
Different ELM regimes at ASDEX Upgrade and their linear stability analysis

Andreas O. Burckhart

München 2013



Different ELM regimes at ASDEX Upgrade and their linear stability analysis

Dissertation zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften (Dr. rer. nat.)
an der Fakultät für Physik der Ludwig-Maximilians-Universität
München

durchgeführt am
Max-Planck-Institut für Plasmaphysik, Garching bei München

vorgelegt von
Andreas O. Burckhart
geboren am 12.05.1985 in Meyrin

München, den 13.09.2013



Erstgutachter:	Prof. Dr. Hartmut Zohm
Zweitgutachter:	Prof. Dr. Gregor Morfill
Tag der mündlichen Prüfung:	16. Dezember 2013

Acknowledgments

I would like to express my appreciation to my committee chair Prof. Dr. Hartmut Zohm, who always expressed great interest in my work and who gave me guidance throughout this thesis.

My deepest gratitude goes to my advisor, Dr. Elisabeth Wolfrum, who always received new results with the greatest enthusiasm. I would like to thank her for being such a great and caring advisor.

Prof. Dr. Sibylle Günter and Prof. Dr. Karl Lackner, as well as Dr. Erika Strumberger helped me through the theoretical parts of this thesis. I am deeply indebted to them for having taken the time to explain to me the meaning and ramifications of the different assumptions made in the MHD simulation codes. Moreover, I thank them for discussing with me the various numerical problems that emerged in the course of this work.

The greatest help toward the end of this thesis probably came from my very good friends Drs. Rachael McDermott and Mike Dunne, who proof-read my work and made many suggestions that strongly contributed to improving it. Mike and I worked very closely together during the last three years, and we had many fruitful discussions about ELMS, MHD, and many other topics. Both of them also helped me to take my mind off things when my work didn't go as planned. Thank you guys, you are the best!

I would also like to thank Dr. Christian Konz and Fuad Osmanlic, who were there during the first part of my thesis and helped me to get started with HELENA and ILSA.

I am grateful to Dr. Peter Lang who contributed to the pellet ELM-triggering section of this thesis, as well as to Drs. Marc Beurskens and Samuli Saarelma for fruitful discussions.

Drs. David Coster and Hans-Joachim Klingshirn, thank you for helping me run ILSA and HELENA more efficiently and providing me with more computing resources in the last weeks of this thesis, which I desperately needed.

I am also thankful to the AUG experiment leaders, diagnosticians and the ECRH team who always welcomed my, sometimes challenging, ideas for experiments.

I had a great office mate, Dr Christian Vorpahl, who also became a good friend. Thank you for the fun we had in that office over the last three years.

I would like to thank all the other members of the AUG team with whom I worked or simply had lively discussions with over lunch during my time at IPP, including Matthias, Thomas, Fabian, Athina, Sylvia, Elli, Ben, Sina, Steffen, Tim, Philip, Diarmuid, Livia, Alex, Laura, Gregor, Leena, Felix and Marco.

A thank you also goes to my friends outside of IPP, who I unfortunately don't see nearly often enough. Thank you to Peter, Alan, Stephanie, Christian, Christoph, Markus, Andreas, Joachim, Stephan, Marc, Robert, Flynn and Till.

Most importantly, none of this would have been possible without the love and patience of my family. All these years, my parents and my sister have been a constant source of love, concern, support and strength. I would like to express my heart-felt gratitude to them.

Zusammenfassung

Edge localised modes (ELMs) sind magnetohydrodynamische (MHD) Instabilitäten, die am Rand von magnetisch eingeschlossenen Fusionsplasmen auftreten, wenn diese sich in der sogenannten high confinement mode (H-Mode) befinden. Sie führen zu periodischen Energie- sowie Teilchenverlusten und begrenzen dadurch die Qualität des Einschlusses. Außerdem verursachen sie eine erhebliche Wärmebelastung der Gefäßwände, die in größeren Maschinen der nächsten Tokamak-Generation möglicherweise nicht mehr zu bewältigen sein wird. Die genaue Natur dieser Instabilitäten ist jedoch noch unklar. Die gängigste Theorie, mit der ELMs üblicherweise beschrieben werden, ist die sogenannte peeling-ballooning Theorie. Diese postuliert, dass von kritischen Druckgradienten und Stromdichten getriebene Moden für ELMs verantwortlich sind. In der vorliegenden Arbeit wird dieses Modell mit experimentellen Daten, die am ASDEX Upgrade Tokamak gemessen wurden, getestet. Zum ersten Mal wird eine breite Auswahl an ELM-Szenarien mit derselben Methodik hinsichtlich ihrer Stabilität untersucht. Der Vergleich zwischen Theorie und Experiment wird in mehreren aufeinander aufbauenden Schritten durchgeführt. Zunächst werden kinetische und magnetische Messungen mit Hilfe des Grad-Shafranov Gleichungslösers CLISTE zu einem Plasma-gleichgewicht kombiniert, das dann mittels des Computerprogrammes HELENA verfeinert wird. Schließlich wird die Stabilität dieses Gleichgewichtes mit ILSA, einem linearen MHD Stabilitätscode, berechnet.

Theoretisch sollte das peeling-ballooning Modell für alle Typ-I ELM-Szenarien gültig sein. Daher wird in dieser Arbeit die Stabilität von einigen verschiedenen H-Moden, in denen Typ-I ELMs vertreten sind, auf peeling-ballooning Moden untersucht. Manche dieser Szenarien sind mit dem Modell konsistent, in anderen treten ELMs jedoch schon weit unter der idealen MHD Grenze auf, oder diese Grenze kann weit überschritten werden. Kurz vor typischen Typ-I ELMs liegt das Gleichgewicht meistens unter oder auf der Stabilitätsgrenze, wobei eine erhebliche Streuung zu beobachten ist. Zusätzlich wird eine H-mode, in die neben Deuterium auch Stickstoff eingeblasen wurde, untersucht. In dieser Entladung wird die Grenze weit überschritten. In einer anderen Entladung liegen Druckgradient und Randstromdichte auf der Stabilitätsgrenze, aber sinken wenn Mikrowellenheizung am Plasmarand angewandt wird. Entgegen der Theorie steigt die ELM-Frequenz, obwohl sich das Gleichgewicht von der Stabilitätsgrenze entfernt.

Andere Szenarien, in denen vermutet wird, dass ELMs nicht durch peeling-ballooning ausgelöst werden, werden auch untersucht. In Entladungen in denen Typ-II und Typ-III ELMs auftreten, befinden sich Druckgradient und Randstromdichte weit unter der Grenze für peeling-ballooning Stabilität. Dies deutet darauf hin, dass andere Mechanismen verhindern, dass sich die kinetischen Profile weiter aufbauen. Das Verhalten der untersuchten Typ-III ELMs ist konsistent mit der Hypothese, dass sie von resistiven Austauschinstabilitäten ausgelöst werden.

Die Möglichkeit, ELMs durch gefrorene Deuterium Pellets auszulösen, wird erforscht. In Entladungen, die durchgeführt wurden, als die Gefäßwände durch Kohlenstoffkacheln geschützt wurden, war es möglich, zu jedem beliebigen Zeitpunkt in einer H-Mode ELMs durch Pellet-injektion auszulösen. Seit wolframbeschichtete Kacheln verwendet werden, ist es jedoch nur dann möglich ELMs auszulösen, wenn eine bestimmte Verzögerungszeit seit dem vorhergehenden ELM verstrichen ist. Diese Verzögerung kann erheblich reduziert werden, indem Stickstoff in das Gefäß eingeblasen wird. In H-Moden mit Wolfram als Material der ersten Wand, mit oder ohne Stickstoffblasen, ist es möglich ELMs durch Pellets auszulösen, wenn sich der Druckgradient und die Stromdichte am Plasmarand 20% unter der Schwelle für spontane ELMs befinden. Da der ELM-induzierte Kollaps des Druckprofils am Plasmarand in der Wolframmaschine und ohne Stickstoffblasen am stärksten ist und die Erholung des Profils

langsamer erfolgt, ist auch die Verzögerungszeit für das erfolgreiche Auslösen von ELMs durch Pellets länger.

Aktuelle Vorhersagen der ITER Randgradienten werden auch unter Verwendung des peeling-ballooning Modells untersucht. Die Ergebnisse zeigen, dass die Vorhersagen sehr optimistisch sind, aber dass das Plasma mit Temperaturen und Dichten, die 30% unter den Werten der Vorhersagen liegen, stabil wäre.

Einige der Ergebnisse in dieser Arbeit stellen die Gültigkeit des peeling-ballooning Modells in Frage. In vielen der dargestellten Fälle treten ELMs deutlich unter dem Schwellwert für peeling-ballooning Stabilität auf. In einem Fall wird dieser Wert wiederum erheblich überschritten. Die vorgestellten Ergebnisse zeigen, dass das peeling-ballooning Modell nicht ausreichend ist, um das Auftreten von ELM-Instabilitäten vollständig zu erklären. Vielmehr ist ein weiterer, noch unbekannter Auslösemechanismus erforderlich.

Abstract

Edge localised modes (ELMs) are magnetohydrodynamic (MHD) instabilities that occur at the edge of magnetically confined fusion plasmas. They periodically expel particles and energy from the confined region. In addition to limiting the confinement, they cause high heat fluxes to the walls of the tokamak which may not be manageable in larger, next-generation devices. However, the exact nature of the instabilities that drive ELMs is still unknown. The most commonly invoked theory to explain the occurrence of ELMs is the peeling-ballooning model which posits a critical edge pressure gradient and current density. In this thesis, this model is tested against experimental data gathered at the ASDEX Upgrade (AUG) tokamak. For the first time, a broad selection of ELM scenarios is analysed with respect to ideal MHD stability using the same methodology. The comparison of experiment and theory is performed using a stability analysis chain, which consists of combining kinetic and magnetic measurements to generate self-consistent plasma equilibria with the Grad-Shafranov solver CLISTE, refining this equilibrium with the HELENA code, and, finally, determining its stability using ILSA, a linear MHD stability code.

In theory the peeling ballooning model should apply to all type-I ELM scenarios. Therefore, the stability of several different type-I ELMy H-mode plasmas is analysed with respect to peeling ballooning modes. While some of them are consistent with the model, in others ELMs occur well below or above the ideal MHD stability limit. The standard type-I ELMy H-mode regime exhibits considerable variations with equilibria both well below and at the stability limit depending on the discharge. In addition, a nitrogen-seeded case in which the edge pressure gradient greatly exceeds the stability limit is identified. In another discharge, the edge pressure gradient and current density, which are on the threshold for marginal stability, relax when edge heating is applied. Contrary to the theory, as the equilibrium becomes more stable against peeling-ballooning modes, the ELM frequency actually increases.

Other scenarios in which peeling-ballooning modes are not thought to be the ELM trigger are also examined. In discharges featuring type-II and type-III ELMs the edge pressure gradient and current density are well below the threshold for peeling-ballooning instabilities, consistent with expectations. This indicates that other mechanisms must be present that prevent the kinetic profiles from building up. The behaviour of the analysed type-III ELMs is consistent with the hypothesis that they are driven by resistive interchange modes.

The feasibility of triggering ELMs via frozen deuterium pellets is also examined in detail. In discharges performed when the walls of AUG were protected by carbon tiles, pellets triggered ELMs whenever they were injected into an H-mode plasma. In the tungsten-walled AUG, however, it is only possible to trigger ELMs after a specific lag time has elapsed since the previous ELM. This lag time can be significantly reduced, though not eliminated, when injecting nitrogen into the divertor. In the tungsten-walled AUG, with and without nitrogen, it is possible to trigger ELMs via pellets when the pedestal pressure gradient and the edge current density are 20% below the threshold for spontaneous ELMs. Since the ELM-induced collapse of the plasma edge pressure profile is strongest in the unseeded tungsten machine and the profile recovery is slower, the lag time for successful triggering is longer.

Recent extrapolations of the ITER pedestal are also tested using the peeling-ballooning model. It is found that the predictions are very optimistic, but that stability could be achieved with temperatures and densities which are 30% lower than in the extrapolations.

Some of the results presented in this thesis question the validity of the peeling-ballooning theory. In many of the presented cases, ELMs occur well below the threshold for peeling-ballooning stability and in one case well above. The presented results indicate that the peeling-ballooning model is insufficient to fully explain the triggering of ELM instabilities. Rather, another as yet unknown trigger mechanism is required.

Contents

1	Introduction	1
1.1	High confinement mode	4
1.2	Edge localised modes	6
1.3	Scope of this work	8
2	Theory and codes	9
2.1	Bootstrap current	12
2.2	Equilibrium, CLISTE and HELENA	13
2.3	ILSA	14
2.4	Stability chain	15
3	Heating systems and diagnostics	19
3.1	Neutral beam injection	20
3.2	Electron cyclotron resonance heating	20
3.3	Magnetic measurements	23
3.4	Charge exchange recombination spectroscopy	23
3.5	Electron cyclotron emission	24
3.6	Thomson scattering	25
3.7	Lithium beam impact excitation spectroscopy	26
3.8	DCN interferometer	26
3.9	Integrated data analysis	27
3.10	Combining the diagnostics	27
4	Optimisation and limitations of the stability chain	31
4.1	Comparison between different codes	31
4.2	Convergence	37
4.3	Choice of Ψ_b	39
4.4	Effect of Ψ_b on the mode width	40
4.5	Ion diamagnetic drift stabilisation	41
4.6	Finite resistivity effects	42
4.7	Effects of a realistic density profile	43
4.8	Influence of measurement errors	44
4.9	Final parameters used in the stability calculations	46
5	Results	49
5.1	Type-I ELMs: peeling-ballooning stability in typical H-modes	49
5.2	Type-III ELMs: small ELMs close to the L-H threshold	56
5.3	Type-II ELMs: strongly fuelled discharges close to double null	58
5.4	Type-I ELM mitigation at high densities	61
5.5	N-seeded ELMs: ELM mitigation by nitrogen seeding	65
5.6	ELM triggering via frozen deuterium pellets	70
5.7	Influence of edge ECRH heating on peeling-ballooning stability	78
5.8	Discussion	84
6	Implications for ITER	87

7	Conclusions and outlook	93
	Appendices	97
A	Derivation of the MHD equations used in ILSA	97
B	Solver	101

1 Introduction

As a result of the rising energy demand and the growing desire to be independent of fossil fuels, it is now more important than ever to explore new energy sources. Nuclear fusion is a promising candidate to take over the base load of energy production and significant progress has been made in the last decades on the path to a working power plant prototype. Present day fusion research mainly looks into the hot fusion of deuterium (D) and tritium (T) [1]. Of all considered reactions, D-T has the highest cross section with a maximum at achievable collision energies. This fusion reaction produces energy according to the equation

$$D + T \longrightarrow {}^4\text{He} + n + 17.59\text{MeV}, \quad (1.1)$$

where the created helium nucleus carries 3.54MeV and the neutron 14.05MeV of the kinetic energy. Reactor designs foresee that the fuel, which at fusion relevant temperatures will take the form of a plasma, will burn in a sealed vacuum chamber. The walls, being bombarded by the hot neutrons created in the fusion reaction, will have to be actively cooled. The water or liquid metal that is used as coolant will then be used to create electric power through the use of turbines as is done in conventional power plants. The use of a liquid metal coolant containing lithium is promising because it allows to breed the tritium required for the operation of the reactor [2]. Current design concepts suggest that lithium could flow through a so-called breeding blanket, absorb neutrons created in the D-T reaction and decay into tritium and Helium.

To maximise the fusion rate three main parameters have to be optimised. The first one is the energy confinement time τ_E , which is defined as the total energy of the plasma divided by the amount of power necessary to sustain it. The density n must also be maximised, because the more particles that reside in a given volume, the higher the chance that two will collide. The optimal value for the third parameter, the temperature T , was determined via power balance calculations to be 150 million degrees centigrade, or 13keV. Ideally, after this temperature is reached via auxiliary heating methods it will be sustained by collisions between the hot alpha particles created by fusion reactions and the thermal plasma particles. We call this situation ignition.

In a fusion reactor, contact between the plasma and the reactor walls must be avoided because at fusion relevant temperatures contact would damage the vacuum vessel and immediately cool down the plasma, shutting down the self-sustained fusion reactions. One of the most promising approaches to contain the plasma is magnetic confinement fusion (MCF), which consists of trapping the charged plasma particles on magnetic field lines. While they can travel freely along the field lines, their perpendicular velocity is limited to a gyro-motion. Present day MCF research mainly explores toroidal configurations because they avoid the end losses that occur in linear devices. However, this configuration gives rise to other particle losses. The general effect of a force on charged particles that gyrate around magnetic field lines can be derived from the guiding centre ansatz [4]. A force perpendicular to the magnetic field will result in a drift of the gyrating particles. This drift is perpendicular to both the force and the magnetic field and is given by

$$\mathbf{v}_D = \frac{\mathbf{F} \times \mathbf{B}}{qB^2}. \quad (1.2)$$




Figure 1.1: The $\vec{E} \times \vec{B}$ -drift results in radial losses [3]

Two drifts, which both point in the vertical direction, arise directly from the toroidal magnetic configuration, the curvature drift and the ∇B drift. The first one is due to the centrifugal force, since the hot particles can move freely on the magnetic field lines, effectively following a circular trajectory. Acting radially outwards, this force causes an upward drift of the ions and a downward drift of the electrons for the geometry shown in figure 1.1 because of the charge dependence in equation 1.2. The second drift is due to the inhomogeneous magnetic field. The magnetic field on the inside of the ring will necessarily be stronger than on the outside, which is why the inboard side is also called the high field side (HFS) and the outboard side the low field side (LFS). Since the gyro-radius depends on the local magnetic field, the orbit of the charged particles in a tokamak is not perfectly circular. This results in a drift in the same direction as the curvature drift. The charge dependence of those two drifts leads to a charge separation (see figure 1.1). The resulting electric field, pointing in the z -direction, creates an $\vec{E} \times \vec{B}$ drift that causes the plasma to drift radially outwards so quickly that the discharge can not be sustained for long. To prevent charge separation it is sufficient to add a magnetic field in the poloidal direction, thereby twisting the magnetic field lines and redistributing the particles evenly in the plasma.

In the tokamak concept, which has been studied since the 1950s, the poloidal field is generated by driving a toroidal plasma current. This current is induced using a transformer coil, a solenoid placed in the centre of the torus. The plasma itself, with a low resistivity of around $10^{-9}\Omega\text{m}$, acts as the secondary winding of a transformer. In present-day devices, a loop voltage of around 1V is usually sufficient to drive a current greater than 1MA. Like a transformer, a tokamak is inherently run in pulsed mode, but new techniques are currently being developed to drive non-inductive current, thereby increasing the discharge length. For example, current can be driven using microwaves as will be described in section 3.2.

As can be seen in figure 1.2a, the twisted field lines in a tokamak plasma form closed surfaces, called flux surfaces. The center of the innermost flux surface is called the magnetic axis. After a certain number of toroidal and poloidal turns, depending on the toroidal magnetic field and the current density profile, a twisted field line closes onto itself. These surfaces are called rational flux surfaces. This is the case when the




Figure 1.2: *Simplified representation of a tokamak (left) [3], cross-section of the ASDEX Upgrade tokamak (right) [IPP database]*

so-called safety factor

$$q = \frac{r}{R} \frac{B_t}{B_\theta} \quad (1.3)$$

is a rational number. R is the major radius (distance of the magnetic axis to the central axis of the transformer coil, see figure 1.1), r is the minor radius (distance of the considered point from the magnetic axis), B_t is the toroidal magnetic field and B_θ is the poloidal magnetic field. Figure 1.2b shows the poloidal cross section of ASDEX Upgrade (AUG) including a plasma. The innermost flux surface is not at the centre of the plasma, but has a slight outward shift called the Shafranov shift. It is due to the kinetic pressure of the plasma and the $\mathbf{j} \times \mathbf{B}$ force from the plasma current, which also exists in a simple current loop [5]. These two forces, called hoop forces, move the plasma radially outward. Applying a homogeneous magnetic field in the z-direction using a pair of Helmholtz coils prevents this radial movement of the plasma (vertical field coils in figure 1.2a).

An actual tokamak plasma like the one sketched in figure 1.2a would extend out to the wall or other limiting components and is, therefore, called a limiter plasma. Such limiter plasmas have the disadvantage that the bulk plasma is in constant contact with the wall or limiter materials, thereby damaging the material structure and accumulating the melted or sublimated materials in the main plasma. The higher the charge Z of these impurities, the more they cool the plasma because of high bremsstrahlung and, if the temperature is not high enough for them to be fully stripped of electrons, line emission. If the plasma cools down too much it can lead to a current quench, disrupting the discharge. While lower Z impurities don't have such a great impact on the plasma temperature, they are also unfavourable since they dilute the D-T fuel, thereby decreasing the fusion reaction rate in a reactor. Uncontrolled impurity accumulation must, therefore, be avoided at all cost. This is especially true for machines like AUG [6],

the recently upgraded Joint European Torus [7] and ITER [8], which all have, or will have in the case of ITER, a full metal wall containing the high Z element tungsten (W). To control the impurity content in the bulk plasma, these three machines, as well as most modern tokamaks, were designed using the divertor principle. Instead of the main plasma being limited by the wall, an additional coil which is usually mounted under the vessel carries a current parallel to the plasma current. The magnetic field from this coil and the poloidal field from the plasma B_θ combine to zero at the so-called X-point (see poloidal cross section of the plasma in figure 1.2b). The flux surface containing the X-point is called the last closed flux surface (LCFS) or separatrix, because it separates the closed flux surfaces inside from the open ones outside. Typically, AUG plasma discharges are run in lower single null (LSN), which means that the X-point is at the bottom. However, it is also possible to run discharges in double null (DN) with the separatrix containing two X-points, one at the bottom and one at the top, or in upper single null (USN). The field lines outside of the separatrix hit vessel components before closing onto themselves, which means that the plasma is not contained, or scraped off. This region is, therefore, called the scrape off layer (SOL). The open flux surfaces in the divertor geometry result in the presence of two separatrix legs under the X-point. Since the velocity of charged particles parallel to the field lines is much higher than their radial velocity, most plasma losses crossing the separatrix into the SOL travel toward the wall parallel to these legs. For this reason, so-called divertor plates are mounted in the path of the separatrix legs (see figure 1.2b). These plates are usually built from carbon or metals with a high melting point like tungsten, and can be actively cooled. Since in a divertor tokamak the bulk plasma is not in contact with vessel components, higher core plasma temperatures can be reached and the impurity content can be kept low. However, all particle and non-radiative energy losses from the plasma impact the divertor plates on a very small surface, which may cause significant problems for future bigger machines like ITER due to high heat fluxes [9].

1.1 High confinement mode

While in a tokamak the individual ions and electrons are confined to magnetic field lines, turbulence and collisions between particles cause radial transport, leading to particle losses. Turbulent transport, which increases with temperature and density gradients, imposes limits to these gradients. As a consequence, density and temperature profiles are constrained by profile stiffness. A profile $f(r)$ is stiff if

$$\frac{d}{dr} \ln(f(r)) = \text{const.} \quad (1.4)$$

The low confinement mode (L-mode) was the first mode of operation run in a tokamak. To illustrate the effect of profile stiffness, a typical L-mode pressure profile is shown in green in figure 1.3.

However, a revolutionary discovery was made at the ASDEX tokamak in 1982 [10,11], when a new operational regime was found that overcomes profile stiffness at the plasma edge. This new regime is called high confinement regime, or H-mode. Figure 1.3 shows representative H- and L-mode pressure profiles, plotted against the minor radius normalised to the separatrix location. While core gradients are still limited by profile stiffness, in H-mode the edge profile is much steeper than in L-mode. Therefore, the




Figure 1.3: *Representative plasma pressure profiles during different modes of operation*

whole profile sits on a so called pedestal, dramatically increasing the confined amount of energy and particles. In most present day tokamaks, H-mode is achieved by increasing the heating power of the plasma over a certain density-dependent threshold [12]. The steeper edge gradient originates from an edge transport barrier (ETB), a region at the edge of the confined plasma in which radial particle and heat transport is significantly reduced. In general, turbulent convection cells limit the pressure gradient by increasing the radial transport of plasma particles. These convection cells can be torn apart if the




Figure 1.4: *Schematic representation for the reduction of turbulent convection cells (a) and experimental evidence of reduced density fluctuations achieved in H-mode operation on AUG (b) [3]. The dashed red line indicates the transition into H-mode.*

$\mathbf{E} \times \mathbf{B}$ rotation velocity of the plasma is sheared. The remaining cells are much smaller, as is shown schematically in figure 1.4a. It has been shown that the $\mathbf{E} \times \mathbf{B}$ rotation velocity shear at the edge is higher in H-mode than in L-mode [13], supporting the idea that this shear is responsible for the ETB and leads to an L-H transition. Experimental evidence confirming the reduction of turbulence when the plasma reaches H-mode is

presented in figure 1.4b. In this spectrogram, the level of density fluctuations measured via Doppler reflectometry is shown. When the plasma reaches H-mode at $t=1.705$ s the turbulence level is drastically reduced.

1.2 Edge localised modes

The steep gradients associated with ETBs provide a source of free energy for magnetohydrodynamic (MHD) instabilities. Therefore, H-mode plasmas are usually accompanied by the appearance of so-called edge localised modes (ELMs) [14]. These periodically occurring instabilities expel particles and energy from the plasma edge into the SOL in short time intervals of around 1ms. While smaller ELMs only affect the pedestal region, lowering the edge gradients and the height of the pedestal top, in some cases the ELM-affected area can extend in to the mid-radius of the plasma. In figure 1.5 ELMs are indicated by vertical dashed lines. Panel a shows the time traces




Figure 1.5: Time traces of divertor shunt measurements (a), the edge line averaged electron density (b), the edge electron temperature (c), the plasma stored energy (d) and the position of the outer separatrix (e) in a typical AUG H-mode discharge.

from shunt measurements in the outer (black) and inner (red) divertor of AUG, which are roughly proportional to the divertor temperature. After every ELM the divertor temperature increases for a couple of milliseconds due to the hot plasma particles arriving along the open field lines. The loss of these hot particles causes the edge density (panel b) and temperature (panel c) to drop. A corresponding drop in the total energy confined in the plasma can also be seen in panel d. ELMs also lead to a contraction of the plasma volume, which means that the separatrix moves inwards at the outboard side (panel e) and outwards at the inboard side (not shown) [15]. ELMs that occur in a stable discharge with constant global plasma parameters are typically very reproducible. Figure 1.6 shows how the maximal edge n_e and T_e gradients evolve over the course of 105 type-I ELM cycles, which is the most common ELM type occurring in AUG H-mode plasmas. The data are synchronised to the ELM onset. After the sudden




Figure 1.6: *Maximal electron temperature (a) and density (b) gradient over a full ELM cycle and selected profiles (c-j) [16]*

collapse of the gradients the recovery takes place over several distinct phases [17]. The origin of the ELM losses becomes obvious when comparing panel c with d and e with f in figure 1.6, which show the edge electron temperature and density profiles at two different times in the ELM cycle. ρ_p is the normalised minor radius and is defined in equation 3.5. Outside of a normalised radius of 0.9 the temperature profile collapses. The density profiles pivots around the separatrix. Between 0.85 and 1.0 the density decreases, while the SOL density increases.

ELMs can typically be categorised into three types [14]:

- Type-I ELMs are the biggest ELMs occurring in tokamaks. They are also the most common ones since they occur in strongly heated H-mode discharges with high confinement, and most experiments are run using scenarios of optimised plasma performance. The induced losses are typically on the order of 5% to 10% of the total plasma stored energy, but can also reach up to 40%. Type-I ELMs can be identified by their frequency, which increases with higher heating power, and by their size.
- Type-II ELMs occur at high densities in plasmas close to double null. Unlike type-I ELMs, which appear as big distinct events in the divertor temperature, edge temperature, and edge density signals, type-II ELMs cannot necessarily be distinguished from one another. Rather, they appear as fluctuation-like events. In addition, type-II ELMs are characterised by broad band magnetic fluctuations between 30 and 50kHz [18, 19]. The per-ELM energy losses induced by type-II ELMs are very small.
- Type-III ELMs occur just above the H-mode power threshold. Starting from very small events with a frequency above 1kHz when the heating power is just above the threshold, they gradually change into bigger, more easily distinguish-

able crashes whose frequency decreases as the heating power is ramped up. Eventually, if the power is increased sufficiently, the type-III ELMs give way to the larger type-I ELMs. Type-III ELM losses typically lie under 2%.

ELMs do have beneficial effects, since they flush out impurities, cleansing the plasma, but the released energy may still pose a danger to the machine. In current fusion experiments the ELM induced heat loads on the divertor target plates are typically under the material limits of the plasma facing components (PFC). Future fusion reactors, however, will contain a much bigger plasma volume than present-day devices and their ELM-induced peak heat fluxes are predicted to greatly exceed the tolerable limit of 10MWm^{-2} [20].

1.3 Scope of this work

Since ELMs will pose a threat to the PFCs of future fusion devices, it is important to find ways to mitigate or suppress them. However, we still don't have a fundamental understanding of what drives them. While several theories attempt to explain the occurrence of ELMs, the most prevalent one being the peeling-ballooning theory which will be presented in chapter 2, they do not always succeed in describing the experimental data. The aim of this thesis is to analyse the main ELM regimes accessible in AUG. The experimental data will be compared to the theoretical peeling-ballooning model in order to benchmark the theory. Several cases in which this widely accepted model successfully predicts the occurrence of ELMs will be discussed. However, other experimental conditions will also be presented in which the model fails, as the ELM occurs far below or above the peeling-ballooning threshold for stability.

In addition to the peeling-ballooning theory, chapter 2 will introduce the concept of a plasma equilibrium, present the stability code ILSA and describe the stability chain used in this work. In chapter 3 the heating systems employed in the presented discharges will be introduced, as well as the diagnostics used for the analysis. The numerical limits of ILSA, a convergence study of different input parameters and the influence of measurement errors on the final result of the code will be discussed in section 4. The differences between typical type-I ELMs and other ELM regimes, as well as the comparison of experimental data to the peeling-ballooning theory will be presented in chapter 5. Important implications of these results for ITER will be discussed in chapter 6. Finally, chapter 7 will give a short summary of the results presented in this thesis.

2 Theory and codes

The theoretical interpretation of ELMs has been the subject of intense research for the last 30 years. The most widespread theoretical approach to explain the occurrence of ELMs is the peeling-ballooning theory. It assumes a coupling between two type of magnetohydrodynamic (MHD) instabilities, the pressure gradient driven ballooning modes and the current driven peeling modes.




Figure 2.1: *Schematic representation of an $n=3$, $m=9$ mode. The left figure is a horizontal cut through the 3D plot in the middle, the right plot is the poloidal cross section.*

In general, modes in a tokamak can be described by their mode numbers n and m , where n is the number of toroidal periods of the perturbation and m is the number of poloidal periods. Figure 2.1 shows the deformation of the magnetic field caused by an $n=3$, $m=9$ instability. The color coding in the 3D picture (b) denotes the amplitude of the perturbation. The deformation shown in the toroidal cross-section (a) is strongly exaggerated for clarity. The poloidal cross-section is shown in panel (c). After 9 toroidal and 3 poloidal turns around the torus, the mode closes onto itself. Since modes typically form standing waves along the field lines, causing specific modes to be resonant on rational q surfaces according to $m = nq$, this mode is most likely to be localised on the $q=3$ surface.

