

The Plasma Response to Applied Magnetic Perturbations for ELM Control in Tokamaks

David Anthony Ryan

Doctor of Philosophy

University of York
Physics

November 2017

Abstract

Edge Localised Modes (ELMs), a repetitive MHD instability, present a risk to ITER machine components and scientific objectives, necessitating a robust ELM mitigation strategy. The application of Resonant Magnetic Perturbations (RMPs) is an effective means of suppressing or mitigating ELMs, and a set of RMP coils will be installed on ITER. Although there are several working theories, a robust predictive theory of ELM control by RMPs is currently lacking. In previous assessments of the efficacy of the ITER RMP system, the response of the plasma to the applied RMP was typically overlooked; however it is well known that the plasma response constitutes a significant correction to the total plasma perturbation. In this thesis, the MARS-F code (Liu *et al* 2000 *Phys. Plasmas* **7** 3681) is used to study the linear plasma response to applied RMPs using realistic geometry and experimental parameters. The pitch aligned components and the amplified peeling response components of the RMP, both previously implicated in the ELM mitigation and suppression mechanisms, are studied. It is shown that although the pitch aligned components are strongly screened by the plasma response, they may be finite near the plasma edge, and may be driven by the amplified peeling response via poloidal harmonic coupling. RMPs can be tuned by adjusting the phase $\Delta\phi_{UL}$ between the toroidal waveforms in the upper and lower rows of RMP coils. A scheme for optimising the coil phase to maximise ELM mitigation for a given plasma equilibrium is devised for ASDEX Upgrade plasmas. A database of ASDEX Upgrade RMP experiments is assembled and used to benchmark the scheme against MARS computations, and validate it against experimental measurements. The scheme is then extended to the ITER RMP coil system.

Contents

Abstract	i
List of Tables	iv
List of Figures	v
Acknowledgements	vii
Declaration of Authorship	viii
1 Introduction	1
1.1 Energy for Human Development	1
1.2 Fossil Fuels	2
1.3 Carbon Free Electricity Generation	4
1.4 Nuclear Fusion	5
1.5 Achieving Terrestrial Fusion	7
1.5.1 The Lawson Criterion	9
1.6 The Tokamak	11
1.7 Edge Localised Modes	14
1.8 Outline	17
2 Background	19
2.1 Magnetohydrodynamics	19
2.2 MHD Equilibrium	21
2.3 Plasma Stability	23
2.4 RMP ELM Control Observations	25
2.5 RMP ELM Control Theories	27
2.6 Plasma Response Model - MARS-F	32
3 Coupling between peeling response and pitch aligned component	35
3.1 Resonant Field Amplification	35
3.2 Amplified Peeling Response	35
3.3 Equilibrium	37
3.4 Benchmark	38
3.5 Plasma Response to Applied Perturbation	41

3.5.1	Pitch Aligned Components	46
3.5.2	Robustness of Peeling Response to X Point Truncation	49
3.6	Poloidal Harmonic Coupling on ASDEX-Upgrade	53
3.7	Experimental Validation of MARS Simulations	55
3.8	Chapter Results Summary	65
4	Coil Phase Optimisation	68
4.1	Alignment of Applied Field with Equilibrium Field	70
4.1.1	Extracting $\Delta\phi_{opt}$ from Plasma Response Computation	70
4.1.2	Effect on Alignment of Equilibrium Parameters β_N and q_{95}	71
4.2	ASDEX Upgrade Optimal Coil Phase	74
4.2.1	Scaled Equilibrium Set	74
4.2.2	MARS-F Computed $\Delta\phi_{opt}$ Across (β_N, q_{95}) Domain	74
4.2.3	Parametrisation of $\Delta\phi_{opt}$	78
4.3	Parametrisation Benchmarking	79
4.3.1	Benchmarking Database	79
4.4	Sources of Uncertainty for $\Delta\phi_{opt,quad}$	83
4.5	Parametrisation Experimental Validation	85
4.5.1	Coil Phase Scan Experiments	86
4.5.2	Comparison of measured $\Delta\phi_{opt}$ with predictions	92
4.6	ITER Coil Phase Parametrisation	93
4.6.1	ITER reference equilibrium and scaled equilibrium set	94
4.6.2	Parametrisation of ITER coil phase	96
4.7	Database Regression Analysis	99
4.8	Choice of Toroidal Waveform	102
4.9	Chapter Results Summary	104
5	Conclusion	106
5.1	Summary and Discussion	106
5.2	Outlook	113
A	Derivation of displacement resonance condition	115
	Bibliography	118

List of Tables

3.1	Equilibrium parameters of 30835	38
3.2	Validation discharge parameters	56
4.1	Fitting RMSE for different polynomial models	78
4.2	2D quadratic coefficients for optimum coil phase	78
4.3	Experimental coil phase scan equilibrium parameters	88
4.4	Coefficients of ITER upper coil phase parametrisation	98
4.5	Coefficients of ITER lower coil phase parametrisation	99
4.6	Coefficients of ITER coil phase parametrisation without middle set	99
4.7	RMSE between ITER scan and parametrisation	99

List of Figures

1.1	Human development	1
1.2	Climate change	3
1.3	Fusion reaction cross sections	8
1.4	Particle motion	8
1.5	Fusion power balance	9
1.6	Tokamak principles	11
1.7	Triple product progress	12
1.8	Plasma exhaust components	13
1.9	Plasma pedestal	15
2.1	Definition of flux label	21
2.2	Definition of safety factor	22
2.3	P-B stability space	24
2.4	Experimental ELM suppression and mitigation	26
2.5	Parameter spaces of suppression and mitigation	26
2.6	Stochastic theory explanation	28
2.7	Magnetic island induced pressure flattening	28
2.8	Geometric and SFL coordinate systems	34
3.1	Spectral nomenclature	36
3.2	Equilibrium profiles of 30835	39
3.3	Sketch of ASDEX Upgrade RMP coils	40
3.4	Applied RMP toroidal spectrum	40
3.5	Resonant vacuum benchmark	41
3.6	Detailed vacuum benchmark	42
3.7	Sketch of real toroidal angle offset to phase difference.	43
3.8	Poloidal spectra of vacuum and total fields	44
3.9	Coil phase dependence of peeling response	45
3.10	Pitch aligned component profiles	46
3.11	Pitch aligned component coil phase dependence	47
3.12	Sideband phase scan	47
3.13	X point truncation scan	50
3.14	Variation of peeling response with edge geometry	51
3.15	Invariance of peeling response relative to safety factor	52
3.16	Single m spectrograms	54
3.17	Single m pitch aligned components	54
3.18	Plasma response diagnostics	55
3.19	30839 $n = 2$ plasma response	57

3.20	31021 $n = 1$ plasma response	58
3.21	Real space magnetic and displacement perturbations	58
3.22	Measured $n = 2$ poloidal structure.	60
3.23	Measured $n = 1$ poloidal structure.	60
3.24	Measured displacement profiles	62
3.25	Phase scan displacement measurement	63
3.26	Equilibrium uncertainty.	64
3.27	Uncertainty in core response.	64
4.1	Analytic formula for optimum coil phase	70
4.2	Analytic formula validation	71
4.3	Effect of beta and q95 on alignment	72
4.4	Pressure distortion of magnetic axis	73
4.5	Scaled equilibrium set for ASDEX Upgrade	75
4.6	Example scaled equilibrium profiles	75
4.7	Removal of phase wraps	76
4.8	Optimum coil phase with beta and q95	77
4.9	Example fit of kinetic profiles to data	80
4.10	Scatter of validation points in beta and q95 space	80
4.11	Benchmarking point plasma boundaries	81
4.12	Benchmarking of 2D quadratic parametrisation against MARS	82
4.13	Coil phase rotation scan	83
4.14	Coil phase pedestal uncertainty	84
4.15	Coil phase pedestal scan	84
4.16	Plasma boundary variation in validation points	86
4.17	Coil phase scan experimental trace	87
4.18	Measurement of experimental optimum coil phase	89
4.19	Example of ambiguous fit to experimental data	90
4.20	Lag in effective coil phase caused by PSLs	90
4.21	Experimental coil phase as measured with density and ELM frequency	91
4.22	Experimental validation of 2D quadratic parametrisation	93
4.23	ITER coil set	95
4.24	Validation of 2D analytic optimal coil phase formula	95
4.25	ITER reference equilibrium profiles	96
4.26	ITER scaled equilibrium set	97
4.27	Optimum coil ITER phase	98
4.28	Pitch aligned components against peeling response	100
4.29	Normalised pressure against pedestal pressure	101
4.30	Toroidal waveform choice	102
4.31	Misalignment penalty	102

Acknowledgements

I would like to thank my PhD supervisor at Culham, Yueqiang Liu, whose insight, patience and encouragement made this work possible. I would also like to thank my PhD supervisor at York, Ben Dudson, for many hours of lectures and interesting discussions. Thanks are also owed to Andrew Kirk and Nick Walkden at the CCFE, Mike Dunne, Mathias Willensdorfer and Manuel Garcia-Munoz at IPP Garching, and Paolo Piovesan at Consorzio RFX, who provided crucial assistance, opportunities and collaborations throughout this work. I would like to express my deepest gratitude to my parents, Jane and Tony Ryan, for so many years of encouragement, advice and much needed support. Finally I would like to thank my wife Sophie, who has been by my side in every moment.

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 20142018 under grant agreement No 633053, and is part-funded by the EPSRC through the Fusion Doctoral Training Network (grant number EP/K504178/1), and part-funded by the RCUK Energy Programme (under grant EP/I501045). To obtain further information on the data and models underlying this paper, please contact PublicationsManager@ccfe.ac.uk. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

Declaration of Authorship

I, David A Ryan, declare that this thesis titled, ‘Plasma Response to Applied Magnetic Perturbations for ELM Control in Tokamaks’ and the work presented in it are my own.

I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- No part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution.
- Where I have consulted the published work of others, this is always clearly attributed.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

This thesis contains work previously presented at conferences, and published in (Ryan *et al* 2015 *Plasma Physics and Controlled Fusion* **57** 095008), (Ryan *et al* 2017 *Plasma Physics and Controlled Fusion* **59** 024005), (Willensdorfer *et al* 2016 *Plasma Physics and Controlled Fusion* **58** 114004), and (Piovesan *et al* 2017 *Plasma Physics and Controlled Fusion* **59** 014027). Credit for these works is clearly attributed herein.

Dedicated to my wife Sophie, my own Star on Earth.

Chapter 1

Introduction

1.1 Energy for Human Development

We live an interesting age. After the devastation of the first half of the last century, through technology, cooperation and economic development, humanity has made huge

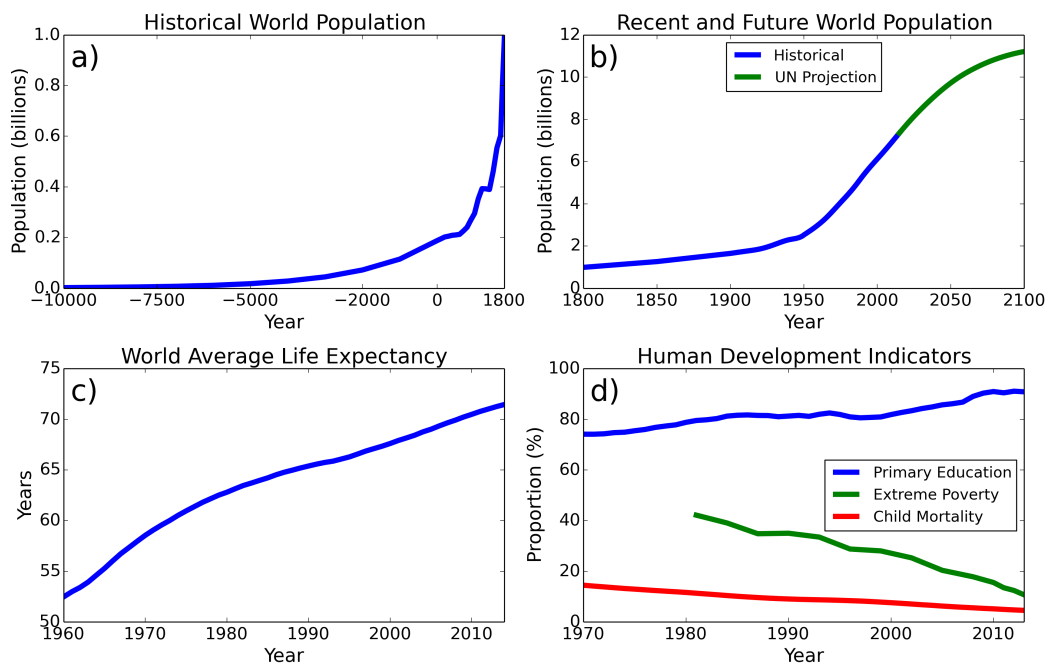


FIGURE 1.1: *The More the Merrier* a) Data from [1]. Until very recently there were precariously few humans. It took tens of millennia for the human population to reach 200 million, and another 1000 years to add another 80 million. b) It now takes a single year to add 80 million. While intuition warns that such a huge population boom must inevitably lead to scarcity and misery, a brief examination of the data tells a different story. c,d) Data from [2]. Even while the population increases, the global poverty and child mortality rates have consistently dropped, and the global average life expectancy and education rate consistently risen.

strides in prosperity and quality of living, as illustrated in figure 1.1. Between 1960 and 2014, the global average life expectancy rose from 52 to 71 years[3]. The global primary education completion rate has risen from 74% to 90% since 1970[4], the global rate of extreme poverty has dropped from 42.2% to 10.7% since 1981[5], and by historical standards inter-state warfare has virtually vanished from the Earth[6]. In the aggregate, humanity has never before been so healthy, prosperous, and peaceful, which is especially remarkable in the context of unprecedented population growth. Having taken tens of thousands of years to reach 1.8 billion in 1915, the population then reached 7.35 billion a mere 100 years later, and is projected to reach 11 billion by 2100[7].

Sufficient per capita energy consumption is fundamentally essential to a high quality of life[8], and so the continuation of the progress of the past century of human development will necessarily entail a large increase in global energy production. It is projected that between 2012 and 2040, total global energy production and electricity production will rise by 48% and 69% respectively[9]. However, faced with the finite availability of fossil fuels and looming threat of anthropogenic climate change, it is becoming ever more apparent that a radical transition in energy production methods will be required.

1.2 Fossil Fuels

In 2014, 81.1% of primary energy consumed globally and 66.7% of electricity generated was produced by burning fossil fuels[11]. Using the ratio of current proven fossil fuel reserves to yearly production, we may crudely estimate how many more years we may safely rely on fossil fuels to provide the bulk of our energy needs. As of 2015, these ratios are 51 years for oil, 53 years for natural gas, and 114 years for coal[12]. These figures are of course underestimates, since proven reserves are regularly revised upwards by hydrocarbon exploration and advances in resource recovery technology[13]. However large the reserves are eventually found to be, fossil fuels are a finite and non-renewable resource which cannot be relied on indefinitely.

Perhaps a more pressing concern than their availability, are the effects of burning fossil fuels on human health. Studies by the World Health Organisation have found that exposure to fossil fuel air pollution increases the risk of numerous debilitating diseases, most notably stroke, heart disease and lung cancer[14]. Because the most polluted areas tend also to be the most densely populated, in 2014 92% of the world's population lived in areas where the air pollution was above WHO recommended limits, and exposure to ambient air pollution currently causes 3.7 million premature deaths per year[15].

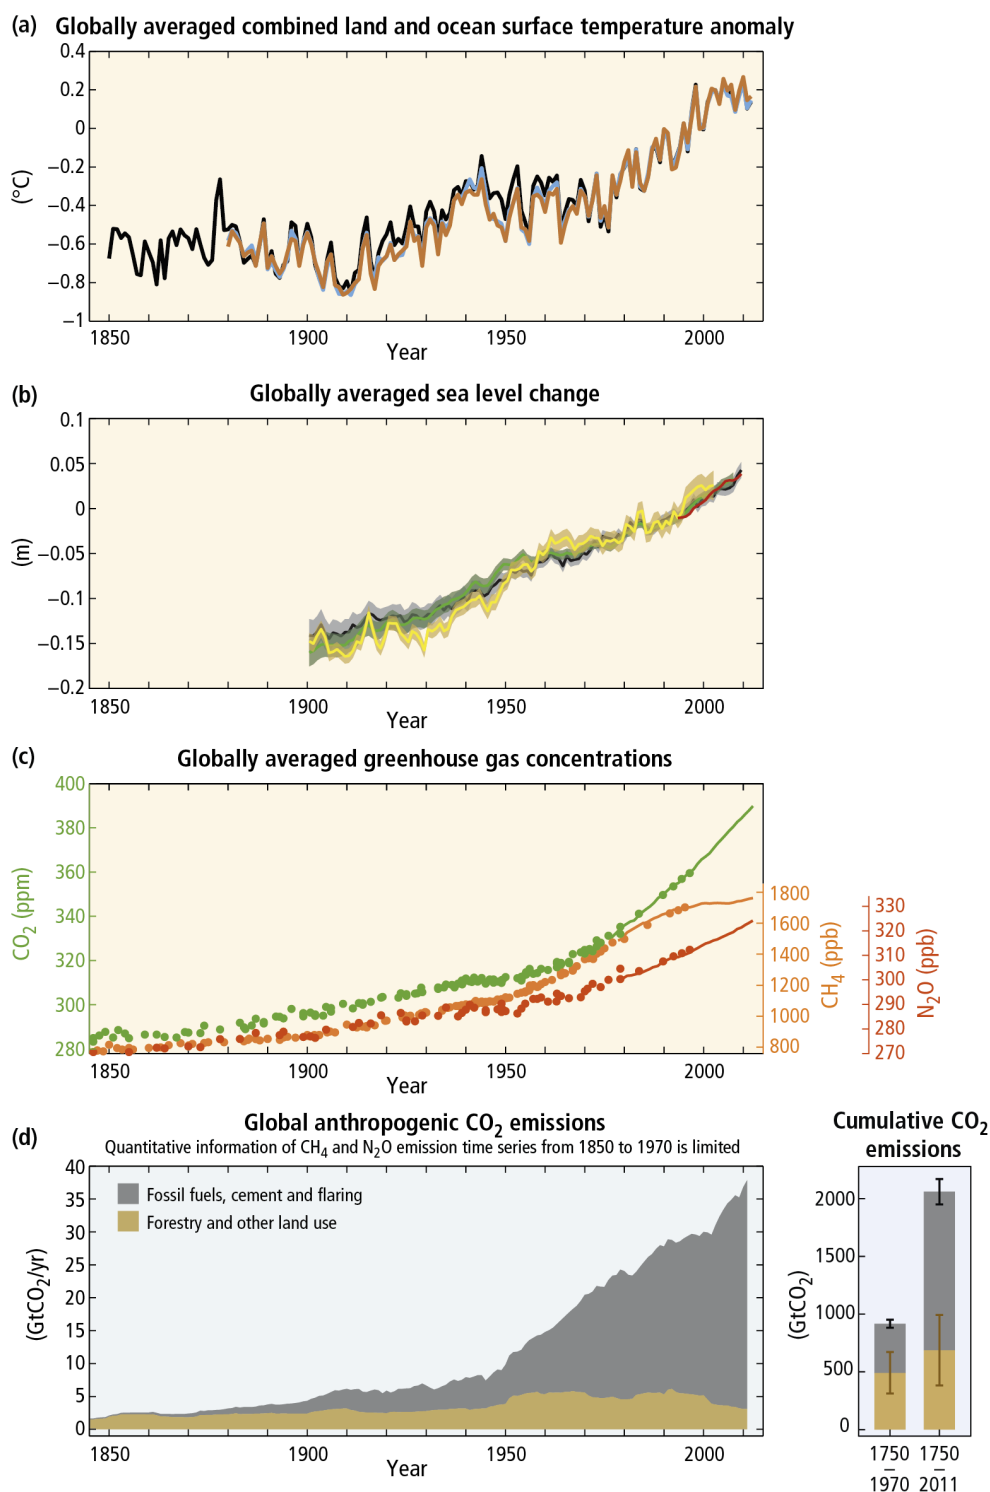


FIGURE 1.2: Figure reproduced from Figure SPM.1 on page 3 of [10]. a) Average global ocean and land surface temperature change relative to 1986-2005 mean. Datasets are HadCRUT4 version 4.1.1.0 (Black), NASA GISS (Blue), and NCDC MLOST version 3.5.2 (Orange). b) Average global change in sea level relative to 1986-2005 mean, aligned to have the same value in 1993. Black dataset reported in Church and White (2011), yellow dataset reported in Jevrejeva et al (2008), green dataset reported in Ray and Douglas (2011), all using tide gauge reconstruction. Red dataset reported in Nerem et al (2010) produced with satellite altimetry. c) Concentrations of various greenhouse gases in the atmosphere, determined from ice core measurements (dots) and direct measurements of atmospheric gas (lines). d) Global CO₂ emissions due to human activity, and cumulative CO₂ emissions since the industrial revolution.

A further side effect of excessive use of fossil fuels is modification of the Earth's atmospheric composition, and consequently the Earth's climate. Numerous independent datasets show a steady increase in global surface temperatures and globally averaged sea levels, since the beginning of the industrial age[10] (see fig 1.2a,b)). This is attributed with high confidence to increasing concentrations of atmospheric CO₂ and other greenhouse gases which are now at an 800,000 year high and still increasing rapidly, caused by the burning of vast quantities of fossil fuels[10] (see 1.2c,d)). In some instances, the effects of anthropogenic climate change can be benign, such as the opening of new shipping lanes through the Arctic ocean by declining Arctic ice[16]. However, these are fairly feeble consolation for the negative effects. These are expected to include an increase in the spread of malaria by increasing the range of mosquitoes[17], a loss of biodiversity due to rapid habitat destruction[18], damage to marine ecosystems caused by ocean acidification[19], and an increase in extreme weather events such as droughts[20], cyclones[21] and wildfires[22].

Considering the finite nature of fossil fuels, their impact on human health and effects on our environment, it seems prudent to research and develop alternative energy sources to reduce our dependence on them, and ultimately to phase out their use when feasible. It is probably impossible to do this quickly; despite non fossil fuel energy sources expanding at a faster rate than fossil fuel sources, it is projected that fossil fuels will still provide 78% of energy in 2040[9]. The largest contributors to fossil fuel consumption are transport, industry and electricity generation. Electrification may contribute significantly towards decarbonising transport and heavy industry; the recent increase in adoption of electric vehicles provides an encouraging example (from a few hundred in 2005 to 1.26 million in 2015[23]). However, electrification does not significantly reduce energy consumption, but merely shifts the onus of decarbonisation to electricity generation.

1.3 Carbon Free Electricity Generation

Of the carbon free generating technologies currently available, solar, wind and nuclear fission are the only methods which have the potential to scale to the very large generating capacities needed to replace fossil fuels[24]. However, these still have significant barriers which are inhibiting their ability to meaningfully replace fossil fuel generation. Since solar and wind provide intermittent and variable power output, they are currently unable to follow the daily demand curve and therefore require conventional fossil fuel sources as backup. Nuclear fission provides a reliable and constant output and can follow demand curves. However, in a scenario in which current nuclear technology were scaled up to provide all the world's electricity, uranium reserves would last at most 185 years[24],


and also nuclear fuel enrichment technology would have to be widely distributed to all states, increasing the likelihood of uncontrollable nuclear weapons proliferation.

Energy technology is currently advancing so rapidly that the above discussion may well be irrelevant at time of reading. The issue of finite uranium fuel may be mitigated or overcome by the development of breeder reactors which would vastly increase viable fission fuel reserves[25], and the development of cheap and scalable energy storage solutions would negate the intermittency of solar and wind. Research is proceeding in these and many other directions, but unfortunately predicting the development and deployment of new technologies is notoriously difficult (although the solutions will of course seem obvious in hindsight). Since we cannot tell today which of the many research branches will bear fruit, we must continue to explore all options.

We live an interesting age. Prosperity, security and progress are not fate, but we may hope that by our best efforts and new technologies we may accomplish and safeguard them.



1.4 Nuclear Fusion

Nuclear fusion is the joining together of light atomic nuclei into heavier nuclei. In an exothermic fusion reaction (one that releases energy), the total mass of the reactants (the 'ingredients' of the reaction) is greater than the total mass of the products. The mass which has apparently gone missing in the reaction, Δm , is released as energy according to Einstein's relation $E = \Delta mc^2$. A ready example of a fusion reactor is our own Sun, which every second burns (in the nuclear sense) 620 million tonnes of hydrogen into 616 million tonnes of helium, with the 'missing' 4 million tonnes of mass released as 3.8×10^{26} Joules of energy[26]. The proton-proton fusion chain which takes place in the Sun is in fact an extremely slow reaction, but the Sun compensates with its astronomically large mass. For terrestrial fusion we require a faster reaction, and of the available candidates, the reaction between Deuterium (2_1D) and Tritium (3_1T) has by far the highest reaction rate at achievable temperatures[27], and so this is the primary reaction of interest to fusion research.



