateful to Dr. Rainer Fischer for always being available to discuss results or ideas. Your critical appraisal of my work is beyond value. I also extend my graditude to Dr. Louis Giannone for his maintenance of and upgrades to the magnetic diagnostics so vital for this work, and to Dr. Christoph Fuchs for managing the CLISTE public scripts. Thanks also go to Dr. Wolfgang Suttrop, Dr. Emiliano Fable, and Dr. Thomas Eich for much valued input and explanations.

My thanks also go to the ASDEX Upgrade Team, without whom none of this work would have been possible. I would also like to thank the experiment leaders who facilitated the experiments I was part of, in particular Thomas Pütterich for putting so much time into the design of the discharges. I also thank the administrative staff in both UCC and the IPP: Margaret Bunce, Alice O'Leary, Gabriele Dörsch, Anja Bauer, and Petra Jordan and to Dr. Behler for computer support.

Thanks must go to the office mates I've had over the years at the IPP: Marco Cavedon, Sina Fietz, Dr. Benedikt Geiger, Alex Lebschy, Dr. Stefen Potzel, Dr. Philip Schneider, and Dr. Bernd Wieland. It was always an interesting office to be in. In addition, thanks also go to the other PhD students and postdocs who have made the last few years a lot of fun: Dr. Leena Aho-Mantila, Matthias Bernert, Dr. Jurrian Boom, Livia Cassali, Dr. Ivo Classen, Dr. Hauke Doerk, Dr. Tim Happel, Athina Kappatou, Dr. Hajo Klingshirn, Dr. Rachael McDermott, Dr. Sylvia Rathgeber, Fabian Sommer, Dr. Eleonora Viezzer, Dr. Christian Vorpahl, and Dr. Matthias Willensdorfer. I also thank you all for putting up with my German and for helping me learn.

Particular thanks go to Diarmuid Curran who eased my transition to Munich and who has always been a good friend, and to Andreas Burckhart with whom I've worked

closely throughout my PhD. Your ideas and friendship are a constant boost.
I also thank my friends in Ireland, particularly Aisling, Bernard, Catriona, Dave, Iain, Mark, Pete, Sinead, and Stephen, who have always been there to welcome me back home, to visit, and for a chat/cup of tea.
To Chris, you've always been there for me and your support and encouragement know no bounds. I thank you for everything, in particular for inspiring me to be creative and curious.
Finally, to my parents. You have always been there for me, no matter what. Your unconditional support has helped me through everything from primary school to finishing this PhD so far. Thank you for that, and thank you for the fun we have.