Ballooning modes are a specific form of pressure driven interchange instability. Interchange instabilities try to lower the energy of the system by interchanging two flux tubes. If the magnetic field line curvature is concave toward the plasma (unfavourable curvature), interchanging a longer flux tube with lower pressure from outside with a shorter flux tube with higher pressure further inside can lower the potential energy. While this is the case at the LFS, such an exchange at the HFS of the plasma where the field line curvature is convex toward the plasma (favourable curvature) would lead to a higher potential energy. In a tokamak, where the twisted field lines pass through regions of favourable and unfavourable curvature, the average curvature is the determining factor. A sheared magnetic field can stabilise interchange instabilities because it would require a minimum amount of field line bending when interchanging two flux tubes. In a tokamak, this stabilisation is given by the Mercier criterion

$$-8 \frac{\mu_0 p'}{B_0^2} (1 - q^2) < r \left(\frac{q'}{q} \right)^2 \quad (2.1)$$

with the prime denoting the derivative to the minor radius r and B_0 is the toroidal magnetic field at the magnetic axis [21]. Since p' is negative, in a tokamak interchange instabilities are always stabilised if q is greater than 1. However, if the amplitude of a perturbation varies along a field line and is stronger in regions of unfavourable curvature, the argument of average curvature is no longer valid. The resulting instabilities, so-called ballooning modes, can exist if the pressure gradient is too high for the field line bending that accompanies the modes to be able to stabilise them. These modes are localised on the low field side of a tokamak, typically have high toroidal and poloidal mode numbers ($n=20$ and higher) and a small radial extent. Figure 2.2 shows

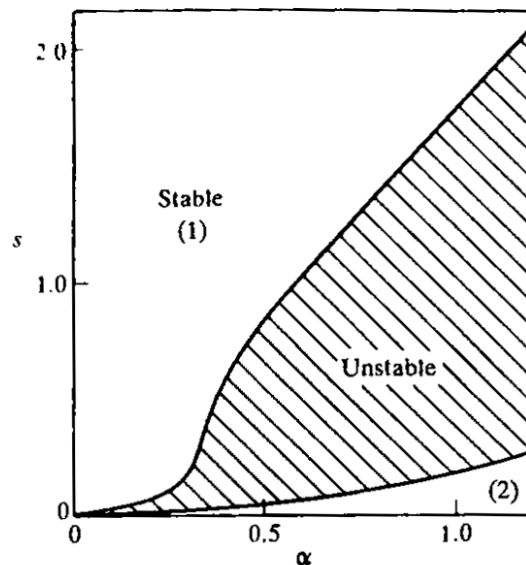


Figure 2.2: s - α diagram showing the stable (white) and unstable (dashed) region for infinite- n ballooning modes [1]

the stable and unstable region for ballooning modes in $s - \alpha$ space, where α is the normalised pressure gradient and s the magnetic shear. Like interchange instabilities, ballooning modes are stabilised by magnetic shear. However, a second stability region exists at very low shear and high pressure gradient in which increasing the shear would destabilise the plasma (region (2) in figure 2.2). In this simple picture the second stability regime is not accessible because the unstable region extends to the origin of the $s - \alpha$ space. However, a strong shaping of the plasma can alter the stability boundary in such a way that the second region is accessible. In this case, a high localised current density can lead to a flat local q -profile and stabilise the ballooning mode.

Peeling modes are another type of instabilities which, at low plasma pressure, are driven by the toroidal current density and its radial gradient [1, 22]. Peeling modes can be unstable if the resonant q surface ($q=m/n$) is located in the vacuum region outside of but close to the LCFS and are stabilised if the resonance resides inside the confined plasma. These modes have a broader radial extent and typically low n and m mode numbers. It is believed that a coupling that takes place between the peeling and the ballooning modes is responsible for the occurrence of ELMs. An overview of this coupling with a focus on the role of the edge current density is given by Snyder *et al.* [23].




Figure 2.3: *Schematic representation of an ELM-cycle in terms of normalised pressure gradient and edge current density [24]*

The main current that contributes to the destabilisation of the peeling modes, the bootstrap current, is introduced in section 2.1. To understand the following description it is important to know that this current is driven by the pressure gradient. Figure 2.3a shows a schematic representation of the theory describing a type-I ELM cycle $j - \alpha$ space, where j is the current density [24]. The red curve marks the peeling-ballooning boundary. In the upper left region the unstable modes are primarily current-driven (peeling) whereas in the lower right corner they are mainly pressure gradient driven (ballooning). The model suggests that an ELM cycle consists of three phases:

1. The pressure gradient builds up until it reaches the ballooning limit, but the current density has a certain resistive delay.
2. The bootstrap current builds up as well, until the joint peeling-ballooning limit is reached.
3. The ELM crash occurs and both the pressure gradient and current density relax, until the cycle starts again.

However, the delay in the recovery of the bootstrap current mainly depends on the current diffusion time scale. It was shown both using a simple theoretical model [17] and experimentally [25] that at AUG this only plays a minor role. The pressure gradient and the edge current density recover nearly simultaneously.

After introducing the bootstrap current in section 2.1, section 2.2 discusses the concept of plasma equilibria and presents the codes that were used in this work to generate them. Computing such equilibria is necessary because the instabilities described in the peeling-ballooning theory do not only depend on local parameters but also on global profiles, and because modes that have their maximum amplitude at different radial locations can still couple. The stability of a plasma against peeling-ballooning modes can be calculated numerically by using codes like ILSA, which was the main tool used in this work. The equations used in ILSA are described in section 2.3 and the solver itself is presented in appendix B. Rather than just calculating the stability of the generated equilibrium it is also interesting to find out where it lies with respect to the peeling-ballooning boundary in $j - \alpha$ space. A description of how to create experimental

diagrams analogous to the sketch shown in figure 2.3 and an overview on the stability chain that was run to produce the results presented in this thesis are given in section 2.4.

2.1 Bootstrap current

All theories that attempt to explain ELMs describe them as a combination of pressure gradient and current driven instabilities. In an H-mode plasma, a high edge pressure gradient results from the ETB (see section 1.1). In turn, this gradient leads to a high edge current density, the so-called bootstrap current, and together they drive the peeling-ballooning modes.

The bootstrap current is due to magnetic mirror effects which occur as a result of the conservation of both the kinetic energy $E = mv^2/2$ and the magnetic moment $\mu = mv_{\perp}^2/2B$ [1, 26], where m is the mass, v the total velocity and v_{\perp} the gyration velocity of the charged particle. Particles with a high v_{\perp}/v_{\parallel} ratio are reflected at the HFS of the plasma when the magnetic field meets the condition

$$B = B_{min} \left[1 + \left(\frac{v_{\parallel 0}}{v_{\perp 0}} \right)^2 \right], \quad (2.2)$$

with v_{\parallel} being the velocity parallel to the magnetic field lines. The index 0 denotes the velocities at the position of minimum magnetic field B_{min} along the particle orbit.




Figure 2.4: *Banana orbits (a) and parallel velocity distribution of charged particles in a tokamak without (b) and with (c) collisions [26]*

In a tokamak the magnetic field is roughly proportional to the inverse of the major radius R and this condition is met for particles whose velocity distribution satisfies $(v_{\parallel 0}/v_{\perp 0})^2 < 2r/(R_0 - r)$. These particles are called trapped or banana particles because they bounce back and forth on an orbit, the poloidal projection of which resembles a banana as can be seen in figure 2.4a. In the presence of a density gradient, more particles will orbit in one direction than in the other at any given point, leading to a net current. The overpopulation of trapped particles with negative v_{\parallel} can be seen in figure 2.4b. However, the resulting current only plays a minor role compared to the total bootstrap current. Rather, the trapped particles give some of their parallel velocity to passing ions and electrons via collisions, which leads to the velocity distribution shown in figure 2.4c. The passing ions and electrons carry the bulk of the bootstrap current. While in first order the bootstrap current is roughly proportional to the pressure gradient, the more complete formula derived by Sauter *et al.* [27, 28] was used to predict the ITER current density necessary for chapter 6.

2.2 Equilibrium, CLISTE and HELENA

The general concept of a tokamak equilibrium is that a force balance is achieved such that the topology of the nested flux surfaces does not change over time. Therefore, no net force can act on the plasma, which means that the kinetic pressure must be balanced by the magnetic force at all points following

$$\mathbf{j} \times \mathbf{B} = \nabla p, \quad (2.3)$$

which implies $\mathbf{j} \cdot \nabla p = 0$ and $\mathbf{B} \cdot \nabla p = 0$. Therefore, the magnetic field lines and the direction of current flow lie on surfaces of constant pressure. At each surface, Ψ is defined as the magnetic flux inside this surface. Since Ψ is constant on a flux surface, this also implies $\mathbf{B} \cdot \nabla \Psi = 0$.

For a cylindrical coordinate system in which the z-axis corresponds to axis of symmetry of the toroidal configuration, an equation can be derived that enables us to calculate the field line topology of an axisymmetric plasma [1, 21]. This equation, known as the Grad-Shafranov equation, takes the form

$$R \frac{\partial}{\partial R} \frac{1}{R} \frac{\partial^2 \Psi}{\partial R^2} + \frac{\partial^2 \Psi}{\partial z^2} = -\mu_0 R^2 p' - \mu_0^2 f f', \quad (2.4)$$

where R and z are the coordinates in the cylindrical system, $f = RB_\phi/\mu_0$ and the prime denotes the derivative with respect to Ψ . The right hand side of the equation contains the toroidal current density

$$j_\phi = Rp' + \frac{\mu_0}{R} f f'. \quad (2.5)$$

Two main types of solvers for the Grad-shafranov equation exist, predictive and interpretative solvers. While predictive codes require the p' and $f f'$ profiles to be known, interpretative solvers infer them from magnetic field measurements from coils mounted in the vessel outside the plasma. In this work, both types of codes were used.

The interpretative equilibrium solver CLISTE [29, 30] is a further development of the Garching equilibrium code written by K. Lackner [31]. It uses data from flux loops and poloidal and radial magnetic field measurements at AUG (see also section 3.3) to generate a basic magnetic equilibrium. While the total normalised pressure β and current I_p are very accurate, the strong edge variation of the pressure gradient and current density profile is only poorly reflected. Constraining the edge pressure profile leads to a much higher accuracy in the final equilibrium. This can be seen in figure 2.5, in which the black profiles were obtained by constraining the edge pressure profile in CLISTE and the red profiles by only using magnetic data. With kinetic constraints the pressure profile (a) and its gradient (b) from the CLISTE equilibrium (solid black line) match the experimental data very well (black crosses). When the pressure profile is not constrained, however, the pressure profile is almost straight (red). Only a very small peaking of the pressure gradient can be seen in the pedestal region. The local toroidal edge current density at the LFS (c, solid lines) and the flux-surface-averaged $\langle \mathbf{j} \cdot \mathbf{B} \rangle$ (c, dotted lines) are also much higher when the pressure profile is used as a constraint to the equilibrium calculation. Without kinetic constraints the safety factor q (d) and the magnetic shear increase monotonically toward the edge, while a region of low shear exists between a normalised radius of 0.98 and 0.99 when the experimental




Figure 2.5: Output profiles of CLISTE with (black) and without (red) kinetic constraints, showing the pressure profile (a), its gradient (b), the edge current density (c) and the safety factor (d)

pressure profile is included. The pressure gradient, the edge current density and the magnetic shear, which are very important quantities in assessing the stability of the plasma, are much more accurate when including the experimental pressure profile as a constraint in CLISTE. In all equilibria presented in this work, the edge pressure profile was, therefore, constrained by using the data from a combination of the diagnostics presented in chapter 3.

In order to achieve the high degree of numerical accuracy required by stability codes, it is necessary to re-calculate the equilibrium by using a predictive equilibrium solver. In addition to the source profiles p' and $f f'$, the fixed boundary code HELENA [32] used throughout this work requires the shape and location of the last closed flux surface. Both HELENA and the stability code ILSA work in field aligned coordinates, which implies that they cannot resolve the separatrix at which the safety factor q and the magnetic shear reach infinity. Therefore, the equilibrium has to be cut off at a certain radius which we call the cutoff parameter. The influence of this parameter on the final stability of the plasma is discussed in sections 4.1 through 4.4. The high resolution HELENA equilibrium serves as an input for the stability calculations to assess whether a plasma is stable against peeling-ballooning modes.

2.3 ILSA

In general, two methods can be used to determine the stability of a system. The energy principle method [33] is based on the change δW in potential energy of the considered system when a small perturbation ξ is introduced. If a perturbation function $\xi(\mathbf{r})$ can be found such that δW becomes negative, the system is unstable. However, this method only calculates whether the system is stable or not, whereas the normal mode technique, which is used in the stability codes CASTOR and MISHKA, also gives a linear growth rate. The normal mode technique solves linearised MHD equations for

small perturbations. If a solution with a finite growth rate is found, the system is considered to be unstable.

Appendix A describes the equations solved by the linear MHD codes CASTOR and MISHKA-1 [34–36]. Starting from the general single-fluid MHD equations, one obtains the following closed set of four equations, describing the nonlinear evolution of the density ρ , the temperature T , the velocity \mathbf{v} and the magnetic field \mathbf{B} :

$$\frac{\partial \rho}{\partial t} + \nabla(\rho \mathbf{v}) = 0, \quad (2.6)$$

$$\frac{\partial T}{\partial t} + (\mathbf{v} \cdot \nabla)T + (\Gamma - 1)T \nabla \cdot \mathbf{v} = 0, \quad (2.7)$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho(\mathbf{v} \cdot \nabla)\mathbf{v} = \frac{1}{\mu_0}(\nabla \times \mathbf{B}) \times \mathbf{B} - \nabla p, \quad (2.8)$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \left(\mathbf{v} \times \mathbf{B} - \frac{\eta}{\mu_0} \nabla \times \mathbf{B} \right). \quad (2.9)$$

When setting the resistivity η to zero, these equations describe the ideal MHD model. In ILSA, Γ can be freely chosen. For instance, $\Gamma = 1$ corresponds to an isothermic process and $\Gamma = C_p/C_V$ to an adiabatic one, with C_p and C_V the heat capacities at constant pressure and volume, respectively. However, since incompressibility was assumed in the stability calculations performed in this thesis, Γ was set to zero. After linearising the problem one obtains four equations, A.24 through A.27, which are solved by the CASTOR code. An extended version of the code also takes into account the viscosity term Π [37].

A faster MHD code, MISHKA-1, uses a reduced set of equations [36]. It assumes an equilibrium velocity $\mathbf{v}_0 = 0$ and takes into account only a pressure perturbation p_1 instead of treating ρ_1 and T_1 separately. After the linearisation, one obtains equations A.32 through A.34. These equations describe the general stability problem in terms of the perturbed variables \mathbf{v}_1 , p_1 and the perturbed vector potential \mathbf{A}_1 . Since in MISHKA-1 the temperature and density perturbations are treated as one, the mass continuity equation A.1 is not explicitly used.

The stability code ILSA [34] can solve the full set of linearised equations A.24–A.27, or it can be run in MISHKA-1 mode (equations A.32–A.34). The solver itself is described in Appendix B.

In this work, ILSA was run in MISHKA-1 mode using the ideal equations. Note that in ideal MHD the eigenvalue λ is real, which means that the mode purely grows and does not oscillate. However, the calculations are still performed using complex numbers and the imaginary part of the result can serve as an indicator for the numerical accuracy of the calculated growth rate γ , according to $\gamma = \text{Real}(\lambda) \pm \text{Im}(\lambda)$.

2.4 Stability chain

The goal of this work is to compare the different ELM regimes accessible at AUG to each other and to the widespread theoretical peeling-ballooning model. Data from the magnetic and kinetic measurements described in chapter 3 are used to generate an equilibrium with self-consistent pressure and current density profiles using CLISTE. The resulting p' and $f f'$ profiles, as well as the calculated plasma boundary, are then

used to generate a highly resolved equilibrium with HELENA. Before running the actual stability code ILSA, the so-called j-alpha workflow is run. This subroutine of HELENA allows us to systematically alter the current density and pressure profiles independently before running HELENA again [38]. In this work, the pedestal pressure width and the edge current density height were multiplied by different scaling factors. Figure 2.6 shows how the pedestal pressure (a) and current density (b) are modified using the j-alpha workflow. The profiles are plotted against Ψ_n , normalised to the value at which the equilibrium was cut off. The scaling factor 1.00 corresponds to




Figure 2.6: Pressure (a) and current density (b) profiles scaled independently using the j-alpha workflow, and corresponding stability diagram (c). The yellow star marks the operational point, whereas the boxes are explained in the text.

the unaltered experimental equilibrium. This equilibrium is typically referred to as the operational point. Since the pedestal pressure width is varied, a low scaling factor leads to a high pressure gradient. The current density height is varied independently, and it is increased by increasing the scaling factor. The core pressure and current density are adjusted such that the total pressure β and current I_p stay the same as in the reference equilibrium. After re-running HELENA using these profiles to obtain an equilibrium for each, ILSA is run in MISHKA-1 mode independently for different toroidal mode numbers to determine the stability of the generated equilibria. Thanks to the j-alpha workflow it is possible to generate $j - \alpha$ diagrams like the one shown in figure 2.6c, which is similar to the sketch shown in figure 2.3, but generated with real experimental data. In this work, the normalisation for the pressure gradient was taken from Miller *et al.* [24]:

$$\alpha = -\frac{2V'}{4\pi^2} \sqrt{\frac{V}{2\pi^2 R_0}} \mu_0 p', \quad (2.10)$$

where the prime denotes the derivative to Ψ and V is the volume. The numbers in figure 2.6c indicate the most unstable toroidal mode number at each grid point in the $j - \alpha$ space. Every number corresponds to an equilibrium re-calculated using a combination of one pressure and one current density profile from figure 2.6a and 2.6b, respectively. The resulting grid allows us to scan the stability boundary, which is shown as black lines. In the red region of the diagram the equilibria are unstable, in the blue region they are stable. The green and orange regions in between are called marginally unstable. The stability criterion is given in terms of the growth-rate γ normalised to

the Alfvén frequency ν_A . The critical value for stability chosen here is $\gamma/\nu_A = 0.06$. While this value is somewhat arbitrary, it accurately separates the region in which ILSA finds clear unstable modes from the one where it does not find any instabilities at all or only numerically unstable solutions. In the high current density and low α region the mode numbers are small (current driven peeling-like instabilities), while in the low current density and high α region high n numbers dominate (pressure driven ballooning modes). Note that because j and α were not scanned down to zero, the origin of the stability diagram is suppressed. Since I_p and β are kept constant when running the j - α workflow, the pressure and density gradients further inside the plasma would become very large when scanning the edge values down to zero and internal modes would become unstable.




Figure 2.7: *Radial mode structure of the toroidal mode with highest growth rate of three cases in figure 2.6c: the one marked by a green box (a, $n=4$), the operational point indicated by the yellow star (b, $n=12$) and the one marked in purple (c, $n=20$)*

The radial mode structures of the poloidal harmonics for the most unstable n of three cases from figure 2.6c are shown in figure 2.7: an equilibrium in the peeling unstable region (2.7a, green box in 2.6c), the operational point (2.7b, yellow star in 2.6c) and a ballooning unstable equilibrium (2.7c, purple box in 2.6c). The color coding, which indicates the poloidal mode number m , is not the same in the three plots because the numbers are very different. While for current driven modes with low n values only a few, but radially very extended, poloidal harmonics play a role, at high toroidal mode numbers many narrow harmonics contribute.

When running ILSA, unstable equilibria converge within minutes or a few hours depending on the radial grid resolution, the number of poloidal harmonics included in the calculation, and how close the prescribed starting eigenvalue is to the actual result. A stable equilibrium, however, will run until the number of allowed iterations is reached. Running the whole grid at once would be very time consuming. To speed up the process and avoid wasting computing resources, for every equilibrium the results from the previously run toroidal mode numbers are extrapolated to compute the starting value for the next run. Additionally, the calculations are started with the most unstable equilibria (high α and high j) and the results from those equilibria are used as inputs

for the surrounding grid points with slightly lower j and α . As soon as the boundary is reached, the calculations are automatically stopped. In practice, this means that the equilibria which, on the $j - \alpha$ grid, are surrounded by stable ones are not run in ILSA. This enables the boundary to be scanned without wasting computing power and disk space. While it seems from figure 2.6c that the gain resulting from this method is very limited, it saved a considerable amount of time and resources when running some of the diagrams shown in section 5 in which the stable region is much larger.

3 Heating systems and diagnostics




Figure 3.1: *Toroidal (a) and poloidal (b) positions of heating systems and diagnostics in the AUG vessel. The yellow arrows show the NBI paths and the ECRH mirrors are shown in black. The path of the DCN lasers are shown in red and the TS lasers in cyan. The lines of sight of the lithium beam (purple), the ECE diagnostic (green) and the toroidal (red) and poloidal (blue) edge charge exchange diagnostic are also shown. The yellow boxes represent the magnetic field measurement coils.*

To attain fusion relevant temperatures auxiliary heating methods are needed in addition to the Ohmic heating that is inherently present in tokamaks. In present day devices such as AUG, these methods are used to not only heat the plasma, but also as tools to alter it such that very specific physics questions can be addressed. Two different AUG heating systems, neutral beam injection (NBI) and electron cyclotron resonance heating (ECRH), were essential in designing the experiments used in this thesis. Figure 3.1a and b show the toroidal and poloidal positions of these heating systems. The path of the NBI lines are shown in yellow and the mirrors of the ECRH system in black. The operation of and physics principles behind these systems will be described in detail in sections 3.1 and 3.2, respectively.

The goal of this thesis is to analyse the stability of the edge pedestal against pressure and current driven instabilities. To achieve this, accurate magnetic equilibrium reconstructions of the plasma, particularly at the edge, are crucial. The methodology of reconstructing the equilibrium was presented in section 2.2. An equilibrium reconstruction requires, at a minimum, magnetic measurements of the plasma. The yellow boxes in figure 3.1b represent the B_θ coils, which are part of the magnetic measurement suite and will be presented in section 3.3. In addition, there is a poloidally distributed array of coils for flux difference measurements that is not shown in figure 3.1 in the

interest of legibility. Together, these sets of coils provide the basis for the magnetic equilibrium of AUG plasmas. However, the quality of an equilibrium construction can be significantly improved if the pressure profile is additionally prescribed. This requires high quality kinetic measurements, which at AUG are made via several different diagnostics at different toroidal and poloidal locations as shown in figure 3.1. Detailed information on all of the individual diagnostics is provided in sections 3.3 through 3.9.

3.1 Neutral beam injection

If highly energetic ions are present in the plasma, they collide with thermal particles and thereby transfer some of their energy. In a future fusion reactor, the 3.5MeV helium nuclei created in the fusion reaction will act as the main heating method. In present day experimental reactors, which carry out experiments with hydrogen or deuterium such as AUG, the same principle is applied using NBI. In AUG, a H^+ or D^+ ion beam is first accelerated to 60-90keV and then passes through neutral hydrogen where the fast beam ions capture electrons from the neutrals [39,40]. The neutralisation process is necessary because a hydrogen or deuterium ion beam would be deflected by the strong toroidal magnetic field of the tokamak and be unable to enter the plasma. As the neutral beam passes through the plasma, more and more of the fast neutrals are ionised by collisions with plasma particles and are trapped on magnetic field lines. While gyrating around the field lines, they thermalise due to further collisions, thereby heating the plasma. At relatively low beam energies like the ones at AUG, ion-ion collisions dominate. At higher energies, as will be the case with the 1 MeV negative-NBI in ITER and the 3.5 MeV alpha-particles from the fusion reaction, the collisions will mainly heat the electrons [1].

Charge exchange processes can occur between the neutrals from the NBI and ions in the plasma. How charge exchange processes with impurity ions can be used to determine the ion temperature is described in section 3.4. Since the NBI produces fast ions in the plasma, it is also used in the investigation of fast ion driven instabilities. These may occur in future reactors because of the high concentration of fast helium ions created during the fusion processes (see equation 1.1).

The eight neutral beams installed in AUG can produce 2.5MW of heating power each in D operation. While six of them have a fixed geometry, two tangential beams can be moved by $\pm 0.5^\circ$ to vary the main power deposition location between 0.4 and 0.6 of the normalised radius [41]. NBI was used as the primary heating system in all discharges analysed in this thesis.

3.2 Electron cyclotron resonance heating

The plasma can also be heated using electromagnetic waves [42]. Electrons gyrate around the magnetic field lines with a cyclotron frequency of

$$\omega_c = \frac{eB}{m_e} \tag{3.1}$$

where e is the elementary charge, B the magnetic field and m_0 the electron rest mass. A microwave with a frequency of ω_c or one of its harmonics can transfer its energy to the electron. At AUG the electrons are heated using the X2 wave, which means that

the electric field of the wave is perpendicular to \mathbf{B}_t and that the resonance occurs at the second harmonic $\nu_{\text{microwave}} = 2 * \omega_c / 2\pi$. Since AUG is typically operated with a magnetic field of 2.5T at the plasma core, the ECRH system is designed for 140GHz microwaves in order to achieve central heating. The microwaves are created in gyrotrons and then led through hollow wave guides into the vacuum vessel. The system is designed such that the incoming waves are reflected on rotatable mirrors before entering the plasma. Since the mirrors can be rotated around two axes, both the toroidal and poloidal deposition location of the heating power can be changed. The radial deposition location can be changed by adjusting the toroidal magnetic field, thereby moving the location of the resonance. This feature enabled the discharges presented in section 5.7 to be performed, in which the influence of edge heating on ELMs is investigated.

By injecting the microwaves tangentially into the plasma, it is also possible to drive a current. This is called electron cyclotron current drive (ECCD). Let us consider a wave whose \mathbf{k} vector has a strong toroidal component in the positive direction. Due to the Doppler shift the incident wave appears to have a slightly lower frequency for electrons with a positive velocity component parallel to the magnetic field line v_{\parallel} . Therefore, these electrons can absorb its energy at a lower magnetic field. Since the microwaves are injected from the LFS, they first interact with the electrons just described and are damped before reaching the cold resonance layer, such that only electrons with a positive parallel velocity are heated. Figures 3.2a and b show the effect of the ECRH heating in velocity space. The half-circles mark regions with the same total kinetic energy. The green boxes schematically show that the electrons with a specific parallel velocity are heated to higher perpendicular velocity v_{\perp} . This asymmetric heating does not directly drive a current because the electrons do not gain any parallel velocity. Rather, it creates a hole in the velocity distribution at a certain v_{\perp} and v_{\parallel} , and an over-population at the same v_{\parallel} but higher v_{\perp} .




Figure 3.2: *Velocity distribution of electrons heated via tangential ECRH waves [42]*

The two competing current drive effects resulting from this asymmetry are illustrated in figures 3.2c and d. The colors show the change of the electron distribution in

velocity space, blue and green being a decrease and red and yellow an increase in the distribution. Since the collision rate decreases with higher energy, particles with lower energies reach an equilibrium more quickly. Therefore, the hole in the velocity distribution created by the heated electrons fills up faster than the over-populated region can de-populate. A higher net number of electrons remain at positive v_{\parallel} . This is known as the Fish-Boozer current drive and generates a negative current. The second current drive mechanism, the Ohkawa current drive, originates from trapped particles (see section 2.1). The electrons in the area above the two solid green lines in figures 3.2b and d are trapped due to their high v_{\perp}/v_{\parallel} ratio. If through microwave heating the passing electrons become trapped particles, as is shown in figure 3.2b and d, they will bounce back and forth on banana orbits and, therefore, will not contribute to the total current. Since the electron density distribution has a hole at low v_{\perp} and positive v_{\parallel} a positive net current is driven. In addition, detrapping of the banana particles can occur through collisions, such that the trapped electrons diffuse into the non-trapped region. At positive parallel velocities, this effect is in equilibrium with the trapping process, but in the negative v_{\parallel} region it leads to an additional positive contribution to the current. Which one of the two mechanisms dominates partly depends on the fraction of passing electrons that become trapped particles and, therefore, whether the resonant magnetic field is located at the LFS or HFS. While performing the discharges presented in section 3.2 an attempt was made to increase (decrease) the edge current density via ECCD to destabilise (stabilise) peeling-ballooning modes and increase (decrease) the ELM frequency. However, the driven currents were too low to have a visible effect on ELMs. In this work, the ECRH power and current deposition profiles were calculated using the beam tracing code TORBEAM [43].




Figure 3.3: *Metal sheet damaged by ECRH radiation (a) and visualisation of the crossing point between non-absorbed ECRH waves and the PFCs (b)*

While being a very versatile heating method, ECRH can also pose a potential danger to the vessel components. If the waves are not fully absorbed they can reach the walls of the vacuum vessel and damage the PFCs or even the vessel itself. The reflection of the microwaves at a cutoff density due to density peaking, incorrect polarisation of the waves and too low electron density can prevent part or all of the power from being absorbed. Figure 3.3a shows a metal sheet that was melted due to incorrect ECRH polarisation in one of the discharges performed in the framework of this thesis. To

prevent serious damage to the vessel in the event of partial absorption, the results from TORBEAM were combined with the 3D AUG visualisation tool AUGDDD [44] to make sure that the waves would intersect the tungsten coated tiles, which are more robust than the other components in the vessel. The AUGDDD screen shot in figure 3.3b shows tiles of the inner divertor in AUG. The red arrow points toward the intersection of the ECRH waves and the tiles as calculated by TORBEAM for one case of only partially absorbed waves.

3.3 Magnetic measurements

In order to calculate a plasma equilibrium some knowledge of the magnetic field outside the plasma is necessary. The yellow boxes in figure 3.1b are poloidally distributed probes, 38 measuring the poloidal magnetic field and 18 measuring the radial component. In addition to those probes, AUG is equipped with 18 flux loops for magnetic flux difference measurements at several poloidal locations. The time resolution of these two diagnostics was recently upgraded from 1kHz to 10kHz. For many of the discharges presented in this work only the low time resolution was available but this is sufficient for the typical time scales considered here.

While the magnetic data is sufficient to produce a plasma equilibrium, knowledge of the pressure profile will greatly improve its quality. Therefore, in all equilibria presented in this work the edge pressure profile was constrained using data from experimental measurements. The diagnostics used to obtain these data are introduced in the following sections.

3.4 Charge exchange recombination spectroscopy

A tokamak plasma does not consist entirely of deuterium: the presence of impurity ions is unavoidable. These impurities limit the performance by radiatively cooling the plasma and via dilution of the D-T fuel necessary for the fusion reaction. However, they can be useful for diagnosing the ion temperature. Since impurities often ionise and recombine in the plasma, line radiation from neutral or partly ionised elements can be measured. Ionised impurities present in the path of the NBI beam, typically carbon, boron, nitrogen and helium, can gain an electron via charge exchange processes with neutral deuterium from the beam. They still have the same energy and velocity as the bulk of the same impurity species. Therefore, the impurity ion temperature can be calculated from the Doppler broadening of the line radiation originating from those impurities due to relaxation of the initially excited state. Since in AUG discharges the ion temperature equilibration time is smaller than the transport time scale, one assumes that the impurities have the same temperature as the main ions [45]. This has also been confirmed by cross-checking the temperature of different ion species [46]. One major benefit of these charge exchange recombination spectroscopy (CXRS) measurements is that they yield localised information because the measured light originates from the intersection of the line of sight with the neutral beam.