The products of the reaction are an α particle (ie, a Helium nucleus 4_2He), a neutron n , and 17.6MeV of energy. The α particle carries 3.5MeV of energy which heats the fuel to maintain the reaction, and 14.1MeV of energy is carried out of the reactor by the

neutron to boil water, which turns a turbine and generates electrical power. Deuterium is abundant in seawater (approximately 1 part in 6400), and so the supply of this reactant is essentially unlimited. Tritium however is radioactive with a short half-life of 12.3 years, and so does not occur in nature and must be bred from Lithium. The two most abundant isotopes of Lithium in nature are ${}^7_3\text{Li}$ and ${}^6_3\text{Li}$ (92.4% and 7.6% respectively of naturally occurring Lithium), which react to neutron bombardment as shown below



Using the above reactions, Lithium is bombarded by the fusion neutrons which breeds Tritium, which is then burnt in the fusion reaction, to produce neutrons which breed more Tritium. So for practical purposes the fuel for early fusion reactors is Lithium, of which there are readily available and widely distributed reserves sufficient for millennia of total world energy consumption, which can be extended to tens of millions of years by extracting Lithium from seawater [28, 29]¹.

A fusion power plant would share all the advantages of conventional nuclear fission: supplying reliable continuous power, requiring very little land compared with solar and wind, would be scalable to meet world demand and would use a negligible quantity of fuel compared with fossil fuels. Unfortunately it would likely also share the disadvantage of high construction costs relative to fossil fuel plants, but this does not preclude fusion from being cost competitive with fission and other low carbon sources[32]. However, fusion has a number of additional crucial advantages over conventional fission. Although the quantities involved are generally small, some products of nuclear fission have long decay chains, which necessitates continually produced spent fission fuel to be safely stored for millennia[33]. The products of the fusion reaction conversely, are radiologically and chemically inert, so there is no continuous production of radioactive waste. Over the course of the fusion reactor lifetime some components will be made radioactive by the neutron bombardment, which must be handled when the reactor is decommissioned. However by careful choice of reactor materials, this small amount of waste will become

¹While D-T is clearly the most practical fusion reaction to use at the current technological level, at sufficiently high temperatures Deuterium will also fuse with other Deuterium nuclei, releasing similar amounts of energy to a D-T reaction. Conditions for D-D fusion are far more technically challenging than D-T fusion, but not prohibitively so[30], and it is possible that once D-T fusion is mastered, later generations of reactors will burn only Deuterium. Hydrogen, which contains concentrations of Deuterium of around 25 ppm[31], is by a wide margin the most abundant and widely available element in the Universe. With an abundance of fuel so incomprehensibly vast, once developed D-D fusion would be the last power source humanity would ever need, capable of providing the energy to remove all other resource limits, and sustaining human civilisations long after the last embers of the last stars.

safe to recycle and reuse on the order of decades, rather than millennia[34]. Furthermore, since fusion reactors have no need of uranium mining, weapon relevant materials or enrichment facilities, the use of a fusion power program to covertly develop nuclear weapons is highly implausible, alleviating concerns of proliferation[35]. Since the fission reaction can take place without active human involvement, and the core of a fission reactor contains many months or years of fuel at any given time, the potential for an uncontrolled or runaway reaction is inherent to fission reactors. Put simply, fission is too easy². In contrast, a terrestrial fusion reaction is incredibly fragile and requires constant intervention to maintain; the loss of any reactor system or containment breach would introduce impurities to the fusion fuel which immediately extinguishes the reaction. Furthermore, at any time the reactor never contains more than a few seconds of fuel. Fusion is therefore inherently safe, as uncontrolled or runaway nuclear fusion reactions are fundamentally impossible. The small quantities of fuel needed and extreme difficulty of using it to create a fusion reaction, also alleviate terrorism concerns relative to fission power.

In summary, a power plant based on nuclear fusion would supply continuous power, would be scalable to meet world demand, would require far less land than conventional renewables, would produce no long lived radioactive waste, would use a universally accessible and practically inexhaustible fuel, would not raise the risk of weapons proliferation, and would be inherently safe from runaway reactions.

1.5 Achieving Terrestrial Fusion

For a reaction to occur, reactants must collide with each other with sufficient kinetic energy to overcome the activation energy of the reaction, which means having a high enough temperature. Atomic nuclei are all positively charged, meaning that for two nuclei to collide they must overcome the strong electrostatic repulsion of 'like' charges. For this reason, nuclear fusion reactions require extreme temperatures to occur (this is not required in fission, since fission reactions are typically initiated by a neutron, which is not subject to the Coulomb force). The number of reactions per unit volume between two reactants A and B is given by $n_A n_B \langle \sigma \bar{v} \rangle_{AB}$, where n_A and n_B are the number densities of reactants A and B, and $\langle \sigma \bar{v} \rangle_{AB}$ is the reaction rate between reactants A and B at a given temperature. Figure 1.3 shows the reaction cross section of D-T and other possible candidate reactions as a function of incident particle energy. The plot

²So easy in fact, that in an Oklo uranium mine in Gabon, a natural fission reactor was discovered which had been assembled by an accident of geology around 1.7 billion years ago, and produced on the order of 100kW for hundreds of millennia[36].

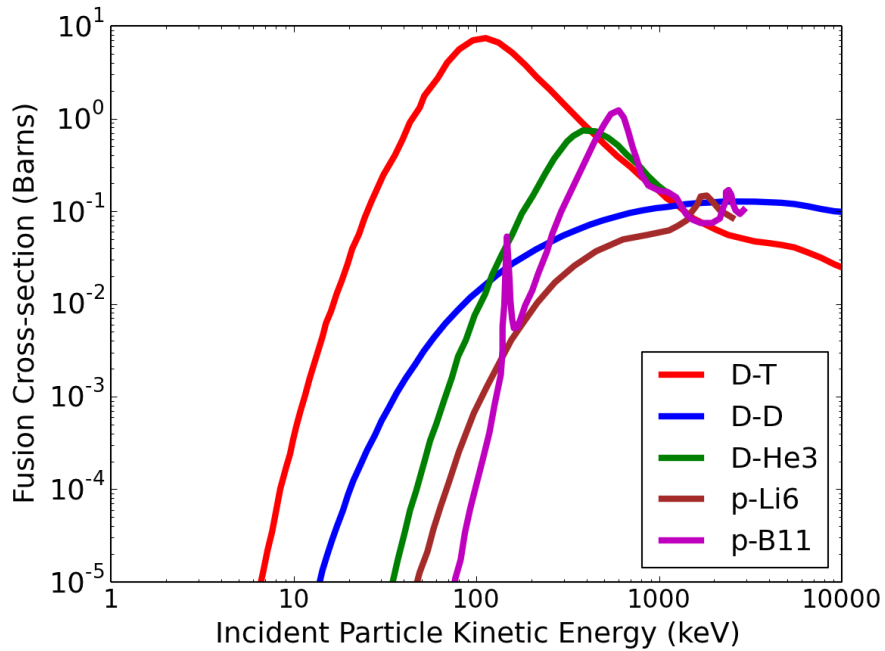


FIGURE 1.3: Cross sections of various fusion reactions with kinetic energy. Figure reproduced from [27]




FIGURE 1.4: Motion of ions (red) and electrons (green) in a magnetic field (blue). Charged particles in a magnetic field follow a helical trajectory through space, moving along magnetic field lines while gyrating around them.

shows that temperatures of at least 10 keV (≈ 120 million $^{\circ}\text{K}$) are required for non-negligible chance of reaction, and also shows why the D-T reaction is chosen: it has the largest cross section (reaction rate) at lowest temperature.

The requirement of these extreme temperatures means that conventional methods of confinement – a material container – clearly could not work, so we must look for methods of confining the fuel without touching it. There are broadly two approaches to fusion fuel confinement: inertial and magnetic. In the inertial confinement approach, small frozen capsules of D-T fuel mix are fired into a vacuum chamber, compressed and

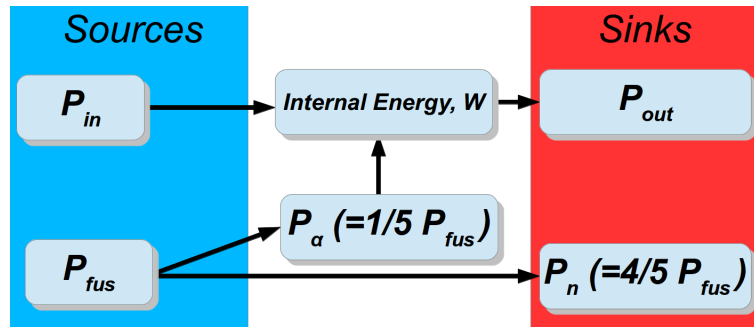


FIGURE 1.5: A simplified sketch of the power balance of a conceptual fusion reactor, from which a criterion for a viable reactor may be derived.

ignited, typically with a powerful laser, launching a thermonuclear burn wave from the pellet centre which consumes the fuel before it disassembles. This process lasts mere microseconds, and the fuel is held in place only by its own inertia, hence the name. Magnetic confinement schemes use strong magnetic fields to confine the fusion fuel, taking advantage of the fact that at fusion relevant temperatures of 10 keV, the fuel is completely ionised. All electrons are stripped from their atoms, and the gaseous fuel becomes a plasma; a mix of ions and electrons. Ionized particles in a magnetic field are subject to the Lorentz force, which strongly restricts their movement perpendicular to magnetic field lines. As sketched in figure 1.4, ions and electrons in a strong magnetic field are free to move parallel to the field, but can only follow circles perpendicular to it. From here we will restrict our discussion to magnetic approaches.

1.5.1 The Lawson Criterion

By considering simple power balance (as sketched in 1.5), we can derive a rough approximation for the conditions which must be satisfied inside a fusion reaction volume in order to extract energy from it. Some power P_{in} must be put into it to heat the plasma to fusion temperatures, power P_{out} leaves the reaction volume through radiation, particle loss and other mechanisms, and some power will be generated by fusion reactions, P_{fus} . Of the fusion power P_{fus} , P_{α} is carried by the α particles and P_n is carried by the neutrons. Since the D-T reaction has only two products, the energy of each reaction is distributed between the products according to their mass, so $P_{\alpha} = P_{fus}/5$. The fusion born neutrons have no charge, so P_n immediately leaves the reaction volume. However the α particles are charged and therefore magnetically confined, so they remain in the reaction volume and add their heat P_{α} to it. Assuming steady state, balancing power entering and leaving the reaction volume gives $P_{in} + P_{\alpha} = P_{out}$. The ratio of fusion power generated to input power is called the power amplification factor $Q = P_{fus}/P_{in}$. When $Q > 1$, the reactor is producing more fusion power than is required to heat the plasma, a condition known as *break-even*. When the fusion power begins to dominate

over the input power, Q becomes large and the plasma is then said to be *burning*. When the fusion power becomes sufficient that the input power may be switched off entirely so $P_\alpha = P_{out}$, $Q \rightarrow \infty$ and the plasma is then said to be *ignited*. Although a burning plasma with moderate $Q \approx 20 - 40$ would be sufficient for a power plant[27], in the following derivation we will set ignition as a minimum condition for a viable fusion reactor, ie, $P_\alpha > P_{out}$.

Fusion α power is the integral over the plasma volume of the reaction rate multiplied by the energy per reaction

$$P_\alpha = \int_V n_D n_T \langle \sigma \bar{v} \rangle E_\alpha dV \quad (1.4)$$

where E_α is the energy of the fusion born α particle. The optimal mix of D-T for fusion is 50:50, ie, $n_D = n_T = n/2$. Substituting this into the above yields

$$P_\alpha = \int_V \frac{1}{4} n^2 \langle \sigma \bar{v} \rangle E_\alpha dV \quad (1.5)$$

The power leaving the reactor $P_{out} = W/\tau_E$, where we define W as the total internal energy, and τ_E as the energy confinement time. The total internal energy can be expressed as

$$W = \int_V \frac{3}{2} (n_e T_e + n_D T_D + n_T T_T) dV \quad (1.6)$$

Assuming thermal equilibrium gives us $T_e = T_T = T_D$, and assuming 50:50 D-T mix gives us $n_D = n_T = n/2$. We may also assume that the plasma is quasi-neutral, ie, the number of positive charges in the plasma is approximately equal to the number of negative charges, since if this were not the case the resulting electric field would immediately correct it. Assuming quasi-neutrality yields $n_e = n$, where n is both the electron number density and ion number density. Making these substitutions above yields

$$W = \int_V 3nT dV \quad (1.7)$$

Applying the ignition condition $P_\alpha > P_{out}$ yields

$$\int_V 1/4 n^2 \langle \sigma \bar{v} \rangle E_\alpha dV > 1/\tau_E \int_V 3nT dV \quad (1.8)$$

Since this is only an approximate calculation, we will restrict the analysis to a small unit of volume, in which density and temperature are approximately constant in space. This allows us to drop the volume integrals to yield

$$n \langle \sigma \bar{v} \rangle E_\alpha > 12T/\tau_E \quad (1.9)$$




FIGURE 1.6: Sketch of the principles of the tokamak plasma confinement device. The toroidal field (blue arrow) is generated by the large D-shaped toroidal field coils (blue coils). A large changing current is driven through the central solenoid (green coils), which drives the plasma current through the plasma which generates the poloidal field. The toroidal and poloidal magnetic fields add together to create a helical field. Figure from [39]

Now in the temperature range of interest, $\langle \sigma \bar{v} \rangle = CT^2 m^3 s^{-1}$ with T in keV and constant $C \approx 1.1 \times 10^{-24} keV^{-1}$ [37]. Using this substitution the above rearranges to

$$n\tau_E T > 3 \times 10^{21} keV m^{-3} s \quad (1.10)$$

Slightly different figures are found in the literature depending on what assumptions are made, but all are within roughly a factor of 2 of this [27, 37, 38]. This is the Lawson criterion, which gives us a useful measure of how close a particular confinement scheme comes to being an energy producing fusion reactor, via the 'triple product' metric $n\tau_E T$.

1.6 The Tokamak

Early investigations into magnetically confined fusion produced a diverse zoo of magnetic field configurations, but none showed the potential to scale to a viable power plant




FIGURE 1.7: The fusion triple product of successive tokamaks has increased at a faster rate than Moore's law, and now is within sight of reactor conditions. Figure from [40]

reactor[27, 41]. That is until results from the T-3 tokamak were presented in 1968[42], showing far superior performance to other contemporary confinement schemes, resulting in the tokamak becoming the main focus of magnetic fusion energy research globally.

While magnetic fields strongly restrict particle motion perpendicular to the field, particles are not constrained in the direction parallel to the field as shown in figure 1.4. To prevent particles from leaving confinement by streaming freely along field lines, toroidal magnetic confinement schemes adopt a 'doughnut shaped' configuration in which field lines form closed loops, so particles freely streaming along field lines do not leave confinement, but loop around the machine like cars on a race track. A purely toroidal field would cause particles to drift out of confinement, and so is insufficient to confine a plasma on its own. However by adding a 'twist' to the field with a magnetic field in the poloidal direction, this drift is negated[37]. Thus a toroidal confinement scheme requires both a toroidal and a smaller poloidal field.

A tokamak consists primarily of a toroidal vacuum chamber with set of magnetic coils arranged around it (the word tokamak derives from the Russian words for "toroidal chamber with magnetic coils"), and a large solenoid filling the hole in the centre of the torus, see figure 1.6. In operation, first the magnetic coils are activated to generate the




FIGURE 1.8: It is unavoidable that at some point the plasma must touch some part of the machine; the component which is used for this is the exhaust system. On the left is a sketch of the original limiter exhaust system used in early tokamaks, a protruding carbon block extended from the vessel wall for the plasma to terminate against. This was effective at the time, but it later became apparent that the limiter was causing too many impurities to enter the plasma, and limiting performance. On the right is a sketch of the divertor exhaust system, so called because the exhaust is 'diverted' away from the main plasma. The use of a divertor decreased the usable plasma volume, but greatly reduced impurity build up, allowing for greatly improved plasma performance.

toroidal magnetic field, and the gaseous D-T fuel is pumped into the vacuum chamber. A continuously ramping current is driven through the central solenoid, inducing an electric field around the torus, which both ionises the fuel into a plasma and drives a large plasma current in the toroidal direction. The induced plasma current produces the poloidal magnetic field required to negate outward particle drifts. The plasma is heated to fusion temperatures by a combination of ohmic dissipation from the plasma current (as in a kettle filament), radio frequency resonance heating (as in a microwave oven), and injection of high energy neutral particles which heat the plasma by collisions.

Since the early success of T-3, the fusion triple product of successive tokamaks has doubled roughly every 1.8 years, exceeding the growth rate over the same period of the famous Moore's Law for the number of transistors in new microprocessors[43], as summarised in figure 1.7. Over this period many crucial advances were made in magnetic fusion technology. To extend pulse lengths, copper magnetic coils which rapidly overheated were replaced with superconducting coils, and the 'limiter' plasma exhaust system was replaced with the divertor system sketched in figure 1.8, which allowed access to improved plasma regimes. By the 1990s tokamaks were approaching the regime of break-even and ignition. In 1996 JT-60U, although not operating with D-T fuel, achieved a fusion triple product of $n_i T_i \tau_E = 1.5 \times 10^{21} keV m^3 s$ (where n_i , T_i are the ion

density and temperature in the core of the tokamak)[44]. The next year, experiments using D-T fuel on the Joint European Torus produced a peak fusion power of 16.1MW for 25.7MW of input power, achieving a triple product of $8.7 \times 10^{20} keVm^{-3}s$ and a measured fusion amplification factor of $Q=0.63$ [45].

Tokamaks are on the brink of proving themselves capable of generating net power, but to cross the threshold into the reactor relevant burning plasma regime, a larger machine is needed: the ITER tokamak. The purpose of ITER, currently under construction in France, is to 'demonstrate the scientific and technological feasibility of fusion energy for peaceful purposes'[46]. Specifically, ITER is designed to produce ≈ 500 MW of fusion power with ≈ 50 MW of input heating power (a fusion gain of $Q \geq 10$) for 300-500 seconds, and to operate in steady state (≈ 3000 s) with $Q \geq 5$ [47]. To spread the cost and distribute fusion expertise widely, the ITER project is an international scientific collaboration on a scale never before attempted. The ITER members are China, India, the European Union, Russia, Japan, Korea and the United States, between them representing over half of the world population and 85% of global GDP. The scientific knowledge gained from ITER will be used to design and operate the first demonstration fusion power plant, DEMO, which will demonstrate tritium self sufficiency, reliability, and send power to the grid.

On the success of ITER depends the fulfilment of over 70 years of fusion research and development, the aspiration of humanity for a clean, safe and unlimited energy source, and the means to forever consign energy scarcity to history.

1.7 Edge Localised Modes

In 1982 experiments on the ASDEX tokamak in Garching yielded a surprising breakthrough[48]. Following the installation of the divertor exhaust system, it was found that once a certain input power was exceeded the plasma would spontaneously enter a regime of much higher confinement, now known as H-mode. In H-mode, a transport barrier forms at the plasma edge which suppresses turbulent transport in that region, resulting in approximately a doubling of the confinement time. It also creates a region of steep pressure gradient called 'the pedestal', which raises the plasma pressure globally as sketched in figure 1.9. Although the current theory is not firmly established, it is thought that the transport barrier is caused by the appearance of highly sheared 'zonal flows' at the plasma edge, which stretch and shred the turbulent eddies which constitute the dominant transport mechanism[49]. H-mode has now been replicated on all other major

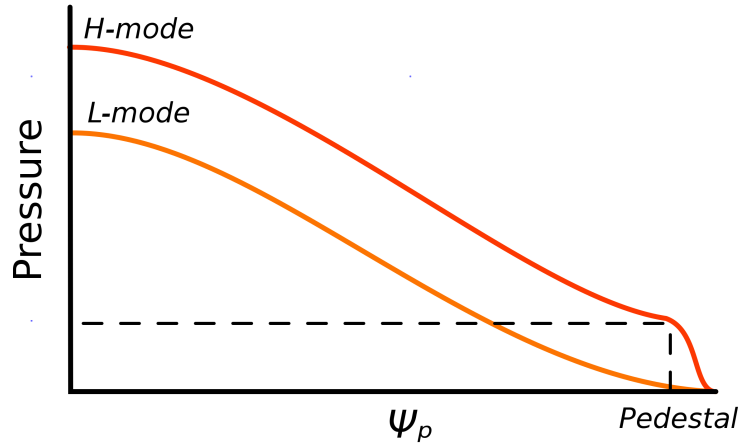


FIGURE 1.9: H-mode (high confinement), which tokamak plasmas enter spontaneously once an input power threshold is exceeded, is characterised by a steep pressure gradient at the edge caused by the appearance of an edge transport barrier. The pedestal boosts the pressure globally, resulting in improved performance relative to L-mode.

tokamaks, and forms the baseline scenario for ITER. While H-mode certainly gives superior performance, it also causes a new plasma instability to appear at the plasma edge, the Edge Localised Mode (ELM).

ELMs occur in a repetitive cycle in H-mode plasmas at rates of 10-100Hz[50], consisting of a steady increase of the pressure pedestal height until a stability boundary is exceeded, followed by a collapse of the pedestal back to a stable state. By collapsing the edge pressure pedestal, the ELM causes the rapid ejection of particles and stored energy ΔW_{ELM} from the plasma. This degrades confinement slightly, but ELMs also perform the useful functions of expelling impurities from the plasma and preventing undesirable density build up, which allows ELMy H-mode to remain in time averaged steady state. The ELM energy ΔW_{ELM} is deposited at the divertor over a very short timescale ($\approx 100 \mu s$) which results in a high transient heat load at the divertor. Although the ELMs of current tokamaks are not large enough to cause any material damage, extrapolations from current machines to ITER indicate that in ELMy H-mode each ELM may contain up to 20MJ of energy, far larger than the $\Delta W_{ELM} \leq 0.66MJ$ limit for material damage at the divertor[51]. Material damage caused by uncontrolled ELMs on ITER would significantly reduce the lifetime of the plasma facing components, impair the endurance and performance of the machine, and possibly threaten the ITER scientific objectives[52, 53]. To avoid this risk, strategies must be developed which either eliminate ELMs, or reduce them in size (energy content) by a factor of at least 30 ($\approx 20MJ/0.66MJ$).

Techniques for controlling ELMs can be divided into passive and active techniques. Passive techniques comprise developing plasma regimes which retain the high confinement of H-mode but without ELMs, and which also have some continuous particle transport

mechanism to prevent density and impurity accumulation. Several such regimes have been identified (I mode[54] and (quiescent) QH mode[55, 56] for example). Currently however these regimes are inflexible and require far more study before it can be determined whether they may be extrapolated to ITER[51], so the standard operating scenario for ITER remains ELMy H-mode[57], which will require active techniques to control ELMs. Active ELM control consists of actively intervening in an ELMy H-mode plasma, in order to prevent ELMs from triggering or reduce their size. During the recovery phase of the ELM cycle as the pressure pedestal height increases, the available stored plasma energy increases accordingly. From this it follows, and is confirmed by experiment[58], that the ELM energy ΔW_{ELM} is roughly proportional to the time since the previous ELM, and so is inversely proportional to the ELM frequency. Specifically[51],

$$\frac{\Delta W_{ELM} \times f_{ELM}}{P_{out}} = 0.2 - 0.4 \quad (1.11)$$

where f_{ELM} is the ELM frequency. This relation shows that the size of individual ELMs may be decreased by intentionally triggering ELMs at a faster rate than their natural frequency, f_{nat} . A factor of 30 increase in the ELM frequency would in principle lead to the factor of 30 decrease in ELM size required to avoid material damage. Two active ELM control systems are to be installed on ITER, a pellet pacing system, and a set of resonant magnetic perturbation (RMP) coils[59].

Pellet Pacing

Injecting pellets of frozen Deuterium was originally developed as a method for fuelling the plasma, but it has been also demonstrated that injecting fuel pellets into ELMy H-mode plasmas can trigger ELMs prematurely, and that the ELM frequency can be synchronised to the pellet launch frequency[60]. Increasing the pellet launch frequency thus increases the ELM frequency, decreasing the ELM size. Pellet pacing has been demonstrated robustly on several tokamaks[61], and on the DIII-D tokamak ELM frequency increases by a factor of over 10 have been achieved with a corresponding 10 fold decrease in the ELM energy[62]. However, uncertainties remain in the effectiveness of the ITER pellet pacing system. Firstly it is not known with certainty whether the relation described in equation 1.11 on which ELM pacing relies, will hold at the ITER scale. Secondly, it has been observed that the area of the divertor over which the ELM energy is distributed, A_{ELM} , increases with ELM energy ΔW_{ELM} [51]. This is advantageous for natural uncontrolled ELMs as it reduces the peak heat flux for large ELMs, however by the same token, it also increases the peak heat flux for small ELMs. If this effect is stronger than expected in ITER, it may reduce the effectiveness of ELM pacing[51].

Most importantly though, ELM pellet pacing has not demonstrated the factor of 30 reduction in ELM size which ITER will require.

Resonant Magnetic Perturbations

It has also been demonstrated that the application of small resonant magnetic perturbations (RMPs) at the plasma edge can increase the frequency of ELMs above the natural frequency, known as ELM mitigation, or even cause ELMs to disappear entirely, known as ELM suppression. While mitigation suffers the same uncertainties in extrapolation to ITER as pellet pacing, these particular uncertainties are not relevant for ELM suppression, which more than satisfies the required factor of 30 reduction in ELM size. ELM suppression by RMPs therefore currently appears the most attractive option for ITER baseline operation, and as such, a flexible set of RMP coils has been added to the ITER design. However, while numerous working theories exist, a robust, predictive and machine independent theory of ELM suppression or mitigation which may be used to assess the ITER ELM coils, is currently lacking.

1.8 Outline

The work presented in this thesis focuses on using simulations of the plasma response to applied RMPs, to further develop current working theories of ELM mitigation and suppression. Simulation results are combined with experimental data to optimise the RMP configuration, this optimisation is tested rigorously on ASDEX Upgrade, and extended to the ITER RMP set.

Chapter 2 introduces and explains background theories and concepts referred to in subsequent chapters, in particular the magnetohydrodynamic description of a plasma, and the fundamentals of plasma equilibrium and stability. Observations of ELM control by RMPs and other RMP effects are summarised, and current prevailing theories of RMP ELM control are explained. The linearised MHD model of MARS-F, used throughout this work, is introduced.

In Chapter 3, the results of an investigation of the plasma response to applied RMPs in an ASDEX Upgrade plasma are presented. The vacuum field is carefully benchmarked against the ERGOS code, and the plasma response to an applied $n = 2$ RMP and $n = 6$ sideband is computed and analysed, with particular emphasis on the pitch aligned components of the field, and the amplified peeling response. Using a contrived single poloidal harmonic boundary condition, it is demonstrated that the pitch aligned

components may be driven by the peeling response via poloidal harmonic coupling. In collaboration with researchers at IPP Garching and Consorzio RFX, the MARS-F computed plasma response is partially validated against experimental measurements of the magnetic perturbation and plasma displacement on ASDEX Upgrade plasmas.

Chapter 4 presents the development and testing of a coil phase optimisation scheme, designed to maximise ELM mitigation for a given plasma equilibrium. By scaling a reference equilibrium in pressure and current, a set of equilibria is produced which spans the ASDEX Upgrade experimental parameter space of edge safety factor and pressure. Using MARS-F, the optimal coil phase at each point is computed, and the optimal coil phase is then parametrised as a function of edge safety factor, pressure and dominant toroidal harmonic of the RMP. This parametrisation is then benchmarked against rigorous MARS-F computations of the optimal coil phase, and validated against experimental measurements of the optimal coil phase from ASDEX Upgrade coil phase scan experiments. The same process is then used to derive a parametrisation of the optimal coil phase for the ITER coil set.

Chapter 5 summarises the results of this thesis in context, and discusses their relevance and potential impact.

Chapter 2

Background

2.1 Magnetohydrodynamics

Many models are used to describe a magnetised plasma. The first principles description of a magnetised plasma is known as the kinetic description, which deals with the evolution of the particle distribution function of particle species j , $f_j(\mathbf{x}, \mathbf{v}, t)$. The distribution function $f_j(\mathbf{x}, \mathbf{v}, t)$ describes the particle density in a 6 dimensional space of position $\mathbf{x} = (x, y, z)$ and velocity $\mathbf{v} = (v_x, v_y, v_z)$, such that the number of particles in a $d\mathbf{x}d\mathbf{v}$ sized box of phase space is $f_j(\mathbf{x}, \mathbf{v}, t)d\mathbf{x}d\mathbf{v}$. Calling the change in f_j due to particle collisions $C(f_j)$ (ie, $\frac{Df_j}{Dt} = C(f_j)$), the Boltzmann equation may be quickly derived (see below) which describes the time evolution of the distribution function

$$\frac{Df_j}{Dt} = \frac{\partial f_j}{\partial t} + \frac{\partial \mathbf{x}}{\partial t} \cdot \frac{\partial f_j}{\partial \mathbf{x}} + \frac{\partial \mathbf{v}}{\partial t} \cdot \frac{\partial f_j}{\partial \mathbf{v}} = \frac{\partial f_j}{\partial t} + \mathbf{v} \cdot \frac{\partial f_j}{\partial \mathbf{x}} + \mathbf{a} \cdot \frac{\partial f_j}{\partial \mathbf{v}} \quad (2.1)$$

$$= \frac{\partial f_j}{\partial t} + \mathbf{v} \cdot \frac{\partial f_j}{\partial \mathbf{x}} + \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_j}{\partial \mathbf{v}} = C(f_j) \quad (2.2)$$

where the chain rule was used to expand the derivative $\frac{D}{Dt}$, and the Lorentz force substituted for the acceleration $\mathbf{a} = \frac{\mathbf{F}}{m}$. While physically accurate, we don't necessarily need all the information contained in f_j , and f_j being a function of 7 variables makes this description very difficult to solve for macroscopic plasma behaviour. In order to make progress towards a simple macroscopic description, we will use the kinetic model as a basis to derive a more appropriate model.

Single fluid Magnetohydrodynamics (MHD) combines Maxwell's equations of electromagnetism with the equations of fluid dynamics, to describe the plasma as an electrically

conducting fluid evolving in response to self-generated and externally applied magnetic and electric fields.

The equations of MHD can be derived from the kinetic description by taking 'moments' of equation 2.2, where taking an n^{th} order moment of arbitrary operator \mathbf{Q} is defined as $\int \mathbf{v}^n \mathbf{Q} d\mathbf{v}$. The zeroth, first and second moments lead to fluid equations expressing conservation of mass, momentum and energy respectively. However, the resulting fluid equation from each moment contains a term which depends on the next highest moment, so we may never produce a closed system of equations with this approach. Therefore after taking the three lowest moments we add an approximation to relate pressure and density, the 'adiabatic approximation', which assumes that the entropy convected with the plasma is conserved. Completing this procedure for both ions and electrons yields separate fluid equations for each, which may be summed to form single fluid equations. The resulting set of single fluid MHD equations is

$$\frac{d\rho}{dt} + \rho \nabla \cdot \mathbf{v} = 0 \quad (2.3)$$

$$\frac{dp}{dt} + \gamma p \nabla \cdot \mathbf{v} = 0 \quad (2.4)$$

$$\frac{d\mathbf{B}}{dt} = (\mathbf{B} \cdot \nabla) \mathbf{v} - (\nabla \cdot \mathbf{v}) \mathbf{B} - \nabla \times (\eta \mathbf{J}) \quad (2.5)$$

$$\rho \frac{d\mathbf{v}}{dt} = -\nabla p + \mathbf{J} \times \mathbf{B} \quad (2.6)$$

$$\mu_0 \mathbf{J} = \nabla \times \mathbf{B} \quad (2.7)$$

where we have used the Lagrangian derivative $\frac{d}{dt} = \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla$

In the above, \mathbf{B} is the magnetic field, p is the plasma pressure, \mathbf{v} is the fluid velocity, ρ is the mass density, γ is the ratio of specific heats, η is the plasma resistivity, and \mathbf{J} is the current density. Equation 2.3 is the continuity equation, describing that the rate of change of mass in a volume element is equal to the difference between the mass entering and leaving the element. Equation 2.4 is the adiabatic equation of state, which describes that the entropy convected with the fluid is constant (ie, $\frac{d}{dt} \left(\frac{p}{\rho^\gamma} \right) = 0$). Equation 2.5 is Faraday's law of electromagnetic induction $\frac{d\mathbf{B}}{dt} = -\nabla \times \mathbf{E}$, combined with Ohm's law for a conducting fluid $\mathbf{E} = \eta \mathbf{J} - \mathbf{v} \times \mathbf{B}$. Equation 2.6 is the equation of motion for a volume element, expressing the acceleration of a volume element due to a pressure gradient, and the Lorentz electromagnetic force on moving charges. Finally equation 2.7 is Amperes law describing the magnetic field induced by a flowing current. Single fluid MHD is widely used to derive and investigate the stability properties of macroscopic magnetised plasma equilibria, which makes it the workhorse model with which plasma geometries and instabilities are described and studied. Plasma confinement schemes which do not




FIGURE 2.1: The flux label ψ_p of a flux surface is defined as the integrated magnetic flux intersecting with the surface S (sketched above in red), which is the toroidal revolution of a curve connecting the magnetic axis with the flux surface. Because lines of B lie in flux surfaces, the poloidal flux intersecting the surface S is independent of the path the curve takes from the magnetic axis to the flux surface. Therefore ψ_p uniquely labels each flux surface.

form stable MHD equilibria will typically only persist for a matter of microseconds before being terminated by some instability, and so are generally considered unsuitable for magnetic fusion.

2.2 MHD Equilibrium

Starting with the equations of MHD 2.3–2.7, and specifying that we search for steady state ($\partial/\partial t = 0$) equilibria with no plasma flow ($\mathbf{v} = 0$), the momentum equation 2.6 immediately reduces to the equation of MHD equilibrium, which is general for any plasma geometry.

$$\mathbf{J} \times \mathbf{B} = \nabla p \quad (2.8)$$

Note that with equation 2.7, \mathbf{J} is eliminated from 2.8, so the equation describes how a pressure gradient can be balanced by a magnetic field, which is the basis of magnetic confinement of plasmas. From the above trivially follows $\mathbf{J} \cdot \nabla p = 0$ and $\mathbf{B} \cdot \nabla p = 0$, which tells us that in MHD equilibria, lines of magnetic field and current must lie on

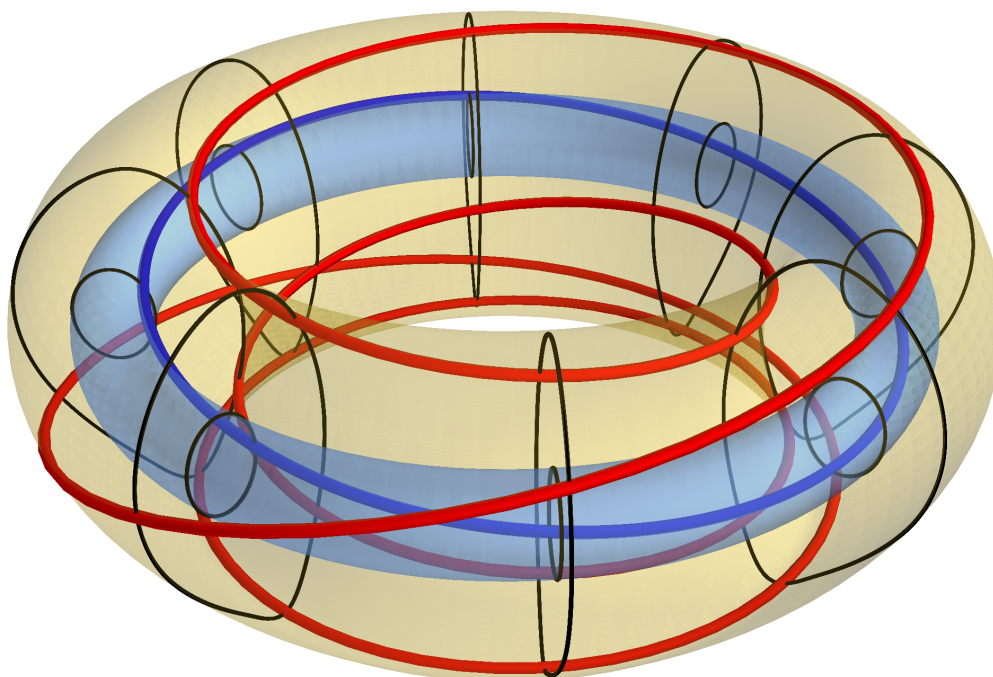


FIGURE 2.2: The 'safety factor' q is the number of toroidal rotations a magnetic field line completes for each poloidal rotation. The above sketches two field lines on two flux surfaces. The blue field line on the blue (inner) surface rotates once toroidally for each rotation poloidally, so $q = 1$. The red field line on the orange (outer) surface rotates 4 times toroidally for each poloidal rotation, so $q = 4$.

surfaces of constant pressure. In tokamaks, we call these surfaces of constant pressure flux surfaces¹.

Since the plasma cross section is not necessarily circular, simply using the minor radius as a radial coordinate is not sufficient. However since the pressure and numerous other quantities are constant on flux surfaces, it makes more sense to use some label of each flux surface as our radial coordinate. By convention, the poloidal magnetic flux ψ_p is usually used to label these flux surfaces, defined as

¹It is interesting to note, that according to equation 2.8, in order for the pressure gradient to be balanced by the magnetic field, both B and J , which lie on flux surfaces, must be non-zero everywhere on each flux surface. According to the amusingly named 'hairy ball theorem' of topology, the only smooth surface for which this is possible is a torus. All other shapes of magnetic flux surface will somewhere have a point where the confining field is zero, for the same reason that if you try to comb the hair on a tennis ball flat, there will always be a point where it sticks up. Therefore only toroidal flux surfaces may balance the pressure everywhere with no 'leaky points'.

$$\psi_p = \int \underline{B} \cdot d\underline{S} \quad (2.9)$$

The flux label ψ_p is the total poloidal flux intersecting a surface S , where S is comprised of a curve connecting the flux surface to the magnetic axis, and a curve which undergoes one toroidal rotation, as sketched in figure 2.1. Therefore, ψ_p serves as a generalised coordinate for the minor radius. Some conventions (and this work) also use s as the generalised coordinate, defined as $s = \sqrt{\psi_N}$ where $\psi_N = (\psi_p - \psi_0)/(\psi_a - \psi_0)$ is the normalised poloidal magnetic flux, in which ψ_0 and ψ_a are the poloidal magnetic flux ψ_p evaluated at the magnetic axis and plasma edge respectively. This definition is convenient, since it ensures that $s = 0$ at the magnetic axis, and $s = 1$ at the plasma edge. Variables which are constant on flux surfaces are called flux functions, since they depend only on ψ_p . The plasma pressure and safety factor q are both flux functions, where the safety factor q is defined as

$$q = \frac{1}{2\pi} \oint \frac{1}{R} \frac{B_\phi}{B_p} ds \approx \frac{r}{R_0} \frac{B_\phi}{B_p} \quad (2.10)$$

and describes how 'twisted' the magnetic field lines are, or their pitch relative to the toroidal direction, as sketched in figure 2.2. The safety factor q is an important concept for understanding plasma stability and interactions of RMPs with plasma equilibria, and will be referred to frequently in this work.

2.3 Plasma Stability

Having defined the plasma equilibrium, we may now investigate its stability; an MHD unstable plasma equilibrium will not last long. Generally to do this the energy principle is invoked, ie, asking for an small arbitrary perturbation to the equilibrium, does the equilibrium lose or gain energy? If it gains energy, then the perturbation has added energy to the equilibrium, and therefore the perturbation shrinks and is stable. However if the equilibrium loses energy, then energy is transferred from the equilibrium to the perturbation, and therefore the perturbation grows. If this is the case for any perturbation, then the equilibrium is unstable.

Starting at the MHD equations, the following equation can be derived for the change in plasma internal energy δW_p (ie, change in energy of the equilibrium) due to an arbitrary perturbation ξ .

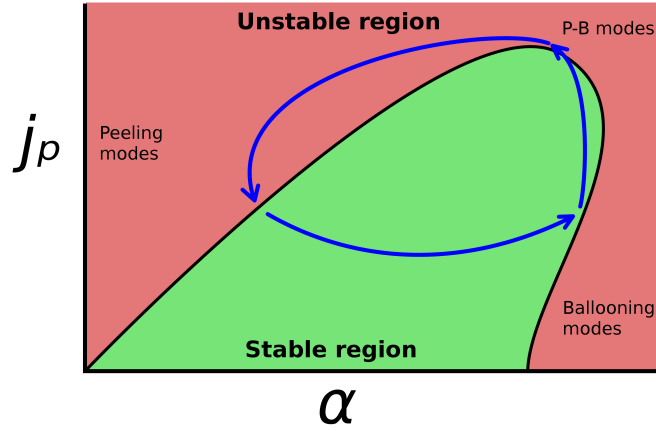


FIGURE 2.3: This sketch shows the stability space of the peeling-ballooning modes which drive ELMs, and a typical ELM cycle which the edge plasma undergoes in ELMy H-mode[50]. P-B modes have a triangular stability space in terms of the normalised edge current density $j_p = Rqsj_{||}/B$, and normalised edge pressure gradient $\alpha = -(2Rq^2\mu_0/B^2)(dp/dr)$. Starting in the stable region, the edge transport barrier causes the pressure pedestal to increase as more heating energy is injected. The pressure gradient driven bootstrap current then rises until the stability limit is exceeded, precipitating an ELM which rapidly collapses the pressure pedestal until the plasma re-enters the stable region, and the cycle begins again.

$$\begin{aligned}
\delta W_p = \frac{1}{2} \int_p & \frac{|B_1|^2}{\mu_0} && \text{Field-line bending} \\
& + \frac{B^2}{\mu_0} |\nabla \cdot \xi_{\perp} + 2\xi_{\perp} \cdot \kappa|^2 && \text{Magnetic compression} \\
& + \gamma p_0 |\nabla \cdot \xi|^2 && \text{Plasma compression} \\
& - 2(\xi_{\perp} \cdot \nabla p)(\kappa \cdot \xi_{\perp}^*) && \text{Pressure term} \\
& - B_1 \cdot (\xi_{\perp} \times b) j_{||} dV && \text{Parallel current term}
\end{aligned} \tag{2.11}$$

In the above, B_1 is the perturbed magnetic field, B the equilibrium magnetic field, κ is the field line curvature, p_0 is the pressure, γ is the ratio of specific heats, and $j_{||}$ is the component of the perturbed current parallel to the magnetic field. For any perturbation ξ , if the associated change of potential energy δW_p is negative (ie, some energy has been transferred from the equilibrium to the perturbation), then that system is unstable. Conversely if δW_p is positive, then the system is stable. It can be seen in equation 2.11, that the first 3 terms are always positive, and so are always stabilising. It always requires energy to bend magnetic field lines, to compress magnetic field lines, or to compress plasma.

However, the last 2 terms, the pressure term and parallel current term, can be either stabilising or destabilising. The pressure term, which contains a factor of ∇p , expresses

how strong pressure gradients such as those found at the edge of an H-mode plasma, can drive the plasma unstable. A ballooning instability is a pressure driven instability, so called because the plasma 'balloons' outwards on the outboard side. Neoclassical effects can also cause large pressure gradients to drive currents, namely, the bootstrap current. The parallel current term, which contains a factor of j_{\parallel} , expresses how high current densities can also drive the plasma unstable. This term can also be large at the edge of an H-mode plasma, where the steep pressure gradients drive large bootstrap currents. These currents can drive peeling instabilities, so called because the outer plasma surface 'peels off' and is ejected.

Peeling and ballooning modes can couple together as 'peeling-ballooning' (P-B) modes. It is thought that peeling-ballooning modes are the MHD instability responsible for triggering explosive ELMs, as evidenced by numerous observations of ELMs occurring near the computed stability boundary for peeling-ballooning modes[50]. In an ELM cycle, the pressure pedestal height slowly builds up, which also increases the edge currents which are driven by strong pressure gradients. With reference to the P-B stability plot in figure 2.3, this moves the plasma equilibrium towards the peeling-ballooning stability boundary, where the destabilising influence of the current and pressure gradient exceeds the stabilising factors, and an ELM is precipitated. This results in a partial collapse of the edge pressure pedestal, moving the plasma back within the P-B stability boundary for the cycle to begin again.

2.4 RMP ELM Control Observations

Since the first observation of ELM suppression by RMPs over a decade ago on DIII-D[64], ELM suppression and mitigation has been replicated on numerous other machines. ELMs are commonly observed as spikes in D_{α} light emission and divertor current, as in figure 2.4 which shows parameter traces and divertor current traces from ASDEX Upgrade experiments in which mitigation and suppression were achieved. The operational space in which ELM mitigation and suppression are achieved is conventionally expressed in terms of pedestal density n_e normalised to the Greenwald density n_{eGW} , and electron pedestal collisionality ν_e^* which is calculated as

$$\nu_e^* = 6.921 \times 10^{-18} \frac{Rq n_e Z_{\text{eff}} \ln \Lambda_e}{\epsilon^{3/2} T_e^2} \quad (2.12)$$

where the Coulomb logarithm $\ln \Lambda_e = 31.3 - \ln(\sqrt{n_e}/T_e)$, R is the plasma major radius in meters, q is the safety factor at the top of the pedestal, ϵ is the inverse aspect ratio (a/R , where a is the plasma minor radius), Z_{eff} is the effective ion charge (ie, including

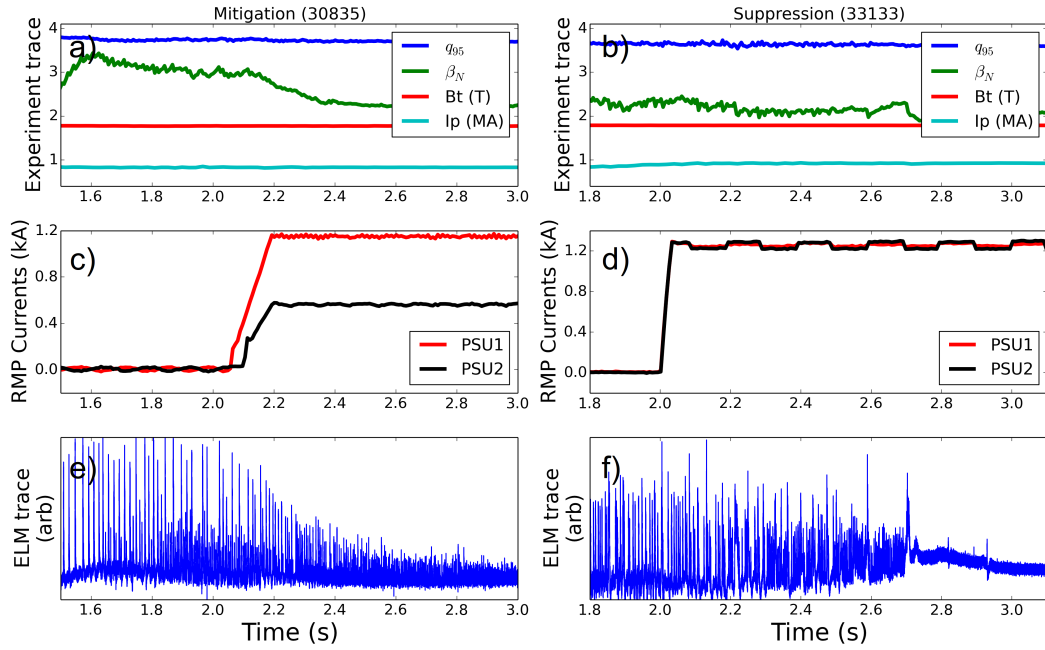


FIGURE 2.4: a,b) Experimental traces of edge safety factor (blue), normalised pressure (green), toroidal magnetic field (red) and plasma current (cyan), from two ASDEX Upgrade plasma discharges. c,d) Currents of the two power supply units of the RMP coils. e,f) The divertor current trace, showing ELM mitigation and suppression respectively, which occur after the RMP coils are energised.

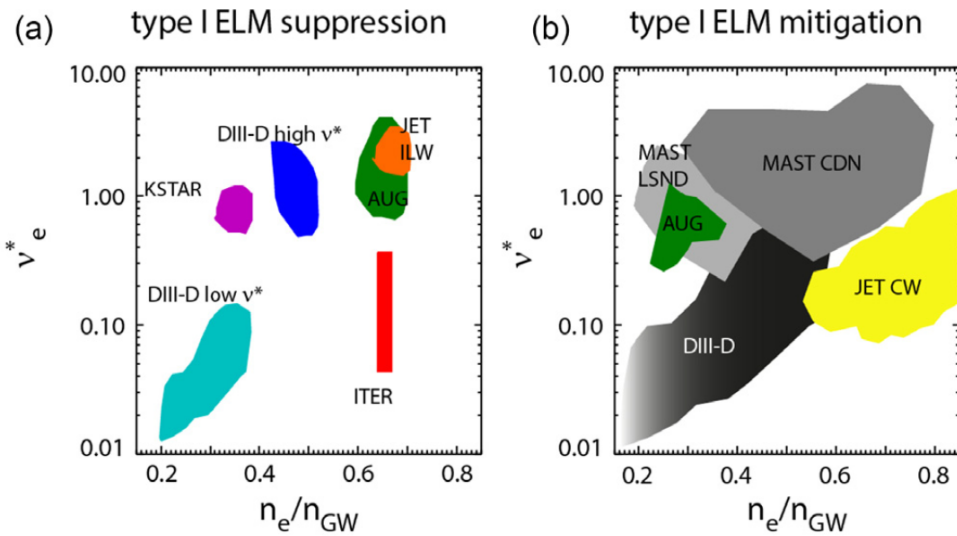


FIGURE 2.5: Parameter space map showing the regions where ELM suppression and mitigation have been achieved, and also the space in which ITER will operate. The mitigation space is far wider than the suppression space, indicating that mitigation is easier to achieve. Figure reproduced from [63].

higher Z impurities), n_e is the electron density at the top of the pedestal, and T_e is the electron temperature at the top of the pedestal. The collisionality and normalised density concisely express the major differences in plasma pedestal regimes between current machines and ITER. Figure 2.5 shows the parameter space of collisionality and density in which ELMs have been suppressed or mitigated on various machines, and the region of operational space which the ITER baseline will occupy. ELM suppression has been achieved within limited operational windows on KSTAR[65], DIII-D[64, 66], EAST[67] and ASDEX Upgrade[68], while ELM mitigation has been achieved in a much wider parameter space on the previously listed machines and also MAST[69].

In addition to collisionality and density, other variables have been identified which influence mitigation and suppression and frame the access windows. Some are equilibrium parameters, such as the edge safety factor q_{95} and the plasma boundary shape, while others are RMP parameters, such as coil current amplitude, dominant toroidal mode number and poloidal spectrum. The degree of ELM mitigation is sensitive to these parameters, while ELM suppression is observed only within particular windows of these parameter spaces. Furthermore, applied RMPs are often seen to induce secondary effects in tokamak plasmas, which may either play a role in the ELM mitigation or suppression mechanisms, or may merely be unintended side effects. Observed secondary effects include a significant drop in the pedestal density known as density pump-out[70], distortions to the plasma boundary shape[71], braking of the plasma rotation[72], and lobe structures appearing around the X point[73] and associated splitting of the divertor strike point.