At the low field side of AUG, two edge CXRS systems are currently installed, the first with toroidal and the second with poloidal lines of sight focused on one of the neutral beams (see figure 3.1a). Those systems have a temporal resolution of up to 2.2ms. They can resolve the AUG pedestal with a radial resolution of 1cm, which is limited by

the channel spacing, and can be improved to 3mm when sweeping the plasma across the lines of sight. Two core systems are also installed in AUG and are routinely evaluated with a 10ms and 20ms time resolution, respectively, and a radial resolution of $\pm 2.5\text{cm}$. In most newer discharges presented in this thesis T_i data from CXRS measurements were used. However, at the time some of the older discharges were performed, the edge CXRS systems were not yet installed. For these discharges the ion temperature was assumed to be equal to the electron temperature.

3.5 Electron cyclotron emission

At AUG, one of the standard techniques to determine the electron temperature is by measuring the electron cyclotron emission (ECE) using a heterodyne radiometer [47]. In a tokamak, the electrons gyrating around the field lines emit cyclotron radiation at different harmonic frequencies $\omega_{ce,n} = n \cdot \frac{eB}{m_e}$, where n is an integer denoting the harmonic, e the elementary charge, B the local magnetic field strength and m_e the electron mass. Since the magnetic field of a tokamak decreases toward the outboard side, every frequency measured corresponds to a specific plasma radius. Microwave radiation travelling through a region that also emits this specific frequency can be absorbed. At a given magnetic field, the microwave spectrum is broadened by the Doppler effect and the relativistic mass increase, which in turn deepens the resonant region for a specific frequency. When the electron density is high enough and the path in which the electron cyclotron radiation can be reabsorbed is long enough such that the resulting radiation intensity corresponds to black body intensity, we call the plasma optically thick. When this is the case, the intensity of the radiation, or radiation temperature T_{rad} , of a certain frequency corresponds to the electron temperature T_e at the location from which the radiation originates. The core of a fusion plasma is typically optically thick, so the assumption $T_e = T_{rad}$ holds. In the pedestal, however, the density is often so low that the plasma is optically thin. The classical ECE evaluation method in the pedestal consists of setting T_e equal to T_{rad} in the optically thick region and extrapolating toward the SOL. However, when using a radiation transport model it is possible to make better use of the radiation temperature. A newly developed method called electron cyclotron forward modelling (ECFM) makes use of Bayesian probability theory to determine the entire electron temperature profile [48]. The solid red line in figure 3.4a shows an electron temperature profile generated using ECFM [48]. Since the plasma is optically thick inside a normalised radius of 0.98, T_e agrees very well with the ECE radiation temperature shown in black. However, further outside the so-called shine through phenomenon can be seen in the radiation temperature. Because the plasma is optically thin at the edge, this enhanced radiation temperature does not originate from the radiation at the cold resonance layer, but from further inside the plasma. The red crosses in figure 3.4a show that ECFM models the radiation temperature very well. It can be seen in figure 3.4b that the residuals all lie below one, which indicates the high quality of the modelling. Since this yields much more detailed pedestal data and gradients, which are especially important for stability calculations, ECFM data was used in this work when it was available.




Figure 3.4: *Electron temperature profile calculated via ECFM (a) and the residuals from the calculation (b) [48]. The normalised radius ρ_p is defined in equation 3.5*

3.6 Thomson scattering

Like CXRS, the Thomson scattering (TS) diagnostic uses Doppler broadening to determine the electron temperature [49]. At AUG, the beams from ten pulsed neodymium-doped yttrium aluminium garnet (Nd:YAG) infrared lasers are guided vertically through the plasma. Part of the light is scattered by individual electrons, being Doppler shifted twice, once during absorption and once during emission. The scattered photons are collected in polychromators through lenses mounted with a vertical view on the laser beam. The local plasma electron temperature along the laser beam can be determined from the Doppler broadening of the collected light. Additionally, the electron density can be calculated from the signal intensity.

The laser generators used in AUG have a pulse duration of 15ns, an energy of the order of 1J per pulse and a repetition frequency of 20Hz. The edge system consists of six beams passing through the vacuum vessel at a radius of 2.13m. The scattering volumes have a height of 2.5cm each and the beam width is 1.5mm. The best radial resolution is achieved at the z-position of the magnetic axis and roughly corresponds to the beam width, since there the laser is tangential to the flux surfaces. At higher and lower locations, the resolution can be as low as 1cm. In order to resolve the whole pedestal, it is necessary to set the outer separatrix position to very high values on the order of $R_{aus} = 2.16\text{m}$. This is often problematic because of shaping constraints: in an elongated plasma the separatrix will often be too close to the limiters which can cause impurity accumulation and might eventually disrupt the plasma. Therefore, edge TS data is not available in most AUG discharges but only in those specially designed for edge measurements. In the dedicated experiments carried out for this work, radial outward sweeps of the plasma were included to obtain good TS measurements of the pedestal.

3.7 Lithium beam impact excitation spectroscopy




Figure 3.5: *Neutral Li-beam injector and LOS of the Li-beam diagnostic*

One of the most accurate methods to measure the electron density in AUG utilises a beam of neutral lithium atoms that are injected with an energy of 30-80keV [50, 51]. When the atoms collide with plasma electrons they can either jump to an excited state and eventually emit line radiation or one of their valence electrons can be stripped off. In the latter case, the resulting Li-ion is confined to a magnetic field line and immediately leaves the volume of the beam. The system of coupled linear equations

$$\frac{dN_i(z)}{dz} = \sum_{j=1}^9 [n_e(z)a_{ij}(T_e(z)) + b_{ij}] N_j(z) \quad (3.2)$$

describes the number of beam Li-atoms N_i in the excited states i , where z denotes the coordinate along the beam. For $i \neq j$, the coefficients a_{ij} describe the excitation and de-excitation processes into the state i due to collisions with ions and electrons. The coefficients of the excitation, de-excitation and ionisation processes from i to any other state are all combined in a_{ii} . b_{ij} denote the Einstein coefficients for spontaneous emission. Equation 3.2 is solved for n_e using Bayesian probability theory [52].

The beam and the two optical heads installed in AUG can be seen in figure 3.5, where CXS and IXS denote the optical heads. The lines of sight only extend to $R_{maj} \approx 2.05$ (ca 5-10cm into the confined plasma, depending on the separatrix position) because at this location the beam attenuation is already so strong that the signal cannot be distinguished from the background radiation. In a medium density discharge the calculated n_e profiles are trustworthy up to the pedestal top (ca 2cm into the main plasma). In the core, the plasma density can be determined via laser interferometry.

3.8 DCN interferometer

The deuterium cyanide (DCN) laser interferometer measures the line integrated electron density [53]. For visible and infrared light ($\omega \gg \omega_{pe}$) the refraction index N of a

plasma changes according to

$$N \approx 1 - \frac{\omega_{pe}^2}{2\omega^2} = 1 - \frac{e^2}{2\epsilon_0 m_e \omega^2} \cdot n_e. \quad (3.3)$$

In AUG, the beam of a DCN-laser is split into three beams. One is guided through the plasma and thus phase-shifted according to

$$\Phi \approx \frac{e^2}{2c\epsilon_0 m_e \omega_0} \int_L n_e dl. \quad (3.4)$$

The second beam is frequency shifted by $\Delta\omega$ using a frequency modulator and the last one serves as a reference. Before the reference and the measurement beam enter the detector, they are both superimposed with the frequency modulated one such that a beat is measured. This method is used because a phase shift on the resulting 10-100kHz signal can be measured relatively easily, whereas it is not possible to directly measure a phase shift of the laser which has a frequency in the THz range.

At AUG, the DCN interferometer diagnostic is sampled at a frequency of 10kHz. In order to obtain full electron density profiles the data from the five lines of sight installed in AUG must be deconvolved, but the strong gradient in the edge cannot be resolved using this method. It is much more advantageous to combine it with Li-beam measurements by using the forward model described in the next section.

3.9 Integrated data analysis

The data from several diagnostic methods can be combined using Bayesian probability theory to form one joint profile [54]. For instance, by combining the data from the DCN interferometer and the lithium beam, one obtains an electron density profile which covers the whole plasma radius, and is very accurate in the pedestal region. Bayesian probability theory allows the inclusion of so-called prior information and weighting of the data from different diagnostics with the corresponding measurement errors. The ECFM method mentioned in section 3.5 is implemented in the integrated data analysis (IDA) framework. For IDA to produce optimal results the data from different diagnostics, primarily ECE and lithium beam, have to be aligned. This is possible using TS data since this diagnostic produces T_e and n_e profiles with the same time- and area-base. Shifts can also be detected by monitoring the IDA residuals.

3.10 Combining the diagnostics

Most discharges presented in this work were heated using a combination of NBI and ECRH. The versatility of the ECRH and ECCD deposition location was exploited in the discharges presented in section 5.7. An overview of the species heated by NBI and ECRH, as well as the quantities measured by the main diagnostics used in this work is given in table 1, where $v_{rot,impurities}$ and $c_{impurities}$ are the rotation velocity and the concentration of the measured impurity species.

As can be seen in figure 3.1b, most diagnostics measure at different toroidal locations. This is not problematic because, aside from slight deviations due to magnetic ripple


Heating method	Heated species
NBI	Mainly ions (at AUG)
ECRH	Only electrons
Diagnostic	Measured quantity
CXRS	T_i , $v_{rot,impurities}$, $c_{impurities}$
ECE	T_{rad} ($\approx T_e$)
TS	T_e and n_e
Li-Beam	n_e (edge)
DCN	n_e (line integrated)

Table 1: *Main heating methods and diagnostics used in this thesis*

effects, all equilibrium quantities in a tokamak are toroidally axisymmetric. The different measurement locations in the vertical plane (figure 3.1b) are a more challenging issue. However, a tokamak consists of nested flux surfaces and most equilibrium quantities are constant on the closed flux surfaces inside of the separatrix. Therefore, it is sensible to map the R and z location of the different measurements to a common area base, which usually takes the form of a normalised plasma radius. When investigating the edge of a tokamak, the normalisation is typically performed using the poloidal flux according to

$$\rho_p = \sqrt{\frac{\Psi(R, z) - \Psi_{axis}}{\Psi_{separatrix} - \Psi_{axis}}}. \quad (3.5)$$

Figure 3.6 shows the kinetic profiles of a typical AUG H-mode discharge. All data,

Figure 3.6: *Electron temperature (a), ion temperature (b) and electron density (c) profiles measured using different diagnostics*

except for the core ion temperature because of its high integration time, were ELM synchronised. In this discharge, all of the data from 2.5s to 4.5s in the discharge that

were collected between 3.5ms and 0.1ms before an ELM crash were overplotted. Fitting such data is also called coherent ELM-averaging. Since the edge ion temperature is very similar to the electron temperature, the assumption $T_i = T_e$ is acceptable in this kind of discharge when no CXRS data is available. While the core temperatures are very different this is not an issue for generating equilibria because only the edge pressure is constrained, whereas the core pressure is allowed to vary. Note that in figure 3.6c the interferometer data appears to have a strong deviation from the other diagnostics because it is a line integrated measurement. The pink crosses show the line-averaged density that should be measured by the different interferometer channels given an electron density profile corresponding to the pink fit. The result agrees very well with the DCN measurements. The Li-beam data (blue in figure 3.6c) has very high uncertainties at the pedestal top because of the strong attenuation of the beam. Therefore, the data collected in this region is typically omitted in the fits to the density profiles.

In this work, the pressure profiles used to generate high resolution equilibria were calculated using the T_e , T_i and n_e data from the diagnostics described in this chapter, after mapping them to ρ_p and aligning them with the information given by the TS system. The ion density was generated from the electron density by using the impurity concentrations determined by the charge exchange diagnostic, or by using the effective charge Z_{eff} value calculated from bremsstrahlung measurements using Bayesian methods [55]. The resulting pressure profiles were combined with magnetic measurements to calculate equilibria using CLISTE. The impact of measurement uncertainties on the final stability calculations is presented in section 4.8

4 Optimisation and limitations of the stability chain

To determine whether a plasma is stable against peeling-ballooning modes, a CLISTE equilibrium is first generated using experimental data, after which this equilibrium has to be re-run in HELENA to increase the resolution and generate a file in the correct format for the stability codes. After running the j-alpha workflow to alter the current density and pressure profiles independently the stability calculations are performed using ILSA. In this chapter, the influence of different input parameters on the results of the stability calculations will be discussed. First, different stability codes will be compared to ensure consistency. The grid density and number of included poloidal harmonics will be checked for convergence. One must differentiate between three different radial grids:

- The internal HELENA grid, where every grid point acts as a node for the Hermite polynomials used to generate the equilibrium.
- The HELENA mapping grid, to which the equilibrium quantities are mapped after the equilibrium has been generated. This is the grid which is used in the equilibrium file that will later be used as an input for the stability code.
- The internal grid of ILSA, CASTOR or MISHKA, on which the stability calculations are performed. The stability code maps the quantities contained in the equilibrium file to this third grid before starting the calculations.

After discussing the so-called cutoff parameter Ψ_b , the effects of ion diamagnetic drift stabilisation, finite resistivity and density profile shape will be discussed. Finally, the influence of kinetic profile uncertainties will be assessed.

4.1 Comparison between different codes

The main stability code used in this work is ILSA, which is a revised version of the CASTOR code [34, 35]. It can either use the full set of equations A.24-A.27 or be run in MISHKA-1 mode, using the reduced set of equations A.32-A.34. In this section, the results from ILSA will be compared to the stand-alone versions of MISHKA-1 and CASTOR.

The first comparison is performed using a synthetic, ballooning-unstable equilibrium with a circular cross-section. The safety factor increases from 1.05 in the core to 1.95 at the edge of the plasma. Whereas it is possible to run ILSA and CASTOR using the same radial grid, MISHKA-1 uses a different mesh accumulation algorithm. To optimise the results, a similar edge grid point density was used in MISHKA-1. This is crucial when comparing the codes at high toroidal mode numbers since their eigenfunctions are very narrow and localised at the far edge. For a meaningful code comparison one needs either a very high grid point density in this region in order to resolve all the modes sufficiently well, or a comparable edge grid point density in all codes. The former is not possible with the stand-alone version of MISHKA-1 since it is limited to a total number of 201 grid points.

Figure 4.1a shows the growth rates of different toroidal modes calculated using CASTOR, MISHKA-1 and ILSA, run in MISHKA-1 and CASTOR mode. The real part of the eigenvalue calculated by the stability codes is the growth rate normalised to the

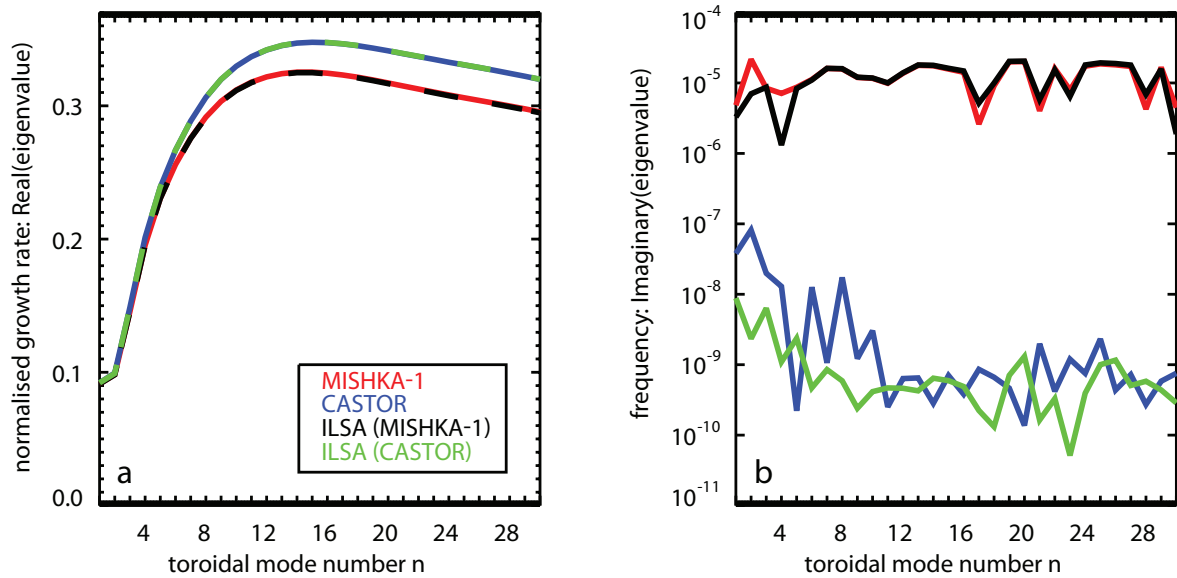


Figure 4.1: *Real (a) and imaginary (b) part of the eigenvalue calculated for a range of toroidal mode numbers with different stability codes*

Alfven frequency, which is generated using the core density. The right plot shows the imaginary part of the eigenvalue. Since all codes were run in ideal mode, the imaginary part should be zero. The deviation from zero can be regarded as an indicator of the numerical accuracy of the solution. The results from CASTOR and those from ILSA in CASTOR mode have much lower imaginary parts than those from the two other codes. Nevertheless, the criterion that the imaginary part should be at least one to two orders of magnitude lower than the real part is valid for all codes. This condition is not necessarily met when using experimental equilibria, especially when they are only marginally unstable. All codes produce the same results at low toroidal mode numbers, whereas at higher mode numbers MISHKA-1 and CASTOR show differences. This is to be expected as they are using a different set of equations. The results from ILSA and from the stand-alone codes match very well when ILSA is run in the corresponding mode. The numerical accuracy is also the same, as can be seen from the imaginary part of the eigenvalue. Figure 4.2 shows the relative contributions of different poloidal harmonics. The mode amplitude is plotted against the s -coordinate, which is the radial coordinate used in the stability codes. It is defined as the square root of the poloidal flux normalised to the value at which the equilibrium was cut off. The poloidal mode structure is almost identical in all four codes. They all produce consistent results for a simple synthetic equilibrium. However, it will be shown in the following that large discrepancies arise when the edge q -profile is very steep, as is the case at the edge of shaped equilibria used on AUG.

Both HELENA and ILSA use the straight field line coordinate system, in which the variables depend on the safety factor q . Since q diverges toward infinity at the separatrix, these coordinates make it necessary to cut off a portion of the plasma at a certain radius Ψ_b (defined in normalised poloidal flux). The closer to the separatrix the equilibrium is cut off, the higher and steeper the edge safety factor profile is, which has two major consequences:




Figure 4.2: *Normalised mode amplitude of different poloidal harmonics for $n=1$ (left) and $n=15$ (right), calculated with different stability codes*

- If the q -profile is too steep, the code runs into numerical errors. To exclude such errors, a convergence study was performed and will be presented in section 4.2.
- Because of the higher safety factor, more resonant surfaces are present at the edge of the plasma, which has a stabilising effect on the peeling mode.

Figure 4.3 shows the eigenvalue calculated with different stability codes for $n=3$ using an equilibrium produced from AUG discharge 23223 just before an ELM crash. A scan of the cutoff parameter Ψ_b was performed. MISHKA-1 (red) is only slightly affected by the cutoff parameter, whereas the other codes show a strong dependence. In most cases the imaginary part of the eigenvalue increases when the equilibrium is cut off closer to the separatrix. The light green curve will be explained later in this section.

The poloidal mode structure calculated by ILSA in MISHKA-1 mode and MISHKA-1 standalone can be seen in figure 4.4. At $\Psi_b = 0.990$, the relative contributions of the different harmonics are similar between the two codes, but at $\Psi_b = 0.996$ the differences are much more pronounced. Changing the codes in such a way that one can import the location of the radial nodes from ILSA into MISHKA-1 and vice-versa confirmed that this is not an effect of the different mesh accumulation.

A scan was performed to explore under which circumstances these discrepancies arise. Figure 4.5 is analogous to a j - α diagram, but shows the whole spectrum of the normalised growth rate. The reference equilibrium (green box) is the same one used for figure 4.3. The different line styles correspond to values of Ψ_b from 0.990 to 0.997 for both MISHKA-1 (red) and ILSA in MISHKA-1 mode (black). The scale factors along the axes represent the multiplication factors that were used to modify the edge current density height and the pressure pedestal width.

In the region of low pressure gradient and high edge current density MISHKA-1 and ILSA agree reasonably well. At higher pressure gradients, however, they start to diverge and this trend becomes even stronger at low edge current density. It can also be seen




Figure 4.3: *Real (left) and imaginary (right) part of the eigenvalue $n=3$, calculated with different stability codes at different cutoff parameters Ψ_b*




Figure 4.4: *Normalised mode amplitude of different poloidal harmonics for $n=3$, calculated with MISHKA-1 (red) and ILSA (black) for $\Psi_b = 0.990$ (left) and $\Psi_b = 0.996$ (right)*

that the cutoff parameter Ψ_b changes the spectrum calculated with MISHKA-1 only slightly, whereas the dependence in ILSA is much stronger. For example, in the case marked by an orange box ILSA finds that for $\Psi_b = 0.993$ the equilibrium is much more stable than $\Psi_b = 0.990$, whereas the MISHKA-1 results are almost identical. Large discrepancies in MISHKA-1 only occur at very low edge current density and high cutoff parameter Ψ_b where ILSA does not find any unstable modes.

The discrepancies between MISHKA-1 results generated with different values of Ψ_b are not due to MISHKA-1 but to the equilibrium file created by HELENA. After calculating

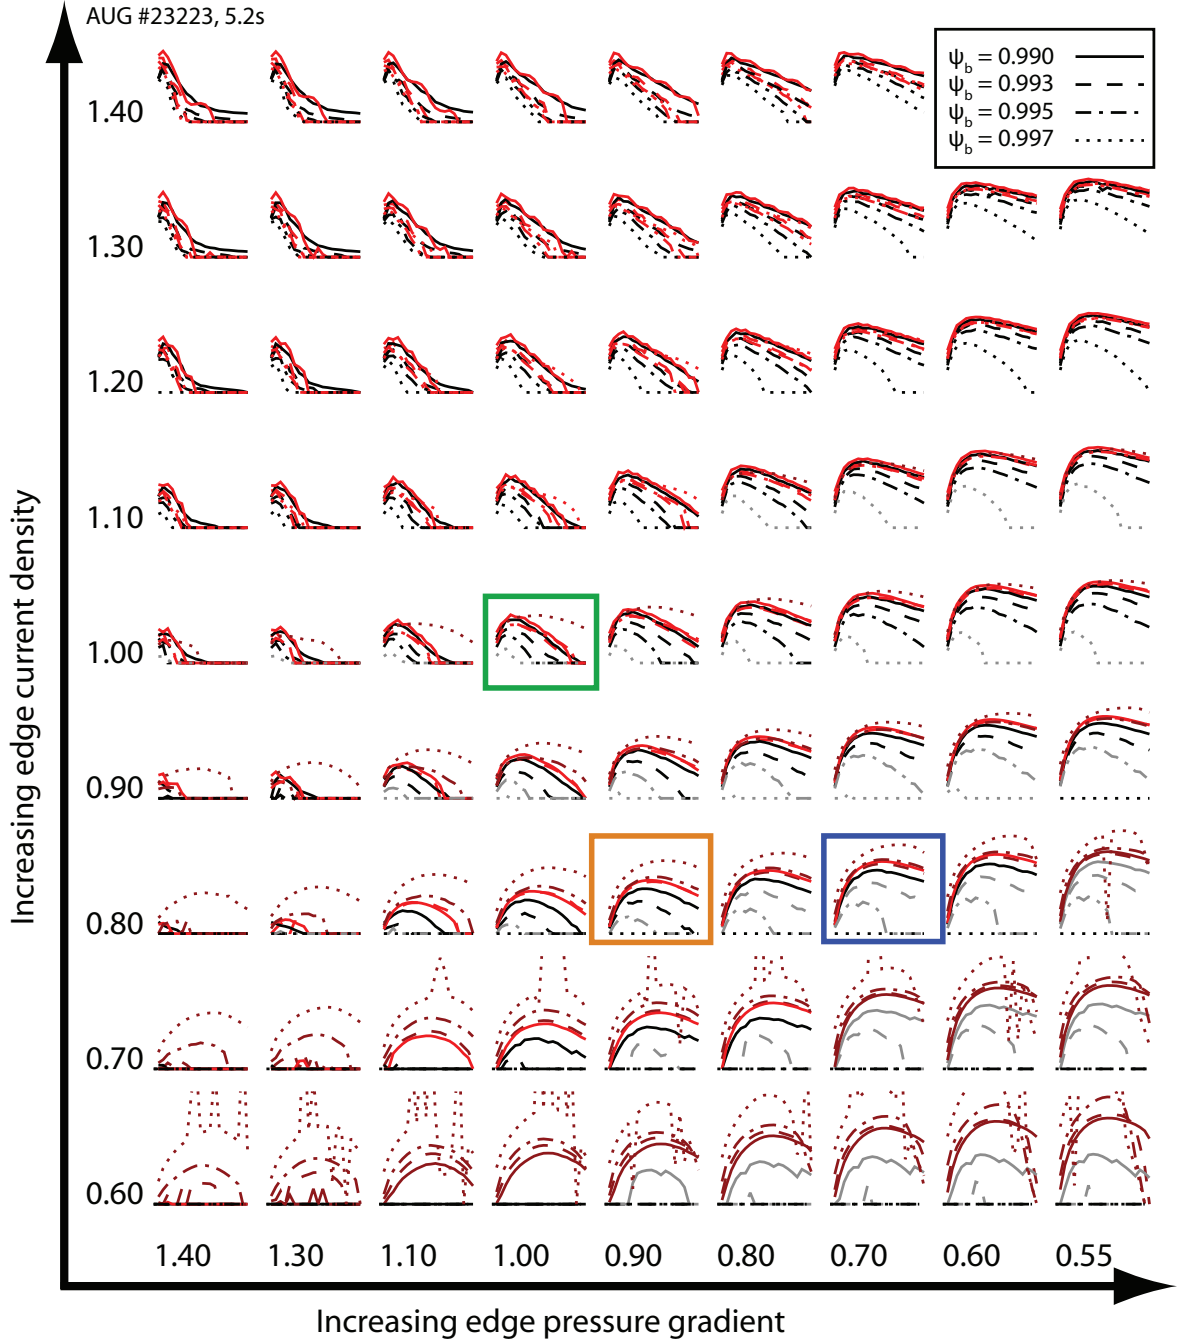


Figure 4.5: Growth rate spectra of different toroidal modes calculated with ILSA (black and grey) and MISHKA-1 (red) for different values of Ψ_b (increasing from solid to dotted). The x-axis of the plots corresponds to the toroidal mode number and ranges from $n=1$ to $n=20$. The y-axis is the normalised growth rate from $\gamma = 0$ to $\gamma = 0.2$. The green box marks the unaltered experimental equilibrium while the orange and blue boxes serve as indicators for equilibria that are described in the text.

the equilibrium using an internal grid, HELENA maps the different quantities to a new grid, which also depends on the safety factor q . If q is too steep, the interpolation to this HELENA mapping grid fails which results in overlapping flux surfaces (red arrow in figure 4.6a). This interpolation problem can be fixed by increasing the HELENA




Figure 4.6: *Part of the HELENA mapping grid when using a low (a) and high (b) radial grid resolution*

mapping grid resolution (figure 4.6b), but when q is too steep increasing the number of radial points even further does not help any more. The MISHKA-1 (ILSA in MISHKA-1 mode) results from unconverged equilibria are plotted in dark red (grey) in figure 4.5 to show that one has to carefully check the HELENA results since in some cases the stability codes still produce results that might, at a first look, appear reasonable. In this comparison, the radial grid resolution was not increased because the standalone version of MISHKA-1 does not allow large equilibrium files to be used as input. The MISHKA-1 results reflect the known behaviour seen in previous works using other stability codes, such as ELITE. Namely, the high pressure gradient region is highly unstable, independent of the edge current density [56].




Figure 4.7: *Derivative of the safety factor q used in MISHKA-1 and ILSA*

A comparison of the radial derivative of the q -profile calculated by the two stability codes shows that the profiles used in ILSA and MISHKA-1 are very different at the edge in the low edge current density cases, even though the input equilibrium was identical (see figure 4.7). The brown curve shows the result of a two-point differentiation performed on the same q -profile after the HELENA result was re-mapped to a mesh with a radial grid size of 2000 instead of 300 points. It can be seen that ILSA calculates more exact values of dq/ds than MISHKA-1. This discrepancy comes from the

mapping routines in the stability codes. Both codes interpolate the equilibrium profiles from HELENA onto their internal grid using a cubic spline function, but with different boundary conditions. For building the spline, MISHKA-1 uses the two-point derivative at $s = 0$ and $s = 1$, whereas ILSA uses the not-a-knot spline method. In this method the third derivative f''' of the function f is matched, according to $f'''(0) = f'''(1)$ and $f'''(N) = f'''(N - 1)$, where N is the index of the last grid point.

As a final test, dq/ds from ILSA was exported and directly read into MISHKA-1 before performing the stability calculations. The result is shown as the light green curve in figure 4.3. When using these more correct values, the MISHKA-1 and ILSA (MISHKA-1) results match extremely well, with discrepancies below 1%. This suggests that the stabilising effect that emerges at low edge current densities and when cutting the equilibrium off further outside is real, and the reason that it is not seen in the MISHKA-1 results is because the code produces incorrect values of dq/ds when the edge q profile is very steep. This result is to be expected because a strong magnetic shear has stabilising effect on ballooning modes (see chapter 2).

4.2 Convergence

Since an accurate q -profile is crucial, convergence was tested by scanning several grid sizes in HELENA. The standard internal HELENA grid used in this work had 300 radial nodes with a higher accumulation at the edge and 257 poloidal nodes. Even in equilibria with the steepest q , increasing the internal HELENA grid density had no effect on the ILSA results. The HELENA mapping grid, however, plays an important role. Figure 4.8 shows the growth rate calculated with ILSA in MISHKA-1 mode




Figure 4.8: Growth rate at different HELENA mapping grid point densities for the reference case (a) and an altered equilibrium (b)

for values of the cutoff parameter Ψ_b between 0.990 and 0.996. The different colors correspond to different HELENA mapping grids. The values given in the legend are the grid point densities of the mapped equilibrium at the very edge of the plasma. In

the left plot, the growth rate of $n = 3$ was calculated for the reference equilibrium (green box in figure 4.5). The right plot was calculated at $n = 15$ using an equilibrium with a q -profile which, when using $\Psi_b = 0.996$, is close to the limit of the capabilities of HELENA (blue box in figure 4.5). These toroidal mode number were used because they roughly have the highest growth rate. The red curves in figure 4.8 correspond to the grid density which was used in figure 4.5, since the stand-alone version of MISHKA-1 does not allow inputs with a larger grid. It can be seen in the steeper q -profile case (right) that this grid density is not sufficiently high to produce accurate results. At $\Psi_b = 0.996$ the HELENA mapping problem described earlier occurs up to $dN/ds = 692$ in the reference case and up to $dN/ds = 2309$ in the altered case. The setting used for the rest of this work is the one shown in blue, in which an edge dN/ds of 2309 was obtained using a total number of 1000 radial grid points. This setting is a good compromise between accuracy, since it is converged at $\Psi_b = 0.994$ for the very steep case, and acceptable file size of 500MB per equilibrium (which corresponds to roughly 60GB per $j - \alpha$ diagram, without counting the data produced by ILSA).