While mitigation does not reduce the divertor heat load as much as suppression, it is far easier to achieve in a much wider parameter space, and the uncertainty about achieving some degree of mitigation on ITER is far lower than achieving full suppression[63].

2.5 RMP ELM Control Theories

There are currently several working theories of ELM suppression and mitigation, none of which is yet supported by a preponderance of evidence. The initial interpretation of the observation of ELM suppression on DIII-D is here called stochastic edge theory. In general the safety factor increases from the plasma centre to the edge, and so at some flux surfaces q will have values which can be expressed as a rational number $q = m/n$, called rational surfaces, at which the field lines complete exactly m poloidal rotations and n toroidal rotations before joining up on themselves. The poloidal spectrum of the applied RMP includes a component b_{res}^1 , which is aligned to the pitch of the equilibrium field lines such that $q = m/n$, where here m is the poloidal harmonic number of the

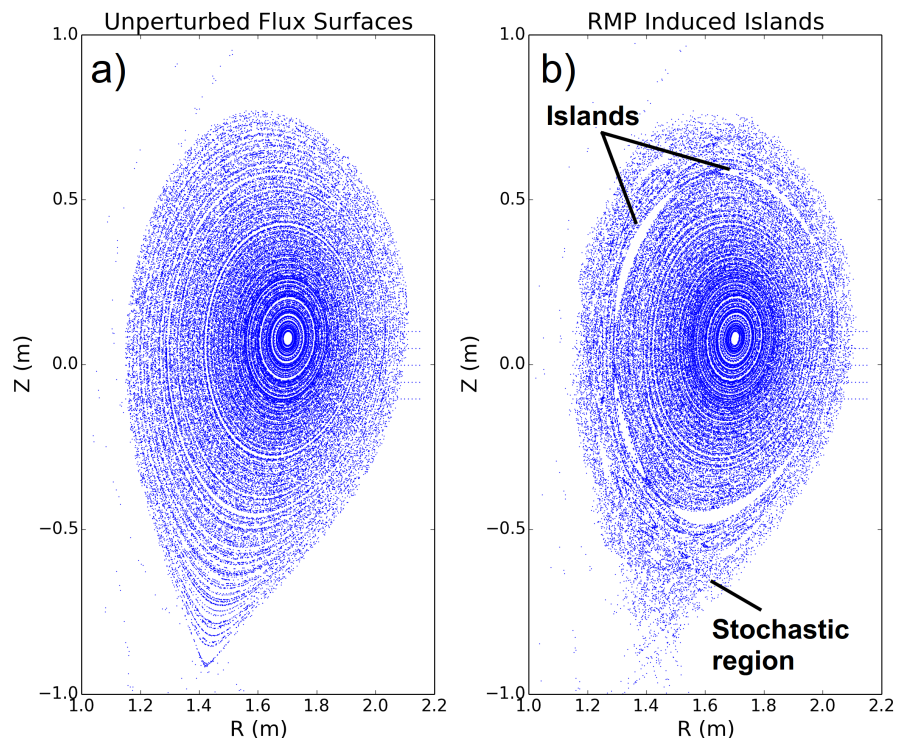


FIGURE 2.6: Poincaré plots show the points where a traced field line intersects with the poloidal plane. These are poincaré plots of an ASDEX Upgrade plasma equilibrium. In a), the magnetic equilibrium is undisturbed so field lines stay on flux surfaces. However in b), an applied RMP has caused islands to open at rational surfaces. At the edge where the rational surfaces are densely packed, the islands have overlapped resulting in a stochastic region, in which field lines may wander away from their flux surface.

Magnetic Island

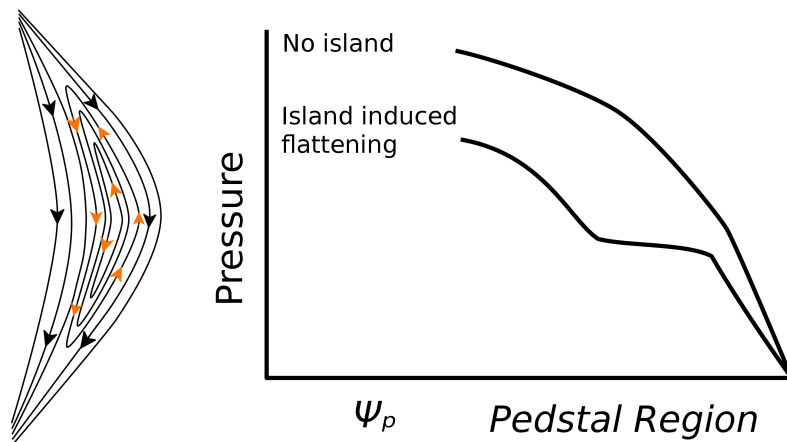


FIGURE 2.7: Magnetic field lines connect one side of magnetic islands to the other, so the pressure across the magnetic island is constant. Therefore the presence of an island causes a localised flattening of the pressure profile (flattening exaggerated for illustrative purposes). If an island is present at the top of the pressure pedestal, it can inhibit the growth of the pedestal, possibly preventing ELMs from destabilizing.

RMP and n the toroidal harmonic number. The pitch aligned component b_{res}^1 , drives the formation of chains of magnetic islands at rational surfaces with width w , which in cylindrical coordinates is

$$w = 4\sqrt{\left|\frac{Rr_s b_{res}^1}{nh B}\right|} \quad (2.13)$$

where R is the major radius of the plasma, r_s is the minor radius of the rational surface, n is the toroidal mode number, $h = d(\log q)/d(\log r)$ is the normalised shear, and B is the equilibrium magnetic field on axis. If these island chains become wide enough and are close enough together, they can begin to overlap. In regions where adjacent island chains overlap, the ordered flux surfaces are destroyed and replaced with a disordered 'stochastic' magnetic field, which has degraded confinement relative to flux surfaces[64]. Figure 2.6 explains this theory visually. The region of degraded confinement at the plasma edge prevents the pedestal pressure from increasing above the MHD P-B boundary, thereby preventing ELMs from occurring. For adjacent island chains of widths w_1 and w_2 a radial distance Δ_{12} apart, the chains overlap if $\frac{w_1+w_2}{\Delta_{12}} > 1$, where $\sigma_{chir} = \frac{w_1+w_2}{\Delta_{12}}$ is known as the Chirikov parameter, and the radial width of the region within which $\sigma_{chir} > 1$ is $\Delta_{\sigma_{chir}}$.

A correlation between $\Delta_{\sigma_{chir}}$ and the appearance of mitigation and suppression has been established on DIII-D[74], and therefore $\Delta_{\sigma_{chir}}$ is used as the guiding figure of merit in designing and assessing the ITER ELM coils[75]. However, stochastic edge theory has some significant deficiencies in its current form. In previous studies[75], the island widths were computed assuming that the total field is simply the linear sum of the equilibrium field and the applied RMP; however it is well known that the plasma response significantly modifies the applied field[76]. The RMP can be amplified by coupling to marginally stable MHD modes, a process known as resonant field amplification[77, 78]. Furthermore, currents are driven at rational surfaces in response to the RMP which act to reduce the pitch aligned component, and in the limit of zero resistivity exactly cancel it. Plasma resistivity is given by the expression[38]

$$\eta \approx \frac{\pi e^2 m_e^{1/2} Z_{eff} \ln \Lambda}{(4\pi \epsilon_0)^2 (k_B T_e)^{3/2}} \quad (2.14)$$

where e is the charge on an electron, m_e the mass of an electron, ϵ_0 the vacuum permittivity, k_B Boltzmann's constant, and T_e the electron temperature (in Kelvin). Due to the high temperatures and $T_e^{-3/2}$ dependence, resistivity in tokamak plasmas is typically negligible, so we would expect the plasma response to greatly reduce the size of islands and hence, the width of the stochastic region $\Delta_{\sigma_{chir}}$. However, towards the bottom of

the plasma pedestal where the temperature is much lower, there is the potential for sufficient resistivity for islands to form[79].

A variant of this theory implicates the rotation zero crossing as a key requirement for ELM suppression. MHD simulations[80–82] have indicated that large islands may form in regions where a rational surface coincides with a zero crossing of the electron perpendicular flow ω_e . This mechanism forms the basis of an alternate theory of ELM suppression. By connecting field lines of inner flux surfaces to outer flux surfaces as shown in figure 2.7, magnetic islands cause a local partial flattening of the pressure gradient. If the ω_e zero crossing coincides with a rational surface at the pedestal top, the island may inhibit the growth in the pedestal width, preventing it from becoming wide enough to trigger an ELM, and thereby suppressing ELMs [83–85].

Another theory implicates RMP induced plasma boundary distortions in the ELM mitigation mechanism. In an undisturbed ELM cycle, following an ELM crash the pressure pedestal gradient and height build up until the P-B stability boundary is again reached, precipitating the next ELM crash. Applied RMPs cause a sinusoidal distortion of the plasma boundary, and the appearance of 'lobe-like' structures near the plasma X point. Stability analyses of the resulting 3D equilibria indicate that these changes to the plasma boundary lower the threshold of P-B stability, so the threshold is reached sooner in each ELM cycle which increases the ELM frequency[86]. An expansion of this theory explains that the edge deformations may also be the cause of the observed RMP induced reduction in density[87], and that the deformations to the plasma boundary may be strongly driven by the coupling of the applied RMP to marginally stable MHD modes near the plasma edge, called edge peeling modes[88]. The amplification of these edge peeling modes was also seen to correlate with the reappearance of ELMs during ELM suppression on DIII-D[89], suggesting that the edge peeling response has a role in driving ELMs prematurely unstable.

Figures of Merit

Despite recent theoretical and experimental progress in understanding of effects of RMPs on ELMs, no theory is yet supported by a preponderance of evidence, and precise conditions required to achieve ELM suppression or a high degree of mitigation are not known. However, by comparing experimental observations with particular aspects of the computed plasma response, here called figures of merit, correlations have been established which may potentially be used for assessing and optimising the ITER ELM suppression and mitigation strategy, even in the absence of a rigorous theoretical understanding of ELM suppression and mitigation.

There are 4 figures of merit in common use. The pitch aligned component b_{res}^1 refers to the component of the total field, including plasma response, aligned to the equilibrium field at the outermost rational surface². The X point displacement ξ_X refers to the maximum absolute value of the distortion of the plasma surface in the vicinity of the plasma X point. The high field side response refers to the amplitude of the total field (including the plasma response) measured at the outboard side midplane, just outside the plasma. The peeling response refers to the maximum amplitude of marginally stable MHD modes near the edge, amplified by the applied RMP.

The figures of merit are of interest to RMP ELM control, because correlations have been established between them and RMP induced effects. b_{res}^1 was shown to correlate with the frequency of mitigated ELMs on ASDEX Upgrade and MAST[70, 72]. ξ_X was shown to correlate with mitigated ELM frequency on ASDEX Upgrade[70], as well as with RMP induced density pump out on MAST[69, 87]. On DIII-D, the high field side response was shown to correlate with density pump out and ELM suppression[90, 91]. The high field side response here refers to the measured or computed total magnetic field including plasma response, at the location of a magnetic probe on the inboard (high field) side, at the tokamak midplane. Finally the peeling response has been shown to correlate with ELM suppression on DIII-D[92, 93], density pump out on MAST[87], and ELM mitigation on ASDEX Upgrade[70, 94]. Modelling works have recently begun to impose some order on these observations, by explaining that the different figures of merit are correlated with each other in certain circumstances[79, 87, 88, 93–96].

The figures of merit may be optimised computationally, and may thereby be used to guide ELM mitigation and suppression experiments on current machines, or to design and assess ELM control systems on future machines. Unfortunately, the set of figures of merit also has some serious deficiencies. Observed correlations between figures of merit and RMP effects are currently not robust or consistent across machines, and many are based on quite limited datasets. Furthermore, the correlations between figures of merit implies at least some level of redundancy in the set. These deficiencies may be addressed by expanding the dataset on which they are based to wider parameter spaces and more machines, and by theoretical investigations of the relationships between different figures of merit.

²At the edge of a diverted plasma q goes to infinity, resulting in an infinite sequence of densely packed rational surfaces. However, as discussed later, for numerical investigations q must be truncated to some finite value, which for fixed n makes the sequence of rational surfaces finite, and determines the m number which defines the outermost rational surface. All results presented in this work are tested for robustness against truncation of the q profile, and found to be insensitive to it.

2.6 Plasma Response Model - MARS-F

In light of the previous discussions, it is clear that methods of computing the plasma response to applied RMPs are required. The varied approaches to simulating the plasma response to RMPs are reviewed and compared in [88, 97]. The general approach to computing the plasma response to applied RMPs used in this work reduces the model to a driven stationary eigenvector problem, with the equilibrium and plasma response as a stationary solution and the applied RMP as the drive force. In this approach, the equations of MHD are first linearised, meaning all quantities are set to be the sum of an equilibrium component and a small perturbation. For example the magnetic field \mathbf{B} is taken to be $\mathbf{B} = \mathbf{B}_0 + \mathbf{b}$, where \mathbf{B}_0 is the equilibrium magnetic field, and \mathbf{b} is the perturbed magnetic field with $|\mathbf{b}| \ll |\mathbf{B}_0|$. Derivation of the linearised equations of MHD are common in textbooks [98], but the MARS model also includes plasma equilibrium flow, which is usually excluded in simple derivations. Derivations of the linearised equations of MHD which include plasma flow can be found in [99, 100]. We now add the assumption that all quantities oscillate in the lab frame at the driving frequency of the applied magnetic field, Ω_{RMP} . With these assumptions, following [100] the linearised equations of resistive MHD including toroidal flow are found as

$$i(\Omega_{RMP} + n\Omega)\xi = \mathbf{v} + (\xi \cdot \nabla\Omega)R\hat{\phi} \quad (2.15)$$

$$i\rho(\Omega_{RMP} + n\Omega)\mathbf{v} = -\nabla p + \mathbf{j} \times \mathbf{B} + \mathbf{J} \times \mathbf{b} \\ - \rho[2\Omega\hat{Z} \times \mathbf{v} + (\mathbf{v} \cdot \nabla\Omega)R\hat{\phi}] - \rho\kappa|k_{||}v_{th,i}|[\mathbf{v} + (\xi \cdot \nabla)\mathbf{V}_0]_{||} \quad (2.16)$$

$$i(\Omega_{RMP} + n\Omega)\mathbf{b} = \nabla \times (\mathbf{v} \times \mathbf{B}) + (\mathbf{b} \cdot \nabla\Omega)R\hat{\phi} - \nabla \times (\eta\mathbf{j}) \quad (2.17)$$

$$i(\Omega_{RMP} + n\Omega)p = -\mathbf{v} \cdot \nabla P - \gamma P \nabla \cdot \mathbf{v} \quad (2.18)$$

In the above, R is the major radius of the plasma, ϕ the toroidal angle with unit vector $\hat{\phi}$, \hat{Z} the vertical unit vector. $\mathbf{V}_0 = \Omega R\hat{\phi}$ is the equilibrium rotation velocity, Ω is the plasma rotation frequency, Ω_{RMP} is the frequency of the applied RMP (zero for static RMP fields), η is the plasma resistivity, and n is the (chosen) toroidal mode number. The final term of equation 2.16 describes landau damping of the ion acoustic wave. κ is a constant numerical coefficient, used to vary the strength of parallel sound wave damping, where $k_{||}$ is the parallel wave number of the sound wave (essentially a geometric factor), and $v_{th,i}$ is the ion thermal velocity. \mathbf{B} , \mathbf{J} , P and ρ are the equilibrium magnetic field, current, pressure and density, which are all known and given as input. The unknown state vector to be calculated is $(\xi, \mathbf{b}, \mathbf{j}, p)$, the plasma displacement, perturbed magnetic

field, perturbed current, and perturbed pressure. The RMP current \mathbf{j}_{RMP} is known and given as input, and the RMP field is calculated from the following relation.

$$\nabla \times \mathbf{b} = \mathbf{j}_{RMP} \quad (2.19)$$

$$\nabla \cdot \mathbf{j}_{RMP} = 0 \quad (2.20)$$

The above equations of linearised MHD including toroidal flow, constitute the MARS-F model formulation [101].

In order to solve for the perturbation quantities, the system of equations (2.15) - (2.20) is discretized, replacing the continuous quantities with their discrete counterparts. MARS uses finite element discretization in the radial direction, and fourier discretization in the poloidal direction, ie, as the sum of a finite number of poloidal modes $e^{(im\chi)}$ where m is the poloidal harmonic number, and χ is the generalised poloidal angle. All quantities in the toroidal direction have a prescribed $e^{(in\phi)}$ dependence, where n is the single toroidal mode number of the applied RMP and ϕ the toroidal angle. By linearisation and then discretization, all time and space derivative operators are replaced by linear multiplicative terms, and so the system of coupled partial differential equations of MHD has been simplified to a linear algebraic problem. The discretized equations can be expressed in matrix form

$$\lambda \mathbf{A} \mathbf{X} = \mathbf{B} \mathbf{X} + \mathbf{X}_0 \quad (2.21)$$

where $\lambda = \Omega_{RMP} + n\Omega$, and \mathbf{X} is the discretized version of the perturbation vector $(\xi, \mathbf{b}, \mathbf{j}, p)$, \mathbf{X}_0 is the discretised version of the source (forcing) term (equation 2.19), and \mathbf{A} and \mathbf{B} are linear operators consisting only of known equilibrium quantities. \mathbf{X} is the unknown to be solved for, which can be done by a matrix inversion.

$$\mathbf{X} = (\lambda \mathbf{A} - \mathbf{B})^{-1} \mathbf{X}_0 \quad (2.22)$$

For a driven steady state system, which is how stable RMP experiments are modelled, inverting the (very large) matrix $(\lambda \mathbf{A} - \mathbf{B})$ is the main computational process of the MARS code. The solution \mathbf{X} is the state vector $(\xi, \mathbf{b}, \mathbf{j}, p)$ at all points of the computational domain: the plasma response.

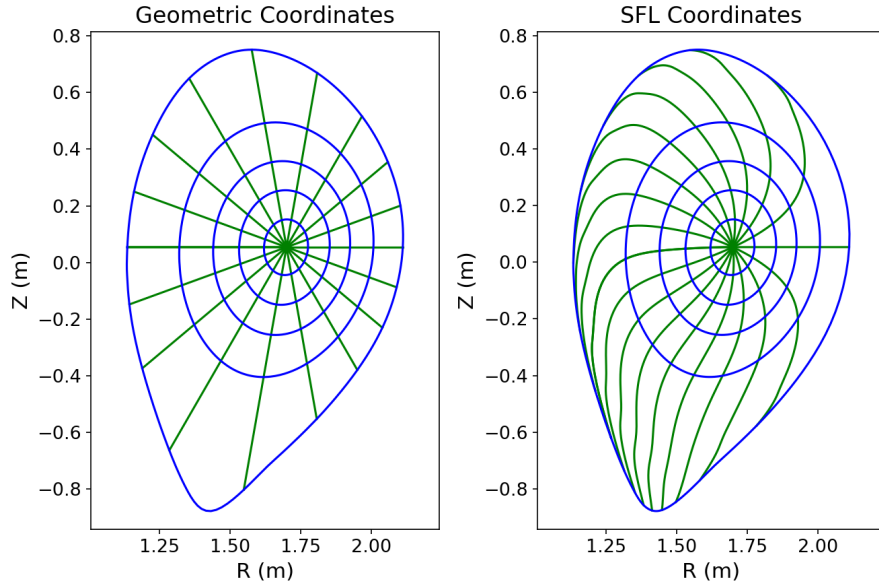


FIGURE 2.8: a) Geometric coordinate system. Poloidal resolution is evenly distributed, allowing good resolution near the RMP coils on the outboard side. b) Straight Field Line (SFL) coordinate system. Poloidal resolution is concentrated near the X point, and is poor near the RMP coils. Both grids show lines of χ in 20 degree increments.

Magnetic Geometry

In this work, two grids with different coordinate systems are usually required: a geometric, and straight field line (SFL) coordinate system. Both use cylindrical coordinates (s, ϕ, χ) , with s the radial coordinate, ϕ the toroidal angle coordinate, and a χ the poloidal angle coordinate, but differ in their definitions of χ . Examples of the two grids for the same equilibrium are plotted in figure 2.8. The blue lines trace flux surfaces at increments of s from 0 to 1 in steps of 0.2, while the green lines trace lines of constant χ from 0° to 360° in increments of 20° . In the geometric coordinate system, χ is simply the intuitive real space poloidal angle, anti-clockwise from the outboard midplane. In the SFL coordinate system however, χ is defined such that field lines in the (ϕ, χ) plane are straight. It is only in SFL coordinates that poloidal fourier harmonics are physically meaningful, and correspond to observable MHD modes in real space, and so poloidal harmonics may only be interpreted in SFL coordinates. However, as demonstrated in figure 2.8, on the outboard side where the RMP field is applied, poloidal resolution is very poor in an SFL coordinate system, making it difficult to properly resolve the applied RMP field. Therefore, MARS plasma response computations are generally performed in geometric coordinates, and interpolated onto a matching SFL grid for interpretation in post process. Both the geometric and SFL grids used in this work are computed for given experimental equilibria using the CHEASE fixed boundary equilibrium solver [106].

Chapter 3

Coupling between peeling response and pitch aligned component

3.1 Resonant Field Amplification

Many systems which support oscillatory behaviour, whether mechanical (swings and bridges), electrical (radios), or magnetohydrodynamic (a toroidal plasma), have resonances: particular frequencies and/or wavelengths at which the system oscillates particularly readily. When a force is applied to these systems at or near a resonant frequency, the corresponding resonant mode is excited to a large amplitude. In this case, the system is a rotating tokamak plasma equilibrium, the modes are marginally stable MHD modes, and the driving force is the applied RMP. It may seem counter-intuitive that a static applied field could illicit this response, but of course since the plasma is rotating, in the frame of the plasma the applied field is not static. Marginally stable MHD modes being coupled to and driven by an applied RMP (or error field) is known as Resonant Field Amplification[77].

3.2 Amplified Peeling Response

One set of modes has received particular attention in the context of applied RMPs. These modes have low toroidal mode numbers, are spatially localised near the plasma edge, and have a peeling mode structure, and are therefore called low n edge peeling modes (hereafter just the edge peeling response). The edge peeling response has been




FIGURE 3.1: Example spectrogram of the field including the plasma response. The spectral regions (ie, regions in m, s space) referred to as 'Kink', 'Core Kink' and 'Peeling' are labelled. More rigorously, the designations 'kink' and 'peeling' refer to specific radial profiles of the plasma displacement. It should be emphasised that m is a discrete quantity; smoothing is applied to the m axis to make modes more easily discernible.

implicated by several studies as having a role in ELM mitigation[70, 94]. Furthermore, a correlation between the computed peeling response and ELM suppression has been detected on DIII-D[92]. It has been suggested that the peeling response may be driving the pitch aligned components by poloidal harmonic coupling[93], thereby increasing the width of the stochastic region which may be necessary for ELM suppression. This motivates the following numerical study into the peeling response, to examine whether the peeling response may drive the pitch aligned components by poloidal harmonic coupling.

At this point it is useful to clarify some nomenclature, illustrated in figure 3.1. The term 'harmonic' refers to a single poloidal fourier harmonic m having poloidal dependence $e^{(im\chi)}$ (where χ is the the straight field line coordinate system poloidal angle). The term 'mode' refers to a particular type of plasma instability, generally comprised of many poloidal harmonics, characterised by its radial displacement profile. The 'kink response' generally refers to the amplification of harmonics just above pitch resonance (ie, $m > qn$). The kink response is further divided into 'core kink' and 'edge peeling'

response, referring respectively to the kink response in the plasma core and near the edge.

This chapter details work previously presented in part in [79]. In this section, the plasma response to an $n = 2$ RMP field including $n = 6$ sidebands applied to an ASDEX-Upgrade plasma equilibrium is computed using the MARS-F code. The sensitivity of the plasma response to the geometry near the X point is also investigated by varying the plasma boundary shape. To investigate the general behaviour of poloidal harmonic coupling in a realistic toroidal geometry, a single poloidal harmonic magnetic perturbation is applied at the plasma edge as a boundary condition. The pitch aligned components resulting from this applied boundary condition were then computed, to investigate whether they could be driven by poloidal harmonic coupling. This chapter also describes work previously presented in [102–104], in which MARS-F simulations of the plasma response are compared with experimental measurements of the resulting magnetic perturbation and plasma displacement, made with external magnetic B_p probes, and plasma emission diagnostics.

This section is organised as follows. Section 3.3 describes the experimental equilibrium used for this study. Section 3.4 describes benchmarking of the applied perturbation against Biot-Savart code ERGOS. Section 3.5 describes the computed plasma response to this applied perturbation, as well as its dependence on the plasma boundary X point truncation. Section 3.6 describes the plasma response to contrived single m applied perturbations in order to investigate poloidal harmonic coupling between the pitch aligned components and the peeling response. In section 3.7 MARS-F predictions of the magnetic and displacement perturbations are compared with experimental measurements of the displacement and magnetic field. Section 3.8 then summarises and discusses these results.

3.3 Equilibrium

The 2D plasma equilibrium used in this study was reconstructed from ASDEX-Upgrade discharge 30835: a low collisionality ELM mitigation experiment. Table 3.1 lists essential information of the plasma equilibrium. The initial equilibrium was constructed using the free boundary equilibrium code CLISTE [105], provided as part of a standardised modelling kit at IPP Garching. Magnetic measurements, the $q = 1$ surface location from measurements of sawtooth instabilities, and the scrape-off layer current were used as constraints on the equilibrium. Using the fixed boundary equilibrium solver CHEASE [106], the equilibrium was then refined and mapped to the MARS-F straight

field line coordinate system. Figure 3.2 a)-c) shows curves fitted to experimentally measured radial profiles of plasma parameters: a) the electron and ion temperature, b) the electron density, and c) the toroidal angular velocity. Plots d)-f) are equilibrium profiles of parameters output from the CHEASE equilibrium code: d) the safety factor profile, e) the plasma pressure, and f) the plasma current density. s is the radial coordinate, defined as $s = \sqrt{\psi_N}$ where $\psi_N = (\psi_{pol} - \psi_0)/(\psi_a - \psi_0)$ is the normalised poloidal magnetic flux, in which ψ_0 and ψ_a are the poloidal magnetic flux ψ_{pol} evaluated at the magnetic axis and plasma edge respectively.

ASDEX Upgrade is equipped with 16 ELM control coils, arranged in two toroidal rings of 8 coils each above and below the midplane. Having 8 coils toroidally allows toroidal mode numbers of $n = 1, 2, 3$ or 4 (although the $n = 3$ inevitably contains a significant $n = 5$ sideband). The coils are powered by 4 independent power supply units, allowing rectangular or sinusoidal toroidal waveforms, and also allowing waveforms to be rotated in the upper and lower rows independently. Figure 3.3 shows the ASDEX Upgrade coil set relative to a typical plasma shape, as well as examples of toroidal current waveforms which may be applied in the upper or lower coil rows for a rectangular $n = 2$ and sinusoidal $n = 1$ applied field.

B_0 (T)	I_p (MA)	q_0	q_{95}	β %	l_i	ν_e^*	S
1.7	0.77	0.8	3.8	2.04	1.01	0.08	3.5×10^8

TABLE 3.1: Parameters of the plasma equilibrium used for this numerical study: ASDEX Upgrade discharge 30835 at 3200ms. B_0 is the vacuum equilibrium field strength at the magnetic axis, I_p is the total plasma current, q_0 and q_{95} are the safety factor at the core and the $\Psi_N = 0.95$ flux surface respectively. $\beta = (2\mu_0 \langle P \rangle)/(B_0^2)$ is the ratio of volume averaged plasma pressure to magnetic pressure on axis, l_i is the normalised plasma internal inductance (see [106] for full definition), ν_e^* is the plasma collisionality evaluated at the pedestal top (following the formula in [107]), and S is the Lundquist number evaluated at the magnetic axis.

3.4 Benchmark

To provide an independent benchmark of the vacuum magnetic perturbation, the vacuum field was computed using MARS-F and also with the Biot-Savart based code ERGOS[108]. Although the toroidal spectrum of the applied perturbation is dominated by a single mode number n_{dom} , it will also inevitably contain sidebands of a finite size. A perturbation with dominant mode number n_{dom} applied using N coils will typically contain large sidebands at mode numbers $(jN \pm n_{dom})$ where j is an integer. Figure 3.4 shows the toroidal waveform and spectrum of the RMP coil currents, applied in 30835 at 3.2s, showing the expected large n_{dom} component and $N \pm n_{dom}$ sidebands. MARS




FIGURE 3.2: Ion and electron temperature, electron density, rotational velocity, safety factor, plasma pressure and current density profiles, from ASDEX-Upgrade discharge 30835, taken at 3200ms into the plasma shot.

computations are limited here to low n , so only the low n ($\lesssim 10$) sidebands can be studied. Not shown in the figure are the negative toroidal harmonics, which by symmetry are the same as the positive harmonics (this is not the case with the poloidal spectrum). In MARS, the symmetry of the toroidal spectrum across $n = 0$ is used to compute both the n and $-n$ response. The lower coil set has a slightly higher n_{dom} component and lower sidebands, due to the slightly smaller toroidal inter coil spacing of the lower coils. Figure 3.5 shows the $n = 2$ pitch aligned components due to the vacuum field, computed both with ERGOS and MARS-F. The figure shows that the pitch aligned components largely agree, so we may be confident that the $n = 2$ components of the MARS and ERGOS fields are in agreement. However, this only examines the n_{dom} component of the field. Figure 3.6a-d) shows a detailed global benchmark between MARS and ERGOS. Previous studies [93] have investigated the largest sideband and found it to be small compared with the dominant n component, and have therefore neglected it. However, the figure shows that the MARS single n global field differs noticeably from the ERGOS full toroidal spectrum field. Since MARS-F is a linear code, and toroidal harmonics do not exhibit the complex coupling behaviour of poloidal harmonics (since they are

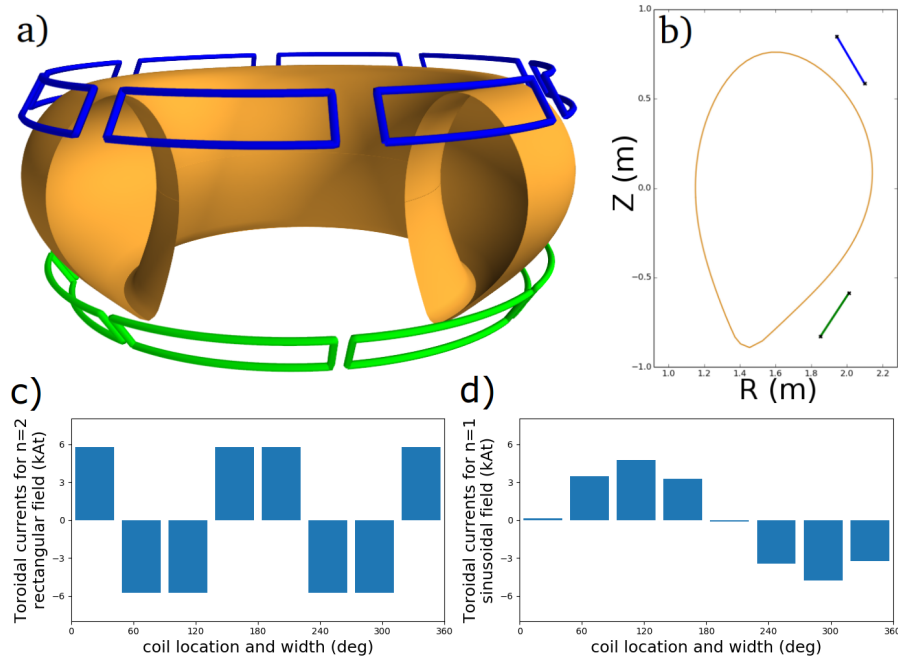


FIGURE 3.3: a,b) Sketch of the ASDEX Upgrade coil set. The coil set consists of an upper and lower row of 8 coils each, and apply non-axisymmetric magnetic perturbations close to the plasma edge. c,d) Example toroidal current waveforms available for use in experiments in the upper or lower row of coils. c) An $n = 2$ rectangular waveform. d) An $n = 1$ sinusoidal waveform.

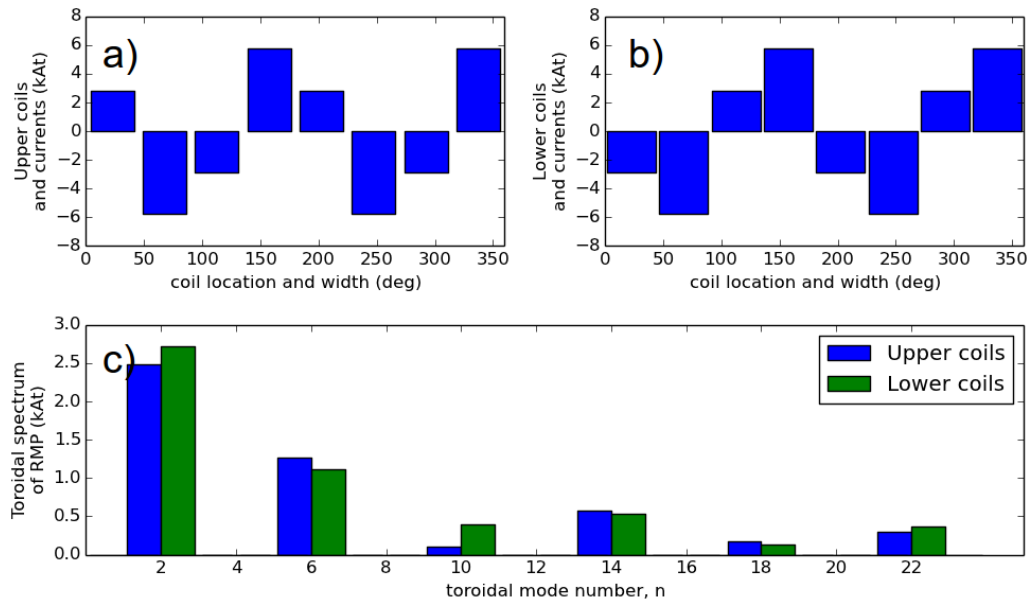


FIGURE 3.4: The toroidal waveform of a) upper and b) lower coil currents applied to ASDEX Upgrade 30835. This toroidal waveform results in the toroidal spectrum plotted in c), where blue is the upper coil toroidal spectrum and green is the lower coil toroidal spectrum. The dominant component is $n = 2$, with a significant $n = 6$ sideband.

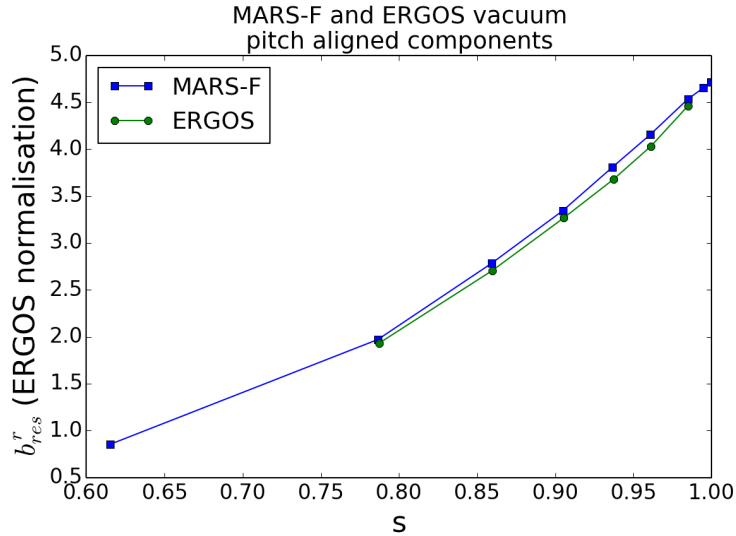


FIGURE 3.5: Vacuum benchmark of the pitch aligned components of the vacuum field, computed with MARS-F and ERGOS. The MARS-F computation has been converted to match the slightly different ERGOS normalisation.

defined on a circle around which all equilibrium quantities are constant), the field due to sidebands can be computed separately and added to the dominant mode number in post process. Figure 3.6e-h) shows the same benchmark, with the inclusion of the $n = 6$ sideband in the MARS computation. The figure shows a good agreement between the MARS vacuum field using $n = 2$ and $n = 6$ with the ERGOS field. This shows that the global field including detailed geometry is well approximated using only the dominant toroidal mode number and largest sideband, therefore, the $n = 2$ and $n = 6$ toroidal harmonics will be considered in this study. The small difference is primarily due to the different representation of the coil currents between the two codes - ERGOS uses a detailed real space description of the coil geometry, and so includes the whole toroidal spectrum. Nevertheless the ERGOS vacuum field provides a useful benchmark for the applied perturbation computed by MARS-F.

3.5 Plasma Response to Applied Perturbation

ASDEX Upgrade has an upper and lower set of RMP coils, allowing the poloidal spectrum to be tuned by varying the toroidal phase difference between the upper and lower coils, $\Delta\phi_{ul}$. It should be noted that the toroidal phase difference $\Delta\phi_{ul}$ refers to the difference in phase between the upper and lower coil sets, rather than the offset between the upper and lower coil current waveforms in real space toroidal angle, $\Delta\phi_{real}$. Rather, the toroidal phase difference $\Delta\phi_{ul}$ is related to the real space offset $\Delta\phi_{real}$ by the relation $\Delta\phi_{ul} = n\Delta\phi_{real}$, as illustrated in figure 3.7. Using MARS-F, the $n = 2$ and $n = 6$

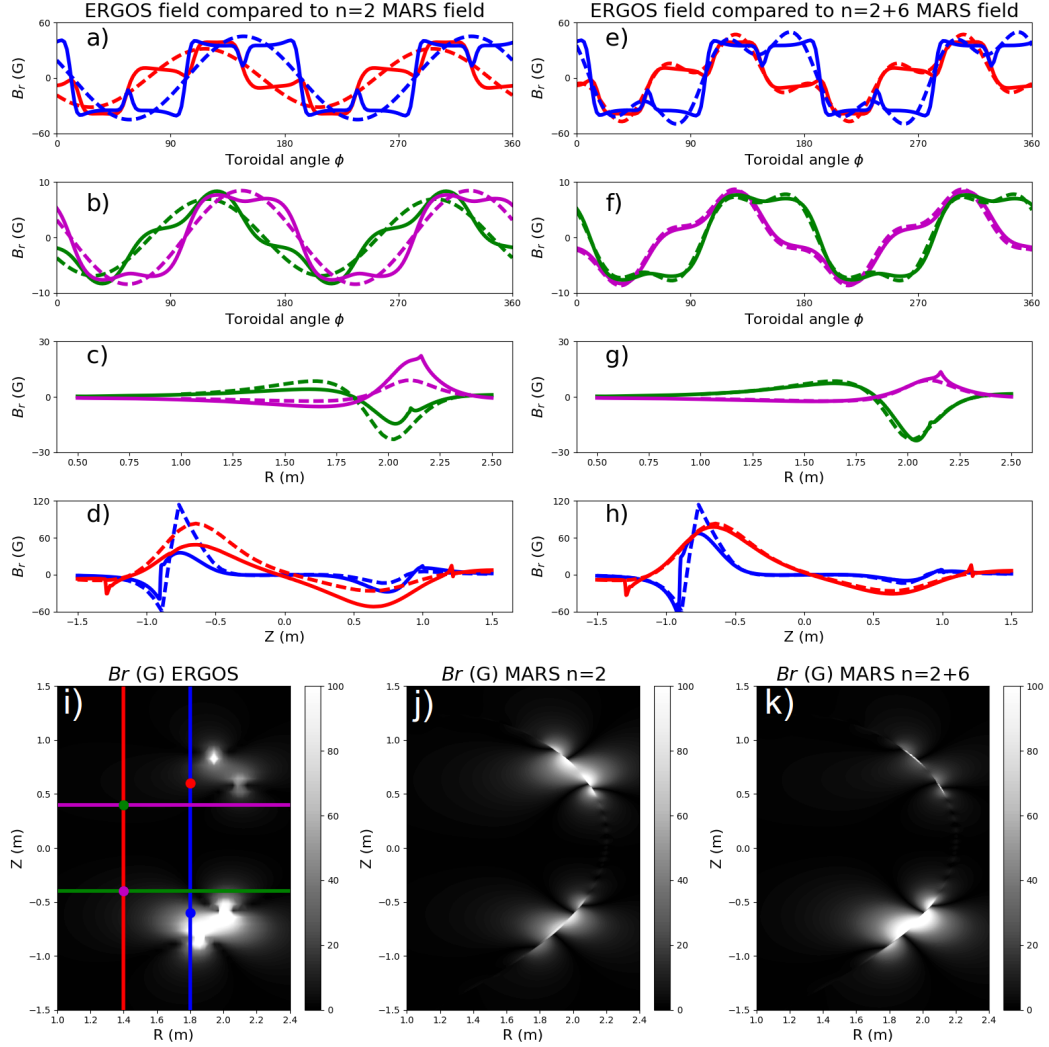


FIGURE 3.6: A benchmark of the global vacuum field between MARS-F and ERGOS. a,b) Vacuum field computed with ERGOS (solid) and MARS $n = 2$ (dashed) with toroidal angle, near to a) and far from b) coils. R,Z locations indicated by the red and blue circles in plot i). c,d) Vacuum field computed with ERGOS (solid) and MARS $n = 2$ (dashed) with height Z c) and major radius R d). Locations of the lineouts are plotted in their corresponding colour in figure i). e,f,g,h) Same results with ERGOS (solid) and MARS $n = 2$ summed with the $n = 6$ component (dashed). i) Contour plot of the ERGOS field in a poloidal (R,Z) plane. j) Contour plot of the MARS $n=2$ field in the poloidal plane. k) Contour plot of the MARS $n=2$ and $n=6$ fields in the poloidal plane. The match with the ERGOS field is greatly improved by the inclusion of the $n=6$ sideband.




FIGURE 3.7: The above sketch demonstrates how the phase difference $\Delta\phi_{ul}$ is achieved in practice, by applying a toroidal angle offset between the lower and upper coils of $\Delta\phi_{ul}/n$. The blue (green) rectangles depict the upper (lower) RMP coils, while the blue (green) sinusoid depicts the toroidal current waveform these coils are applying.

plasma response to an applied $n_{dom} = 2$ perturbation with experimental values of the coil currents was computed. MARS-F is a linear model, which allows the response due to the upper and lower coils to be computed separately, and then summed in post process using the principle of superposition. The plasma response due to both coils is computed in post-process for arbitrary $\Delta\phi_{ul}$ using the relation $b_{\Delta\phi_{ul}} = b_{upper} + b_{lower} \times e^{(-i\Delta\phi_{ul})}$, as in previous plasma response studies [93]. In the plasma bulk and most of the edge region where the electron temperature is high, the Spitzer model for resistivity is sufficient. However, in the limit approaching the plasma edge where the electron temperature can tend to zero, the Spitzer resistivity would tend towards infinity. Therefore to avoid a numerical singularity, the maximum value of the resistivity is fixed for numerical stability; in this study the maximum value of resistivity was fixed at 204 nOhm m (corresponding to 370eV), its value at $s = 0.99$. Figure 3.8 shows the vacuum perturbation, and the total magnetic perturbation including plasma response, ie, $b_{total} = b_{vac} + b_{response}$, both for $n = 2$ and $n = 6$ fields. In this work, 'total' is taken to mean 'including the plasma response'. The figure plots the absolute value of the normal component of the perturbed magnetic field, in the dimensionless unit defined as $|b^1| = \left| \frac{\mathbf{b} \cdot \nabla \psi}{\mathbf{B}_{eq} \cdot \nabla \phi} \frac{q}{R_0^2 B_0} \right|$. There are several noteworthy features. Firstly, for both $n = 2$ and $n = 6$ the total pitch aligned components ($b_{m=qn}^1$) at the inner rational surfaces are close to zero. Due to the high temperature in the plasma bulk, the resistivity is negligibly small, so the plasma response is close to ideal. This causes the pitch aligned components to be almost perfectly screened by induced current sheets at rational surfaces. Closer to the plasma edge however, the temperature is lower and the resistivity is higher, which allows the pitch aligned components near the plasma edge to be finite. Secondly, there are some regions of the spectrogram (ie, in m, s space) which are amplified above their vacuum values by resonant field amplification. As figure 3.8 shows, the $n = 2$ core kink mode is moderately amplified, but the dominant response for both $n = 2$ and $n = 6$ is the edge localised low n peeling mode. There is no core kink response apparent in the $n = 6$ response. Thirdly, both the applied and total fields of the $n = 6$ sideband are much smaller than the dominant $n = 2$ component, since the $n = 6$ component of the toroidal

current waveform is far smaller than the $n = 2$. The amplified peeling response of the $n = 6$ field also has a much smaller radial extent than the $n = 2$, and is also confined to fewer poloidal harmonics.




FIGURE 3.8: a) The $n=2$ poloidal spectrum of the applied vacuum magnetic perturbation, with $\Delta\phi_{ul} = 180^\circ$. The pitch aligned components are highlighted with white circles, and the white dashed line follows the $m = nq(s)$ contour. b) The $n=2$ poloidal spectrum of the magnetic perturbation including the plasma response (ie, the total field). While there is moderate amplification of the core kink, the edge localised peeling response is more appreciable. The pitch aligned components far from the edge are totally screened by the plasma response, but the components close to the edge can remain finite. c),d) $n=6$ vacuum and total poloidal spectra.

To quantify how the $n = 2$ and $n = 6$ peeling response varies with $\Delta\phi_{ul}$, the field amplitude at a representative spectral point (ie, a point in m, s space) is computed while $\Delta\phi_{ul}$ is scanned. Figure 3.9 shows the total and vacuum magnetic field perturbation $|b^1|$ at spectral location $(m, s) = (11, 0.99)$ for the $n = 2$ field, and $(m, s) = (26, 0.99)$ for the $n = 6$ field. The figure shows that the field in this spectral region is amplified above its vacuum value, and that the amplification is strongly dependent on the coil phase shift $\Delta\phi_{ul}$. These results agree qualitatively with a similar study of a DIII-D plasma, which also predicted amplification of the edge peeling response with strong $\Delta\phi_{ul}$ dependence [93]. The figure also shows that the amplification effect (size of the total field relative to the vacuum field) appears greater for $n = 6$, although the total field is still small compared to the $n = 2$ component.

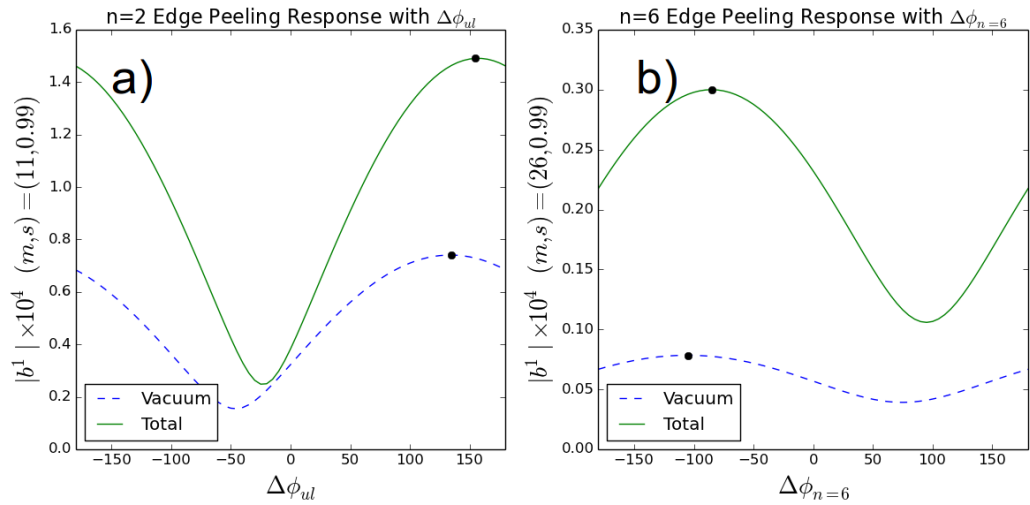


FIGURE 3.9: The magnitude of the perturbation at spectral location $(m, s) = (11, 0.99)$, which is representative of the edge peeling response, with a numerical scan of $\Delta\phi_{ul}$. The dashed line shows the vacuum perturbation, and the solid line is the total perturbation.

3.5.1 Pitch Aligned Components

The white circles in figure 3.1 mark the pitch aligned field components: the spectral points (in m, s space) where $m = nq(s)$. In the straight field line coordinate system, at these spectral points the magnetic perturbation is aligned with the equilibrium magnetic field. The pitch aligned field components are of particular interest to RMP studies, because they determine the widths of islands which are driven at rational surfaces, and thereby the widths of any stochastic regions which may form in response to the RMPs. Figure 3.10 shows the absolute value of the $n = 2$ and $n = 6$ vacuum and total pitch aligned field components, at a coil phase difference of $\Delta\phi_{ul} = 80^\circ$ for $n = 2$, and $\Delta\phi_{ul} = -110^\circ$ for $n = 6$. In Ideal MHD, applied RMPs drive current sheets at rational surfaces which completely screen the magnetic pitch aligned components. Therefore in the bulk plasma, where the resistivity is low and the plasma is approximately ideal, the pitch aligned components of the total field are negligibly small. However, as the figure shows, outside the temperature pedestal close to the plasma edge (roughly $s > 0.95$) the resistivity is much higher, so the pitch aligned components are not completely screened but remain finite. The figure also shows that the $n = 6$ pitch aligned components are far smaller than the $n = 2$. As before, this is primarily due to the smaller size of the $n = 6$ sideband of the toroidal current waveform, but it is also due to the $m = nq$ spectral line being separated from the main spectral lobes of the applied RMP when $n = 6$. Specifically, near the edge the applied RMP has poloidal harmonics mostly $m < 10$, while the $n = 6$ spectral line in the edge region is above $m = 20$.