Figure 4.9: *Growth-rate with different numbers of poloidal harmonics for the reference case (a) and an altered equilibrium (b)*

A second convergence test was performed in order to assess how many poloidal harmonics are necessary. The same two cases were used as in figure 4.8, also with $n=3$ for the reference and $n=15$ for the case altered in j and α . The edge grid density of the mapped equilibrium was $dN/ds = 2309$. Figure 4.9a shows that for $n=3$ in the shallow q case 30 poloidal harmonics are sufficient for convergence. For $n=15$ and a steep q profile, however, full convergence is only reached when between 120 and 140 harmonics are used. Convergence was also tested with respect to the ILSA grid, and 300 radial nodes with a high edge grid accumulation of $dN/ds \approx 20000$ which was used throughout this work was found to be sufficient in all cases.

4.3 Choice of Ψ_b

Once most code parameters were optimised, the task of finding a sensible value for the cutoff parameter Ψ_b remained. It has been shown that the presence of a separatrix has a stabilising effect on peeling modes and that the equilibrium becomes more unstable when more flux is cut off [57], so the choice of Ψ_b will have an impact on the final result.

In linear ideal MHD, a mode and its growth rate can be visualised as the displacement of a magnetic field line at a certain rate. Since such a change in the magnetic field will induce a current, one sensible cutoff criterion can be calculated using the resistive skin depth

$$\delta = \sqrt{\frac{2}{\mu_0 \sigma \omega}}, \quad (4.1)$$

where σ is the conductivity in the plasma layer we cut off. Above the critical frequency ω , a perturbation does not diffuse through the layer δ [1]. If Ψ_b is chosen such that the




Figure 4.10: Growth rate of $n=15$ (black) and critical frequency (blue) at different Ψ_b . The dashed lines show half of the diamagnetic drift frequency for different n 's.

plasma layer being cut off is so small that $\omega = 2/\mu_0 \sigma \delta^2$ is higher than the growth rate of the mode, the current induced by the mode would diffuse faster into the vacuum region than the mode can grow. Since ideal MHD does not allow current diffusion, cutting off this layer is justified. A comparison of ω and the growth rate at different Ψ_b for the $n = 15$ case (figure 4.9b, grey, 200 poloidal harmonics) is shown in figure 4.10. Here, Ψ_b is scanned up to 0.999 and the growth rates are plotted in SI units. The blue curve shows the critical frequency ω , if the resistive skin depth is set to the cut

off layer and σ is the Spitzer conductivity at the cutoff location, calculated from the experimental temperature profile. Ideal MHD is only valid if the plasma is cut off at $\Psi_b = 0.996$ or further inside. Additionally, the sudden drop in the growth rate between $\Psi_b = 0.996$ and $\Psi_b = 0.999$ suggests that HELENA or ILSA encounters numerical problems in this region. This was confirmed after analysing the HELENA equilibrium in more detail. As already mentioned earlier in this chapter, the interpolation to the mapping grid is problematic at steep q and results in overlapping flux surfaces (figure 4.6a). These numerical errors that occur at $\Psi_b \geq 0.998$ are inherent to such steep q -profiles and could not be avoided by changing parameters in the code or by increasing the grid resolution. The cutoff criterion calculated using the resistive skin depth is at the limit of the capabilities of the stability chain. In cases with lower edge current density HELENA already runs into problems at $\Psi_b = 0.996$. Therefore, most stability diagrams shown in this work were all generated twice, using $\Psi_b = 0.990$ and $\Psi_b = 0.994$ to show possible differences.

4.4 Effect of Ψ_b on the mode width




Figure 4.11: *Normalised amplitude of different poloidal harmonics for $n=15$ at $\Psi_b = 0.990$ (top) and $\Psi_b = 0.996$ (bottom) plotted versus the internal ILSA coordinate s (a) and the major radius (b). The red dashed vertical line in (b) represents the position of the separatrix.*

The impact of Ψ_b on the mode width will now be assessed. It was postulated that the radial extent of the mode calculated using linear MHD, could be correlated with the ELM affected area [58]. If the cutoff parameter Ψ_b were to impact the extent of the mode, this would be problematic for the verification of this hypothesis. Figure 4.11 shows the radial structure of the different poloidal harmonics for $n=15$ for two different choices of Ψ_b . The sign of the amplitude is irrelevant, the colors represent the number of the harmonic. When they are plotted against the internal ILSA coordinate s (left), the locations of the different harmonics appear shifted depending on Ψ_b . This is to

be expected since their resonances are at different locations due to the relative shift of the q -profile. However, when plotted in real space the locations of the harmonics are similar and the total extent of the mode into the plasma is the same. For a better comparison of the harmonics $m = 71$ is shown in black and $m = 80$ in purple. While it has not yet been verified whether the extent of the mode envelope into the plasma does hold information about the ELM affected area, it is not affected by the choice of Ψ_b .

4.5 Ion diamagnetic drift stabilisation

At high mode numbers the wavelengths of peeling-ballooning modes become comparable to the ion gyro-radius. At these high mode numbers an approach using kinetic theory is more sensible than MHD. It was shown that the ion diamagnetic drift stabilises the ballooning mode [59,60]. This is due to the magnetic field being frozen into the electron fluid while the mass density of the plasma is determined by the ions. This means that the density perturbation ρ_1 in equations A.24 and A.26 moves in the ion diamagnetic drift direction while the magnetic field perturbation \mathbf{B}_1 in equations A.26 and A.27 moves in the electron diamagnetic drift direction. If the diamagnetic drift velocity, normalised to the size of the mode, is high compared to the growth rate of the mode, the mode is torn apart. The diamagnetic drift frequency ω^* can be written as

$$\omega_{*pi} = \frac{n}{e_i n_i} \frac{dp_i}{d\Psi}, \quad (4.2)$$

where n is the toroidal mode number [61]. Hastie *et al.* extended the stabilisation model, which initially assumed a constant ω_{*pi} , to a radius dependent model because the ion diamagnetic drift frequency varies strongly in the pedestal region [61]. In their example, the stabilising effect of a radially varying $\omega_{*pi}(x)$ was roughly 50% lower than for a constant value calculated from the maximal edge pressure gradient. Since 50% is small compared to the differences that will now be discussed, the present work sticks to the commonly used criterion that a growth rate lower than $\omega_{*pi}/2$ is considered stable. The horizontal dashed lines in figure 4.10 indicate this stability criterion for $n = 1, 2$ and 3 . Since ω_{*pi} is directly proportional to the toroidal mode number and the equilibrium used for figure 4.10 lies in the high pressure gradient part of the j - α diagram, the growth rate of $n = 15$ is almost one order of magnitude lower than $\omega_{*pi}/2$. Because of this linear n and α dependence, normalising to the diamagnetic drift frequency changes the shape of the $j - \alpha$ diagram significantly. Figure 4.12 shows the same stability diagram twice, with the stability criterion being $\gamma/\nu_A = 0.06$ (left), and $2\gamma/\omega_{*pi} = 1$ (right). While in the former case the most unstable mode numbers around the experimental equilibrium are between 7 and 14, diamagnetic drift stabilization suggests that much lower mode numbers between 1 and 4 dominate. It appears that the diamagnetic drift would completely stabilise the plasma at only a slightly lower edge current density, even if the pressure gradient was much higher. This effect comes from the strong stabilisation of higher mode numbers, and because the mode spectrum strongly shifts toward higher toroidal mode numbers. In this region, low mode numbers (n 's around 1 to 4) have a zero (or close to zero) growth rate. However, this is just a rough estimation of the stabilising effect of the diamagnetic drift and analysing this further would require using a two-fluid model. In this work, γ/ν_A contours are shown

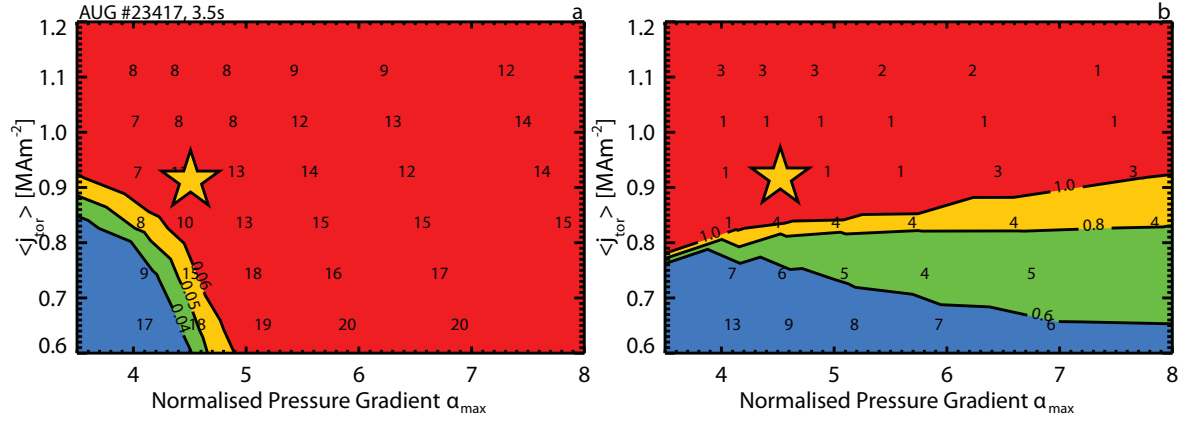


Figure 4.12: *Stability diagram of discharge 23417, showing ν_A (left) and $\omega_{*pi}/2$ (right) contours*

in the stability diagrams, with the caveat that FLR effects have a stabilising effect on high- n modes.

4.6 Finite resistivity effects




Figure 4.13: *Left: Growth rate of $n=5$ for different resistivity values. Right: normalised mode amplitude for $\eta = 10^{-10}$ Ωm*

While in this work the stability codes were run in the ideal-MHD mode, a comparison with results from calculations including a finite resistivity was made. The steeper ∇p case from discharge 23223 used earlier (blue box in figure 4.5) was run using ILSA in resistive CASTOR mode. Figure 4.13 shows the growth rate for different values of the resistivity η . The actual resistivity of the plasma in the pedestal region lies between 10^{-8} and 10^{-7} Ωm, which is two orders of magnitude higher than the values shown in

figure 4.13. While it appears that the growth rate of the modes increases when moving toward a higher resistivity, it can be seen in the right plot that the radial mode structure shows a strong oscillatory behaviour at the edge, indicating numerical problems. Under $\eta = 10^{-10}\Omega\text{m}$ the resistivity is so low that the contribution from the resistive terms in ILSA is negligible. Starting at $\eta = 10^{-9}\Omega\text{m}$ the resistivity starts playing a role, but ILSA cannot resolve the strong oscillations at the rational flux surfaces that occur in resistive MHD. Because of the difficulties the stability code appears to have with a finite resistivity, the remaining calculations presented in this work were all performed using ideal MHD.

4.7 Effects of a realistic density profile




Figure 4.14: *a and b: Density profiles used. c: normalised growth rate for $n=5, 10$ and 15 . d: radial mode structure of $n=10$ for a realistic (positive values) and constant (negative values) density profile.*

In general, the equilibrium density is assumed to be constant when performing stability calculations, even though a density term appears in equations A.24, A.26 and A.32. Most codes allow for a radius dependent density, but this feature is seldom used. In the pedestal region, where the density varies by one order of magnitude within 2cm, the inertia of the plasma changes quite significantly, which might affect the contribution of the different poloidal harmonics, thereby changing the eigenvalue. To assess the influence of the profile shape, a realistic density profile was included in the calculations,

and then gradually changed to a constant profile. To obtain comparable results, the profile was normalised such that the density was the same at $s = 0.99$, which is approximately the position of steepest pressure gradient. Additionally, the density was set to the same value at the plasma centre, such that all cases are normalised to the same value. The profiles used can be seen in figure 4.14a and with a focus on the edge region in 4.14b. It can be seen in figure 4.14c that the growth rate changes slightly between the different density profiles. At a constant density, the growth rate is slightly higher than when using a realistic profile, but a scan has shown that this depends on the exact location at which all cases were set to have the same edge density. The radial mode structure for $n = 10$ is plotted in figure 4.14d. The poloidal harmonics normalised to positive values correspond to a realistic density profile (black curve in figure 4.14 a and b), while the ones normalised to negative values correspond to a constant density (pale blue curve in figure 4.14 a and b). It is clear that the inertia has an effect on the relative mode amplitude. In the realistic case, the contribution from inside (outside) of $s = 0.99$ is lower (higher) than in the case with constant density. This agrees with the higher density and, therefore, inertia inside of $s = 0.99$. While including a realistic density profile does have an effect on the relative contribution of the different poloidal harmonics, the impact on the growth rate is small. Since the main interest is to see if and at which toroidal mode number the plasma is unstable a constant density profile was used in the rest of this work. This choice was also motivated by a significant increase of the imaginary part of the eigenvalue with steeper density profiles, indicating poorer numerical accuracy.

4.8 Influence of measurement errors

The two main systematic uncertainties in the experimental measurements are the ion temperature, which is not available for older discharges and, therefore, assumed to be equal to T_e , and the electron temperature gradient, because of the ECE shine through [48]. The effect of a systematic error in the T_e gradient on the stability was assessed for discharge #23417 for which ECFM results were available (see section 3.5). The ELMs in this discharge occur at two distinct frequencies, with some ELMs occurring directly after the pedestal gradient recovery around 9-10ms after the previous crash, and others after a several millisecond long phase of almost constant pressure gradient. This behaviour has already been described for a similar discharge [17].

The four stability diagrams in figure 4.15 represent the equilibria 3ms before the ELM crash in discharge #23417 for the fast (pedestal still recovering, top) and slow (fully recovered pedestal, bottom) ELM cycles, using the classical IDA analysis (left) and ECFM (right). They were all generated using $\Psi_b = 0.990$. In the slow ELM cycles the classical IDA profiles lead to a pressure gradient and edge current density that are 20% lower than the stability boundary, whereas when using the ECFM method the operational point is in the marginally unstable region. Note that since the assumption $T_i = T_e$ was used, and recent highly resolved edge CXRS measurements suggest that ∇T_i is slightly lower than ∇T_e , the correct operational point might be just outside of the unstable region [62]. In the equilibrium 3ms before a fast ELM crash, the operational point is not quite on the stability boundary, which is to be expected since the profiles have not yet fully recovered. In this case, the difference between the equilibrium obtained using ECFM and the one generated with classical IDA is also very large. Nev-




Figure 4.15: T_e profiles (left) and stability diagrams (right) of AUG #23417

ertheless, the position of the stability boundary itself barely changes between the four cases. This suggests that the stability calculations are rather robust against errors in the kinetic data used for computing the input equilibrium. However, accurate current density and kinetic profiles are needed for identifying the exact position of the operational point, which determines whether the plasma is unstable or not. When available, Thomson scattering data was used to complement the ECE or ECFM temperature in this work.

A second potential systematic error is the location of the steep gradient region in relation to the equilibrium separatrix. In order to assess this, the kinetic profiles were shifted by $\pm 5\text{mm}$ before the CLISTE equilibrium was generated. The resulting perturbed equilibria were then analysed with ILSA.




Figure 4.16: Toroidal mode spectra for three different shifts between kinetic and magnetic data

A very strong effect of this shift on the toroidal mode spectrum can be seen in figure 4.16. An outward shift of the pressure profile leads to higher growth rates, while an

inward shift appears to be stabilising. Figure 4.17 shows the pressure gradient (a),




Figure 4.17: *Pressure gradient (a), current density (b) and q -profile (c) calculated using different shifts of the pressure profile*

the current density (b) and the q -profile (c) of the reference equilibria calculated using the three shifts. The pressure and current density profiles from the equilibrium with the negative shift (black) are lower and broader than in the correct equilibrium (red), which results in the lower growth rates seen in figure 4.16. The broader current density peak explains the lower mode numbers. It is clear that the shift used in the blue case (+5mm) is wrong because the edge current density is very low. Nevertheless, this example shows that a high pressure gradient at the separatrix strongly destabilises the high- n ballooning modes. In this case, even the high shear (see blue q -profile in figure 4.17c) is not sufficient to stabilise them.

In this work, great care has to be taken when aligning the kinetic data to minimise systematic errors in the stability calculations. The correct profile location can be detected via carefully monitoring the residuals from CLISTE [25].

4.9 Final parameters used in the stability calculations

- A HELENA equilibrium mapped to a radial mesh with $dN/ds > 2300$ at the edge.
- An ILSA grid with a total of 300 radial nodes and with an edge grid accumulation of $dN/ds \approx 20000$.
- A total number of poloidal harmonics between 30 and 180 depending on the toroidal mode number n (to save computing power and memory).
- Most $j - \alpha$ diagrams were run with two values of the cutoff parameter Ψ_b , 0.990 and 0.994, to observe possible differences.
- ILSA was run in ideal mode using the MISHKA-1 equations.
- A constant density profile was assumed.
- The kinetic profiles were aligned such that the residuals from CLISTE are minimised.
- When CXRS data was available, T_i profiles were used rather than setting $T_i = T_e$.

- To determine T_e either the ECFM method was used, or the ECE radiation temperature was combined with Thomson scattering data.

5 Results

In the previous chapters the peeling-ballooning theory and the methodology used to compare it to experimental data were introduced. The stability chain presented in section 2.4 will now be applied to several different types of ELMs. While many publications already exist on the peeling-ballooning stability of type-I ELMs [23, 63–68], two type-I cases will still be presented in section 5.1 to build a basis for comparing further scenarios. Type-III ELMs, smaller events which arise just above the L-H power threshold, will be explored in section 5.2. Type-II ELMs and a type-II like ELM-free H-mode regime emerging at high densities will be discussed in sections 5.3 and 5.4. Two scenarios designed to mitigate the power flux to the walls of future fusion devices, namely nitrogen seeding in the divertor and ELM-triggering via pellets will be discussed in sections 5.5 and 5.6. It was found at the TCV tokamak that heating the plasma edge with ECRH can also increase the ELM frequency and decrease the per-ELM losses. Such discharges were repeated at AUG and are presented in section 5.7.

5.1 Type-I ELMs: peeling-ballooning stability in typical H-modes

As already mentioned in section 1.2, type-I ELMs are the most commonly occurring ELM type. They occur in strongly heated H-mode scenarios with high confinement. Time traces taken from a typical type-I ELMy H-mode discharge are shown in figure 5.1. To ease the comparison between the different ELM regimes these same time traces will




Figure 5.1: *Type-I ELMs: time traces of the divertor current from shunt measurements (black), the power to the outer divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature just inside the pedestal top (red)*

be presented for all discharges discussed in this chapter. The black curve shows shunt measurements in the divertor. The signal is named *Ipolsola* because it corresponds to the poloidal SOL current measurement in the outer (außen) divertor. This signal is

roughly proportional to the temperature in the outer divertor and is, therefore, a very good indicator of hot particles arriving along the open field lines. At AUG, this signal is typically used to determine the onset of an ELM. The power reaching the outer divertor plates measured with infrared cameras is shown in the second frame. The total kinetic energy confined in the plasma, W_{mhd} , is shown in green. The data from the outermost channel of the DCN interferometer can be seen in blue. The electron density is integrated along the line of sight and, therefore, given in m^{-2} . The radiation temperature in the last frame is taken from an ECE channel at the pedestal top where the assumption $T_e = T_{rad}$ is typically valid.

The discharge 23417 shown in figure 5.1 was run with a 1MA plasma current and -2.5T toroidal magnetic field, which are the typical values used in H-mode experiments at AUG. It was fuelled with deuterium gas at a rate of $3.8 \cdot 10^{21}s^{-1}$ and heated with 5MW NBI and 1.3MW ECRH power. While most ELMs had a frequency on the order of 80Hz, some occurred after an inter-ELM time of only 7-8ms. The ELMs induced 6MW power loads on the outer divertor and the average amount of energy lost per ELM was 10% of the total W_{MHD} . The pedestal top temperature typically dropped by 10% during each ELM, while the drop in edge density was slightly lower.




Figure 5.2: Pre type-I ELM profiles of the electron and ion temperature (a), the electron density (b), the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)

The black curves in figure 5.2 show the pedestal profiles of the electron temperature, the electron density, the total pressure and its gradient, the current density and the safety factor q from the same discharge. The red curves are from a different discharge which will be discussed later in this section. Since discharge 23417 was carried out before the edge CXRS system was installed, the assumption $T_i = T_e$ was made. The T_e and n_e curves are averaged over many ELMs from data taken in 1ms time intervals directly before the ELM crashes. The crosses in panels c and d are the total pressure and its gradient. They are calculated from the temperature and electron density profiles. The ion density was calculated from n_e assuming a radially constant Z_{eff} value of 1.6. The solid lines are the CLISTE fits to the experimental data. In all equilibria presented in this section the pressure profile in CLISTE was only constrained with

kinetic data between the pedestal top and the separatrix. Further inside, it was more loosely parametrised to account for the fast ion pressure. The total W_{MHD} , which can be calculated very accurately from magnetic measurements alone, then acted as a constraint on the magnitude of the core pressure profile. The edge current density shown in panel e is the local toroidal current density at the low field side of the plasma, while the q-profile in panel f is flux surface averaged.




Figure 5.3: *Stability diagram of type-I ELMy discharge 23417*

The stability diagram in figure 5.3 shows the result from the stability chain presented in section 2.4 using the profile in figure 5.2. Stability is shown as a function of two destabilising quantities, the pressure gradient and the current density. α_{max} is the maximum edge pressure gradient normalised according to equation 2.10. The maximal toroidal edge current density is flux-surface averaged and also serves as a proxy for the magnetic shear. The positions of the numbers mark the grid points of the $j - \alpha$ grid while the numbers themselves correspond to the dominant toroidal mode. The diagram is typical for type-I ELMs. In the high j , low α region the most unstable modes are low n peeling modes while the low current, high α region is dominated by high n ballooning modes. The mode numbers were only scanned up to $n = 20$ but the toroidal mode number spectra calculated for this region are almost flat which indicates that the dominant mode number is only slightly larger than 20. The operational point, indicated by a yellow star, lies in a region of intermediate n . It is rather far inside the unstable region, while according to peeling-ballooning theory the ELM crash should already have occurred when the operational point reaches the marginally unstable region. However, the calculations do not take into account the stabilising effects of diamagnetic drifts and sheared rotation. Additionally, the assumption $T_i = T_e$ might not be valid as the ion temperature pedestal is typically slightly wider than the electron temperature [15], and the stability diagram was only generated using a low cutoff parameter $\Psi_b = 0.990$. When using $\Psi_b = 0.994$ and an experimental T_i profile, the stability boundary would move slightly into the red region of the plot, while the operational point would move toward the boundary.

As the diagnostics available at AUG have greatly improved since the discharge was performed, it was repeated in order to analyse the ELM cycle in greater detail. Edge T_i profiles are now available and the time resolution of the magnetic measurements has increased. However, the hardware from the ECE diagnostic is older which leads to more scattered T_{rad} data. The time traces from the repeat of discharge 23417 can




Figure 5.4: *Type-I ELMs: time traces of the divertor current from shunt measurements (black), the power to the outer divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature at the pedestal top (red)*

be seen in figure 5.4. While the ELM-induced divertor power loads and the amount of confined energy were comparable to the reference, the ELM frequency was clamped between 30 and 40 Hz with only a very few outliers. The energy and density losses were much higher than in the reference, as was the density itself. The temperature at the pedestal top was lower than in the reference.

The pre-ELM profiles are shown in red in figure 5.2 to compare them to the reference discharge. While the pedestal top electron temperature is similar, the pedestal itself is slightly wider. However, the uncertainties at the pedestal top are relatively large in this discharge and the electron temperature in this region is strongly influenced by the constraints imposed by the modified tanh fit that was used. The T_i pedestal is lower and wider than T_e and the electron density is 15% higher than in the reference. The resulting total pressure gradient and edge current density are 35% lower, leading to a less pronounced low shear region and a higher edge q than in the reference. Rather than only calculating the peeling-ballooning stability of the pre-ELM case, the evolution of the operational point with respect to the stability boundary will be assessed by analysing different time points in the ELM cycle.

The evolution of the kinetic profiles during type-I ELM cycles in AUG was analysed in great detail in reference [17]. It was found that the recovery after the crash can be divided into several distinct phases. Directly after the crash the electron temperature and density profiles stay relaxed for roughly 1.5 ms. Strong oscillations picked up by magnetic diagnostics indicate that there is still MHD activity in this phase. The T_e profile then briefly starts to recover but this recovery halts after 1ms, at which point the density recovers. After roughly 2.5ms, when the density profile has reached its pre-ELM shape, the second T_e recovery phase sets in. Finally, the n_e and T_e gradients and, therefore, ∇p_e either stay clamped at their maximal values before the next ELM crash occurs, or the instability arises immediately. The same behaviour was observed




Figure 5.5: *Recovery of the electron temperature gradient (top) and temperature profiles corresponding to the different phases (bottom) [62]*

in the reference discharge 23417 just discussed. The evolution of the edge electron and ion temperature profiles of the newer discharge 27963 throughout the ELM cycle is shown in figure 5.5 [62]. The large top frame shows the recovery of the maximal edge T_e gradient and the subsequent plots the temperature profiles during the phases marked by the colored bars. The black T_e profiles were calculated using ECFM and the red curve is a fit to the T_i data shown in blue. Note that the pedestal top electron temperature shown here is lower than in figure 5.2 because there TS data was included in the modified tanh fit in figure 5.2 and because the fit function assumes symmetry between the pedestal top and bottom. Additionally, only one ECE channel is available in the region slightly inside the pedestal top, and it is lower than one would assume from the shape of the remaining channels (see kink in the ECFM T_e profiles between $\rho_p = 0.92$ and 0.96 in figure 5.5). It is not clear whether this is due to a wrong calibration of the ECE channel or an instability localised in this region prevents the electron temperature from recovering. While in this discharge the phases in the recovery of the gradient are not as clear as described in reference [17], they can still be distinguished. Phase (ii) in figure 5.2 corresponds to the collapsed phase after the crash. Phase (iii) is the phase in which the T_e profile recovery briefly halts. It is followed by the final ∇T_e recovery phase, which lasts from the end of the blue region to the end of the purple region. Finally, the electron temperature gradient stays clamped to its maximum, while the pedestal width continues to increase. The question arises why the temperature and density profiles can reach their final gradient long before the next ELM crash occurs. To explore this, selected data from the equilibrium evolution presented in reference [62] were chosen for a stability analysis, namely the equilibria corresponding to phases (i),

(iii) (iv) and (v) in figure 5.5.




Figure 5.6: *Stability diagrams at different times in the type-I ELM cycle of discharge 27963*

The stability diagrams corresponding to these phases are shown in figure 5.6. While the diagrams were calculated using $\Psi_b = 0.990$ the dashed black curve in the last diagram represents the stability boundary (red to orange region) calculated using $\Psi_b = 0.994$. Regardless of which boundary is used, the occurrence of the ELM crash can not be explained from peeling-ballooning theory. Figure 5.5a shows the diagram corresponding to the beginning of the temperature recovery. It is far inside the stable region, as one would expect from a post-ELM equilibrium. Toward the end of the recovery phase of ∇T_e , marked by (iv) in figure 5.5, the edge current has almost reached its pre-ELM value while α is still 20% lower than just before the ELM. Diagram c corresponds to the beginning of the phase with clamped ∇T_e and d to the pre-ELM equilibrium. While the location of the operational point is similar, the stability boundary moves inside the formerly stable region. This is due to the pedestal broadening at constant pressure gradient. However, just before the ELM the equilibrium is still clearly stable and according to peeling-ballooning theory an increase of almost 50% is necessary in the edge pressure gradient and current density to render it unstable. This is well above the experimental errors.

The operational points of phases (i), (iii), (iv), (v) and (vi), as well as their respective stability boundaries are overplotted in figure 5.7 using the color coding from figure 5.5. Aside from not reaching the stability boundary toward the end, the plasma evolution during the ELM cycle behaves as one would expect. At the beginning of the recovery (phase iii, blue), the pressure gradient and edge current density are still very low,




Figure 5.7: *Evolution of the operational point during a type-I ELM cycle of discharge 27963 in $j - \alpha$ space*

before they start to move in the direction of the high j and high α region of the peeling-ballooning stability boundary. Toward the end of the recovery phase (purple) the stability boundary is close to the pre-ELM boundary because the pedestal is still relatively wide. In phase v (cyan equilibrium), the boundary is at much higher values of the pressure gradient because of the narrower pedestal at the beginning of the saturated gradient phase. It appears that the current density and the pressure gradient are slightly higher than in the pre-ELM cases, but it can be seen in the work published by M. Dunne (figure 5.7 in reference [15]) that this is within the experimental uncertainties. Note the pre-ELM case (green) and the equilibrium toward the end of the ELM cycle (yellow) are very close. However, the boundary still moves toward the operational point since the pedestal width continues to increase. This is consistent with former investigations of the evolution of the electron pressure profile, which found that the pedestal top continues to increase toward the end of the ELM cycle while its gradient stays clamped to a maximum value [17]. This is also consistent with the assumptions of the EPED model [69], which hypothesises that the pedestal gradient is first limited by a transport instability while the pedestal height and width continue to evolve, leading to an unstable peeling-ballooning mode.

A gyrokinetic analysis of the same equilibria as the ones presented here was performed with the GENE code [70] and preliminary results indicate that kinetic ballooning modes may be near marginal stability in the saturated gradient phases [71]. While the results presented here do not support the triggering of the ELM by an ideal peeling-ballooning mode, the stability boundary is still expected to evolve between the pre-ELM time point and the ELM crash as the pedestal width continues to increase. This evolution could not be analysed because of the high integration time of the CXRS diagnostic. Assuming a linear evolution of the boundary in time and a constant inter-ELM time of 30ms, we can extrapolate the pre-ELM boundary from the cyan boundary in figure 5.7 (16ms before the ELMs), the yellow (12ms before the ELM) and the purple one (2-3ms before the ELM). The result (black dotted line) is still far away from the green operational point, which is not expected to change. The pressure gradient of the pre-ELM equilibrium in discharge 27963 is 30% lower than the stability boundary. In this discharge, type-I ELMs can, therefore, not be explained by ideal MHD stability.

5.2 Type-III ELMs: small ELMs close to the L-H threshold

Type-III ELMs, which occur just above the L-H threshold, are smaller than type-I ELMs and can be identified by their frequency which increases with heating power. Type-III ELMs could be relevant for the initial operational phase of ITER. Distinct type-III ELMs are rare at AUG: most discharges only feature type-I events or small filament-like structures which can not clearly be identified as ELMs. However, they do appear in H-mode plasmas close to the H-mode density limit [72]. Discharge 27138 featured large ELMs causing more than 10% W_{MHD} losses each, which gradually disappeared to give way for type-III ELMs as the density increased.