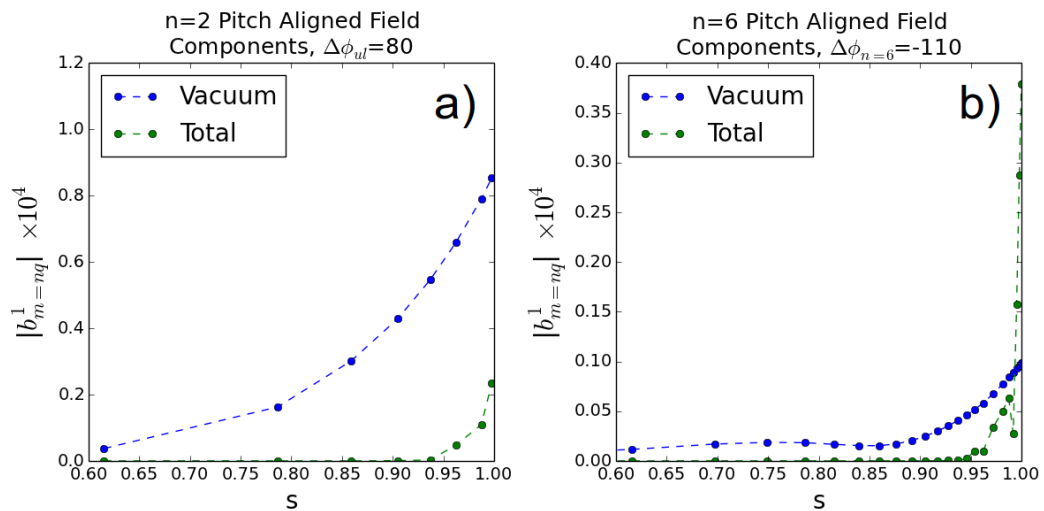


FIGURE 3.10: The pitch aligned components of the applied vacuum perturbation and total perturbation including plasma response. In the plasma bulk far from the edge the pitch aligned components are well screened, but in the edge region can be finite as the resistivity is higher.

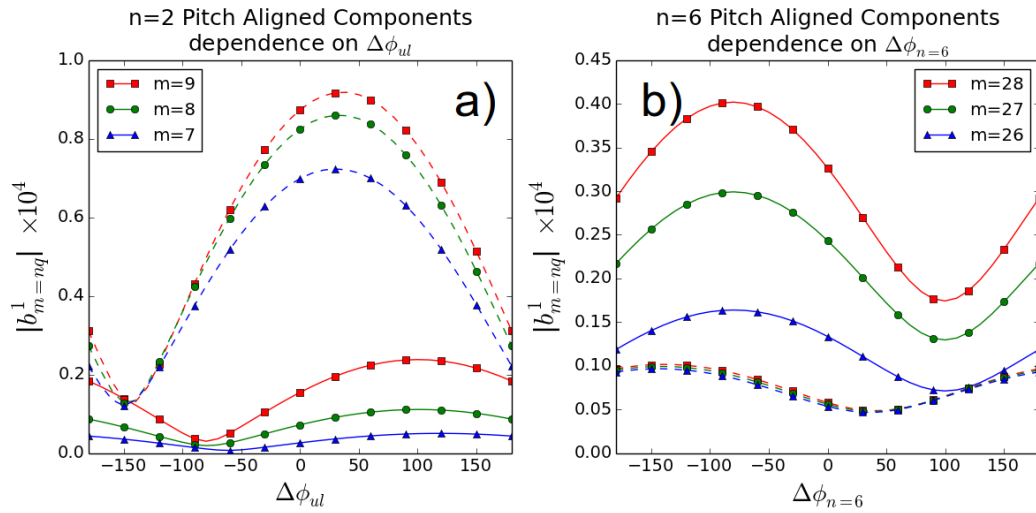


FIGURE 3.11: Absolute values of the outermost 3 pitch aligned components, in the vacuum approximation (dashed lines) and the total field (solid lines). The maximum pitch aligned component including the plasma response is offset 60° from its vacuum value. This provides further evidence that vacuum modelling alone is insufficient to predict the coil phase for optimum pitch alignment.

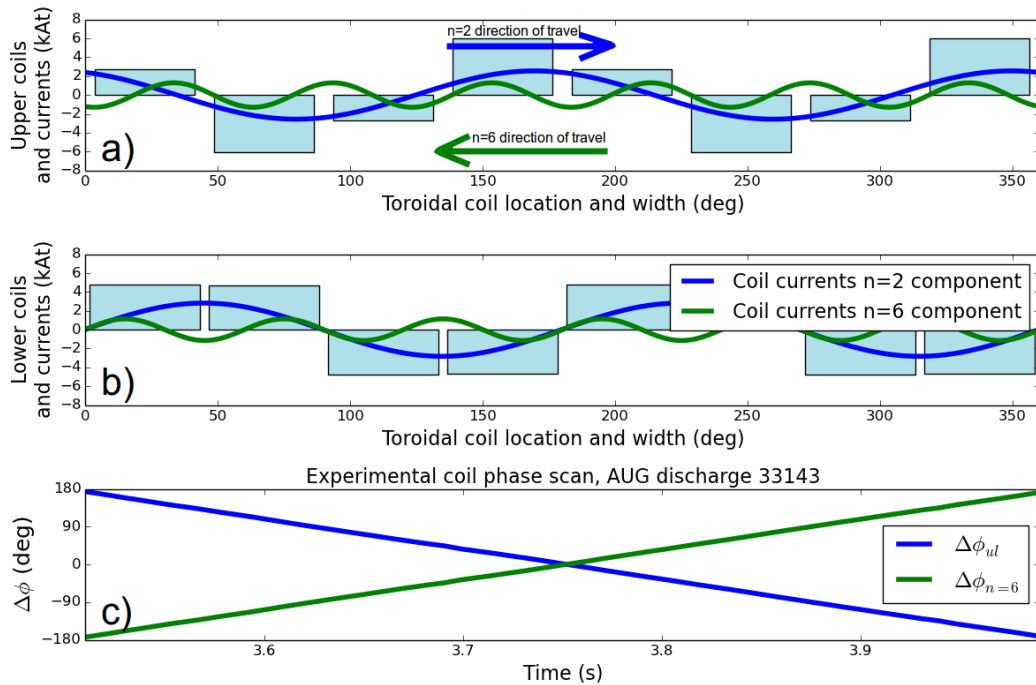


FIGURE 3.12: a,b) Experimentally applied coil currents in the a) and b) lower coils at an instant of time during a scan of the coil phase $\Delta\phi_{ul}$, are represented by the cyan bars. To generate the phase scan, the upper coil currents form a travelling wave which rotates the $n = 2$ field (blue line) toroidally. However, the $n = 6$ component (green line) travels at the same rate in the opposite direction. c) $\Delta\phi_{ul}$ of the $n = 2$ and $n = 6$ components of the applied RMP. For the case of an $n_{dom} = 2$ field, $\Delta\phi_{ul}$ for the $n = 6$ is simply the negative of $\Delta\phi_{ul}$ for the $n = 2$.

Figure 3.11 plots the absolute values of the outermost 3 pitch aligned components (at the $m = 7, 8,$ and 9 rational surfaces for $n = 2$, or the $m = 26, 27,$ and 28 rational surfaces for $n = 6$) as a function of n_{dom} coil phase difference $\Delta\phi_{ul}$. The solid lines show the pitched aligned components of the total field, while the dashed lines show the vacuum field. The maximum value of the $n = 2$ vacuum pitch aligned components (ie, vacuum alignment) occurs at $\Delta\phi_{ul} = 30^\circ$, whereas the maximum value of the $n = 2$ total pitch aligned components occurs at $\Delta\phi_{ul} = 90^\circ$. It is apparent that the coil phase which maximises the total pitch aligned components is offset from its vacuum value by 60° . This offset is also apparent in the $n = 6$ pitch aligned components. This demonstrates that a vacuum field aligned applied field does not maximise the total pitch aligned components, and therefore to optimise the coil configuration to maximise total pitch aligned components, the plasma response must be accounted for. It should be noted that this phase shift is not exclusive to the pitch aligned components; other spectral regions also experience a shift in their dependence on $\Delta\phi_{ul}$ (see figure 3.9), but this will not be examined here. The $\Delta\phi_{ul}$ dependence of the total and vacuum fields is robust with respect to changes in q_a caused by changes in plasma shape (of the type discussed in section 3.5.2), and therefore the 60° offset of the total field from the vacuum is also robust to these changes. The magnitudes of the outermost pitch aligned components, however, were found to be slightly sensitive to q_a . This may be caused by changes to the rotation and resistivity at outermost rational surfaces, which results from movement of the rational surfaces caused by changes to q_a . The sensitivity of the plasma response to plasma rotation and resistivity is described in previous studies [76].

A noteworthy feature of figure 3.11 is that for certain ranges of $\Delta\phi_{ul}$, the total pitch aligned components are of the same order or larger than their vacuum value. This is unexpected, given the strong screening effects which act to reduce the pitch aligned components. This is particularly evident for the outermost $n = 6$ pitch aligned components, possibly due to their higher n number allowing them to be packed more closely to the plasma edge in the region of relatively high resistivity. In the context of this linear single fluid resistive MHD model, resonant field amplification is the sole method by which parts of the total perturbation spectrum could exceed their vacuum values. It was previously suggested [92], that since the amplified edge localised peeling response is spectrally close to the outermost pitch aligned components, poloidal harmonics of the amplified peeling response may couple to and drive the pitch aligned components by poloidal harmonic coupling. This possibility motivated the investigation into poloidal harmonic coupling on ASDEX-Upgrade, described in section 3.6.

Experimental scans of $\Delta\phi_{ul}$ are typically performed by rotating the upper toroidal waveform of currents relative to the lower. An example of this is plotted in figure 3.12, using the coil currents of discharge number 33143 (since 30835 had a static $\Delta\phi_{ul}$). It is useful

to note that, for the case of $n_{dom} = 2$ discharges, which with 8 coils toroidally will always have a significant $n = 6$ component, the phase difference of the $n = 2$ component $\Delta\phi_{ul}$ is simply the negative of the phase difference of the $n = 6$ component $\Delta\phi_{n=6}$. With reference to figure 3.11, which plots the outermost pitch aligned components with experimentally applied $\Delta\phi_{ul}$, applying $\Delta\phi_{n=6} = -\Delta\phi_{ul}$ the figure shows that pitch aligned components of the $n = 2$ and $n = 6$ fields are maximised at roughly the same experimental coil phase difference. It has been suggested[109] that stochasticity may be enhanced by sidebands, which 'fill in' the gaps between islands of the dominant harmonic. The observation that the $n = 2$ and $n = 6$ islands are maximised at the same $\Delta\phi_{ul}$ supports this theory.

3.5.2 Robustness of Peeling Response to X Point Truncation

In tokamaks with divertor exhaust systems, the plasma boundary has at least one 'X point' where the poloidal field becomes zero. At an X point, the safety factor q tends to infinity and so is undefined. In the MARS-F flux-based coordinate system this corresponds to a numerical singularity, and so the X point must be excluded by truncation of the plasma boundary. Truncation effectively imposes a finite q at the plasma edge q_a , by approximating the divertor configuration as a limiter configuration with an otherwise similar shape. Previous works [93] suggest that the predicted amplification of the low n edge localised peeling response may be dependent on the truncation around the X point and the value of the edge safety factor q_a .

To test how robust the predicted peeling response is with respect to changes in X point plasma boundary shape, the plasma response was recomputed as the plasma boundary shape was incrementally changed towards including an X point. Figure 3.13 shows the plasma boundaries resulting edge q profiles of 5 equilibria. The equilibria are indistinguishable apart from their differing levels of truncation around the X point, and resulting different edge q profiles. Figure 3.14 shows the edge spectrograms computed for different levels of X point truncation and consequent values of q_a . The spectrograms show a radial distortion of the spectrum, but the amplitude of the peeling response does not appear to change. Varying the edge q profile shifts the $m = nq(s)$ curve, and the spectrogram is radially distorted to follow its movement. This suggests that although the response at a particular radial position s may change when the extent of truncation is varied, the response at a particular value of q does not. Figure 3.15 shows the edge radial profile of the $m = 12$ harmonic of the total magnetic perturbation, for the 5 values of q_a used in the study. The figure shows that when $q(s)$ is used as the radial coordinate, varying the edge truncation has no effect on the amplified peeling response. This suggests that altering the edge q profile does not affect the amplitude of the peeling response, but only

radially distorts its spectral structure. This suggests that the amplified peeling response would not disappear in the limit of including the X point. It is emphasised that this is a study of the resonant field amplification of *stable* edge peeling modes, not the stability of these modes. Therefore this finding is not in conflict with previous works [110] which find that the stability of edge peeling modes can be sensitive to the inclusion or exclusion of an X point.




FIGURE 3.13: Variation of plasma boundaries and edge q profiles while incrementally sharpening the X point, approaching a separatrix. A sharper (ie, less truncated) X point results in higher values of q_a in the equilibrium construction.




FIGURE 3.14: Edge regions of spectrograms from the 5 geometries, corresponding to 5 values of q_a . Increasing q_a shifts the $m = nq$ line, not altering the peeling response amplitude, but 'compressing' it radially.

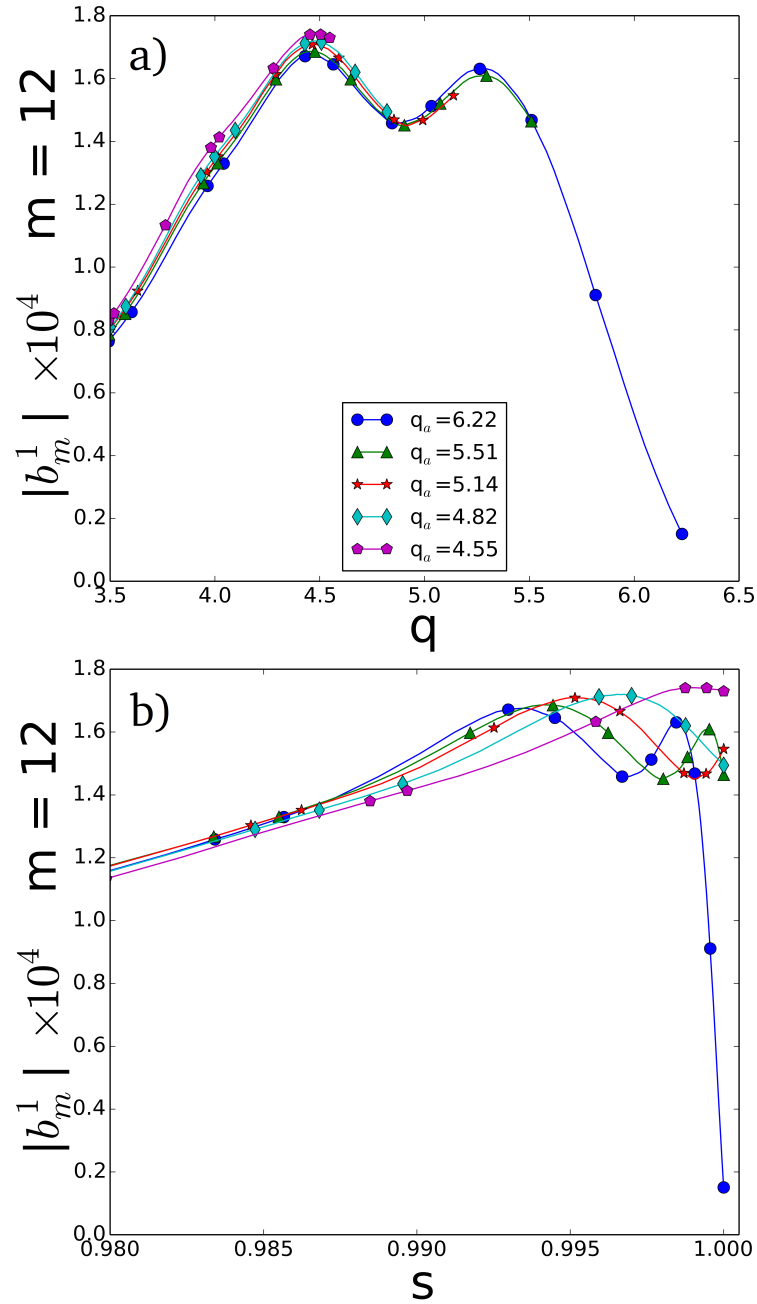


FIGURE 3.15: a) The radial profile near the edge of the $m = 12$ harmonic of the total magnetic perturbation, using q as radial coordinate. b) The radial profile near the edge of the $m = 12$ harmonic of the total magnetic perturbation, using s as radial coordinate. These plots demonstrate that when s is used as a radial coordinate, the peeling response appears sensitive to the extent of X point truncation. Conversely when safety factor q is used as the radial coordinate, the edge profiles coincide up to each value of q_a .

3.6 Poloidal Harmonic Coupling on ASDEX-Upgrade

Analytic theory indicates 3 primary sources of poloidal harmonic coupling in tokamak plasmas[111]: toroidicity causes coupling of adjacent poloidal harmonics with $\Delta m = |m - m'| = 1$, ellipticity couples harmonics with $\Delta m = 2$, and triangularity couples harmonics with $\Delta m = 3$. In the Straight Field Line (SFL) coordinate system which MARS and other spectral codes use, which is necessary for a meaningful definition of a poloidal harmonic, geometrical and physical quantities are not meaningfully separable, so poloidal coupling can in principle occur between any pair of poloidal harmonics. Poloidal coupling manifests both in the vacuum field, and the plasma response.

To investigate the generic behaviour of poloidal coupling, MARS-F was used in an atypical manner. Instead of the applied perturbation being due to RMP coils, the perturbation was applied with a prescribed magnetic perturbation b_{BC}^1 , applied at the plasma boundary. This is a valid procedure, because the perturbation at any closed surface (eg, the plasma boundary) completely determines the perturbation inside that surface, subject to the same physics as a conventional RMP coil calculation. A magnetic perturbation consisting of single poloidal harmonic m_{BC} with unit amplitude ($b_{BC}^1 = 1.0$ for $m = m_{BC}$) with toroidal mode number $n = 2$, was applied at the plasma boundary, and the resulting vacuum and total magnetic fields in the plasma bulk were computed. Figure 3.16 shows spectrograms of a vacuum field and total field, produced by applying unit amplitude $m_{BC} = 7$ and $m_{BC} = 13$ magnetic perturbations to the plasma edge. The perturbation is purely single m at the plasma edge; however, in the plasma bulk, the poloidal spectrum is broadened by poloidal harmonic coupling, and also shifted to lower $|m|$. This behaviour appears to be general for any applied m_{BC} . Figure 3.17 shows the absolute values of the pitch aligned components, which resulted from applied boundary perturbations with $m_{BC} = 11, 12$ and 13 . For these applied perturbations, m_{BC} were chosen to have no corresponding rational surface in the plasma (ie, $m_{BC} > n * q_a$). Therefore, the pitch aligned components can be non-zero only by poloidal harmonic coupling. These results show that a single m perturbation of unit amplitude applied at the plasma boundary, can drive the pitch aligned components to be appreciable relative to the size of the applied perturbation (in this work, up to $|b_{m=nq}^1|/|b_{BC}^1| = 0.2$). This result suggests that poloidal harmonic coupling is a viable mechanism by which the amplified peeling response may drive the pitch aligned components.

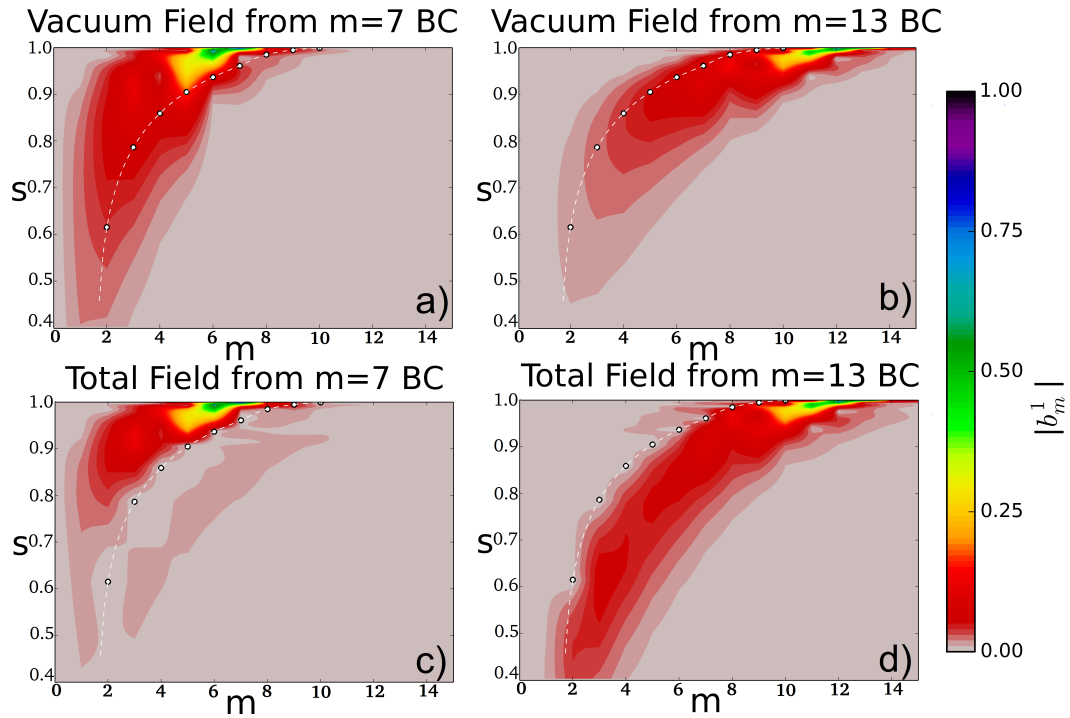


FIGURE 3.16: Spectrograms showing the perturbation which results from the application of a unit amplitude single harmonic perturbation, with $m_{BC} = 7$ and $m_{BC} = 13$, as a boundary condition at the plasma edge. Both the total and vacuum fields are shown. At the edge the perturbation is prescribed to be single m , but in the plasma bulk the spectrum broadens and shifts towards lower $|m|$ by poloidal harmonic coupling. Note that the colormap in these plots is reversed with respect to previous spectrograms.




FIGURE 3.17: The pitch aligned components of the total fields resulting from applied $m_{BC} = 11$, 12, and 13 unit amplitude perturbations at the plasma boundary. The toroidal mode number used was $n = 2$. The absolute value of the pitch aligned components is up to 20% of the applied field in this case. Since there are no $m = 11$, 12 or 13 rational surfaces in the plasma, these pitch aligned components may only be non-zero due to poloidal harmonic coupling.

3.7 Experimental Validation of MARS Simulations

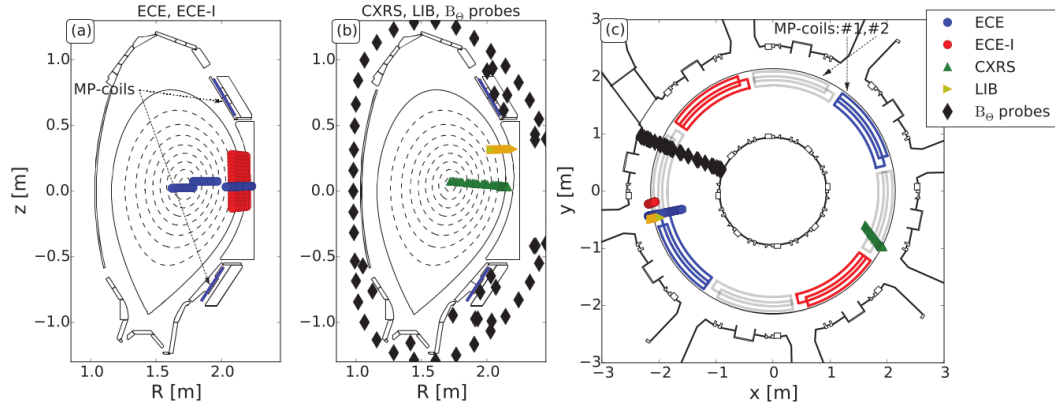


FIGURE 3.18: Figure reproduced with permission from [102]. Sketch of the locations of which various diagnostics on ASDEX Upgrade may measure the plasma displacement and magnetic perturbation in response to applied RMPs.

The work detailed in this section was previously presented in [102–104], where the background and implications of these results are discussed in detail. This section focuses only on the results as further benchmarking and a partial experimental validation of simulation results from the MARS-F code. My contribution to these works was all MARS-F simulations, and investigations of the effects of equilibrium uncertainty on the predicted plasma displacement. All results from other codes such as VMEC and M3D-C1, as well as experimental procedures, analysis and results, were the work of collaborators.

The plasma response manifests both as a displacement of plasma surfaces (ξ), and as a magnetic perturbation (b), both of which may be measured in experiments. The magnetic plasma response is measured using arrays of magnetic pickup coils, as sketched in figure 3.18. Since pickup coils measure the rate of change of magnetic field enclosed in the loop, which must then be integrated, the signal of the RMPs and the plasma response is optimised when the applied RMPs are rotated rigidly in the toroidal direction, which is achieved by sinusoidal variation of the RMP coil currents. Toroidal rotation of the RMP also allows the toroidal structure of the plasma response to be measured with other static diagnostics. However, requiring a rigid rotation can be disadvantageous, because the requirement of sinusoidal toroidal waveform reduces the maximum value which can be applied to the coils, and also the rotation causes significant attenuation by the passive stabilisation loops installed on ASDEX Upgrade. Pickup coils measure the magnetic field outside the plasma edge, but the magnetic perturbation inside the plasma bulk cannot be directly measured. However, by collecting radiation emitted from the plasma bulk, the plasma displacement may be measured globally. If the displacement is known then the magnetic perturbation may be deduced from it (neglecting resistivity), via the ideal linearised induction equation below.

$$\mathbf{b} = \nabla \times (\boldsymbol{\xi} \times \mathbf{B}) \quad (3.1)$$

The plasma displacement is measured using the electron cyclotron emission (ECE) diagnostic, soft x-ray (SXR) diagnostic, and charge exchange recombination spectroscopy (CXRS) diagnostic. The ECE, SXR and CXRS lines of sight and locations are sketched in figure 3.18. These diagnostics are used to make fine measurements of the plasma temperature profiles, and from these measurements small movements of flux surfaces are deduced, based on the approximation that temperatures are constant on flux surfaces. The displacement may be computed as $\xi_r = \left(\frac{\partial T_{rad0}}{\partial r} \right)^{-1} \delta T_{rad}$, where δT_{rad} is the perturbation of the radiation temperature due to the plasma response, and $\frac{\partial T_{rad0}}{\partial r}$ is the slope of the average radiation temperature profile, both measured by the ECE (for electrons) and CXRS (for ions) diagnostics[104]. Accurate determinations of the plasma displacement from the ECE, CXRS and SXR measurements are generally non-trivial and are not reported here, but are described in [102, 104] and references therein. The MARS-F model has been extensively validated against experiments by previous works. In [112, 113], MARS-F predictions of the total magnetic field and plasma displacement in response to applied RMPs, are compared with measurements on DIII-D. Good agreement was found with MARS-F until the pressure approached or exceeded the no-wall limit, at which point the kinetic MARS-K was required to recover the experimental measurements. Continued experimental validation of MARS-F on other machines is vital to build and maintain confidence in its predictions.

Plasma Response Measurement Experiments

Shot	n	q_{95}	β_N	$n_e \times 10^{19} m^{-3}$
30839	2	5.46	1.5	7.2
31021	1	4.33	1.6	10.0
32138	1	5.12	2.1	5.9
31034	1	4.17	2.5	9.3

TABLE 3.2: Equilibrium parameters of the discharges used for experimental measurements of the plasma response.

In [102], a rigidly rotating $n = 2$ RMP is applied to ASDEX Upgrade discharge number 30839, while in [103, 104] a rigidly rotating $n = 1$ RMP is applied to discharge numbers 31021 and 32138. To investigate experimentally the dependence of the $n = 1$ plasma response on the poloidal spectrum, a scan of coil phase $\Delta\phi_{ul}$ was performed in discharge 31034, which was otherwise a repeat of 31021. Global plasma parameters of these discharges are found in table 3.2. The plasma displacement in response to the

applied RMP is predicted using MARS-F, VMEC and M3D-C¹, and from these simulation results synthetic diagnostic data are produced. In [102], synthetic ECE data are compared with experimental ECE measurements of the plasma displacement in the edge region. Both the predicted and measured displacements exceed vacuum computations, which indicates an amplification of the displacement by the plasma response, and supports the existence of the edge peeling response. In [103, 104], ECE, CXRS and SXR measurements are combined for a global measurement of the plasma displacement as described in section 3.7. Magnetic probes around the plasma edge also detect a strong amplification of the applied RMP, further evidencing a strongly amplified edge peeling response.




FIGURE 3.19: Figure reproduced with permission from [102]. Poloidal spectra of the $n = 2$ magnetic and displacement plasma response for ASDEX Upgrade discharge 30839, computed with MARS-F and VMEC. a,d) The applied vacuum magnetic perturbation. b,e) The total magnetic perturbation including the plasma response. The expected magnetic edge peeling response is clearly visible in the $m > nq$ region. c,f) Plasma displacement due to the applied RMP. The displacement is also concentrated near the edge, but unlike the magnetic response is mostly resonant ($m = nq$).

Computational predictions of the plasma response to the applied $n = 2$ RMP in discharge 30839 are plotted in figure 3.19. The figure shows the radial profiles of the poloidal harmonics of the predicted vacuum magnetic field, total magnetic field including plasma response and plasma displacement, computed with VMEC and separately with MARS-F. The figure shows reasonable agreement in the magnetic perturbation between the two codes, and a close quantitative agreement in displacement. Some difference in predicted magnetic field between the two codes is expected, given the radical differences in approach of the two codes. Originally designed to solve for stellarator equilibria, VMEC[114] solves for a 3D equilibrium by minimising the potential energy of the plasma. Conversely, MARS-F applies a magnetic perturbation to a given 2D equilibrium, and




FIGURE 3.20: Poloidal spectra of the $n = 1$ magnetic and displacement plasma response for ASDEX Upgrade discharge 31021, computed with MARS-F. a) Applied vacuum field. b) Total field including plasma response. There is an edge peeling response, but also a core kink response is apparent. c) The plasma displacement is resonant, and significant core modes are apparent.

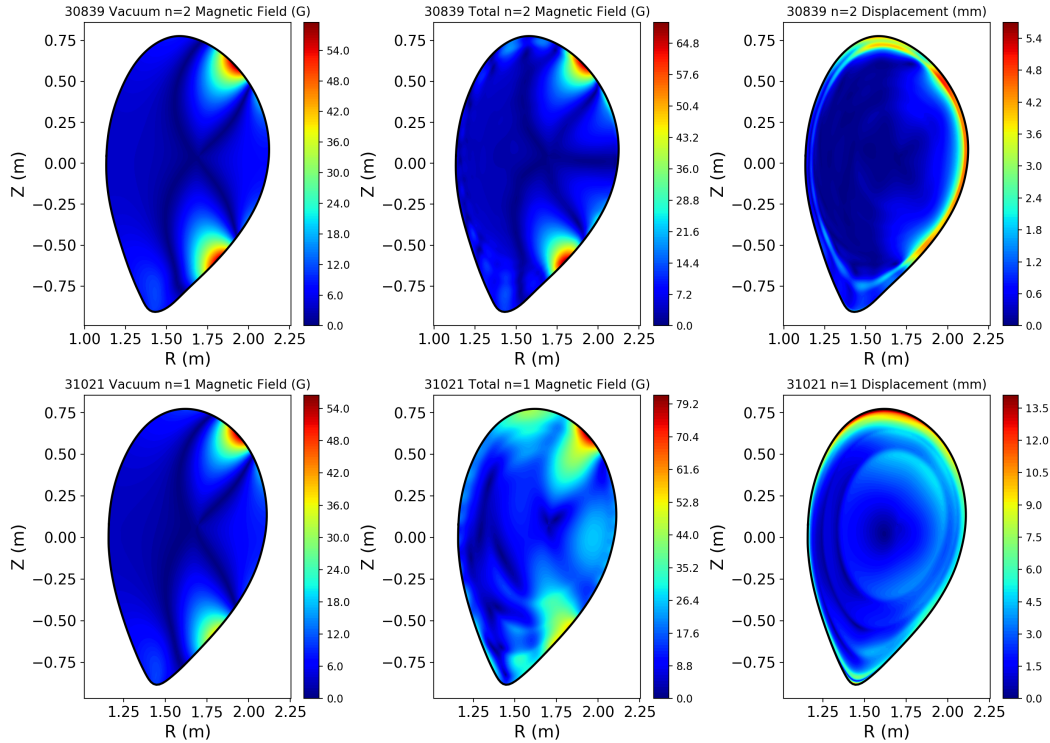


FIGURE 3.21: Real space (R,Z) structure of the magnetic and displacement response to applied $n = 1$ and $n = 2$ RMPs. For 30839 the applied field was an $n = 2$, whereas for 31021 the field was $n = 1$. a,d) Vacuum fields applied to discharges 30839 and 31021. b,e) Total field including plasma response of the $n = 1$ and $n = 2$ experiments. c,f) Plasma displacement of the $n = 1$ and $n = 2$ experiments. The $n = 2$ displacement is mostly concentrated close to the edge, while the $n = 1$ displacement is more global. Because the displacement is resonant it is concentrated on rational surfaces, which are apparent in the figure.

solves the linearised equations of MHD for the total perturbation. Furthermore, VMEC does not include resistivity or toroidal plasma rotation, while MARS-F does.

Figure 3.20 shows the radial profiles of the poloidal harmonics of the computed vacuum and total magnetic fields and plasma displacement for $n = 1$ RMP experiments 31021. The magnetic vacuum field and plasma response appear qualitatively quite similar in the $n = 1$ and $n = 2$ experiments, the main differences being that the resonant condition $m = nq$ follows a different spectral curve, and that the amplification effect is larger in the $n = 1$ case. In the displacement spectrograms, both the $n = 1$ and $n = 2$ have large peaks close to the plasma edge; however, the $n = 1$ also features a prominent core response which is absent in the $n = 2$ displacement spectrogram. This is because the q profile of the $n = 1$ discharge is close to 1 and quite flat in the core, so the resonant condition $q \approx m/n$ (discussed below) is approximately satisfied over a larger radius.

Both figures 3.19 and 3.20 show the non-resonant ($m > nq$) edge peeling response in the magnetic spectrogram, and a strongly resonant response ($m = nq$) of the plasma displacement. To reveal the cause of the differences in poloidal structure between the magnetic field and plasma displacement, we may take the linearised ideal induction equation 3.1, and apply it to the case of a general tokamak plasma. The result (a complete derivation is included in Appendix A) is a relation between b_1 and ξ_1 , the magnetic and displacement quantities plotted in figure 3.19.

$$\xi_1 \approx b_1 / (m - nq) \quad (3.2)$$

It can be seen from the above relation that the displacement ξ_1 is expected to be dominated by the resonant components, where $m = nq$. Figure 3.21 shows the bulk vacuum and total magnetic fields and plasma displacement of the $n = 1$ and $n = 2$ experiments, computed with MARS-F. Comparing the $n = 1$ and $n = 2$ magnetic perturbations shows that the $n = 1$ plasma response has a larger amplification effect, and is more global, due to the slightly higher plasma pressure, and having a flat core q profile where the above resonance is close to satisfied.

Poloidal Mode Number Measurement

In [102] using spatially resolved Electron Cyclotron Emission Imaging (ECE-I) measurements, the poloidal structure of the displacement at the $s = 0.97$ flux surface was experimentally measured, and also predicted using VMEC and MARS-F modelling. Using an approximate Fourier representation of the displacement (approximating the displacement at a flux surface to be dominated by a single poloidal harmonic), $\xi = \xi_a e^{i(m\chi - n\phi)}$




FIGURE 3.22: Figure reproduced with permission from [102]. Generalised straight field line poloidal angle χ of ECE-I channels on a single flux surface ($q = 5.35$), plotted against the phase of the measured displacement perturbation, derived from the toroidal rotation of the perturbation by angle ϕ . Blue and green points are the measured data from the ECE-I diagnostic, which demonstrate close agreement with synthetic data produced using VMEC. Linear fit lines are used to determine the dominant poloidal harmonic, which matches expectations within uncertainties.

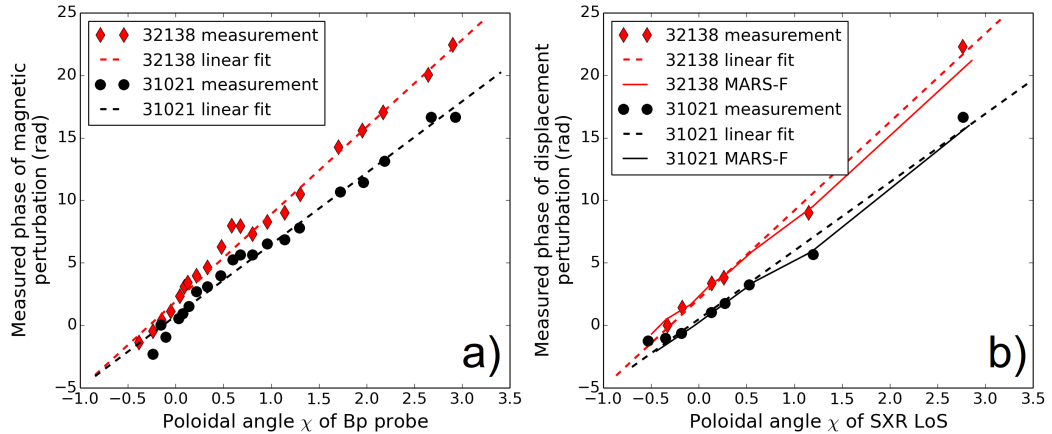


FIGURE 3.23: Figure reproduced with permission from [104]. a) Generalised straight field line poloidal angle χ of intersections of SXR lines of sight with the plasma edge, plotted against the relative phase of the $n = 1$ plasma displacement. b) Poloidal angle χ of B_p magnetic probes, plotted against the relative phase of the $n = 1$ measured magnetic field, which includes both the vacuum field and plasma response. In both plots the phase of the $n = 1$ displacement and total magnetic field at the corresponding SFL poloidal angle is also plotted, showing good agreement with experimental measurements.

where ϕ is the toroidal angle and χ the SFL poloidal angle, it is apparent that the dominant poloidal mode number of the perturbation at a particular flux surface may be found as $m = \Delta\phi/\Delta\chi$, where $\Delta\phi$ is an increment in toroidal phase corresponding to an increment in poloidal angle of $\Delta\chi$. Several channels of the ECE-I diagnostic localised around the $s = 0.97$ surface ($q \approx 5.35$) at different poloidal locations allow the displacement to be measured at varying χ with approximately constant s and q , which provides the $\Delta\chi$. The rotation of the applied RMP induces the resulting displacement to rotate toroidally in synchrony, allowing the perturbation to be measured at varying ϕ . In figure 3.22, the toroidal phase shift of the displacement from some fixed toroidal offset $\Delta\phi$ measured at distinct ECE-I channels, is plotted against the generalised poloidal angle χ of each ECE-I channel. The plot includes measured data from the ECE-I diagnostic, as well as synthetic data based on the predictions of the displacement made with VMEC and MARS-F plotted in figure 3.19. The plot shows a roughly constant $\Delta\phi/\Delta\chi$, indicating a single dominant poloidal harmonic at the $q \approx 5.35$ surface. A resonant poloidal harmonic of an $n = 2$ perturbation at the $q \approx 5.35$ surface, leads us to expect $m = nq \approx 10.7$. The synthetic ECE-I data produced using a VMEC simulation of the plasma displacement predicts $m = 10.72 \pm 0.63$, and the measured ECE-I data shows $m = 9.83 \pm 0.98$, in agreement within uncertainties.

In [103, 104], the poloidal structure at the plasma edge of both the $n = 1$ plasma displacement and total magnetic field are measured, using poloidal arrays of B_p probes and SXR channels with lines of sight tangent to the plasma edge at varying poloidal locations. Two B_p arrays separated toroidally by 180 degrees were used for clean extraction of the $n = 1$ component of the magnetic perturbation. Figure 3.23 a) shows the toroidal phase delay of the experimentally measured magnetic field against poloidal location of the B_p probes in SFL coordinates. As previously, the gradient of the trend line indicates the dominant poloidal mode number of the perturbation near the plasma edge. The measured dominant mode numbers for both shots closely match the dominant poloidal mode numbers predicted by MARS-F simulations. Figure 3.23 b) plots the experimentally measured phase of the displacement at the plasma edge as a function of poloidal angle in SFL coordinates, and a linear fit to the data. The solid lines are equivalent measurements from synthetic diagnostics generated using MARS-F computed displacement profiles. The figure shows that the experimentally measured dominant poloidal harmonic closely matches the MARS-F computed dominant poloidal harmonic in both the magnetic and displacement perturbations, and for two distinct discharges.

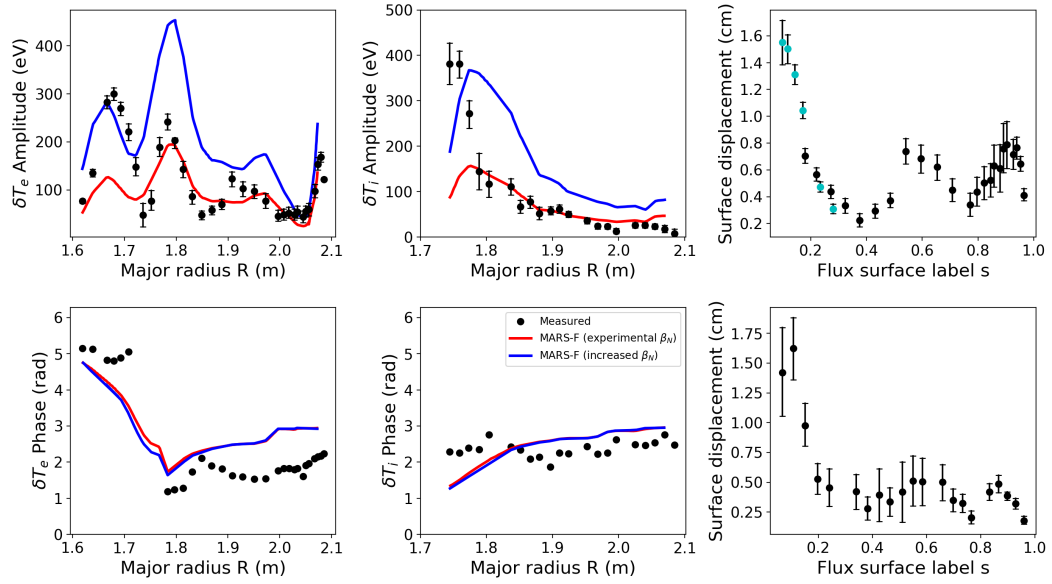


FIGURE 3.24: Figure reproduced with permission from [104], showing the radial profile of the plasma displacement in discharge 32138. a, d) Amplitude and phase of the electron temperature perturbation δT_e caused by the plasma displacement ξ_r . Black circles show directly measured δT_e , while the red line shows δT_e resulting from the MARS-F predicted displacement. b, e) Amplitude and phase of the ion temperature perturbation δT_i with MARS-F prediction. c, f) Plasma displacement profiles measured using ion (CXRS) and electron (ECE) perturbations.

Displacement Radial Profile Validation

Figure 3.24 shows the experimentally measured profiles of the radiation temperature perturbation δT_{rad} , measured with spatially resolved ECE and CXRS channels. The ECE diagnostic measures the electron temperature profile and perturbation, while the CXRS measures the ion temperature profile and perturbation. Measurements shown in the figure are taken during discharge 32138 between 2.2-3.6s. Using MARS-F the displacement profile of the perturbation ξ_r is computed, and the simulated ion and electron temperature perturbations are computed using $\delta T_{rad} = \xi_r \left(\frac{\partial T_{rad0}}{\partial r} \right)$. The synthetic $\delta T_{e,i}$ data derived from MARS-F predicted displacements agree qualitatively for both the ion and electron temperatures, and match quite well quantitatively outside the core, with the exception of the phase of δT_e which deviates significantly from the measured phase. In the core region, the large measured displacement peak is not reproduced by synthetic data.

Displacement Coil Phase Dependence Validation

In discharge 31034, the poloidal spectrum of the applied $n = 1$ RMP is varied by scanning the coil phase difference between the upper and lower rows $\Delta\phi_{ul}$. Figure 3.25 shows the




FIGURE 3.25: Figure reproduced with permission from [104]. Variation of the $n = 1$ plasma displacement with $\Delta\phi_{ul}$ at a) the plasma core and b) the plasma edge. Also plotted are predictions from MARS-F and M3D-C¹. The experimental result and two codes agree fairly well in coil phase dependence. Predictions of the two codes are not drastically dissimilar in amplitude, but they both overestimate the experimental measurements. This may be a fault with the comparison, since the experimental measurements are both line averaged quantities, while the numerical predictions are made locally.

measured amplitude of the core ($s=0.2$) and edge ($s=0.95$) plasma displacement as $\Delta\phi_{ul}$ is experimentally scanned, measured using the SXR diagnostic as reported in [103]. The experimental measurement is averaged over several scans of $\Delta\phi_{ul}$. These scans are also performed using codes MARS-F and M3D-C¹[115]. Both codes MARS-F and M3D-C¹ closely match each other in amplitude and $\Delta\phi_{ul}$ dependence, as well as matching the $\Delta\phi_{ul}$ dependence of the experimental measurement. The amplitude of the experimental measurement is not recovered however. It is suggested in [104] that this may be because the experimental measurement is a line averaged measurement, while the numerical predictions are for the displacement local to a flux surface, making a direct comparison not fully rigorous.

Computed Displacement Uncertainty

Figure 3.24 demonstrates a significant underestimate of the core displacement perturbation by MARS-F. In this section, it will be investigated whether this discrepancy may be caused by uncertainty in the equilibrium reconstruction. Two of the distinguishing features of these $n = 1$ discharges is a q profile which is close to 1 and quite flat in the core, and also relatively high plasma pressure, with β_N generally a large fraction of the no-wall limit. The plasma response and resonant field amplification are known to be sensitive to the q profile and plasma pressure[112, 116, 117], but the core safety factor and core pressure profile shape are determined rather vaguely. Due to the resonance




FIGURE 3.26: Profiles of a) plasma pressure, b) current density and c) safety factor, as the core safety factor was scanned using CLISTE, keeping other quantities fixed. All the resulting equilibria met the constraints imposed by the edge magnetic measurements, and so are all plausible equilibrium profiles.

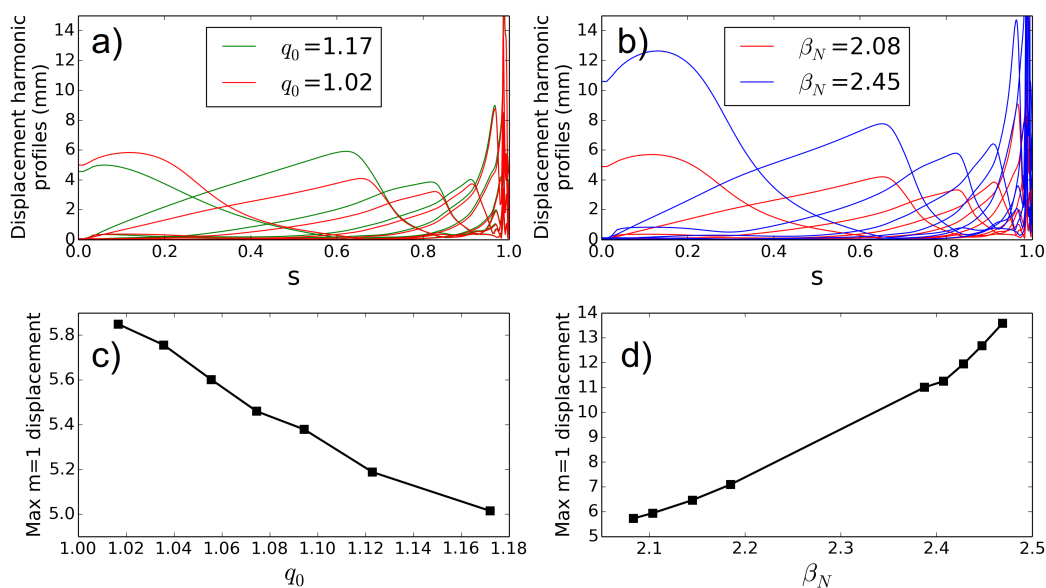


FIGURE 3.27: Figure reproduced with permission from [104]. The core safety factor was scanned within its uncertainty, and the pressure scanned through the range of values it took during the discharge. The core response is shown to be sensitive to both parameters. The above plots demonstrate a scan in core safety factor q_0 with constant β_N . Although not shown here, the plasma β_N was also scanned by raising the pressure profile and adjusting the other parameters self consistently.

condition expressed in equation 3.2, the displacement is expected to be especially sensitive to small changes in q in this instance, where q is only slightly above 1 for large regions of the plasma bulk. The uncertainty in the core q and pressure is not wholly due to measurement uncertainty, but also because for a given set of magnetic measurements (which comprise the primary constraints for equilibrium reconstruction) there is no unique solution for the equilibrium near the core. Further uncertainty is added to the core pressure by lack of data of the fast ion population, which even in current tokamaks may constitute a non-negligible correction to the core plasma pressure.

In order to quantify the effect of the uncertainty in equilibrium reconstruction on the computed plasma displacement, the core safety factor and pressure of equilibrium 32138 were varied within plausible uncertainties ($q_0=1.0-1.2$, $\beta_N=2.08-2.45$), and the plasma displacement recomputed. Figure 3.26 summarises the scaling of the equilibrium, which was performed using the CLISTE equilibrium code for scaling q_0 , and CHEASE for β_N . Figure 3.27 shows the variation in computed plasma displacement resulting from the changes to the equilibrium. As expected, the maximum core displacement increases as q_0 approaches 1, and also increases with β_N . It appears that uncertainty in q_0 alone is insufficient to explain the discrepancy between the computed and observed core displacement. By increasing β_N however, the computed core displacement can be made to approach the experimentally observed value. In figure 3.24 the computed temperature perturbation for $\beta_N = 2.45$ (blue line) is shown, demonstrating an improvement in regions which match poorly when the experimental value of $\beta_N = 2.08$ (red line) is used. However, since the pressure profile is scaled self-similarly in this study, the edge localised displacement harmonics are also amplified resulting in the loss of agreement between simulation and experiment at the edge. It is suggested in [104] that global agreement may be improved with localised modifications to the plasma pressure profile, such as an increase in only the core pressure rather than a global pressure increase. However, quantifying and reducing the uncertainties in core pressure and safety factor, for example by including MSE constraints in the equilibrium reconstruction[118], would be more rigorous.