Figure 5.8: *Type-III ELMs: time traces of the divertor current from shunt measurements (black), the power to the outer divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature just inside the pedestal top (red)*

Figure 5.8 shows time traces of discharge 27138 after the transition to type-III ELMs. It was run at 800kA with a magnetic field of -2.5T, 7.5MW of NBI heating and no ECRH. The amount of fuelling, $6 \cdot 10^{21} \text{s}^{-1}$, was relatively high because the discharge was used to explore the density limit of H-mode operation. The type-III ELM frequency lay between 100Hz and 400Hz. The divertor shunt measurements and the divertor power load time traces show that compared to the inter-ELM signal, type-III ELMs are much smaller than type-I ELMs (figs. 5.1 and 5.4). While type-I ELM losses are generally around 10%, figure 5.8 indicates energy losses below 2% and a drop of the edge density between 2% and 3%. The ECE time trace, which represents the radiation temperature slightly inside the pedestal top at $\rho_p = 0.93$, shows a drop by roughly 100eV after each ELM. Half of the temperature loss is recovered within 0.3ms after the crash. Note that this could also be an artifact due to a fast ELM-induced contraction of the plasma which might re-expand very quickly. Since the ECE diagnostic measures at a fixed radial position, the ELM-induced contraction of the plasma causes a given edge ECE channel to measure the lower temperatures at larger values of ρ_p . When the plasma re-expands, the channel returns to its initial radial position relative to the separatrix, such that an apparent increase in temperature is observed. However, the

time resolution of the standard plasma equilibrium is not high enough to resolve such fast movements.




Figure 5.9: *Pre type-III ELM profiles of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)*

The blue curves in figure 5.9 show that the kinetic pedestal profiles prior to type-III ELMs are much shallower and wider than in the type-I ELM My H-mode cases presented earlier (cf. figure 5.2). While it is not sensible to compare the absolute pressure gradient to the type-I ELM My H-modes discussed in the previous section because they were run at 1MA plasma current, a type-I-ELMy H-mode run at 800kA is presented in section 5.7. Discharge 29294 (black curves in figure 5.9) has a slightly lower pressure gradient and a much narrower pedestal, even though it was only heated with 2.5MW NBI and 2.3MW ECRH. The wide pressure gradient profile of the type-III ELM My H-mode 27138 drives a broad local edge toroidal current density slightly above 1MAm^{-2} . Because of the wide edge current, the region of low shear around $\rho_p = 0.98$ is also very broad. If the plasma is in the second stability region, this low shear region will have a stabilising effect on the ballooning modes. In the type-I pre-ELM case the toroidal edge current density is slightly higher, but its flux-surface average slightly lower (see figure 5.10). The $j - \alpha$ diagrams in figure 5.10 shows that in type-III ELM My H-mode the plasma is stable against peeling-ballooning modes (left), while the type-I case is very close to the stability boundary. The diagrams were generated using $\Psi_p = 0.994$, and the boundaries for $\Psi_b = 0.990$ are indicated by the dotted black curves. The operational points are at similar locations, but the stability boundary of the type-I case is at lower j and α values. Both a 50% higher edge current density and pressure gradient are necessary to drive the plasma unstable. This is not surprising because unlike in the model for type-I ELMs, finite resistivity is believed to play a non negligible role in the trigger mechanism of type-III ELMs [73]. However, the resistivity mainly depends on the temperature, which is very similar in the two cases. The boundary in the type-I case is probably lower because of a slightly lower triangularity than in the type-III case and because the current density is more localised at the plasma edge.

In the type-I ELM My phase of discharge 27138 the pedestal pressure gradient decreases




Figure 5.10: *Stability diagram of the plasma before a type-III ELM*

as the density increases because of a lower temperature. This stabilises the type-I ELMs, but it was shown for the similar discharge 28726 that the radial electric field well at the edge weakens with increasing density [74]. Type-III ELMs are thought to be driven by resistive interchange instabilities which are stabilised by shear in the radial electric field [73]. As this shear decreases with the depth of the E_r well when the density increases, type-III ELMs appear.

5.3 Type-II ELMs: strongly fuelled discharges close to double null

Type-II ELMs are small fluctuation like events which do not induce large crashes in the kinetic profiles like type-I ELMs. Rather, the losses are most likely caused by frequent small filaments detaching from the plasma. In AUG, type-II ELMs occur in high triangularity discharges close to DN and with a high density. The triangularity depends on the radial positions of the upper and lower X-points compared to the position of the geometric axis of the plasma, the geometric axis being equidistant from the outer and inner separatrices at the horizontal midplane. Unlike the magnetic axis, it is independent of the Shafranov shift. The higher the triangularity, the more D-shaped the plasma.

The type-II ELM discharge analysed in this section was run with a plasma current of 800kA at -2.5T, 5MW NBI heating and 0.75MW ECRH. The fuelling was $9.3 \cdot 10^{21} \text{s}^{-1}$ because a high gas fuelling rate is necessary to prevent impurity accumulation in discharges with high triangularity. Figure 5.11 shows the poloidal topology of the presented discharge. In the first phase (black), the triangularity was already high, but the upper separatrix was still far from the second X-point. This phase was still dominated by type-I ELMs with a frequency of 40-80Hz. However, when the upper separatrix was gradually moved toward the second X-point, the regime changed to a type-II ELM H-mode (red).

Characteristic time traces from the two phases are shown in figure 5.12. While in the type-I phase the ELM-induced heat loads reach up to 7MW in the outer divertor, in the type-II phase they fluctuate around 1-3MW, similar to the inter-ELM phase in the type-I ELM H-mode. During type-II ELMs the signal from the shunt measurements in the divertor (top panel) shows much more activity than during the inter-ELM phases.




Figure 5.11: *Shape of the equilibrium during the type-I ELM phase (black) and the type-II phase (red) in discharge 25740*

The type-I ELMs induce energy losses of 10%-20% and the drop in edge density is in the same range. In the second phase the confined energy is higher, which is mainly due to the higher density despite the slightly lower temperature. The type-II ELMs, which have a frequency between 200Hz and 1kHz, induce energy losses up to 3% and density fluctuations in the range of 1.5%. While the type-I ELMs are characterised by large drops in the pedestal top temperature of more than 50%, the type-II ELMs cause strong temperature fluctuations of around 50eV. These fluctuations occur over a broad band of frequencies between 20 and 65kHz just inside the pedestal top [18, 19]. They can also be observed in the signals from magnetic pickup coils. While it is possible to correlate some of the larger events in the divertor with drops in T_{rad} and n_e , most of the fluctuations seem uncorrelated. This suggests that type-II ELM losses are very localised in poloidal and toroidal angle and the heat and particles travel on a toroidally localised bundle of open field lines to the divertor. Therefore, they reach the divertor plates at a specific toroidal location that depends on the pitch angle of the field lines and where they were released from the confined plasma, rather than being a global phenomenon like type-I ELMs. Such localised losses would explain why the events observed by the ECE diagnostic cannot necessarily be seen in the time traces from the divertor diagnostics shown in figure 5.12 as both sets of measurements correspond to distinct toroidal locations.

Figure 5.13 shows the pre-ELM kinetic profiles in the type-I phase (black) and the averaged profiles in the type-II phase (red). Before type-I ELMs the pedestal top temperatures of the ions and electrons are similar. The ion temperature is less affected by the change to type-II ELMs than the electron temperature, which drops by 20%. While the pedestal top and core density increase in the second phase, the pedestal itself widens only slightly. The lower temperature and wider pedestal cause the pedestal top pressure and the pressure gradient to drop when the type-I ELMs disappear. In turn, this drives a lower but broader edge current density. Therefore, the q-profile has higher




Figure 5.12: *Type-I (left) and type-II (right) ELMs: time traces of the divertor current from shunt measurements (black), the power to the outer divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature at the pedestal top (red)*

values at the plasma edge during type-II ELMs and the region of low shear is not as pronounced.

The stability diagrams for both phases, calculated with $\Psi_b = 0.994$, can be seen in figure 5.14. The type-I pre-ELM equilibrium is 10% under the stability boundary in current and pressure gradient, which is within the uncertainties of the measurements. To generate the equilibria used here, the ECFM was ELM-averaged to a single profile. If, however, one would fit T_e with the commonly used modified tanh function, the pressure gradient would be 10% steeper and the equilibrium would be marginally unstable. The type-II equilibrium, however, is far from the unstable region of the diagram. Both the pressure gradient and the edge current density must increase by roughly 60% to reach the stability boundary, which is well above the experimental uncertainties. Note that the stability boundary in the right figure is at higher values of α than in the left one. This is because higher pressure gradients become more stable as the triangularity is increased and the equilibrium gets closer to DN, which is consistent with JET results [67, 75]. Unlike at JET, the low current region also becomes more stable. Previous results suggest that the radial mode structure of type-II ELMs is narrower than that of type-I ELMs, which may indicate that the extent of the peeling-ballooning mode is linked to the ELM affected area [76]. As the type-II equilibrium is so far in the stable region and ILSA does not find any unstable modes, such a comparison is not possible. However, comparing the radial mode structure of the dominant toroidal mode at $\alpha = 6$ and $\langle j_{tor} \rangle = 0.8 \text{ MA m}^{-2}$ in the two diagrams showed that the mode width does not decrease in the type-II case (see figure 5.15, the colors change toward red with increasing m).

The stability analysis of type-II ELMs indicates that these instabilities have a too low pressure gradient and edge current density to be driven by peeling-ballooning modes. Previous results that indicate that the radial extent of the mode driving type-II ELMs is smaller than in type-I ELMs could not be confirmed. However, in the framework




Figure 5.13: *Pre-type-I ELM profiles (black) and profiles during a type-II ELMy phase (red) of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)*




Figure 5.14: *Stability diagram of the plasma before type-I ELMs in the reference phase of discharge 25740 (left) and during the type-II phase*

of the EPED model, an interruption of the pedestal width evolution could cause a small ELM event [77]. Microtearing modes that are present slightly inside the pedestal top region and cause the broad band fluctuations in T_e could prevent the pedestal top electron temperature from reaching the same values as in the type-I phase [78]. This would explain why the plasma pressure does not reach the peeling-ballooning boundary and the type-I ELMs disappear. However, this explanation is highly speculative.

5.4 Type-I ELM mitigation at high densities

When increasing the fuelling in type-I ELMy H-mode discharges, type-I ELMs disappear and smaller, type-II like events start to appear. As this regime is more easily accessible at low plasma current, the 600kA, 2.2T discharge 28017 is presented. It was heated with 5MW NBI and 0.6MW central ECRH, and fuelled at $2.8 \cdot 10^{21} \text{s}^{-1}$.




Figure 5.15: *Radial mode structure of the dominant toroidal mode number at $\alpha = 6$ and $\langle j_{tor} \rangle = 0.8 \text{ MA m}^{-2}$ from the two diagrams shown in figure 5.14*

The type-I ELM frequency was very scattered between 20Hz and 60Hz and dropped to only sporadically occurring ELMs with increasing density. Above a certain edge density, type-I ELMs stopped occurring altogether. It was planned that the edge of the discharge presented in this section would be heated with ECRH, but due to an incorrect polarisation setting the majority of the microwaves was not absorbed. To ensure that the disappearance of the ELMs is not correlated with the small amount of edge-deposited ECRH or the glowing PFCs resulting from the incorrect polarisation, the discharge is compared to a similar one without edge ECRH heating. Figure 5.16




Figure 5.16: *Line integrated core (blue) and edge (black) electron density in two discharges. The red crosses mark the change into ELM-free H-mode*

shows the core and edge line-integrated core densities of the two similar discharges. The red crosses mark the time and density at which the last type-I ELM appears. The density threshold for reaching the ELM-free H-mode is the same in both discharges, but it is not clear whether this is a core or an edge effect. However, since the second discharge did not have any ECRH power deposited at the edge of the plasma or on the

PFCs, ECRH cannot be responsible for the suppression of the ELMs in the reference.




Figure 5.17: *Type-I ELMs (left) and ELM-free H-mode (right): time traces of the divertor current from shunt measurements (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature just inside the pedestal top (red)*

Figure 5.17 shows the time traces of the type-I (left) and ELM-free phase (right) of discharge 28017. In the left panel, the magnitude of the divertor current is a factor of two lower than for typical type-I ELMs, while the W_{MHD} drops are of the order of 7%, slightly lower than the typical 10% drop. In addition to the ELMs themselves, strong inter-ELM fluctuations can be observed in the divertor shunt measurements. As the density increases, the fluctuations become the dominant energy loss mechanism and the type-I ELMs disappear. Directly after the transition to the ELM-free H-mode the stored energy exceeds the pre-ELM level in the type-I ELM phase. The confinement soon starts to degrade after the time window shown in figure 5.17, while the density stays high. Although the time traces from the ELM-free phase in figure 5.17 are reminiscent of the type-II ELMs presented in the last section, a Fourier analysis of the pickup coil signals showed that they do not feature the 20-65kHz broad band fluctuations which are characteristic of type-II ELMs.

The kinetic profiles, shown in figure 5.18, change only a little when the ELM-free phase sets in. The electron density and total pressure pedestal top values are slightly higher, but this has little effect on the pressure gradient because the pedestal width slightly increases. The edge current density, however, is 10% lower in the ELM-free phase.

Both the pre-ELM equilibrium and the ELM-free case are in the stable region for peeling-ballooning modes (figure 5.19). While the pre-type-I ELM case is slightly closer to the unstable region and the boundary itself slightly moves, these trends are within the uncertainties of the experimental measurements. Note, however, that the dominant toroidal mode numbers at the stability boundary are relatively small compared to the high triangularity type-I and type-II ELM plasmas discussed in section 5.3, indicating that in discharge 28017 a different kind of instability dominates, which has a peeling like behaviour. Such a current driven instability can be expected from the high relative




Figure 5.18: *Pre-type-I ELM profiles (black) and profiles during an ELM-free phase (blue) of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)*




Figure 5.19: *Stability diagrams of the plasma before a type-I ELM (left) and in the ELM-free phase (right) of discharge 28017*

current density in this discharge. With a maximum of 0.9 to 1 MA m^{-2} it is only slightly lower than in some of the 1 MA discharges discussed earlier, while the total plasma current is only 600 kA . The small events arising with the suppression of the type-I ELMs at high densities could be more closely related to the type-III ELMs discussed in section 5.2 which also have small dominant mode numbers.

Figure 5.20 shows the radial mode structure of the dominant toroidal mode number of the pre-ELM equilibrium and the ELM-free phase. The modes shown here correspond to the equilibria with 10% increased current density and pressure gradient. The broad harmonics confirm the peeling-like behaviour of the instability. However, the modes extend far into the plasma, which contradicts a possible correlation between the width of the linear MHD mode and the ELM size.

Resistive effects could also be important with the dominant current driven modes. This would also account for the small size of the ELM crash, where the resistivity




Figure 5.20: *Radial mode structure of the dominant toroidal mode number of the type-I pre-ELM equilibrium and the ELM-free phase. In both cases the modes shown here correspond to the equilibria with 10% increased current density and pressure gradient.*

dampens significant changes in the current density. The small fluctuations would then correspond to a “soft” peeling limit, analogous to the soft ballooning limit described for type-II ELMs by Saarelma *et al.* [75]. A resistive MHD analysis and non-linear simulations are required to further analyse this hypothesis.

5.5 N-seeded ELMs: ELM mitigation by nitrogen seeding

As already mentioned, ELM-induced heat loads might exceed the heat flux limits of the PFCs of future larger fusion devices. To mitigate the power loads both from the ELMs and in the inter-ELM periods new scenarios have been developed in which low-Z impurities are seeded in the divertor, in addition to the normal D fuelling. The power radiated by these impurities is much higher than that radiated by deuterium, cooling the divertor efficiently. The hot particles leaving the confined plasma and arriving in the divertor along the open field lines are cooled down before reaching the plates, causing a lower heat flux than they otherwise would. It was found that seeding nitrogen is ideal for this purpose [79]. As a beneficial side-effect, nitrogen seeding in the divertor also increases the stored energy in the plasma in a full metal tokamak. In a machine with carbon walls, however, no improvement was observed [80, 81]. It appears that nitrogen seeding recovers the confinement degradation that occurred when changing from carbon PFCs to a full metal wall. Since introducing a low-Z impurity into the plasma edge dilutes the deuterium ions, and thus leads to a lower ion density, the ion temperature can increase while the total pedestal pressure stays constant. If one assumes stiff core T_i profiles, this effect can propagate to the plasma centre, leading to a higher stored energy content. However, the confinement improvement observed when seeding nitrogen is too strong to be caused by this effect alone and no conclusive explanation has yet been found [82]. In addition to the confinement increase, it was found that the presence of nitrogen mitigates the ELMs and, in some cases, increases the ELM frequency.




Figure 5.21: *Pick up coil signal and divertor current during non-seeded ELMS (a and b) and N-seeded ELMS (c and d), as well as the evolution of the n_e profile in the two phases (e) [83]*

In the full-W AUG, the actual crash of typical type-I ELMS has been hypothesised to consist of two phases, which can be seen in the time traces in figure 5.21a and b [83]. These frames show the signal from a magnetic pickup coil and the divertor current. In the first phase, which lasts roughly 0.5ms, the pedestal top pressure is eroded but the steep gradient region remains unaffected (phase 1a). For better visualisation, this erosion is sketched for the electron density in panel e. In the second phase, which is typically 1ms long but can also last up to 5ms, the pressure gradient drops and more energy is lost (phase 1b). In nitrogen-seeded discharges, however, the second phase does not occur (see panels c and d). The recovery from the ELM begins directly after the pedestal top is eroded. While this could be the reason for the smaller ELM size and higher frequency, it does not explain the higher pre-ELM energy stored in the plasma. During the 2013 experimental campaign of AUG, a reproducible discharge was performed in which nitrogen seeding led to a dramatic increase in stored energy of 40%. This 1MA, -2.5T discharge was strongly heated using 10MW NBI, 1.5MW ECRH and 1.8MW ion cyclotron heating, which is radio-frequency heating affecting both the ions and the electrons. The discharge was also very strongly fuelled at $23 \cdot 10^{21} \text{s}^{-1}$ deuterium and an additional $22 \cdot 10^{21} \text{s}^{-1}$ nitrogen seeding in the second phase. Note that the fuelling rate is expressed in electrons per second. While when fuelling hydrogen or deuterium this also corresponds to the neutral fuelling rate, when injecting impurities the rate has to be divided by the charge number of the impurity species. The ELM frequency, between 120 and 150Hz in the reference phase, dropped slightly when seeding was applied, in contrast to typical observations [83]. Figure 5.22 shows the evolution of the plasma over several ELM cycles in the reference phase (left) and with nitrogen seeding (right). The ELMS, which are much smaller in the seeded phase, do not cause any power loads to the inner divertor, which is fully radiatively cooled. The per-ELM energy loss is slightly lower and 40% of this energy is recovered within 0.5ms after the end of the crash. Note that the y-axis of the W_{MHD} signal is different for the non-seeded and seeded phases because of the dramatic increase of the confined energy. However, the scale was kept the same to enable a better comparison of the




Figure 5.22: *Type-I ELMs without (left) and with (right) N-seeding: time traces of the divertor current from shunt measurements (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature around $\rho_p = 0.9$ (red). Note that the plotted W_{MHD} range is different in the left and right plots.*

losses. While the drops in temperature are comparable in both phases, the particle losses are much lower when nitrogen is seeded. The improved confinement appears to originate mainly from an increase in temperature.




Figure 5.23: *Pre-type-I ELM profiles without (black) and with (blue) N-seeding of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)*

This can also be seen in the temperature profiles shown in figure 5.23. Compared to the reference phase (black), both the ion and electron temperatures strongly increase in the seeded phase (blue). It appears that the T_e pedestal broadens with similar gradient, while the T_i gradient increases. However, the ion temperature is only poorly diagnosed in the reference phase because the CXRS systems were set to measure nitrogen. This is sensible for the seeded phase in which the intensity of the nitrogen signal is very

high, but not in the reference phase. When this experiment is repeated in the future it should be performed in two separate discharges, such that the CXRS systems can measure boron in the reference discharge. While the edge density gradient is similar




Figure 5.24: *Stability diagrams of two pre-ELM plasmas without (left), and with (right) N-seeding. In the top ones the growth rate is normalised to ν_A , in the bottom ones to $\omega_{*,pi}/2$*

in both phases, the difference in the pedestal top density lies within the uncertainty of the DCN and TS measurements, which disagree in the first phase. According to the data from the DCN interferometer, the edge density should be similar in both phases. Because of the inaccuracies in T_i and n_e , the pedestal pressure in the reference phase might be higher than the one shown in figure 5.23. Even accounting for these uncertainties, the pressure profile in the presence of nitrogen is much steeper with comparable width. The almost 100% increase in pressure gradient leads to a 30% increase in edge current density at the outboard side of the plasma. In addition to the local toroidal current density at the LFS, figure 5.23 shows the parallel current density averaged over each flux surface (dashed lines). Interestingly, $\langle \mathbf{j} \cdot \mathbf{B} \rangle$ is the same in both phases. This means that the peeling drive relative to the ballooning drive will be lower since the pressure gradient strongly increases.. It was recently shown that when seeding nitrogen the flux surface-averaged edge current density can even drop despite an increasing pressure gradient [15]. This was found to be consistent with neoclassical calculations of the bootstrap current. The local current density on the outboard side is higher because it includes the Pfirsch-Schlüter currents, which increase with the pressure gradient but average to zero on a flux surface.

It can be seen in the top panels in figure 5.24, which were generated using $\Psi_b = 0.994$, that both the unseeded (left) and the N-seeded (right) pre-ELM equilibria are unstable. The unseeded reference case is close to the stability boundary and, since it was not very well diagnosed, its uncertainties are relatively high. The quality of the T_i and n_e profiles collected in the N-seeded plasma, however, is very good. Only the T_e gradient might be slightly too steep, since a very steep function was fit to the relatively scattered data. However, this is not sufficient to explain why the operational point is so far in the unstable region. The dominant toroidal mode numbers are very high, which means that the stabilising effect from the diamagnetic drift will be stronger than in the discharges presented earlier. The bottom panels show the same diagrams with the growth rate normalised to the ion diamagnetic drift frequency divided by two, where a mid-pedestal density of $4 \cdot 10^{19} \text{m}^{-3}$ was used in the normalisation. Since high toroidal mode numbers are strongly stabilised, relatively low mode numbers become dominant, $n=8$ in the reference and $n=4$ in the seeded case. With this normalisation the unseeded plasma is now marginally stable, while the seeded case is still 10% in the seeded region. The radial mode structure of the $n=19$ mode, which is dominant when using the ν_A




Figure 5.25: *Radial mode structure of the dominant toroidal modes in figure 5.24*

normalisation, is shown in the left frame of figure 5.25. The structure is similar in both phases, with the radial extent in the seeded phase being slightly larger. In the right frame the modes which are dominant when the diamagnetic drift stabilisation is taken into account are compared. The mode extends further inside in the seeded phase, as was the case without $\omega_{*,pi}$ stabilisation. This confirms that the reduced ELM size is not due to a narrower instability and indicates that there are two distinct phases to the ELM crash. The first, then, would be linked to linear stability, while the second is likely a nonlinear process. Due to this nonlinearity, the second phase cannot be investigated using a linear MHD code like ILSA.

To summarise, when seeding nitrogen into the divertor of a full metal machine reproducible improvement in confinement and a smaller ELM size are observed. The pressure gradient and local toroidal current density at the LFS of the plasma strongly increase while $\langle \mathbf{j} \cdot \mathbf{B} \rangle$ stays constant or even drops. The resulting equilibrium is highly unstable against the same peeling-ballooning modes that are dominant in the unseeded reference H-mode plasma while the mode structure from the linear calculations and the ELM phenomenology of the first phase in the crash are identical. This

suggests that a stabilising mechanism must exist that is not included in ideal MHD. Since the dominant mode numbers are very high, the diamagnetic drift stabilisation is very strong. When taking it into account both the unseeded and seeded cases are within 10% of the stability boundary. The first phase of the ELM crash is most likely linked to linear peeling-ballooning stability, but the second phase, which is suppressed when seeding nitrogen, is a nonlinear phenomenon and cannot be described using linear codes.

5.6 ELM triggering via frozen deuterium pellets

At AUG, cubic frozen deuterium pellet with a side length of 1.5-2.0mm can be injected into the plasma with velocities up to 1km/s by using a centrifuge. In future large fusion reactors these pellets will constitute the main source of particle fuelling. Type-I ELMs are typically triggered as a side effect of injecting pellets into an H-mode plasma [84]. If the pellet repetition frequency is high enough, it is even possible to artificially set the ELM frequency. Recent experiments at DIII-D achieved an increase of the natural ELM frequency by a factor of 12 [85]. Since a higher ELM frequency at constant input power reduces the energy losses per ELM, and thus leads to a lower peak power flux to the divertor, ELM pacing has long been considered to be a possible solution to the too high ELM-induced heat fluxes that threaten to melt the plasma facing components of future devices like ITER.




Figure 5.26: *Pellet triggered ELMs in C-AUG: time traces of the divertor D-alpha radiation (black), the pellet monitor (green), the power to the outer divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature just inside the pedestal top (red)*

Figure 5.26 shows the time traces for discharge 19829. In addition to the same signals as shown in the previous sections, a monitor signal that indicates the arrival of pellets in the plasma is plotted in green. Since the divertor currents were not available at the time of the discharge, the D- α radiation is shown instead. During the plotted time interval, four spontaneous and two triggered ELMs occur. The pellet monitor shows the

two fragments of a broken pellet arriving into the plasma within 1.5ms of each other, where the first fragment arrives 4.3ms after the previous spontaneous ELM. Since both pellet fragments trigger an ELM, it appears that with a fast enough pellet launcher ELM-pacing would be possible at frequencies of around 500Hz. Both triggered ELMs cause a lower power load to the divertor than the spontaneous ELMs, which is crucial for ELM mitigation via pellets. However, the discharge shown in figure 5.26 was carried out when some of the plasma facing components of AUG were still made of non-coated graphite. It was recently shown at JET [86] and later confirmed at AUG [87] that ELM triggering is much more difficult in a full metal machine. It appears that in a metal device ELM triggering via pellets is only possible after a specific lag time.




Figure 5.27: *ELM energy losses in spontaneous and pellet-triggered ELMs in the C AUG (a), with full tungsten walls (b) and in N-seeded discharges (c) [87]*

Figure 5.27 shows the relative W_{MHD} losses of spontaneous (blue) and pellet-triggered ELMs (red) in AUG discharges run in the carbon AUG (a), the full tungsten machine without (b) and with (c) N-seeding [87]. The green triangles in (b) are triggered ELMs from an ELM-pacing experiment at 70Hz. In carbon, energy losses induced by triggered ELMs are lower than the losses from spontaneous ones. ELMs can already be triggered immediately after a previous crash; no lag time is observed. However, the natural ELM frequency is very scattered and a spontaneous ELM can already occur 2-4ms after the previous event (panel a). The spontaneous ELMs occurring in discharges run with a full tungsten wall which are shown in panel b have a minimum inter-ELM time of 12ms. However, only a subset of the ELMs occurring in the analysed discharges are plotted, and in some cases a few events already occur after 8ms. Successful ELM-triggering is only possible 8-10ms after the previous event, as can be seen from the zero W_{MHD} losses before 10ms in panel b. However, when nitrogen is introduced into the full tungsten machine, successful triggering can already be achieved after 2ms (panel c). Unlike in the carbon machine, this is possible despite a minimal inter-ELM time of spontaneous ELMs of 8ms. The possibility of faster triggering could be connected to the lack of the second phase in the ELM crash with N-seeding and the faster recovery of the pedestal discussed in section 5.5 [83]. If the pedestal top is eroded but the steep gradient region is only slightly affected, a sudden local density increase induced by the pellet might cause the plasma to become unstable against peeling-ballooning modes. In this section one discharge of each of the three scenarios introduced in figure 5.27 will be analysed. In the C-AUG and nitrogen seeded cases two equilibria will be presented, one taken just before a spontaneous ELM and one taken in the recovery phase when ELMs can already be triggered. In the tungsten case, three equilibria will be discussed, one taken just before a crash and two from the recovery phase, just before and after the triggering limit.




Figure 5.28: *Pellet triggered ELMs in W-AUG: time traces of the divertor current from shunt measurements (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature just inside the pedestal top (red)*

The W-AUG discharge without N-seeding was run at 1MA, 2.5T and heated using 5MW NBI and no ECRH heating. The gas fuelling level was set to $6.7 \cdot 10^{21} \text{s}^{-1}$ and the resulting frequency of spontaneous ELMs was scattered between 30 and 130Hz. Figure 5.28 shows three typical responses of the plasma following a pellet. The left panel shows a pellet triggering an ELM shortly after the previous event. The inter-ELM time, 7ms, corresponds to the inter-ELM time of the spontaneous ELMs occurring with the highest frequency in this discharge, 140Hz. In the second panel a typical triggered ELM occurring 20ms after the previous one can be seen. This time corresponds to the lower range of the natural ELM frequency. The last panel shows the arrival of a pellet 5ms after the previous crash which does not trigger an ELM. It does, however, seem to trigger small filaments than can be seen in the divertor shunt measurement and power signals. In general, the ELMs triggered in this discharge are not smaller and do not induce a lower heat flux in the divertor than spontaneous ones. At the arrival of each pellet the density first strongly increases, as expected from a deuterium pellet whose primary function is to fuel the plasma. After successful ELM triggering, it drops from this high density to the value prior to a spontaneous ELM, rather than dropping to the typical post-ELM value. The latter is only reached after the next spontaneous event. When triggering is not successful, the density remains high until the next ELM occurs. The sharp drop in the radiation temperature after each pellet is an artifact from the ECE measurements because the microwaves can not pass through the layer of high density plasma caused by the pellet.

The edge temperature profiles in this discharge were not very well diagnosed. The ECE data from the channels in the steep gradient region are very scattered and no radial sweep of the plasma was performed to gather TS data and improve the quality of edge T_i measurements. However, it is still the best diagnosed discharge of this type and for technical reasons it was not possible to repeat it with a stronger focus on edge diagnostics. The temperature profiles presented in figure 5.29 are, therefore, only




Figure 5.29: Profiles at different time points in the ELM cycle in W-AUG of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)

preliminary until the experiments can be repeated in the next experimental campaign of AUG. Three time slices are presented in this analysis. The first one is taken just before the ELM crash (black). For the second one (blue), the data from 4ms to 7ms after the ELM crash was fitted using coherent ELM-averaging. Such a large interval had to be chosen to achieve a minimum data quality and because the edge T_i measurements are integrated over 2.2ms. These profiles correspond to a plasma in which ELM triggering via pellets is not yet possible. The last set of profiles shown in figure 5.29 (red) was taken between 7ms and 10ms after the ELM crash when triggering is possible. While all red profiles almost reach the high pre-ELM values of the black profiles, 5.5 ms after the ELM the pedestal is still very shallow. The pressure gradient is at less than 50% of its maximum value and the edge current density is also very low. Therefore, it is not surprising that it is only possible to trigger ELMs in the later time interval. While it would be interesting to correlate the profile evolution with the possibility of successful ELM triggering, the quality of the data is not good enough to perform a detailed ELM cycle analysis like the one presented in reference [17].




Figure 5.30: Stability diagrams of a W-AUG plasma before a spontaneous type-I ELM (left), 5.5ms after the ELM crash (middle) and 8.5ms after the ELM crash (right)

However, a preliminary stability analysis could be performed at the three time points in the ELM cycle and is shown in figure 5.30 ($\Psi_b = 0.994$). It reflects what was expected

from the profiles in figure 5.29: the pre-ELM time point is on the stability boundary, while the two other equilibria are stable. However, the equilibrium from 7-10ms in the recovery phase only lies 20% below the stability boundary. Since an ELM can be triggered, it appears that a pellet provides a high enough local density increase to obtain the pressure gradient necessary to reach the stability boundary. Between 4ms and 7ms after the ELM crash, however, a pressure gradient increase of more than 100% and, more importantly, a large current density increase are necessary to reach the boundary. It is unlikely that an incoming pellet can have such a dramatic impact on the pressure gradient that the edge current density instantaneously increases by 40%. Therefore, it is not possible to trigger an ELM in this phase.




Figure 5.31: *Pellet triggered ELMs in W-AUG with N-seeding: time traces of the divertor current from shunt measurements (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature around $\rho_p = 0.9$ (red)*

When seeding nitrogen in the divertor the lag time necessary for successful ELM triggering decreases significantly. The nitrogen-seeded discharge chosen for performing a stability analysis was run at 1MA, -2.6T and heated with 10MW NBI, 1.5MW ECRH and 1.5MW ion cyclotron heating. The fuelling rate was $9.3 \cdot 10^{21} \text{s}^{-1}$ deuterium with $9.9 \cdot 10^{21} \text{s}^{-1}$ nitrogen seeded in the divertor. As mentioned before, the nitrogen fuelling rate has to be divided by seven to obtain the fuelling rate of the neutrals, because the measurements are always expressed in electrons per second. The resulting ELMs occurred with a frequency of 80 to 130Hz. Three pellets, all triggering an ELM can be seen in figure 5.31. The first and third pellets both trigger an ELM very shortly after the onset of the previous spontaneous one, 2.3ms and 2.5ms respectively. In the first panel, a further spontaneous ELM occurs shortly after the triggered one, which might be due to the surplus of particles caused by the pellet. While in this discharge the ELM frequency does not typically exceed 130Hz, a few sporadic spontaneous ELMs do occur only 2ms after the previous one. This does not happen only after the density was increased by a pellet as is the case in the left panel of figure 5.31. As expected, ELM-triggering is also possible later in the ELM cycle (middle panel). The ELMs triggered shortly after the previous event induce a lower power load to the divertor than

spontaneous ones. While this is also reflected in the energy losses, it is difficult to state this for the temperature since there is a gap in the ECE channels between $\rho_p = 0.9$ and the centre of the pedestal, and because the plotted ECE channel at $\rho_p = 0.9$ is in high density cut-off after a pellet is injected. While the density increases after each pellet as in the unseeded case, it needs more time to revert to its original level. This is due to the lower ELM-induced particle losses when nitrogen is seeded (cf. section 5.5).




Figure 5.32: Profiles at different time points in the ELM cycle in W-AUG with N-seeding of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)

The black profiles shown in figure 5.32 were generated from data taken just before an ELM, while the red ones were ELM-averaged between 2ms and 4.5ms after an ELM crash, when ELM-triggering is already possible. The kink at the T_e pedestal top is due to the gap in the ECE channels already mentioned. Because of this gap, the CLISTE equilibrium was generated using a low pedestal top pressure constraint. The general behaviour of the profiles reflect what was observed in the unseeded reference case. The profiles at 3ms after the previous event, when ELM-triggering is possible, are almost fully recovered. The pressure gradient and edge current density only lie 10% below the pre-ELM values. As in the unseeded case, this is also consistent with an arriving pellet triggering an ELM.

However, stability calculations indicate that both equilibria should be stable. The diagrams in figure 5.33 were generated using $\Psi_b = 0.994$, the dotted black lines using $\Psi_b = 0.990$. Peeling-ballooning theory does not describe the ELMs occurring in this discharge since the equilibrium prior to spontaneous events would need a 15% higher edge current density and pressure gradient to be marginally unstable. The post-ELM case lies 20% under the threshold. However, one has to keep in mind that because of the missing ECE channels at the pedestal top, the real experimental temperature and, therefore, pressure profile, might be steeper or wider. This possibility is likely since the T_e and T_i gradients in figure 5.32 are similar despite the typical observation that the T_e pedestal is usually narrower and steeper than T_i [15]. Whether the actual pre-ELM equilibrium is 25% steeper and lies on the calculated stability boundary or




Figure 5.33: *Stability diagrams of a N-seeded plasma before a spontaneous type-I ELM (left) and 3ms after the ELM crash (right)*

a different stability limit exists close to the measured pre-ELM value of ∇p , a similar pressure gradient increase as after 8.5ms in the unseeded case (roughly 15%) would be sufficient to trigger an ELM 3ms in the recovery phase of the N-seeded discharge. Such an increase could arise from the strong local density increase caused by an arriving pellet.




Figure 5.34: *Pellet triggered ELMs in C-AUG: time traces of the divertor D-alpha radiation (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature in the pedestal (red)*

To present ELM-triggering in the uncoated carbon device, a different discharge was chosen than the one that was shown in figure 5.26 because it was better diagnosed. It had a plasma current of 1MA, -2.6T, 5MW NBI, 1.3MW ion cyclotron heating and no ECRH. It was fuelled at only $0.3 \cdot 10^{21} \text{s}^{-1}$ and had an ELM frequency between 45-55Hz, with only very few ELMs slightly exceeding 60Hz. Unlike the nitrogen-seeded discharge in tungsten, no sporadic spontaneous ELMs occurred immediately after a previous event. It can be seen in figure 5.34a and b that ELM-triggering was already successful 1.6ms after the previous ELM crash, much quicker than the fastest occurring natural

ELM. However, the power loads induced by triggered ELMs were not necessarily smaller than in the case of spontaneous ones (first panel). When analysing additional pellet-induced ELMs no clear indication of lower energy losses can be observed either. After an ELM is triggered the density reverts to the same level as before the arrival of the pellet: if the density is still low from the previous ELM, after a short increase it will return to this level. However, if the density has recovered since the previous event, the triggered ELM has little effect on it (right panel). In some cases, the density after the triggered ELM is even higher than before (not shown). This behaviour is consistent with the W-AUG and N-seeded cases shown earlier. The temperature, however, drops independent of the elapsed time. Here, the drops in the T_{rad} time trace are real as the ECE diagnostic was not affected by high density cut off after the pellet injection.




Figure 5.35: Profiles at different time points in the ELM cycle in C-AUG of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)

The corresponding kinetic profiles are presented in figure 5.35. Again, the black profiles were taken prior to an ELM and the red ones were ELM-synchronised between 1.5 and 3ms after the crash. Since the edge CXRS system was not yet installed when the PFCs were still partly made of uncoated graphite, the assumption $T_i = T_e$ was made. While the ELM-induced drops in the temperature are much lower than in the W-coated reference case, the pedestal density 2ms after a spontaneous ELM crash is still much lower than the pre-ELM value. The pressure profile is not yet recovered and the local edge current density still needs to increase by roughly 25% before it reaches its pre-ELM value. However, as it can be seen from the position of the operational point in figure 5.36, the flux-surface averaged toroidal current density is similar just before and 2ms after the ELM crash. A closer inspection of the experimental inputs to the equilibrium reconstruction and re-running CLISTE using different parameters did not change this trend in the flux surface averaged current density, which could be due to a change in the Z_{eff} profile after the ELM. However, in discharges in which $\langle j_{tor} \rangle$ drops after the ELM crash such an ELM-induced change in Z_{eff} should also occur. The stability calculations shown in figure 5.36 were run with $\psi_b = 0.994$, while the black curve is the boundary at $\Psi_p = 0.990$. The pre-ELM equilibrium is 10% in the unstable region,




Figure 5.36: *Stability diagrams of a plasma in C-AUG before a spontaneous type-I ELM (left) and 2ms after the ELM crash (right)*

but it should be kept in mind that the assumption $T_i = T_e$ was made, implying that the actual pressure gradient might be slightly lower. Additionally, the toroidal mode numbers are quite high, which means that the diamagnetic drift stabilisation effect will be relatively strong. The post-ELM equilibrium is in the stable region, but very close to the stability boundary. A small increase of the pressure gradient or the edge current density would be sufficient to destabilise it.

To conclude, ELM triggering via pellets is possible directly after the previous ELM in a carbon machine, while it is only successful after a certain lag time in full metal machines. However, this lag time can be shortened by seeding nitrogen in the divertor. The results presented in this section indicate that ELM triggering is only possible if the plasma is close to the stability boundary or at least to the pre-ELM values of the equilibrium, if the latter is not unstable. A pressure gradient of roughly 20% under the pre-ELM value is sufficient for successful ELM triggering, but the exact value remains to be determined. This is consistent with results from DIII-D, where the pressure gradient drops by 20% during the ELM crash, at which point ELM triggering is already successful [85]. Therefore, the recovery phase after the ELM is the decisive factor in determining whether triggering is possible or not. The smaller ELM crash and the faster recovery of the pressure profile when seeding nitrogen explains the shorter lag time. While it would be interesting to perform a detailed analysis of the ELM cycle to determine the exact shape of the critical kinetic profiles necessary for ELM triggering, the data quality of the analysed discharges is not sufficient. However, further experiments aiming for a comparison of the threshold for ELM triggering in unseeded and N-seeded discharges in W-AUG will be performed in the next experimental campaign. While it is not possible to repeat discharges in C-AUG, it may be possible to simulate carbon walls by seeding methane.

5.7 Influence of edge ECRH heating on peeling-ballooning stability

Experiments carried out at AUG showed that it is possible to pace ELMs by applying modulated ECRH power to the plasma edge [88]. However, this was only possible if the modulation frequency was close to the natural ELM frequency. More recently, it

was found at the TCV tokamak that depositing continuous wave (CW) ECRH power at the edge of a plasma can increase the natural ELM frequency, while lowering the expelled per-ELM energy [89]. ELM pacing using modulated ECRH power was also achieved. The interpretation of these results is based on a 0D power integrator model in which, after an ELM crash, the pedestal energy accumulates depending on the amount of ECRH power and heating location, and an ELM is triggered when a certain energy threshold is reached. However, a detailed analysis of the evolution of kinetic profiles was not possible due to technical limitations of the experimental setup.

In the framework of this thesis a scenario was designed for AUG discharges in which the deposition location of ECRH heating was gradually shifted from the core to the edge to reproduce the experiments carried out at TCV. Several similar discharges were performed to achieve this. First, test discharges were carried out with modulated ECRH power to verify the ECRH power deposition location as predicted by TORBEAM. The response of the ECE channels at the corresponding radius confirmed the predicted deposition location within a few cm.

All discharges were carried out at 800kA, -2.5T, 2.5MW NBI and a fuelling level of $4.5 \cdot 10^{21} \text{s}^{-1}$. In the reference discharge with 2.37MW central ECRH heating the ELM frequency lay between 40 and 80Hz, but a few sporadic ELMs occurred already 6-8ms after the previous event. Figure 5.37 shows that the discharge behaves like a




Figure 5.37: *Type-I ELMs with central ECRH: time traces of the divertor current from shunt measurements (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature in the pedestal (red)*

typical type-I ELMy H-mode in AUG. The ELMs themselves have the typical length of ELMs occurring in the full metal AUG, which indicates the presence of the two phases discussed in section 5.5 [83]. The ELM-induced energy losses are approximately 10%, while the drop in edge density and in pedestal top temperature are higher.

After gathering the reference data, a final discharge was performed in which the ECRH power was gradually shifted toward the edge. However, a minimum amount of central ECRH heating was used at all times to prevent impurity accumulation. Figure 5.38 shows the 2-D ECRH power deposition location (top) in the reference discharge and






Figure 5.38: Poloidal cross section of AUG showing the path and deposition location of the ECRH beams (top) and deposited power density (bottom) in the four considered phases

the three phases of the edge heated one, as well as the deposited power density as a function of the normalised radius (bottom). Note that while it appears that edge heating deposits less power, this is because the power density per unit volume is plotted and the plasma volume per flux surface is much larger at the edge. In the last phase of the discharge, 1.6MW of power was deposited around $\rho_p = 0.90$.

The plasma can have different responses to edge ECRH heating. In a test discharge performed several weeks before the actual experiment presented here, edge ECRH increased the ELM frequency as observed at TCV. However, in the subsequent discharges analysed the ELM frequency, formerly between 40 and 80Hz, did not simply increase. Rather, the few sporadic ELMs already occurring after 6-8ms after the previous crash already present in the reference discharge grow in number to form a second ELM frequency band between 100 and 130Hz (see figure 5.39). As the ECRH heating was shifted toward the edge, this second band became more populated.

Figure 5.40 illustrates this trend. While the ELM frequency is still low when mid-radius heating is applied (left), it increases when the power from one gyrotron is shifted to the edge. However, one ELM seems to be 'missing', which results in the two frequency bands already mentioned. In the last frame, which shows a time interval in which 1.6MW of ECRH heating was deposited around $\rho_p = 0.9$, only ELMs from the high

Figure 5.39: *ELM frequency dependence of the ECRH power deposition location*Figure 5.40: *Type-I ELMs at different ECRH deposition locations: time traces of the divertor current from shunt measurements (black), the power to the inner divertor (red), the total amount of energy confined in the plasma (green), the line integrated edge electron density (blue) and the radiation temperature in the pedestal (red)*

frequency band can be seen. While these short intervals were specifically chosen to illustrate the two frequency bands, both bands are still partly populated in the mid-radius and in the edge heating cases. However, the trend toward increasing population of the higher band with edge deposition is very clear (see figure 5.39). The ELMs from the upper frequency band typically induce less power loads to the divertor and cause lower energy losses, while the particle losses are comparable. The crash of the pedestal temperature is considerably lower. The W_{MHD} signals in figures 5.37 and 5.40 show that the stored energy significantly decreases as heating is shifted toward the edge. This is mainly due to the core density becoming peaked in the presence of central ECRH and then flattening when part of the power is removed [90]. This effect also propagates to the edge (see edge density traces in figures 5.37 and 5.40).

The kinetic profiles corresponding to three phases with different ECRH power deposition locations are shown in figure 5.41. The phase with mid-radius heating at $\rho_p = 0.4$ and 0.6 was omitted for legibility, but its profiles are very similar to the core-heated case (black). The T_i profiles from discharge 29295 are only shown up to $\rho_p = 0.94$ because the data from the core CXRS system was not available at the time of the analysis. However, since the pressure profiles were only constrained in CLISTE up to




Figure 5.41: *Pre-ELM profiles at different ECRH deposition locations of the electron and ion temperature (a), the electron density (b) the total pressure (c), its gradient (d), the toroidal current density at the LFS (e) and the q -profile (f)*

the pedestal top, this is not cause for concern. The pedestal top values of all kinetic profiles, T_e , T_i and n_e , drop when shifting the ECRH deposition location to the edge. While the pressure gradient reaches 430 kPa m^{-1} in the core heated phase, it drops by 30% when moving the power deposition. The edge current density peak shows the same trend. With edge ECRH heating ELMs are triggered faster than with core heating, even though the pressure gradient and current density are significantly lower.

The stability diagrams shown in figure 5.42 were generated using the profiles in figure 5.41 and with $\Psi_b = 0.994$ (again, the dotted black lines represent the border between the unstable and marginally stable region for $\Psi_b = 0.990$). As observed in the profiles in figure 5.41, the operational point moves toward lower pressure gradient and current density values as the heating power is shifted toward the edge. The stability boundary, however, does not move. While with core heating only the operational point is close to the marginally unstable region, the edge heated cases are stable. The decrease of the pressure gradient and the increasing population of the higher ELM frequency band with edge ECRH cannot be explained by ideal MHD. However, the ECRH power deposited inside the pedestal top could drive turbulent instabilities that prevent the temperature and density from fully recovering.

Since the loss in W_{MHD} is mainly due to lower density peaking because of the loss of core-deposited ECRH power a similar experiment was performed in which the core ECRH power was kept constant but edge ECRH was added later in the discharge. The ELM frequency increased from 25Hz-90Hz initially to 70Hz-200Hz in the phase with additional edge ECRH. The W_{MHD} dropped by 5% and the line integrated core density by 2% in the second phase, while no change in the ELM losses or the edge density was observed. While the T_e pedestal top value slightly increased with edge heating, T_i decreased by 10%. This also supports the hypothesis that ECRH might drive turbulent transport at the pedestal top and indicates that it mainly affects the ions. No stability analysis could yet be performed for this discharge.

As described in section 3.2, it is possible to drive a current with ECCD using the mi-




Figure 5.42: *Stability diagrams of pre-ELM equilibria at different ECRH deposition locations*

crowave heating system. This feature was used in the design of experiments aimed at driving an additional current at the edge of the plasma at the same location as the bootstrap current, thereby influencing the pedestal stability. This would have been the first direct experimental evidence for the role of the edge current density in the ELM crash. However, since the efficiency of ECCD current drive strongly depends on the local electron temperature, it is very difficult to increase the edge current density enough to make a significant difference. Furthermore, achieving the correct localisation is challenging. The ECRH deposition location, current drive efficiency, and the corresponding ECRH mirror angles were calculated using TORBEAM. When the angles are set such that the main deposition is near the bootstrap maximum around $\rho_p = 0.98$, the absorption of the wave strongly decreases and reflections occur because the incident wave is almost tangential to the steep density pedestal. It is not safe to perform experiments at such angles, since the reflections of microwave radiation could damage the machine. Figure 3.3 shows a metal plate melted due to ECRH radiation. Though this was caused by an incorrect polarisation of the wave, reflections would have similar consequences and could also damage diagnostics. Additionally, during an ELM crash, the outer separatrix jumps about 1cm inwards while the X-point moves upwards. Therefore, a safety margin must be included when calculating the ECRH mirror angles, such that the wave does not miss the plasma and directly hit vessel components during an ELM. ECRH deposition is possible and allowed in AUG up to a maximum radius of $\rho_p = 0.90$. Two discharges were carried out with opposite toroidal

angles of the ECRH beam to drive current in the co- and the counter directions. The power deposition was kept the same in order to disentangle heating from current drive effects. No change in the ELM frequency was observed between the two discharges. However, the ELM frequency was very scattered between 70 and 200Hz. As such, a small change would not be visible. The power and current deposition profiles can be




Figure 5.43: Pre-ELM profiles with co- (blue) and counter- (red) edge ECCD of the deposited power (a) and deposited current density (b), the electron and ion temperature (c), the electron density (d) the total pressure (e), its gradient (f), the toroidal current density at the LFS (g) and the q -profile (h)

seen in figure 5.43. Because of the low temperature at the plasma edge, TORBEAM calculations indicate that a total current of only 3.3kA was driven around $\rho_p = 0.9$. All profiles, including the edge current density are the same within the measurement errors and the CLISTE confidence bands. Therefore, no stability analysis is shown. In summary, not all results from the TCV experiments could be reproduced. The ELM frequency did not increase linearly with the amount of edge heating power deposited at the edge. Rather, a second, higher ELM frequency band became more populated. However, the trend toward lower ELM energy losses with increasing ELM frequency was observed, but the total W_{MHD} decreased. From the point of view of peeling-ballooning theory, there is no reason why edge ECRH heating should lead to an earlier ELM crash. As the ECRH power deposition is shifted outwards, the operational point moves deeper into the stable region of the $j - \alpha$ diagram while the stability boundary stays the same. Edge ECRH seems to drive turbulence that mainly affects the ions and limits the pedestal pressure. The attempt to directly test the peeling-ballooning theory by increasing the edge current density via ECCD was not successful because of the low current drive efficiency in the relatively cold plasma edge, the difficulties in depositing the current in the same region as the bootstrap current, and the extreme scatter in the ELM frequency when applying ECRH to the plasma edge.

5.8 Discussion

In this chapter, several different types of ELMs were analysed and used to test the widely accepted peeling-ballooning model. One type-I ELM case was presented that agrees very well with the theory since the operational point was on the stability boundary, but in a second analysed type-I ELMy H-mode, the occurrence of ELMs is incon-

sistent with this theory. Around the middle of the ELM cycle, the maximum edge current density and pressure gradient stop increasing, but the pedestal still widens, which, at the same pressure gradient, increases the pressure at the pedestal top. This pedestal broadening destabilises the plasma and the stability boundary moves toward the operational point. This is also consistent with the general picture described by the peeling-ballooning model [69]. However, at the onset of the ELM crash the pressure gradient is still roughly 30% under the stability boundary. In this discharge type-I ELMs cannot be explained by the peeling-ballooning theory.

Type-III ELMs are believed to be driven by resistive interchange instabilities and, if so, cannot be described using ideal MHD [73]. The results presented in section 5.2 are consistent with this hypothesis, as the pre-type-III equilibrium discussed is very stable. The phenomenology of the type-III ELMs analysed also agrees with such an explanation. They appear as the density increases and it was shown that a higher density leads to a lower shear in the edge radial electric field [74]. Within the picture, in the type-I ELM phase this shear is still high enough to stabilise the resistive interchange instabilities, but under a certain threshold the stabilising effect becomes insufficient and type-III ELMs occur at pedestal values which are under the type-I threshold.

The type-II ELM equilibrium presented in section 5.3 is much more stable than the type-I reference case, which is within 10% of the stability boundary. This is partly due to the lower pressure gradient and edge current density in the type-II phase, and partly to the stability boundary moving toward higher α values as a consequence of the stronger shaping and closeness to DN configuration. One can speculate that microtearing modes that are present slightly inside the pedestal top region cause the broad band fluctuations observed in T_e and prevent the pedestal top electron temperature from reaching the same values as in the type-I phase, thereby suppressing the type-I ELM crash.

It was observed that type-I ELMs disappear when increasing the density of the plasma above a certain threshold. Type-II like fluctuations occur in the ELM-suppressed phase, which could constitute the dominant transport mechanism in the pedestal. However, the equilibrium in this phase is closer to type-III than to type-II ELMs, as the dominant mode numbers are rather low and the edge current fraction to the total current is relatively high. The radial extent of the most unstable peeling-ballooning mode is similar in the ELM-suppressed phase and in the type-I ELM reference. This contradicts a possible correlation between the width of the linear MHD mode and the ELM size.

It is possible to improve the confinement of an H-mode plasma by seeding nitrogen into the divertor. A discharge in which a dramatic confinement improvement of 40% was achieved was discussed in section 5.5. The ELMs in nitrogen-seeded discharges differ from non-seeded ELMs in that they do not feature the two phases that are typical for type-I ELMs occurring in a full metal machine. The N-ELM crash is very similar to the first phase of a non-seeded ELM [83]. The radial extent of the dominant modes is the same in the seeded and non-seeded phases, consistent with this result. As the ELM itself is a nonlinear phenomenon, the second phase of non-seeded ELMs cannot be described using linear MHD, so the lower ELM losses cannot be explained with the codes used in this thesis. Typically, $\langle \mathbf{j} \cdot \mathbf{B} \rangle$ is lower in nitrogen seeded discharges despite the higher pressure gradient [15]. However, in the discharge discussed in section 5.5 the much steeper pressure gradient that accompanies the improved confinement in

the seeded phase drives a similar $\langle \mathbf{j} \cdot \mathbf{B} \rangle$ and a slightly higher $\langle j_{tor} \rangle$ than in the non-seeded reference phase. Because of the higher pressure gradient and current density values, the plasma becomes much more unstable against peeling-ballooning modes. Even though the diamagnetic drift stabilises the highest mode numbers, $n=3-8$ are still unstable when including this stabilising effect. While this is the only case presented in this section in which the equilibrium is well above the stability threshold for peeling-ballooning modes, it does question the general validity of the peeling-ballooning model. In the framework of ideal MHD, an ELM should be triggered before the pedestal pressure gradient and current density can reach such high values.

The ability to trigger ELMs using frozen deuterium pellets differs between N-seeded and non-seeded discharges in a full metal machine [86]. The lag time which must elapse after an ELM crash to achieve successful triggering is much lower when seeding nitrogen. In a carbon machine, it is possible to trigger an ELM immediately after the previous one. Different time points of an ELM cycle in each of these three scenarios were analysed in section 5.6. While the N-seeded pre-ELM equilibrium is stable against peeling-ballooning modes, the non-seeded one and the one from C-AUG are (marginally) unstable. However, in all three cases a pressure gradient 20% below the pre-ELM gradient is sufficient for a pellet to trigger an ELM. When the edge pressure gradient is only at half the value of the pre-ELM gradient, however, ELM-triggering is not successful. Therefore, the amplitude of the ELM crash and the recovery phase of the pedestal profiles are the decisive factors in determining whether triggering is possible or not. As those two factors are different in C-AUG ELMs, W-AUG ELMs and N-seeded ELMs, the lag time necessary before triggering differs depending on the scenario. While it would be interesting to analyse the recovery of the kinetic profiles in more detail to determine the exact triggering threshold, the quality of the data was not sufficient in the discharges presented in this thesis. However, it is planned to repeat these discharges in the next experimental campaign of AUG with a stronger focus on edge diagnostics.

ELMs can also be influenced by depositing localised ECRH power at the plasma edge. It was shown at TCV that the ELM frequency increases and the per-ELM losses decrease when shifting the power deposition to the edge [89]. In the experiments discussed in section 5.7, however, a different behaviour was observed. Rather than the ELM frequency increasing, a second, higher frequency band formed and became more populated as more power was moved to the edge. A peeling-ballooning analysis of the profiles showed that the core-heated reference equilibrium is close to the stability boundary, while the operational point moves deeper into the stable region when shifting the power to the edge. This is mainly due to a relaxation of the kinetic profiles while the stability boundary stays the same. One can speculate that ECRH drives turbulence at the deposition location, thereby enhancing transport and limiting the pedestal top and gradients, thereby preventing the plasma from reaching the peeling-ballooning stability limit. Since the relaxation of the gradients does not suppress the ELMs but, on the contrary, increases their frequency, this is another indication that peeling-ballooning theory is missing a key ingredient to explain the ELM trigger.

6 Implications for ITER

The next big step in tokamak research is the commissioning of ITER. With a plasma volume of 840m^3 , a toroidal magnetic field of 5.3T and a 15MA plasma current, it is expected to produce 500MW of fusion power in D-T operation, 10 times more power than is needed from auxiliary heating systems to sustain the plasma [8]. Recent extrapolations applying a linear regression on the performance parameters of present-day mid-sized and large devices provide predictions for the ITER pedestal [83]. The analysis presents predictions for the temperature pedestal height and width and the density pedestal width. Since no scaling for the density pedestal height is available, it was assumed to be $7 \cdot 10^{19}\text{m}^{-3}$, as it was in previous predictions of the ITER pedestal [91]. In this section these extrapolations are used to calculate the neoclassical current density profile from the formula derived by Sauter *et al.* [27, 28] in order to generate self-consistent equilibria. A peeling-ballooning analysis is then presented.




Figure 6.1: Pre type-I ELM profiles of the electron and ion temperature (a), the electron density (b), the total pressure (c), its gradient (d), the flux surface averaged $\mathbf{j} \cdot \mathbf{B}$ (e) and the q -profile (f) of an AUG discharge.

The validity of the bootstrap current formula from Sauter *et al.* was confirmed for the edge of AUG discharges [25]. Figure 6.1 shows the profiles of the first discharge discussed in section 5.1. However, instead of the local current density at the LFS, $\langle \mathbf{j} \cdot \mathbf{B} \rangle$ is shown in panel e, which can be compared with the result from the bootstrap current calculations more easily because it does not contain the Pfirsch-Schlüter current contribution. The result from the neoclassical calculation (red curve) matches the CLISTE result extremely well at the position of the maximal edge current density.

According to the scaling from Schneider *et al.* [83], at 74MW of total heating power the ITER pedestal top electron temperature could reach 4.7keV , the electron pressure roughly 50kPa and the ion pressure 73kPa or 25kPa , depending on whether the DIII-D data is included or not. The analysis presented in the following uses the predicted pedestal top electron temperature of 4.7keV . Given the large uncertainties in the ion pressure prediction, the ion temperature is assumed to be equal to T_e . The pedestal top electron density is set to $7 \cdot 10^{19}\text{m}^{-3}$ and Z_{eff} to 1.2 with boron as the main impurity species. In the prediction the temperature pedestal width varies between 0.03 and 0.043, depending on which scaling is used, while the density pedestal width

is between 0.012 and 0.014 Ψ_n [83]. Instead of only using these values, a large range of temperature and density pedestal widths was scanned independently in the present analysis. However, the pedestal top values were kept constant in all cases. Realistic synthetic profiles were modelled after the shape of experimental AUG profiles and are shown in figure 6.2. While in the left frames the data is plotted over the whole radius,




Figure 6.2: *Synthetic ITER electron density (top) and temperature (bottom) profiles. The magnified pedestal is shown in the middle frames for clarity. The ITER boundary is shown in panel (e).*

the pedestal is magnified in the middle ones to point out the different widths. An almost flat core electron density profile and a slightly peaked temperature profile were chosen, as is the case in typical AUG discharges. Both the temperature and electron density widths were scanned from 0.008 to 0.17 in terms of Ψ_n . Self-consistent equilibria were generated using a combination of the temperature and density profiles shown in figure 6.2 to build a 2D $\Delta T - \Delta n$ grid for the stability calculations. Nine selected input profiles used in HELENA can be seen in figure 6.3. The blue curves are pressure input profiles for HELENA, calculated from all possible combinations of the three temperature and the three density profiles shown in red in figure 6.2. The resulting flux-surface-averaged neoclassical current density, a combination of the bootstrap and the Ohmic current, can be seen in black. In each case the Ohmic current density was scaled by changing the loop voltage in order to match the total of 15MA plasma current foreseen in ITER. Since both the density and the temperature gradients drive a part of the bootstrap current, the current density can exhibit a double peak when the T and n pedestal widths are too different (see figure 6.3f). To smooth out the profiles, two Gaussian curves were fitted to the bootstrap current profile, one to the left and one to the right side of the maximum, before adding the result to the Ohmic current density. Finally, a soft smoothing filter was applied. The resulting current density profiles that were used as an input to HELENA are shown in red. The equilibria were then generated using the realistic ITER plasma boundary in figure 6.2e.

Figure 6.4a shows the stability diagrams generated from the self-consistent equilibria. Instead of α and $\langle j_{tor} \rangle$, the T and n_e pedestal widths are used as the x and y axis, respectively. Note that the axes are logarithmic. The operational point predicted by




Figure 6.3: Selected ITER pressure and current density profiles consistent with the red profiles in figure 6.2

Schneider *et al.*, represented by a star, is far in the unstable region. The minimum value of ΔT_{ped} that is still stable is 0.07 in terms of Ψ_n , given $\Delta n_{ped} = 0.012$. In this region of the diagram the stability of the plasma is almost independent of the density pedestal width. It is only at much higher values of Δn_{ped} that the broadening of the density pedestal starts to become stabilising. With $\Delta T_{ped} = 0.04$, Δn_{ped} must be 0.1 for the equilibrium to be marginally stable. Figure 6.4b shows the same stability diagram in $j - \alpha$ space. Since self-consistent bootstrap current density and kinetic pressure profiles were used, and the bootstrap current is roughly proportional to α , only the diagonal is shown. The grid points in some regions are very dense because the same values for j and α can be achieved with different combinations of the temperature and density profiles. The y axis corresponds to the maximum current density outside of $\Psi_n = 0.8$. In the cases with very low bootstrap contribution like the one shown in figure 6.3g this corresponds to the Ohmic current at $\Psi_n = 0.8$, which explains why so many points are at $\langle j_{tor} \rangle = 0.6 \text{ MA m}^{-2}$. The pressure gradient of the predicted ITER pedestal is far above the stability threshold. The channel that leads to the stable island at $\alpha = 5$ and $\langle j_{tor} \rangle = 1.2 \text{ MA m}^{-2}$ corresponds to the stable region at very high Δn_{ped} in figure 6.4a. Such a wide density pedestal is thought to be unlikely because in present-day devices, the density pedestal is narrower than the temperature pedestal [83].

One of the main results of the scaling presented in reference [83] is the different width of the ITER n_e and T_e pedestal. Therefore, the stability calculations were repeated using the predicted ITER pedestal widths of 0.04 (T) and 0.012 (n_e) in terms of Ψ_n but varying the pedestal top value. The corresponding edge electron density and temperature profiles can be seen in figures 6.5a and b. The electron density pedestal top was scanned from $2.6 \cdot 10^{19} \text{ m}^{-3}$ to $14.6 \cdot 10^{19} \text{ m}^{-3}$, the temperature from 1keV to 7keV. Panels c through k show the edge current density profiles (black) calculated from the red T and n_e profiles and the fits that were used as inputs to HELENA (red), as well as the pressure profiles (blue). The edge bootstrap contribution to the total




Figure 6.4: *Stability diagram of the ITER width scan as a function of T_{ped} and n_{ped} (left), and as a function of $\langle j_{tor} \rangle$ and α*




Figure 6.5: *ITER height scan*

current ranges from 1% (panel g) to 25% (panel e).

The axes on the stability diagram shown in figure 6.6a are the pedestal top temperature and electron density. The ITER prediction from Schneider *et al.*, represented by the yellow star, is in the unstable region. The stability analysis suggests that the ITER prediction is too optimistic. If the hypothesis of fixed pedestal density and temperature widths of 0.04 and 0.012 is true for ITER and if peeling-ballooning modes limit the pedestal, the pedestal top density will only reach $4 \cdot 10^{19} \text{ m}^{-3}$ at $T_{ped} = 4.7$. At lower pedestal temperatures the density could be higher. However, the assumption $T_i = T_e$ was made and the edge T_i gradient is typically lower than ∇T_e . In AUG, using experimental T_i profiles instead of assuming it equal to the electron temperature can lower the pressure gradient up to 20%. Figure 6.6b shows the stability of the ITER height scan as a function of α and $\langle j_{tor} \rangle$. The pressure gradient of the predicted ITER pedestal is almost a factor of two above the stability threshold. This means that a broader ion temperature profile that leads to a 20% lower pressure gradient is not sufficient to stabilise the plasma.

The total pedestal top pressure resulting from $n_{ped} = 7 \cdot 10^{19} \text{ m}^{-3}$ and $T_{ped} = 4.7 \text{ keV}$ is 90kPa. The pressure pedestal width, $\Delta p = 0.04$ is set by the temperature profile. The marginally stable grid point at 4keV and $4.6 \cdot 10^{19} \text{ m}^{-3}$ in figure 6.6 has a pedestal top




Figure 6.6: *Stability diagram of the ITER height scan as a function of ΔT_{ped} and n_{ped} (left), and as a function of $\langle j_{tor} \rangle$ and α*

pressure of 60kPa, while the one at 3keV and a density of $6.6 \cdot 10^{19} \text{m}^{-3}$ has 66kePa. Both values are lower than the one predicted by Snyder *et al.*, $p_{ped} = 92\text{kPa}$, with the same pressure pedestal width of 0.04 in terms of Ψ_n [69]. However, they agree with the most pessimistic pedestal top pressure of 66kPa extrapolated from present-day devices, which was obtained when including data from DIII-D in the regression [83]. Since the triangularity of the plasma has a stabilising effect on peeling-ballooning modes, a higher pedestal top pressure might be achieved via stronger shaping of the plasma boundary.

7 Conclusions and outlook

In this thesis, different types of ELMy H-modes were analysed, compared to each other and used to test the peeling-ballooning model: the most commonly invoked theory to explain the occurrence of ELMs. The stability chain, a sequence of four steps, was used to determine whether the experimental plasmas were stable against peeling-ballooning modes, or if specific unstable ideal modes could be found. First, the data acquired during plasma discharges from temperature, density and magnetic diagnostics were combined to generate self-consistent equilibria using CLISTE. These equilibria were then refined with HELENA to achieve a higher degree of numerical accuracy. The third step was to create a $j - \alpha$ grid by independently varying the edge current density height and the pressure pedestal width by using the j-alpha workflow. Finally, the ILSA code was run to determine the stability of the plasma, the dominant toroidal mode number and the position of the stability boundary in $j - \alpha$ space.

The limits of the stability chain were explored and inconsistencies between ILSA and the standalone version of MISHKA-1, which arise when the q-profile is very steep, were found. These inconsistencies are caused by the different interpolation methods used in MISHKA-1 and ILSA and they have a non-negligible influence on the results of the stability boundary in the low current density region of the $j - \alpha$ diagrams. The influence of Ψ_b , the cutoff parameter which determines at which flux surface the equilibrium is cut off such that it is possible to describe it in field-aligned coordinates, was explored. As cutting off further outside has a stabilising effect and no satisfactory reason was found to favour one value of Ψ_b over another, many of the stability diagrams presented in this thesis were generated twice, using $\Psi_b = 0.990$ and 0.994 . The difference in the calculated stability boundaries typically lay on the order of 10%, which is within the measurement errors of the pressure gradient. Furthermore, it was found that systematic measurement errors have only a small influence on the location of the stability boundary. However, such errors have to be kept as low as possible since they will distort the position of the operational point with respect to the boundary. It was also shown that the correct alignment between the magnetic and kinetic measurements is crucial because a shift strongly influences the growth rate spectra of the peeling-ballooning modes.

In the main part of this thesis, the peeling-ballooning model was applied to different types of ELMy H-mode plasmas. Discharges in which ideal MHD should be valid, namely type-I ELMy H-modes, were compared to discharges featuring ELM-types that are believed to be stable against peeling-ballooning modes, type-II and type-III ELMs, as well as an ELM-suppressed regime at high densities. Furthermore, other type-I ELMy scenarios, including discharges with nitrogen seeding, pellet ELM triggering and with ECRH heating at the plasma edge were tested against peeling-ballooning theory.

The results from the type-II and type-III ELMy H-modes were partly in agreement with previous findings. Both of them were stable against peeling-ballooning modes, which is expected as they are believed to be driven by other instabilities. The appearance of type-III ELMs at high densities is consistent with the hypothesis that they are driven by resistive interchange modes which are stabilised by a strong shear in the radial electric field, and this shear decreases with increasing density. The type-II H-mode was very stable against peeling-ballooning instabilities, because the stability boundary

moved to a region of higher pressure gradient and the experimental pedestal pressure gradient relaxed compared to the type-I reference phase. It appears that a different transport mechanism than a peeling-ballooning mode limits the pedestal top pressure such that the threshold for type-I ELMs cannot be reached. This is consistent with the broad-band fluctuations that are observed in the electron temperature inside the pedestal top, which indicates turbulence in this region.

Another regime in which type-I ELMs are suppressed was analysed. This suppression occurs above a certain density threshold. While the plasma edge was dominated by type-II like fluctuations, the equilibrium was closer to the stability boundary than is observed in type-II ELMy plasmas and the fraction of edge current to the total plasma current was higher. The toroidal mode numbers of the dominant instabilities calculated from the peeling-ballooning model were relatively low and the radial extent of the poloidal harmonics was similar to the type-I reference case. This contradicts a correlation between the ELM size and the mode extent. However, given the low temperature pedestal, Resistive effects could also be important with the dominant current driven modes. This would also account for the small size of the ELM crash, where the resistivity dampens significant changes in the current density. The small fluctuations would then correspond to a “soft” peeling limit, analogous to the soft ballooning limit described for type-II ELMs by Saarelma *et al.* [75]. A resistive MHD analysis and non-linear simulations are required to further analyse this hypothesis.

Several type-I ELMy H-modes, operational regimes in which ELMs are believed to be driven by peeling-ballooning modes, were also analysed. Two similar discharges, which were carried out several years apart, were compared to each other and to the model. While the pre-ELM equilibrium of the reference discharge was unstable, consistent with ideal MHD theory, the operational point of the repeated experiment was deep in the stable region. 15-20ms before the ELM crash, the pedestal gradients had already reached their final value, but the pedestal top continued to increase slowly because of the increasing pedestal width. As expected from the peeling-ballooning and the EPED models [69], the stability boundary moved toward lower pressure gradients as the pedestal expanded at constant gradient. However, the pre-ELM equilibrium was still 30% under the stability threshold. Reaching this threshold is also a necessary (and sufficient) condition for the occurrence of an ELM within the EPED model. While this result alone might not be sufficient to question the applicability of the peeling-ballooning theory, two additional pieces of evidence were presented in this thesis: an edge-heated H-mode in which the ELM frequency increased as the edge pedestal gradients decreased and a N-seeded case in which a pedestal exists which is deep inside the unstable region of the $j - \alpha$ diagram.

A type-I ELMy H-mode with central ECRH and NBI heating was discussed, that was, within error bars, on the stability boundary for peeling-ballooning modes. When the ECRH power was moved toward the edge of the plasma, the T_e , T_i and n_e pedestals relaxed, possibly due to turbulence caused by the ECRH power deposited near the pedestal top. As a consequence, the pressure gradient and the edge current density also decreased while the stability boundary stayed at the same location. One might expect that an operational point more than 30% below the stability boundary would cause a transition to an ELM-suppressed regime. However, the opposite was the case: a second, higher band of the ELM frequency formed and became more populated as the ECRH power deposition location was shifted to the plasma edge.

A third case was presented which could also not be explained by the peeling-ballooning model. When seeding nitrogen into the divertor during a type-I ELMy H-mode discharge, the confined energy increased by 40%. This increase was mainly due to a steepening of the pedestal gradients and an increase of the pedestal top temperature and density. The pre-ELM equilibrium in the unseeded phase was close to the stability boundary, slightly in the unstable region. When introducing nitrogen into the divertor, the plasma became highly peeling-ballooning unstable because the pressure gradient increased while the stability boundary remained unchanged. Including the stabilising effect of the diamagnetic drift stabilised the modes with high n numbers, but the lower $n=3-8$ modes remained highly unstable. This indicates that ideal peeling-ballooning theory was not the driving mechanism for the ELMs occurring in either of the two phases, or that an additional stabilising effect exists, which was much stronger in the N-seeded phase and shifted the stability boundary to higher values of j and α .

Another question that was addressed in this thesis is the feasibility of triggering ELMs via frozen deuterium pellets. In discharges performed when the walls of AUG were still made from carbon, pellets would trigger an ELM whenever they were injected into an H-mode plasma. Since the walls were coated with tungsten it is only possible to trigger them after a specific lag time has elapsed since the previous ELM. This lag time can be significantly reduced when injecting nitrogen into the divertor. This is due to two connected factors: the lower ELM losses, and the faster recovery from the ELM crash. The ELM-induced collapse of the pedestal in C-AUG discharges was so low that the perturbation or the density increase created by a pellet was sufficient to trigger an ELM. In the W-AUG and the N-seeded discharges analysed in this work the pedestal pressure profile just after the triggering limit was also close to the pre-ELM pedestal, and the lag time for successful triggering was close to the minimum inter-ELM time of spontaneous ELMs. As spontaneous ELMs can occur after such an interval, this suggests that it is the perturbation created by the pellet that is responsible for ELM triggering, rather than a local density increase that drives the pressure gradient above the stability boundary. While it is unclear if peeling-ballooning modes are the driving mechanism in the occurrence of type-I ELMs, the results presented here show that it is possible to trigger them via pellets when the pedestal pressure gradient and the edge current density are 20% below the typical threshold for spontaneous ELMs in the corresponding scenario. It is important to understand the nonlinear evolution of ELMs because the ELM losses and the recovery phase after the ELM determine the maximum frequency with which they can be triggered via pellets. The exact triggering threshold could not be determined due to the insufficient quality of the experimental data. However, it is planned to repeat the discharges in the next experimental campaign of AUG with a stronger focus on edge diagnostics.

In some of the discharges discussed in this thesis the peeling-ballooning model is consistent with the experimental results, while in others it fails to explain the occurrence of ELMs. A possible explanation of this could be a hypothesis discussed by Wilson *et al.*, which was dubbed “Circumstantial evidence” [58]. The authors propose that the peeling-ballooning mode could exist, but cannot grow because of the sheared flow in the plasma edge. In this case, the mode would limit the pedestal but not result in an ELM crash. The ELM itself would then be triggered by a non-MHD event. This would be consistent both with the AUG ELMs that occur when the equilibrium is well below the peeling-ballooning stability boundary and those that exist on or near it. However,

this hypothesis is highly speculative and no alternative is proposed to an ideal MHD mode as the trigger mechanism.

Recent extrapolations of the ITER pedestal were tested against the peeling-ballooning model [83]. The extrapolations attempt to predict the T_e and n_e pedestal width, as well as the T_e value at the pedestal top. The pedestal top electron density is assumed to be $7 \cdot 10^{19} \text{m}^{-3}$. In this thesis, two scans were performed on the basis of these extrapolations. In the first one, the n_e and T_e pedestal widths were scanned independently assuming a pedestal top temperature of 4.7keV and density of $7 \cdot 10^{19} \text{m}^{-3}$. In the second one, the predicted pedestal widths were used and the pedestal top values were scanned. For each combination of temperature and density profiles the bootstrap current was calculated and a self-consistent equilibrium was generated. A stability analysis was then performed using ILSA. It was found that if the peeling-ballooning stability represents an upper limit then the predicted values are very optimistic. To obtain a pedestal top temperature of 4.7keV at a density of $7 \cdot 10^{19} \text{m}^{-3}$, the density pedestal must be almost ten times wider than predicted. This is thought to be unlikely since in present-day devices the density pedestal is narrower than the temperature pedestal. Another possibility is a 75% wider temperature pedestal than predicted. A stable plasma at the predicted pedestal width can be achieved if the pedestal top temperature is only 3keV or both n_e and T_e have pedestal top values which are 30% lower than in the prediction. These cases would lead to a total pedestal top pressure of 60-66kPa, which is consistent with the worst case scenario presented in the extrapolations, but it is 30% lower than the results presented by Snyder *et al.* [69]. The effects of plasma shaping remain to be tested with this model.

Some of the results presented in this thesis question the validity of the peeling-ballooning theory. In many of the presented cases, ELMs occurred well below the threshold for peeling-ballooning theory and in one case well above. In general, ideal MHD should be more applicable to hotter plasmas, because the resistivity decreases. However, the hottest plasma, the high power nitrogen seeded case, was well above the threshold for ideal peeling-ballooning stability. While the number of discharges presented in this thesis is hardly sufficient for a definitive assertion, the results indicate that the peeling-ballooning model is insufficient to fully explain the triggering of ELM instabilities. Rather, another as yet unknown trigger mechanism is required.

Appendices

A Derivation of the MHD equations used in ILSA

In this appendix the derivation of the equations solved by the linear MHD codes CASTOR and MISHKA-1 is presented. It is mainly based on the ILSA manual written by E. Strumberger *et al.* [34], the description of the CASTOR code [35] and the description of MISHKA-1 [36]. The approximations used in the following are:

- A small perturbation
- A single-fluid plasma, which implies:
 - A net charge density of zero
 - A small displacement current compared to the conduction current
- Isotropic plasma pressure
- Scalar electric resistivity

Such a plasma can be described with the mass continuity equation

$$\frac{d\rho}{dt} + \rho \nabla \cdot \mathbf{v} = 0, \quad (\text{A.1})$$

the momentum equation

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla) \mathbf{v} = \mathbf{j} \times \mathbf{B} - \nabla p, \quad (\text{A.2})$$

the equation of state

$$\frac{d}{dt}(p\rho^{-\Gamma}) = 0, \quad (\text{A.3})$$

and Ohm's law

$$\mathbf{E} + \mathbf{v} \times \mathbf{B} = \eta \mathbf{j}, \quad (\text{A.4})$$

where ρ is the mass density, \mathbf{v} the fluid velocity, p the scalar pressure, \mathbf{j} the current density, \mathbf{E} the electric field, \mathbf{B} the magnetic field, η the scalar resistivity and Γ the polytrope index. The derivative d/dt can be expressed as

$$\frac{d}{dt} = \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla. \quad (\text{A.5})$$

In addition to these, the Maxwell equations are necessary to build a set of self-consistent equations describing a resistive MHD plasma. We ignore the displacement current $\epsilon_0 \partial \mathbf{E} / \partial t$, since the temporal variations are relatively slow, and we assume a zero charge density ρ_q . Therefore, the Maxwell equations simplify to Gauss' law

$$\nabla \cdot \mathbf{E} = 0 \quad (\text{because } \rho_q = 0), \quad (\text{A.6})$$

Gauss's law for magnetism

$$\nabla \cdot \mathbf{B} = 0, \quad (\text{A.7})$$

Faraday's law

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E}, \quad (\text{A.8})$$

and Ampère's law

$$\mu_0 \mathbf{j} = \nabla \times \mathbf{B} \quad (\text{because displacement current } \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \text{ is negligible}), \quad (\text{A.9})$$

where μ_0 is the vacuum permeability.

Equation A.3, written as

$$-\Gamma p \rho^{-\Gamma-1} \frac{d\rho}{dt} + \rho^{-\Gamma} \frac{dp}{dt} = 0, \quad (\text{A.10})$$

can be combined with equation A.1 to form

$$\Gamma p \rho^{-1} \rho \nabla \cdot \mathbf{v} + \frac{dp}{dt} = 0. \quad (\text{A.11})$$

When applying the Lagrangian derivative (equation A.5), one obtains

$$\frac{\partial p}{\partial t} + (\mathbf{v} \cdot \nabla)p + \Gamma p \nabla \cdot \mathbf{v} = 0. \quad (\text{A.12})$$

If we assume that the plasma behaves like an ideal gas, according to $p = \rho kT/m$, where m is the particle mass, k the Boltzmann constant and T the plasma temperature, equation A.12 transforms to

$$T \frac{\partial \rho}{\partial t} + \rho \frac{\partial T}{\partial t} + T(\mathbf{v} \cdot \nabla)\rho + \rho(\mathbf{v} \cdot \nabla)T + \Gamma \rho T \nabla \cdot \mathbf{v} = 0 \quad (\text{A.13})$$

and since $\partial \rho / \partial t = -\rho \nabla \cdot \mathbf{v} - (\mathbf{v} \cdot \nabla)\rho$ (eq. A.1 & A.5), it can be simplified to

$$\frac{\partial T}{\partial t} + (\mathbf{v} \cdot \nabla)T + (\Gamma - 1)T \nabla \cdot \mathbf{v} = 0. \quad (\text{A.14})$$

We, therefore, obtain the following closed set of four equations, describing the nonlinear evolution of ρ , T , \mathbf{v} and B :

$$\frac{\partial \rho}{\partial t} + \nabla(\rho \mathbf{v}) = 0 \quad (\text{eq. A.1 \& A.5}), \quad (\text{A.15})$$

$$\frac{\partial T}{\partial t} + (\mathbf{v} \cdot \nabla)T + (\Gamma - 1)T \nabla \cdot \mathbf{v} = 0 \quad (\text{eq. A.14}), \quad (\text{A.16})$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho(\mathbf{v} \cdot \nabla)\mathbf{v} = \frac{1}{\mu_0}(\nabla \times \mathbf{B}) \times \mathbf{B} - \nabla p \quad (\text{eq. A.2 \& A.9}), \quad (\text{A.17})$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \left(\mathbf{v} \times \mathbf{B} - \frac{\eta}{\mu_0} \nabla \times \mathbf{B} \right) \quad (\text{eq. A.8 with } E \text{ from eq. A.4 and } \mathbf{j} \text{ from eq. A.9}). \quad (\text{A.18})$$

The nonlinear description of the plasma by equations A.15 through A.18 can be linearised by expanding around an equilibrium state. Using this expansion, a quantity $\mathbf{C}(\mathbf{x}, t)$ can be written as

$$\mathbf{C}(\mathbf{x}, t) = \mathbf{C}_0(\mathbf{x}) + \text{Re} \{ e^{\lambda t} \mathbf{C}_1(\mathbf{x}) \} \quad (\text{A.19})$$

where the subscript 0 denotes the equilibrium quantity and 1 the first order perturbation. When introducing this approach to equation A.15, one obtains

$$\frac{\partial}{\partial t} \rho_0 + \frac{\partial}{\partial t} (e^{\lambda t} \rho_1) + \nabla \cdot ((\rho_0 + e^{\lambda t} \rho_1)(\mathbf{v}_0 + e^{\lambda t} \mathbf{v}_1)) = 0$$

$$\lambda e^{\lambda t} \rho_1 + \nabla(\rho_0 \mathbf{v}_0) + e^{\lambda t} \nabla(\rho_0 \mathbf{v}_1) + e^{\lambda t} \nabla(\rho_1 \mathbf{v}_0) + \cancel{e^{2\lambda t} \nabla(\rho_1 \mathbf{v}_1)} = 0.$$

$\partial \rho_0 / \partial t$ is zero and the crossed out term is very small since it contains the product of two perturbed quantities. In this section terms marked in blue only contain equilibrium quantities and are zero since the initial equations from which they are derived also have to be valid for the equilibrium state. The equation simplifies to

$$\lambda \rho_1 = -\rho_0 \nabla \cdot \mathbf{v}_1 - \mathbf{v}_1 \cdot \nabla \rho_0 - \rho_1 \nabla \cdot \mathbf{v}_0 - \mathbf{v}_0 \cdot \nabla \rho_1. \quad (\text{A.20})$$

The same ansatz on equation A.16 yields

$$\frac{\partial}{\partial t} (\cancel{T_0} + e^{\lambda t} T_1) + [(\mathbf{v}_0 + e^{\lambda t} \mathbf{v}_1) \cdot \nabla] (T_0 + e^{\lambda t} T_1) + (\Gamma - 1)(T_0 + e^{\lambda t} T_1) \nabla \cdot (\mathbf{v}_0 + e^{\lambda t} \mathbf{v}_1) = 0,$$

$$\lambda e^{\lambda t} T_1 + \mathbf{v}_0 \cdot \nabla T_0 + e^{\lambda t} \mathbf{v}_0 \cdot \nabla T_1 + e^{\lambda t} \mathbf{v}_1 \cdot \nabla T_0 + \cancel{e^{2\lambda t} \mathbf{v}_1 \cdot \nabla T_1} + (\Gamma - 1) \left[T_0 \nabla \cdot \mathbf{v}_0 + e^{\lambda t} T_0 \nabla \cdot \mathbf{v}_1 + e^{\lambda t} T_1 \nabla \cdot \mathbf{v}_0 + \cancel{e^{2\lambda t} T_1 \nabla \cdot \mathbf{v}_1} \right] = 0,$$

$$\lambda T_1 = -\mathbf{v}_0 \cdot \nabla T_1 - (\Gamma - 1) T_1 \nabla \cdot \mathbf{v}_0 - \mathbf{v}_1 \cdot \nabla T_0 - (\Gamma - 1) T_0 \nabla \cdot \mathbf{v}_1. \quad (\text{A.21})$$

Equation A.17 simplifies to

$$(\rho_0 + e^{\lambda t} \rho_1) \frac{\partial}{\partial t} (\cancel{\mathbf{v}_0} + e^{\lambda t} \mathbf{v}_1) + (\rho_0 + e^{\lambda t} \rho_1) [(\mathbf{v}_0 + e^{\lambda t} \mathbf{v}_1) \cdot \nabla] (\mathbf{v}_0 + e^{\lambda t} \mathbf{v}_1) = \frac{1}{\mu_0} [\nabla \times (\mathbf{B}_0 + e^{\lambda t} \mathbf{B}_1)] \times (\mathbf{B}_0 + e^{\lambda t} \mathbf{B}_1) - \frac{k}{m} \nabla(\rho_0 T_0 + e^{\lambda t} \rho_1 T_0 + e^{\lambda t} \rho_0 T_1)$$

$$\lambda \rho_0 e^{\lambda t} \mathbf{v}_1 + \rho_0 (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_0 + e^{\lambda t} \rho_1 (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_0 + e^{\lambda t} \rho_0 (\mathbf{v}_1 \cdot \nabla) \mathbf{v}_0 + e^{\lambda t} \rho_0 (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_1 = \frac{1}{\mu_0} (\nabla \times \mathbf{B}_0) \times \mathbf{B}_0 + e^{\lambda t} \frac{1}{\mu_0} (\nabla \times \mathbf{B}_0) \times \mathbf{B}_1 + e^{\lambda t} \frac{1}{\mu_0} (\nabla \times \mathbf{B}_1) \times \mathbf{B}_0 - \frac{k}{m} [\nabla \rho_0 T_0 - e^{\lambda t} \nabla(\rho_1 T_0 + \rho_0 T_1)].$$

In the second step, the second order terms have already been removed. After removing the (blue) equilibrium terms, which zero out according to equation A.17, the equation becomes

$$\lambda \rho_0 \mathbf{v}_1 = -\rho_1 (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_0 - \rho_0 [(\mathbf{v}_1 \cdot \nabla) \mathbf{v}_0 + (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_1] + \frac{1}{\mu_0} [(\nabla \times \mathbf{B}_0) \times \mathbf{B}_1 + (\nabla \times \mathbf{B}_1) \times \mathbf{B}_0] + \nabla(\rho_1 T_0 + \rho_0 T_1). \quad (\text{A.22})$$

The last one, equation A.18 becomes

$$\begin{aligned} \frac{\partial}{\partial t}(\mathbf{B}_0 + e^{\lambda t} \mathbf{B}_1) &= \nabla \times \left((\mathbf{v}_0 + e^{\lambda t} \mathbf{v}_1) \times (\mathbf{B}_0 + e^{\lambda t} \mathbf{B}_1) - \frac{\eta}{\mu_0} \nabla \times (\mathbf{B}_0 + e^{\lambda t} \mathbf{B}_1) \right) \\ \lambda \mathbf{B}_1 &= \nabla \times \left(\mathbf{v}_0 \times \mathbf{B}_1 + \mathbf{v}_1 \times \mathbf{B}_0 + \cancel{e^{\lambda t} \mathbf{v}_1 \times \mathbf{B}_1} - \frac{\eta}{\mu_0} \nabla \times \mathbf{B}_1 \right) \end{aligned}$$

Using the vector identity $\nabla \times (\mathbf{A} \times \mathbf{B}) = \mathbf{A}(\nabla \cdot \mathbf{B}) - \mathbf{B}(\nabla \cdot \mathbf{A}) + (\mathbf{B} \cdot \nabla) \mathbf{A} - (\mathbf{A} \cdot \nabla) \mathbf{B}$, and because $\nabla \cdot \mathbf{B}_1 = 0$ (equation A.7) we get

$$\lambda \mathbf{B}_1 = -(\mathbf{v}_0 \cdot \nabla) \mathbf{B}_1 - \mathbf{B}_1(\nabla \cdot \mathbf{v}_0) + (\mathbf{B}_1 \cdot \nabla) \mathbf{v}_0 + \nabla \times \left(\mathbf{v}_1 \times \mathbf{B}_0 - \frac{\eta}{\mu_0} \nabla \times \mathbf{B}_1 \right) \quad (\text{A.23})$$

The set of four equations solved by the CASTOR code are, therefore,

$$\lambda \rho_1 = -\rho_0 \nabla \cdot \mathbf{v}_1 - \mathbf{v}_1 \cdot \nabla \rho_0 - \rho_1 \nabla \cdot \mathbf{v}_0 - \mathbf{v}_0 \cdot \nabla \rho_1, \quad (\text{A.24})$$

$$\lambda T_1 = -\mathbf{v}_0 \cdot \nabla T_1 - (\Gamma - 1) T_1 \nabla \cdot \mathbf{v}_0 - \mathbf{v}_1 \cdot \nabla T_0 - (\Gamma - 1) T_0 \nabla \cdot \mathbf{v}_1, \quad (\text{A.25})$$

$$\begin{aligned} \lambda \rho_0 \mathbf{v}_1 &= -\rho_1 (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_0 - \rho_0 [(\mathbf{v}_1 \cdot \nabla) \mathbf{v}_0 + (\mathbf{v}_0 \cdot \nabla) \mathbf{v}_1] - \nabla(\rho_0 T_1 + \rho_1 T_0) + \\ &\quad \frac{1}{\mu_0} [(\nabla \times \mathbf{B}_0) \times \mathbf{B}_1 + (\nabla \times \mathbf{B}_1) \times \mathbf{B}_0], \quad (\text{A.26}) \end{aligned}$$

and

$$\lambda \mathbf{B}_1 = -(\mathbf{v}_0 \cdot \nabla) \mathbf{B}_1 - \mathbf{B}_1(\nabla \cdot \mathbf{v}_0) + (\mathbf{B}_1 \cdot \nabla) \mathbf{v}_0 + \nabla \times \left(\mathbf{v}_1 \times \mathbf{B}_0 - \frac{\eta}{\mu_0} \nabla \times \mathbf{B}_1 \right). \quad (\text{A.27})$$

While CASTOR solves equations A.24-A.27, MISHKA-1 uses a reduced set of equations [36]. It assumes an equilibrium velocity $\mathbf{v}_0 = 0$ and takes into account only a pressure perturbation p_1 instead of treating ρ_1 and T_1 separately. After the linearisation, equation A.12 becomes

$$\begin{aligned} \frac{\partial}{\partial t}(\cancel{p_0} + e^{\lambda t} p_1) + ((\cancel{\mathbf{v}_0} + e^{\lambda t} \mathbf{v}_1) \cdot \nabla)(p_0 + e^{\lambda t} p_1) + \Gamma(p_0 + e^{\lambda t} p_1) \nabla \cdot (\cancel{\mathbf{v}_0} + e^{\lambda t} \mathbf{v}_1) &= 0, \\ \lambda p_1 &= -(\mathbf{v}_1 \cdot \nabla) p_0 - \Gamma p_0 (\nabla \cdot \mathbf{v}_1). \end{aligned} \quad (\text{A.28})$$

Equation A.17 linearises to

$$\begin{aligned} \lambda e^{\lambda t} \rho_0 \mathbf{v}_1 + \cancel{e^{2\lambda t} \rho_0 (\mathbf{v}_1 \cdot \nabla) \mathbf{v}_1} &= \frac{1}{\mu_0} (\nabla \times \mathbf{B}_0) \times \mathbf{B}_0 + e^{\lambda t} \frac{1}{\mu_0} (\nabla \times \mathbf{B}_0) \times \mathbf{B}_1 \\ &\quad + e^{\lambda t} \frac{1}{\mu_0} (\nabla \times \mathbf{B}_1) \times \mathbf{B}_0 - \nabla(p_0 + e^{\lambda t} p_1) \quad (\text{A.29}) \end{aligned}$$

$$\lambda\rho_0\mathbf{v}_1 = \frac{1}{\mu_0}(\nabla \times \mathbf{B}_0) \times \mathbf{B}_1 + \frac{1}{\mu_0}(\nabla \times \mathbf{B}_1) \times \mathbf{B}_0 - \nabla p_1. \quad (\text{A.30})$$

Finally, with zero equilibrium velocity, equation A.27 can be simplified to

$$\lambda\mathbf{B}_1 = \nabla \times \left(\mathbf{v}_1 \times \mathbf{B}_0 - \frac{\eta}{\mu_0} \nabla \times \mathbf{B}_1 \right). \quad (\text{A.31})$$

Expressing the perturbed magnetic field as a function of the perturbed vector potential $\mathbf{B}_1 = \nabla \times \mathbf{A}_1$ yields the following set of equations:

$$\lambda\rho_0\mathbf{v}_1 = \frac{1}{\mu_0}(\nabla \times \mathbf{B}_0) \times (\nabla \times \mathbf{A}_1) - \frac{\mathbf{B}_0}{\mu_0} \times (\nabla \times \nabla \times \mathbf{A}_1) - \nabla p_1, \quad (\text{A.32})$$

$$\lambda p_1 = -(\mathbf{v}_1 \cdot \nabla)p_0 - \Gamma p_0(\nabla \cdot \mathbf{v}_1), \quad (\text{A.33})$$

$$\lambda\mathbf{A}_1 = \mathbf{v}_1 \times \mathbf{B}_0 - \frac{\eta}{\mu_0} \nabla \times \nabla \times \mathbf{A}_1. \quad (\text{A.34})$$

The equations A.32-A.34 describe the general stability problem solved by MISHKA-1 in terms of the variables \mathbf{A}_1 , \mathbf{v}_1 and p_1 .

B Solver

To solve the linearised equations A.24-A.27 or A.32-A.34, they are converted to a large non-hermitian eigenvalue problem

$$\lambda\mathbf{S}\mathbf{w} = \mathbf{R}\mathbf{w}, \quad (\text{B.1})$$

where the perturbed quantities are expressed by the vector $\mathbf{w}^T = (\rho_1, \mathbf{v}_1, T_1, \mathbf{A}_1)$, or $\mathbf{w}^T = (p_1, \mathbf{v}_1, \mathbf{A}_1)$ in MISHKA-1 mode. The matrix \mathbf{S} only contains diagonal elements and the matrix \mathbf{R} contains the differential operators and equilibrium quantities from equations A.24-A.27 or A.32-A.34. The components of \mathbf{w} are expanded in Fourier series in toroidal and poloidal angle, and described by finite quadratic and cubic Hermitian elements in the radial direction. The trial functions then have the form

$$w_k(s, \theta, \varphi) = \sum_{m,j,p} \alpha_{m,j}^{k,p} h_{p,k}^j e^{im\theta} e^{in\varphi}, \quad (\text{B.2})$$

where s , θ and φ are the radial, poloidal and toroidal coordinates. The indices k , m , j and p label the variables, the poloidal Fourier harmonics, the radial nodes and the interpolating functions, respectively. The explicit form of the Hermite polynomials $h_{p,k}^j$ can be found in appendix A of the ILSA manual [34]. To solve the eigenvalue problem the Galerkin method is applied [92], in which the weighting function that is multiplied to the equation is the same as the trial function:

$$\begin{aligned} & \lambda \int \int \int (h_{p,k}^j e^{-im\theta} e^{-in\varphi} \sum_{\bar{p}, \bar{m}, \bar{k}, \bar{j}} \sum_l e^{il\theta} (\sqrt{g} S^{k, \bar{k}})_l \alpha_{\bar{m}, \bar{j}}^{\bar{k}, \bar{p}} h_{\bar{p}, \bar{k}}^{\bar{j}} e^{i\bar{m}\theta} e^{i\bar{n}\varphi}) ds d\theta d\varphi \\ &= \int \int \int (h_{p,k}^j e^{-im\theta} e^{-in\varphi} \sum_{\bar{p}, \bar{m}, \bar{k}, \bar{j}} \sum_l e^{il\theta} (\sqrt{g} R^{k, \bar{k}})_l \alpha_{\bar{m}, \bar{j}}^{\bar{k}, \bar{p}} h_{\bar{p}, \bar{k}}^{\bar{j}} e^{i\bar{m}\theta} e^{i\bar{n}\varphi}) ds d\theta d\varphi \end{aligned} \quad (\text{B.3})$$

After integrating over the poloidal and toroidal angles, the whole problem can be written as

$$\lambda \mathbf{B} \mathbf{x} = \mathbf{A} \mathbf{x}, \tag{B.4}$$

where $\mathbf{x} = \alpha_{\bar{\mathbf{m}}, \bar{\mathbf{j}}}^{\bar{k}, \bar{p}}$ denotes the vector of the expansion coefficients. The explicit form of the matrices \mathbf{A} and \mathbf{B} can be found elsewhere [34, 35]. The eigenvalue equation is solved by ILSA using the inverse vector iteration method.

References

- [1] J Wesson. *Tokamaks*. Oxford University Press, 1987.
- [2] AR Raffray, M Akiba, V Chuyanov, L Giancarli, and S Malang. Breeding blanket concepts for fusion and materials requirements. *Journal of Nuclear Materials*, 307-311:21–30, 2002.
- [3] H Bartels, H Bosch, R Brakel, M Drevlak, U Fantz, H Hartfuss, D Hartmann, D Hoffmann, R Kleiber, A Könies, K Krieger, E Poli, B Scott, W Suttrop and H Zohm. *IPP Summer University for Plasma Physics*, 2007.
- [4] T Boyd. *The Physics of Plasmas*. Cambridge University Press, 2003.
- [5] P Bellan. *Fundamentals on Plasma Physics*. Cambridge University Press, 2008.
- [6] R Neu, V Bobkov, R Dux, JC Fuchs, O Gruber, A Herrmann, A Kallenbach, H Maier, M Mayer, T Pütterich, V Rohde, ACC Sips, J Stober, K Sugiyama, and the ASDEX Upgrade Team. Ten years of W programme in ASDEX Upgrade – challenges and conclusions. *Physica Scripta*, 2009(T138):014038, 2009.
- [7] G F Matthews, M Beurskens, S Brezinsek, M Groth, E Joffrin, A Loving, M Kear, M-L Mayoral, R Neu, P Prior, V Riccardo, F Rimini, M Rubel, G Sips, E Villedieu, P de Vries, ML Watkins, and EFDA-JET contributors. JET ITER-like wall – overview and experimental programme. *Physica Scripta*, 2011(T145):014001, 2011.
- [8] M Shimada, DJ Campbell, V Mukhovatov, M Fujiwara, N Kirneva, K Lackner, M Nagami, VD Pustovitov, N Uckan, J Wesley, N Asakura, AE Costley, AJH Donné, EJ Doyle, A Fasoli, C Gormezano, Y Gribov, O Gruber, TC Hender, W Houlberg, S Ide, Y Kamada, A Leonard, B Lipschultz, A Loarte, K Miyamoto, V Mukhovatov, TH Osborne, A Polevoi, and ACC Sips. Progress in the ITER physics basis, Chapter 1: Overview and summary. *Nuclear Fusion*, 47(6):S1, 2007.
- [9] T Eich, B Sieglin, A Scarabosio, A Herrmann, A Kallenbach, GF Matthews, S Jachmich, S Brezinsek, M Rack, and RJ Goldston. Empirical scaling of inter-ELM power widths in ASDEX Upgrade and JET. *Journal of Nuclear Materials*, 438, Supplement(0):S72 – S77, 2013.
- [10] F Wagner, G Becker, K Behringer, D Campbell, A Eberhagen, W Engelhardt, G Fussmann, O Gehre, J Gernhardt, G v Gierke, G Haas, M Huang, F Karger, M Keilhacker, O Klüber, M Kornherr, K Lackner, G Lisitano, GG Lister, HM Mayer, D Meisel, ER Müller, H Murmann, H Niedermeyer, W Poschenrieder, H Rapp, H Röhr, F Schneider, G Siller, E Speth, A Stäbler, KH Steuer, G Venus, O Vollmer, and Z Yü. Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak. *Phys. Rev. Lett.*, 49:1408–1412, 1982.
- [11] F Wagner. A quarter-century of H-mode studies. *Plasma Physics and Controlled Fusion*, 49(12B):B1–B33, 2007.

- [12] F Ryter, T Pütterich, M Reich, A Scarabosio, E Wolfrum, R Fischer, M Gemisic Adamov, N Hicks, B Kurzban, C Maggi, R Neu, V Rohde, G Tardini, and the ASDEX Upgrade Team. H-mode threshold and confinement in helium and deuterium in ASDEX Upgrade. *Nuclear Fusion*, 49(6):062003, 2009.
- [13] RM McDermott, B Lipschultz, JW Hughes, PJ Catto, AE Hubbard, IH Hutchinson, RS Granetz, M Greenwald, B LaBombard, K Marr, ML Reinke, JE Rice, D Whyte, and Alcator C-Mod Team. Edge radial electric field structure and its connections to H-mode confinement in Alcator C-Mod plasmas. *Physics of Plasmas*, 16(5):056103, 2009.
- [14] H Zohm. Edge localized modes (ELMs). *Plasma Physics and Controlled Fusion*, 38:105–128, 1996.
- [15] M Dunne. *Inter-ELM evolution of the edge current density profile on the ASDEX Upgrade tokamak*. PhD thesis, University College Cork, 2013.
- [16] A Burckhart. *Temporal evolution of the electron temperature and density profiles during type I ELM cycles at ASDEX Upgrade*. Diplomarbeit, Universität München, 2009.
- [17] A Burckhart, E Wolfrum, R Fischer, K Lackner, and H Zohm. Inter-ELM behaviour of the electron density and temperature pedestal in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 52(10):105010, 2010.
- [18] E Wolfrum, M Bernert, JE Boom, A Burckhart, IGJ Classen, GD Conway, T Eich, R Fischer, A Gude, A Herrmann, NC Luhmann Jr, M Maraschek, R McDermott, HK Park, T Pütterich, J Vicente, B Wieland, M Willensdorfer, and the ASDEX Upgrade Team. 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, 2011.
- [19] JE Boom. 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, 2012.
- [20] G Federici, CH Skinner, JN Brooks, JP Coad, C Grisolia, AA Haasz, A Hassanein, V Philipps, CS Pitcher, J Roth, WR Wampler, and DG Whyte. Plasma-material interactions in current tokamaks and their implications for next step fusion reactors. *Nuclear Fusion*, 41(12):1967, 2001.
- [21] JP Freidberg. *Ideal Magnetohydrodynamics*. Plenum Publishing Company Limited, 1987.
- [22] JW Connor. A review of models for ELMs. *Plasma Physics and Controlled Fusion*, 191, 1998.
- [23] PB Snyder and HR Wilson. ELMs and the role of current-driven instabilities in the pedestal. *Contributions to Plasma Physics*, 42(2-4):258–271, 2002.

-
- [24] RL Miller, MS Chu, JM Greene, YR Lin Liu, and RE Waltz. Noncircular, finite aspect ratio, local equilibrium model. *Physics of Plasmas*, 5(4):973–978, 1998.
- [25] MG Dunne, PJ McCarthy, E Wolfrum, R Fischer, and L Giannone. Measurement of neoclassically predicted edge current density at ASDEX Upgrade. *Nuclear Fusion*, 52(12), 2012.
- [26] AG Peeters. The bootstrap current and its consequences. *Plasma Physics and Controlled Fusion*, 42:B231, 2000.
- [27] O Sauter, C Angioni, and YR Lin-Liu. Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime. *Physics of Plasmas*, 6(7):2834, 1999.
- [28] O Sauter, C Angioni, and YR Lin-Liu. Erratum: Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime, *Physics of Plasmas* 6 2834 (1999). *Physics of Plasmas*, 9(12):5140, 2002.
- [29] PJ McCarthy. Analytical solutions to the Grad-Shafranov equation for tokamak equilibrium with dissimilar source functions. *Physics of Plasmas*, 6(9):3554–3560, 1999.
- [30] PJ McCarthy. 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.
- [31] K Lackner. Computation of ideal MHD equilibria. *Computer Physics Communications*, 12(1):33–44, 1976.
- [32] GTA Huysmans, JP Goedbloed, and W Kerner. Isoparametric bicubic Hermite elements for solution of the Grad Shafranov equation. *International Journal of Modern Physics C*, 2(1):371–6, 1991.
- [33] IB Bernstein and EA Frieman. An energy principle for hydromagnetic stability problems. *Proceedings of the Royal Society A*, 244(1236):17–40, 1958.
- [34] E Strumberger, C Konz, V Ingochine, and C Tichmann. ILSA code: Linear MHD stability studies. Manual, Max-Planck-Institut für Plasmaphysik, Garching, 2005.
- [35] W Kerner, JP Goedbloed, GTA Huysmans, S Poedts, and E Schwarz. CASTOR : Normal-Mode Analysis of Resistive MHD Plasmas. *Plasma Physics*, 303:271–303, 1998.
- [36] AB Mikhailovskii and SE Sharapov. Optimization of computational MHD normal-mode analysis for tokamaks. *Plasma Physics Reports*, 23(10):844–57, 1997.
- [37] E Strumberger, S Gunter, P Merkel, S Riondato, and E Schwarz. Numerical MHD stability studies: toroidal rotation, viscosity, resistive walls and current holes. *Nuclear Fusion*, 45(9):1156–1167, 2005.
-

- [38] F Osmanlic. *Ideal linear MHD stability analysis of the edge plasma in ASDEX Upgrade*. Diplomarbeit, Technische Universität München, 2011.
- [39] A Stäbler, J-H Feist, E Speth, JL Dunne, S Goetz, B Heinemann, A Krauss, R-C Kunze, H Lohnert, J Sielanko, W Szyszek, O Vollmer, and K Wittenbecher. Design of the neutral beam injection system for ASDEX Upgrade. In *Fusion Technology 1988*, pages 620 – 624. Elsevier, Oxford, 1989.
- [40] B Streibl, PT Lang, F Leuterer, JM Noterdaeme, and A Stäbler. Chapter 2: Machine design, fueling, and heating in ASDEX Upgrade. *Fusion Science and Technology*, 44(3):578–592, 2003.
- [41] A Stäbler, J Hobirk, F Leuterer, F Meo, and JM Noterdaeme. Chapter 13: Current drive in ASDEX Upgrade. *Fusion Science and Technology*, 44(3):730–742, 2003.
- [42] R Prater. Heating and current drive by electron cyclotron waves. *Physics of Plasmas*, 11(5):2349, 2004.
- [43] E Poli, AG Peeters, and GV Pereverzev. TORBEAM, a beam tracing code for electron-cyclotron waves in tokamak plasmas. *Computer Physics Communications*, 136(1–2):90 – 104, 2001.
- [44] T. Lunt, J.C. Fuchs, K. Mank, Y. Feng, F. Brochard, A. Herrmann, V. Rohde, and N. Endstrasser. A new 3D viewer as an interface between the ASDEX Upgrade CAD models and data from plasma modelling and experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 623(2):812 – 814, 2010. 1st International Conference on Frontiers in Diagnostics Technologies.
- [45] E Viezzer, T Pütterich, GD Conway, R Dux, T Happel, JC Fuchs, RM McDermott, F Ryter, B Sieglin, W Suttrop, M Willensdorfer, E Wolfrum, and the ASDEX Upgrade Team. High-accuracy characterization of the edge radial electric field at ASDEX Upgrade. *Nuclear Fusion*, 53(5):053005, 2013.
- [46] R McDermott. Private communication.
- [47] HJ Hartfuss, T Geist, and M Hirsch. Heterodyne methods in millimetre wave plasma diagnostics with applications to ECE, interferometry and reflectometry. *Plasma Physics and Controlled Fusion*, 39(11):1693, 1997.
- [48] SK Rathgeber, L Barrera, T Eich, R Fischer, B Nold, W Suttrop, M Willensdorfer, E Wolfrum, and the ASDEX Upgrade Team. Estimation of edge electron temperature profiles via forward modelling of the electron cyclotron radiation transport at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 55(2):025004, 2013.
- [49] H Murmann. The thomson scattering systems of the ASDEX Upgrade tokamak. *Review of Scientific Instruments*, 63(10):4941–4943, 1992.
- [50] J Schweinzer. Reconstruction of plasma edge density profiles from LiI(2s-2p) emission profiles. *Plasma Physics and Controlled Fusion*, 34(7):1173–1183, 1992.

-
- [51] M Willensdorfer, E Wolfrum, R Fischer, J Schweinzer, M Sertoli, B Sieglin, G Veres, F Aumayr, and the ASDEX Upgrade Team. Improved chopping of a lithium beam for plasma edge diagnostic at ASDEX Upgrade. *Review of Scientific Instruments*, 83(2):023501, 2012.
 - [52] R Fischer. Probabilistic lithium beam data analysis. *Plasma Physics and Controlled Fusion*, 50(8):085009, 2008.
 - [53] C Demichlis. Equipe TFR - tokamak plasma diagnostics. *Nuclear Fusion*, 18(5):647–731, 1978.
 - [54] R Fischer, CJ Fuchs, B Kurzan, and W Suttrop. Integrated data analysis of profile diagnostics at ASDEX Upgrade. *Fusion Science and Technology*, 58(2):675–684, 2010.
 - [55] SK Rathgeber, R Fischer, S Fietz, J Hobirk, A Kallenbach, H Meister, T Pütterich, F Ryter, G Tardini, and E Wolfrum. Estimation of profiles of the effective ion charge at ASDEX Upgrade with Integrated Data Analysis. *Plasma Physics and Controlled Fusion*, 52(9):095008, 2010.
 - [56] PB Snyder, RJ Groebner, AW Leonard, TH Osborne, and HR Wilson. Development and validation of a predictive model for the pedestal height. *Physics of Plasmas*, 16(5):056118, 2009.
 - [57] GTA Huysmans. External kink (peeling) modes in X-point geometry. *Plasma Physics and Controlled Fusion*, 47(12):2107–2121, 2005.
 - [58] HR Wilson, SC Cowley, A Kirk, and PB Snyder. 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, 2006.
 - [59] WM Tang, RL Dewar, and J Manickam. Influence of diamagnetic drifts on critical beta in tokamaks. *Nuclear Fusion*, 22(8):1079–1081, 1982.
 - [60] WA Cooper and KT Tsang. Diamagnetic drift stabilization of ballooning modes in tokamak geometry. *Nuclear Fusion*, 21(11):1477–1479, 1981.
 - [61] RJ Hastie, PJ Catto, and JJ Ramos. Effect of strong radial variation of the ion diamagnetic frequency on internal ballooning modes. *Physics of Plasmas*, 7(11):4561–4566, 2000.
 - [62] M Dunne. Sensitivity of plasma edge equilibrium reconstruction to experimental uncertainties at ASDEX Upgrade. *To be submitted*, 2013.
 - [63] S Saarelma. ELM phenomenon as an interaction between bootstrap-current driven peeling modes and pressure-driven ballooning modes. *Plasma Physics and Controlled Fusion*, 42(5A):A139, 2000.
 - [64] S Saarelma, V Parail, Y Andrew, E De La Luna, A Kallenbach, M Kempenaars, A Korotkov, A Loarte, JS Lönnroth, P Monier-Garbet, JK Stober, and W Suttrop. MHD stability analysis of diagnostic optimized configuration shots in JET. *Plasma Physics and Controlled Fusion*, 47(5):713–731, 2005.
-

- [65] C Konz, PB Snyder, and LD Horton. Comparative Linear MHD Stability Analysis for ASDEX Upgrade and DIII-D Power Scan Studies. *Proc. 36th EPS Conf.*, (June):2–5, 2009.
- [66] PB Snyder. ELMs and constraints on the H-mode pedestal: peeling-ballooning stability calculation and comparison with experiment. *Nuclear Fusion*, 44(2):320, 2004.
- [67] P. B. Snyder, H. R. Wilson, J. R. Ferron, L. L. Lao, A. W. Leonard, T. H. Osborne, A. D. Turnbull, D. Mossessian, M. Murakami, and X. Q. Xu. Edge localized modes and the pedestal: A model based on coupled peeling–ballooning modes. *Physics of Plasmas*, 9(5):2037–2043, 2002.
- [68] LL Lao. Dependence of edge stability on plasma shape and local pressure gradients in the DIII-D and JT-60U tokamaks. *Nuclear Fusion*, 41(3):295, 2001.
- [69] PB Snyder. A first-principles predictive model of the pedestal height and width: development, testing and ITER optimization with the EPED model. *Nuclear Fusion*, 51(10):103016, 2011.
- [70] Frank Jenko. Particle pinch in collisionless drift-wave turbulence. *Physics of Plasmas*, 7(2):514, 2000.
- [71] D Hatch. private communication.
- [72] W Suttrop, V Mertens, H Murmann, J Neuhauser, J Schweinzer, and ASDEX-Upgrade Team. Operational limits for high edge density H-mode tokamak operation. *Journal of Nuclear Materials*, 266-269:118–123, 1999.
- [73] Y Igitkhanov, O Pogutse, G Janeschitz, and J Cordey. Physics of the L-H transition and type III-ELMs phenomena (scaling properties and dimensionless analysis). *Contributions to Plasma Physics*, 40(3-4):368–374, 2000.
- [74] M Bernert. *Analysis of the H-mode density limit in the ASDEX Upgrade tokamak using bolometry*. PhD thesis, Universität München, 2013.
- [75] S Saarelma, A Alfier, MN Beurskens, R Coelho, HR Koslowski, Y Liang, and I Nunes. MHD stability analysis of small ELM regimes in JET. *Plasma Physics and Controlled Fusion*, 51(3):035001, March 2009.
- [76] S Saarelma, S Günter, LD Horton, and ASDEX Upgrade Team. MHD stability analysis of type II ELMs in ASDEX Upgrade. *Nuclear Fusion*, 43:262, 2003.
- [77] PB Snyder, TH Osborne, KH Burrell, RJ Groebner, and AW Leonard. 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), 2012.
- [78] H Doerk, F Jenko, T Gorler, D Told, MJ Pueschel, and DR Hatch. Gyrokinetic prediction of microtearing turbulence in standard tokamaks. *Physics of Plasmas*, 19(5):055907, 2012.

-
- [79] A Kallenbach, R Dux, JC Fuchs, R Fischer, B Geiger, L Giannone, A Herrmann, T Lunt, V Mertens, R McDermott, R Neu, T Pütterich, S Rathgeber, V Rohde, K Schmid, J Schweinzer, W Treutterer, and the ASDEX Upgrade Team. Divertor power load feedback with nitrogen seeding in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 52(5):055002, 2010.
- [80] J Schweinzer. Confinement of 'improved H-modes' in the all-tungsten ASDEX Upgrade with nitrogen seeding. *Nuclear Fusion*, 51(11):113003, 2011.
- [81] MN Beurskens, J Schweinzer, C Angioni, A Burckhart, CD Challis, I Chapman, R Fischer, J Flanagan, L Frassinetti, C Giroud, J Hobirk, E Joffrin, A Kallenbach, M Kempenaars, M Leyland, P Lomas, G Maddison, M Maslov, R McDermott, R Neu, I Nunes, F Ryter, S Saarelma, PA Schneider, P Snyder, G Tardini, E Viezzer, and E Wolfrum. The effect of a metal wall on confinement in JET and ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, submitted, 2013.
- [82] A. Kallenbach, J. Adamek, L. Aho-Mantila, S. Äkäslompolo, C. Angioni, C.V. Atanasiu, M. Balden, K. Behler, E. Belonohy, A. Bergmann, M. Bernert, R. Bilato, V. Bobkov, J. Boom, A. Bottino, F. Braun, M. Brüdgam, A. Buhler, A. Burckhart, A. Chankin, I.G.J. Classen, G.D. Conway, D.P. Coster, P. de Marne, R. D'Inca, R. Drube, R. Dux, T. Eich, N. Endstrasser, K. Engelhardt, B. Esposito, E. Fable, H.-U. Fahrback, L. Fattorini, R. Fischer, A. Flaws, H. Fünfgelder, J.C. Fuchs, K. Gál, M. García Muñoz, B. Geiger, M. Gemisic Adamov, L. Giannone, C. Giroud, T. Görler, S. da Graca, H. Greuner, O. Gruber, A. Gude, S. Günter, G. Haas, A.H. Hakola, D. Hangan, T. Happel, T. Hauff, B. Heinemann, A. Herrmann, N. Hicks, J. Hobirk, H. Höhnle, M. Hölzl, C. Hopf, L. Horton, M. Huart, V. Igochine, C. Ionita, A. Janzer, F. Jenko, C.-P. Käsemann, S. Kálvin, O. Kar-daun, M. Kaufmann, A. Kirk, H.-J. Klingshirn, M. Kocan, G. Kocsis, H. Kollotzek, C. Konz, R. Koslowski, K. Krieger, T. Kurki-Suonio, B. Kurzan, K. Lackner, P.T. Lang, P. Lauber, M. Laux, F. Leipold, F. Leuterer, A. Lohs, N.C. Luhmann Jr, T. Lunt, A. Lyssoivan, H. Maier, C. Maggi, K. Mank, M.-E. Manso, M. Maraschek, P. Martin, M. Mayer, P.J. McCarthy, R. McDermott, H. Meister, L. Menchero, F. Meo, P. Merkel, R. Merkel, V. Mertens, F. Merz, A. Mlynec, F. Monaco, H.W. Müller, M. München, H. Murmann, G. Neu, R. Neu, B. Nold, J.-M. Noter-daeme, H.K. Park, G. Pautasso, G. Pereverzev, Y. Podoba, F. Pompon, E. Poli, K. Polochiy, S. Potzel, M. Prechtel, M.J. Püschel, T. Pütterich, S.K. Rathgeber, G. Raupp, M. Reich, B. Reiter, T. Ribeiro, R. Riedl, V. Rohde, J. Roth, M. Rott, F. Ryter, W. Sandmann, J. Santos, K. Sassenberg, P. Sauter, A. Scarabosio, G. Schall, K. Schmid, P.A. Schneider, W. Schneider, G. Schramm, R. Schrittwieser, J. Schweinzer, B. Scott, M. Sempf, F. Serra, M. Sertoli, M. Siccino, A. Sigalov, A. Silva, A.C.C. Sips, F. Sommer, A. Stäbler, J. Stober, B. Streibl, E. Strumberger, K. Sugiyama, W. Suttrop, T. Szepesi, G. Tardini, C. Tichmann, D. Told, W. Treutterer, L. Urso, P. Varela, J. Vincente, N. Vianello, T. Vierle, E. Viezzer, C. Vorpahl, D. Wagner, A. Weller, R. Wenninger, B. Wieland, C. Wigger, M. Willensdorfer, M. Wischmeier, E. Wolfrum, E. Würsching, D. Yadikin, Q. Yu, I. Zammuto, D. Zasche, T. Zehetbauer, Y. Zhang, M. Zilker, and H. Zohm. Overview of ASDEX Upgrade results. *Nuclear Fusion*, 51(9):094012, 2011.
-

- [83] P Schneider. *Characterization and scaling of the tokamak edge transport barrier*. PhD thesis, Universität München, 2012.
- [84] P.T. Lang, G.D. Conway, T. Eich, L. Fattorini, O. Gruber, S. Günter, L.D. Horton, S. Kalvin, A. Kallenbach, M. Kaufmann, G. Kocsis, A. Lorenz, M.E. Manso, M. Maraschek, V. Mertens, J. Neuhauser, I. Nunes, W. Schneider, W. Suttrop, H. Urano, and the ASDEX Upgrade Team. ELM pace making and mitigation by pellet injection in ASDEX Upgrade. *Nuclear Fusion*, 44(5):665, 2004.
- [85] L.R. Baylor, N. Commaux, T.C. Jernigan, N.H. Brooks, S.K. Combs, T.E. Evans, M.E. Fenstermacher, R.C. Isler, C.J. Lasnier, S.J. Meitner, R.A. Moyer, T.H. Osborne, P.B. Parks, P.B. Snyder, E.J. Strait, E.A. Unterberg, and A. Loarte. Reduction of edge-localized mode intensity using high-repetition-rate pellet injection in tokamak H-mode plasmas. *Phys. Rev. Lett.*, 110:245001, 2013.
- [86] P.T. Lang, D. Frigione, A. Géraud, T. Alarcon, P. Bennett, G. Cseh, D. Garnier, L. Garzotti, F. Köchl, G. Kocsis, M. Lennholm, R. Neu, R. Mooney, S. Saarelma, B. Sieglin, and JET-EFDA Contributors. ELM pacing and trigger investigations at JET with the new ITER-like wall. *Nuclear Fusion*, 53(7):073010, 2013.
- [87] P.T. Lang, A. Burckhart, M. Bernert, L. Casali, R. Fischer, O. Kardaun, G. Kocsis, M. Maraschek, A. Mlynek, B. Plöckl, M. Reich, F. Ryter, J. Schweinzer, B. Sieglin, W. Suttrop, T. Szepesi, G. Tardini, E. Wolfrum, D. Zasche, and H. Zohm. Pellet ELM trigger and pacing investigations in the all-metal wall tokamak ASDEX Upgrade. *To be submitted*, 2013.
- [88] LD Horton, GD Conway, AW Degeling, T Eich, A Kallenbach, PT Lang, JB Lister, A Loarte, YR Martin, PJ McCarthy, H Meister, J Neuhauser, J Schirmer, ACC Sips, W Suttrop, and the ASDEX Upgrade Team. ITER-relevant H-mode physics at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 46(12B):B511, 2004.
- [89] JX Rossel. *Edge localized mode control in TCV*. PhD thesis, Ecole Polytechnique Federale de Lausanne, 2013.
- [90] C. Angioni, R.M. McDermott, E. Fable, R. Fischer, T. Pütterich, F. Ryter, G. Tardini, and the ASDEX Upgrade Team. Gyrokinetic modelling of electron and boron density profiles of H-mode plasmas in ASDEX Upgrade. *Nuclear Fusion*, 51(2):023006, 2011.
- [91] PB Snyder. Ideal magnetohydrodynamic constraints on the pedestal temperature in tokamaks. *Plasma Physics and Controlled Fusion*, 45(9):1671, 2003.
- [92] L Hageman. *Applied iterative Methods*. Dover Publications, 2004.

Abbreviations

ASDEX AxialSymmetrisches DivertorEXperiment

CXRS Charge Exchange Recombination Spectroscopy

DCN Deuterium cyanide

DN Double Null

ECCD Electron Cyclotron Current Drive

ECE Electron Cyclotron Emission

ECRH Electron Cyclotron Resonance Heating

ELM Edge Localised Mode

ETB Edge Transport Barrier

H-mode High confinement mode

HFS High Field Side

IDA Integrated Data Analysis

JET Joint European Torus

L-mode Low confinement mode

LCFS Last Closed Flux Surface

LFS Low Field Side

Li-beam Lithium beam impact excitation spectroscopy

LSN Lower Single Null

MCF Magnetic Confinement Fusion

MHD MagnetoHydroDynamic

NBI Neutral Beam Injection

Nd:YAG Neodymium-doped yttrium aluminium garnet

PFC Plasma Facing Component

SOL Scrape Off Layer

TS Thomson Scattering

USN Upper Single Null