3.8 Chapter Results Summary

The single fluid resistive plasma response to a static $n_{dom} = 2$ applied RMP field on ASDEX-Upgrade was investigated using the MARS-F code. A benchmark of the vacuum field was first performed by comparing the $n = 2$ vacuum field computed by MARS with the full toroidal spectrum vacuum field computed by ERGOS. It was found that while the $n = 2$ components agree, the bulk field showed that the global full spectrum

ERGOS field differed significantly from the single n MARS field. This was remedied by computing the largest toroidal sideband of the field - the $n = 6$ sideband - and summing the $n = 2$ and $n = 6$ fields. The resulting MARS field compared well with the ERGOS field globally, except in the very close vicinity of the coils. The MARS-F code was used to compute the plasma response to the applied $n_{dom} = 2$ field, and a strongly amplified low n edge peeling response is predicted, localised near the edge in the $s > 0.7$, $m > qn$ region of the magnetic spectrum. It is worth clarifying that the designation of 'low n peeling response' is used to identify this classification of plasma response with previous works [92, 93, 119], not to imply that these results are general for all low n numbers. Both the $n = 2$ and $n = 6$ response were computed, and both show an edge localised peeling response, although the $n = 2$ response also comprises an amplified core kink mode which is not present in the $n = 6$ response. Consistent with previous findings[93], the $n = 6$ response was also much smaller than the $n = 2$, primarily due to the smaller size of the $n = 6$ sideband of the applied field. It is also observed that when finite resistivity is included, resistive dissipation prevents the outermost pitch aligned components near the plasma edge from being completely screened, and therefore they can be finite and even comparable to their vacuum values. Both the amplified peeling response and pitch aligned components have a strong dependence on the phase difference between the upper and lower coil sets $\Delta\phi_{ul}$. Furthermore, the value of $\Delta\phi_{ul}$ which maximises the total pitch aligned components is offset from optimum vacuum pitch alignment by 60° . This result reaffirms the necessity of including the plasma response when calculating the optimum coil phase for stochasticity in RMP experiments; a vacuum prediction would result in a 60° misalignment. It is also noted that the coil phase difference for the $n = 6$ component only was in practise merely the negative of the phase difference for the $n = 2$ component, ie, $\Delta\phi_{n=6} = -\Delta\phi_{ul}$. When accounting for this, it is apparent that the $n = 2$ and $n = 6$ components are maximised at roughly the same values of $\Delta\phi_{ul}$, supporting a recent proposal that sidebands enhance stochasticity induced by RMPs by filling in gaps between the islands of the dominant harmonic[109].

Reducing the plasma boundary truncation such that the boundary shape approached an X point, and thereby increasing q_a , showed no tendency to reduce the amplified peeling response. The peeling response was observed to be distorted radially, following the movement of the spectral line $m = nq(s)$ which is moved by changes to q_a . To rephrase, when q is used as the radial coordinate rather than s , the amplified peeling response is insensitive to numerical truncation of the X point.

A study of poloidal harmonic coupling was undertaken in realistic tokamak geometry, to investigate the suggestion that the amplified peeling response could drive the pitch aligned components by poloidal harmonic coupling. Single m perturbations were applied

at the plasma edge as boundary conditions, and the resulting bulk perturbation computed using MARS-F. The resulting magnetic spectrograms demonstrated that even when perturbations have only a single poloidal mode number at the plasma surface, poloidal harmonic coupling causes the spectrum to broaden in the plasma bulk, and also shift towards lower $|m|$. It is also demonstrated computationally that perturbations spectrally above pitch alignment ($m > nq(s)$) can drive the pitch aligned components by poloidal harmonic coupling, suggesting that the amplified peeling response may drive the outermost pitch aligned components.

The plasma displacement induced by applied $n = 1$ and $n = 2$ RMPs was measured experimentally using ECE, CXRS and SXR diagnostics, and the magnetic perturbation measured using a poloidal array of magnetic probes. The experimentally measured poloidal structure of the total magnetic field and plasma displacement, the radial profile of the displacement, and the dependence of the displacement on $\Delta\phi_{ul}$, are all well recovered by MARS-F simulations, providing partial validation of the MARS-F code. However the amplitude is not, possibly due to uncertainties in equilibrium reconstruction. Using constraints on the core q profile from MSE measurements may reduce this in the future. Furthermore, as part of this study, MARS-F is successfully benchmarked against VMEC and M3D-C¹.

Chapter 4

Coil Phase Optimisation

A robust correlation exists between the mitigated ELM frequency and outermost pitch aligned component b_{res}^1 [70]. As previously described, it is crucial that the ITER ELM control strategy be able to robustly reduce ELM sizes below material damage thresholds. These considerations motivate a study into the coil phase for optimal b_{res}^1 and hence ELM mitigation, on ASDEX Upgrade and ITER. The purpose is to develop and test a robust yet simple parametrisation for the optimum coil phase $\Delta\phi_{opt}$, which may form the basis of a feedback optimisation scheme to maximise the mitigated ELM frequency throughout ITER plasma discharges, particularly during ramp-up and termination during which the plasma parameters, and so possibly the optimum alignment, vary rapidly.

In this chapter, the dependence of $\Delta\phi_{opt}$ on global plasma parameters normalised beta β_N and q_{95} , is investigated and quantified numerically for general ASDEX Upgrade and ITER plasmas, in order to derive and test a simple parametrisation for the optimal coil phase for best ELM mitigation. Normalised beta β_N is here defined as $\beta_N = \beta(\%)a(m)B_0(T)/I_p(MA)$, where B_0 is the equilibrium magnetic field at the magnetic axis, I_p is the total plasma current, a is the plasma minor radius, and β is the normalised plasma pressure $\beta = 2\mu_0\langle p\rangle/\langle B^2\rangle$, where $\langle..\rangle$ denotes an average over the plasma volume.

This chapter is organised as follows. In section 4.1, by scaling the pressure and plasma current of a reference equilibrium of the ASDEX Upgrade tokamak, a set of equilibria spanning a wide parameter space in (β_N, q_{95}) is created, covering the typical ASDEX Upgrade parameter space. It is explained and demonstrated how alignment is affected by both β_N and q_{95} . In section 4.2, at each point in the set of (β_N, q_{95}) points, the plasma response to applied $n = 1 - 4$ perturbations is computed using MARS-F, and the optimal coil phasing $\Delta\phi_{opt}$ derived from these results. Consistent with previous studies [91, 93, 120], it is found that $\Delta\phi_{opt}$ increases smoothly with q_{95} . However, this

work also finds a weak decrease in $\Delta\phi_{opt}$ with β_N , in contrast with a previous experimental and computational study which found that $\Delta\phi_{opt}$ did not vary with β_N [91]. The dependence of $\Delta\phi_{opt}$ on (β_N, q_{95}) is parametrised with a simple 2D quadratic function, the coefficients of which are computed by linear regression and included here for researchers to use for RMP experiment planning on ASDEX Upgrade. In section 4.3, the efficacy of the 2D quadratic parametrisation of $\Delta\phi_{opt}$ is evaluated, by comparing the predictions of the 2D quadratic parametrisation with rigorous MARS-F computations for several 'benchmarking points'. A benchmarking point consists of a plasma equilibrium reconstruction, kinetic profiles fitted to experimental data, and applied RMP coil currents, all from ASDEX Upgrade experiments. The set of benchmarking points contained significant variation in β_N and q_{95} , as well as plasma shape, rotation and kinetic profiles. For each benchmarking point, the plasma response and optimal coil phase are computed using MARS-F, and by comparing the MARS-F results with the 2D quadratic parametrisation, it is shown that the 2D quadratic parametrisation predicts $\Delta\phi_{opt}$ accurately relative to a MARS-F computation to within 34.2 degrees of a plasma response computation for $n = 1$ RMPs, and within 21.3 degrees for $n = 2$ RMPs, for typical experimental (β_N, q_{95}) values. The benchmarking procedure for $n = 3, 4$ is left for future work due to limited experimental data with these n numbers. In section 4.4, possible sources of uncertainty in the 2D quadratic parametrisation are then investigated. By computing $\Delta\phi_{opt}$ for scans of the pedestal width and amplitude of toroidal rotation, it is found that $\Delta\phi_{opt}$ is insensitive to these parameters. The plasma shape is therefore suggested as the primary source of uncertainty in the 2D quadratic parametrisation of $\Delta\phi_{opt}$. In section 4.5, the optimal coil phase was derived from experimental measurements of the ELM frequency during experimental scans of $\Delta\phi_{opt}$, and compared to predictions of the (β_N, q_{95}) parametrisation and rigorous MARS-F computations. The measured optimum coil phase agreed with the quadratic parametrisation to within 40 degrees, and with MARS-F computations to within 35 degrees. In section 4.6, having thoroughly benchmarked and validated the optimal coil phase parametrisation, the procedure for deriving it is repeated for an ITER equilibrium and coil set, and the coefficients of the resulting parametrisation left to form the basis of a feedback coil optimisation system. In section 4.7, the set of aforementioned benchmarking points is statistically analysed, to search for relationships between figures of merit and equilibrium quantities. A correlation is found between b_{res}^1 , and peeling response and ξ_X , supporting previously developed theory[88], however no correlations were detected between figures of merit and equilibrium quantities. Section 4.8 consists of a short study of the optimisation of the RMP field by judicious choice of toroidal waveform. Section 4.9 summarises and discusses the results of the chapter.

4.1 Alignment of Applied Field with Equilibrium Field

4.1.1 Extracting $\Delta\phi_{opt}$ from Plasma Response Computation

As in the previous chapter, the linearity of the MARS model is exploited and the plasma response due to the upper and lower coils (b^u and b^l respectively) are computed separately, then the field due to both coils is reconstructed in post process as $b^t = b^u + b^l e^{-i\Delta\phi_{ul}}$. From this prescription of the field due to both coils, an analytic formula for the optimal coil phase may be derived from simple geometrical arguments in the complex plane. Let the outermost pitch aligned component due to the lower coils be $b_{res}^{1,l}$, and the outermost pitch aligned component due to the upper coils be $b_{res}^{1,u}$, then the total outermost pitch aligned component is $|b_{res}^1| = |b_{res}^{1,u} + b_{res}^{1,l} e^{-i\Delta\phi_{ul}}|$. Since $b_{res}^{1,l}$ and $b_{res}^{1,u}$ are simply complex valued scalars, $|b_{res}^1|$ is maximised when $b_{res}^{1,u}$ and $b_{res}^{1,l}$ are parallel in the complex plane. Therefore the optimal coil phase is simply the angle between $b_{res}^{1,u}$ and $b_{res}^{1,l}$ in the complex plane, given by equation 4.1 below.

$$\Delta\phi_{opt} = \pm \arccos \left(\frac{b_{res}^{1,l} \cdot b_{res}^{1,u}}{|b_{res}^{1,l}| |b_{res}^{1,u}|} \right) \quad (4.1)$$

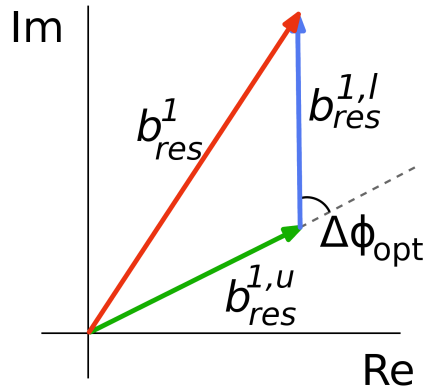


FIGURE 4.1: $b_{res}^{1,l}$ and $b_{res}^{1,u}$ are complex scalars, so the pitch aligned component $|b_{res}^1|$ is maximised when $b_{res}^{1,u}$ and $b_{res}^{1,l}$ are parallel in the complex plane.

This concept is depicted in figure 4.1. A numerical scan of $\Delta\phi_{ul}$ to determine $\Delta\phi_{opt}$ as used in previous works [95] is therefore redundant. There is sign uncertainty in Equation 4.1 however, so the sign of $\Delta\phi_{opt}$ is determined by comparing $|b_{res}^{1,u} + b_{res}^{1,l} e^{-i\Delta\phi_{opt}}|$ and $|b_{res}^{1,u} + b_{res}^{1,l} e^{+i\Delta\phi_{opt}}|$, and choosing the larger. Figure 4.2 shows that computing $\Delta\phi_{opt}$ using a scan of $\Delta\phi_{ul}$, yields the same result as Equation 4.1.




FIGURE 4.2: This figure shows $\Delta\phi_{opt}$ computed using Equation 4.1 and a $\Delta\phi_{ul}$ scan of $|b_{res}^1|$, as used in previous works [79, 95, 120]. It shows that the Equation 4.1 matches the $\Delta\phi_{opt}$ predicted using a scan of $\Delta\phi_{ul}$, which is a more time consuming approach which also limits the precision of $\Delta\phi_{opt}$ to the scan step size.

4.1.2 Effect on Alignment of Equilibrium Parameters β_N and q_{95}

The parameters q_{95} , β_N and n are the chosen variables for this dependence study, because the spectral alignment of the applied perturbation with the equilibrium field, and therefore optimum coil phase $\Delta\phi_{opt}$, is expected to be altered by changes to these parameters. Spectral alignment in this context refers to the extent to which the applied perturbation aligns with the $m = nq$ line in (m, s) space, where $s = \psi_N^{1/2}$ is the radial coordinate and m is the poloidal harmonic number. Figure 4.3 shows the effect of changing plasma β_N and q_{95} on the applied vacuum spectrum and equilibrium resonance line (which satisfies $m = nq$). Figure 4.3a) and b) show that increasing q_{95} or n (with constant β_N) shifts the $m = nq$ line in (m, s) space towards higher m , thereby moving the resonant component relative to the applied spectrum. Figure 4.3c) and d) show that changing β_N (for constant q_{95}) distorts the spectrum of the applied field relative to the $m = nq$ line as further explained below. These plots demonstrate why q_{95} and β_N are used as independent variables for this study.

Figure 4.4 explains the dependence of field alignment on β_N . It is only possible to define physically relevant poloidal harmonics in a straight field line (SFL) coordinate system, in which the poloidal angle is not the geometric poloidal angle θ , but the generalised poloidal angle χ . Generalised poloidal angle χ is defined such that field lines are straight in the (χ, ϕ) plane, where ϕ is the toroidal angle. Therefore, the poloidal coordinate χ and poloidal harmonic number m are defined relative to the magnetic equilibrium. In figure 4.4a,b) the SFL based grid is plotted for low and high β_N cases. In the high β_N

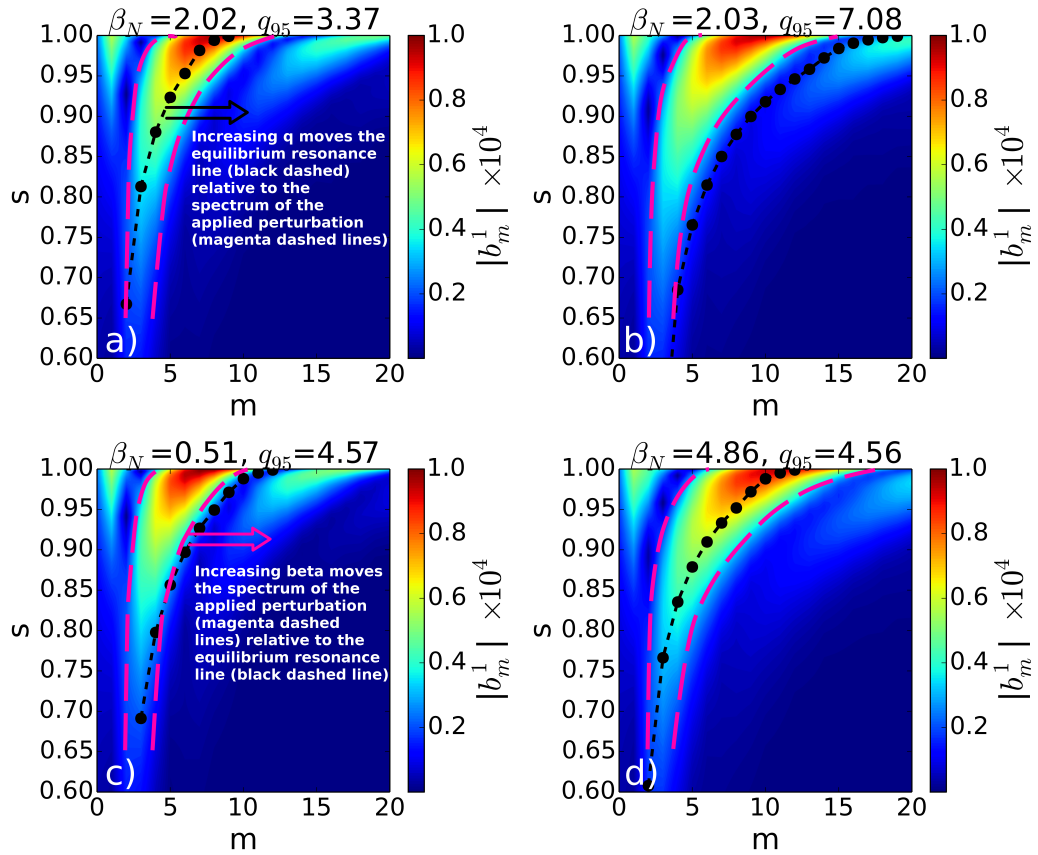


FIGURE 4.3: In a) and b) the $nq(s) = m$ line is plotted relative to the applied vacuum magnetic spectrum for a low (a) and high (b) q_{95} case, with the same value of β_N , demonstrating that changing q_{95} or n shifts the $nq(s) = m$ line relative to the spectrum of the applied perturbation, which changes the RMP alignment. In c) and d) the $nq(s) = m$ line is plotted relative to the vacuum spectrum for a low (a) and high (b) β_N case, with the same q_{95} . The plots show that increasing the plasma pressure β_N distorts the equilibrium magnetic geometry such that the vacuum spectrum is shifted to higher m .

case Shafranov shift causes the magnetic axis to be shifted outwards. Lines of $\chi = \pm 15$ degrees and the plasma boundary enclosed by this angular range are highlighted in black solid lines. They show that in the high β_N case a given range of χ encloses a larger arclength of the low field side plasma boundary than in the low β_N case. Since the mapping from geometric angle θ to SFL angle χ changes with β_N , the field in SFL coordinates is also distorted by the change in β_N , even though the applied fields in real (R, Z) space are identical in both cases. Figure 4.4c) plots the geometric angle $\hat{\theta}$ (where $\hat{\theta}$ is measured from an origin common to both the high and low β_N cases) against the generalised poloidal angle χ . The figure shows that increasing β_N changes the mapping of $\hat{\theta}$ to χ . Figure 4.4d) shows the resulting effect on the vacuum magnetic perturbation at the plasma edge in terms of χ . The figure shows that in the high β_N case, the main features of the applied perturbation are compressed into a smaller range of χ , and the

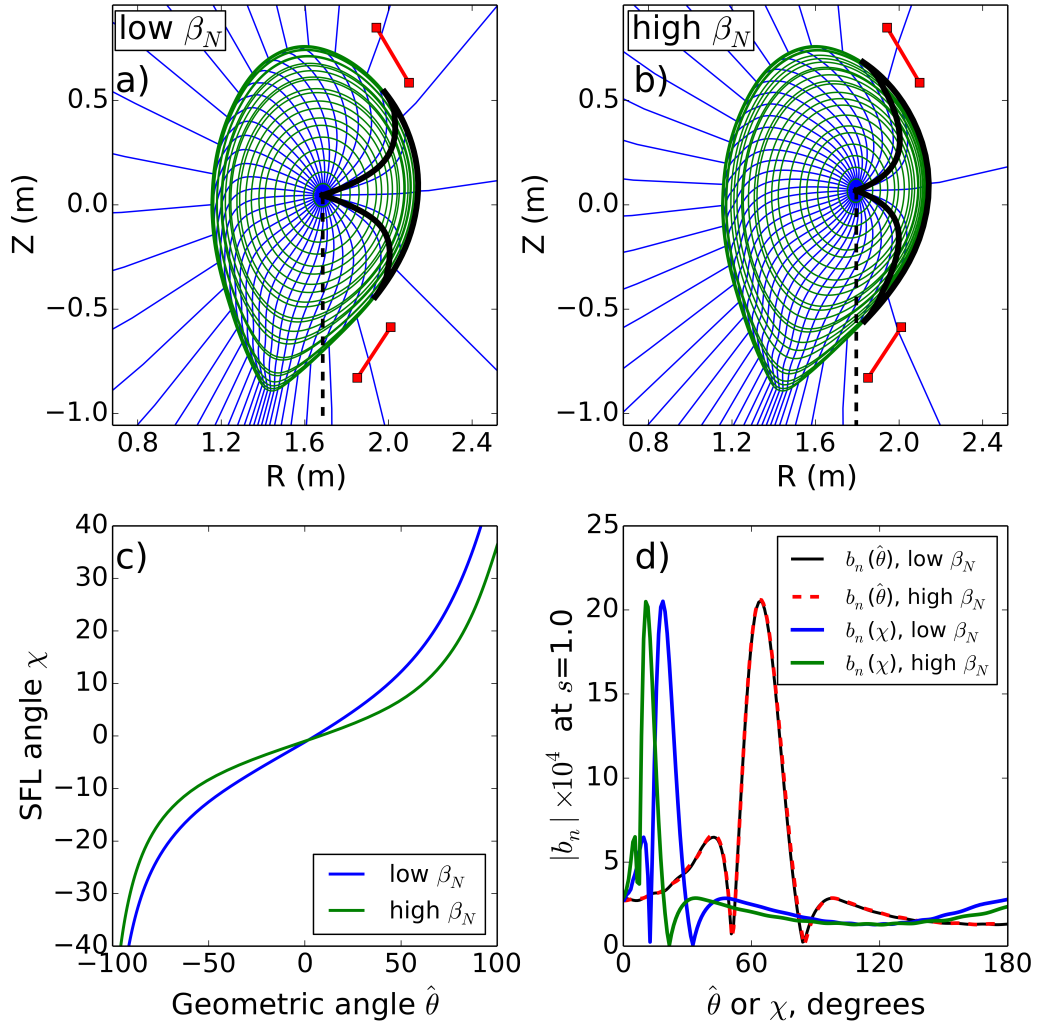


FIGURE 4.4: a,b) Grids of the straight field line coordinate systems based on the equilibrium for low (0.51) and high (4.86) β_N cases. The solid black lines are lines of $\chi = \pm 15$. In the high β_N case, the magnetic axis is Shafranov shifted outwards distorting the magnetic geometry on the outboard side, and the lines of $\chi = \pm 15$ encompass a larger arclength of the plasma boundary. c) The generalised poloidal angle χ is plotted against real space geometric angle $\hat{\theta}$. In the high β_N case there is a lower $d\chi/d\theta$ on the low field side, therefore a smaller extent of χ represents the same geometric features. d) The normal component of the applied vacuum perturbation evaluated at the plasma edge for a low and high β_N case, plotted against geometric coordinate $\hat{\theta}$ which shows that in real space both applied fields are the same, and plotted against generalised poloidal angle χ which shows that in SFL coordinates the low and high β_N fields differ, due to the redistribution of χ with changing β_N . In particular, in the high β_N case the main peaks of the applied field (close to the RMP coils) are compressed into a narrower range of χ . As a result the poloidal spectrum shifts towards higher m in the high β_N case.

perturbation is therefore shifted towards higher poloidal mode numbers with increasing β_N , as is apparent in figure 4.3c,d).

The plasma boundary shape also strongly influences the equilibrium geometry near the edge, and should therefore also have a significant impact on alignment, as explored in [120]. However this is not included in this study, since treating it rigorously would substantially increase the dimensionality of the problem, obstructing the original objective of a simple parametrisation.

4.2 ASDEX Upgrade Optimal Coil Phase

4.2.1 Scaled Equilibrium Set

A dense set of equilibria spanning a wide range of β_N and q_{95} was produced by scaling a reference equilibrium using the CHEASE fixed boundary equilibrium solver [106]. The reference equilibrium was produced using CLISTE from ASDEX Upgrade experiment 30835 at 3.2s, which is used as a standard equilibrium in many works [79, 95]. To avoid a numerical singularity, the plasma boundary was manually smoothed to exclude the X point. The smoothed boundary was held constant in the equilibrium scaling. In order to scan q_{95} , the q profile was rigidly shifted including the edge safety factor q_a , and β_N was scanned by scaling the pressure profile. The current and pressure were adjusted self consistently to maintain valid equilibria. Figure 4.5a) shows the (β_N, q_{95}) grid of the scaled equilibrium set relative to the reference equilibrium, as well as three illustrative cases for which profiles are plotted in figure 4.6. To assess the ASDEX Upgrade experimental parameter space, 20 samples of experimental q_{95} and β_N were taken from 4600 ASDEX Upgrade plasma discharges. Figure 4.5b,c) shows histograms of these samples, demonstrating that the range of values in the scaled equilibrium set is sufficient to encompass typical ASDEX Upgrade global plasma parameters. In Figure 4.6, profiles of q , plasma pressure and current are plotted for 3 of the scaled equilibria, which are annotated in Figure 4.5a), to demonstrate the scaling of the equilibrium profiles.

4.2.2 MARS-F Computed $\Delta\phi_{opt}$ Across (β_N, q_{95}) Domain

For each equilibrium in the set of (β_N, q_{95}) points, the vacuum field and plasma response to the applied RMP field were computed for $n = 1 - 4$ perturbations using the MARS-F code, specifically the outermost resonant field component due to the upper coil set $b_{res}^{1,u}$, and lower coil set $b_{res}^{1,l}$. Using these results, the optimal coil phase was computed with




FIGURE 4.5: a) The scaled equilibrium set (black dots) was produced by scaling the reference equilibrium (cyan pentagon) in plasma current and pressure. The coloured squares denote the 3 illustrative cases plotted in Figure 4.6. b) Histogram of experimental β_N values and c) Histogram of experimental q_{95} values. Values are sampled from 20 timepoints from the flattops of 4600 recent ASDEX Upgrade plasma discharges. It should be noted that q_{95} is here defined as being always positive.




FIGURE 4.6: a) q_{95} was scanned by rigidly shifting the q profile. b) β_N was scaled by scaling the pressure profile self similarly. c) The current density varies to remain consistent with the safety factor and plasma pressure. The designations Case 1,2,3 refer to the coloured points in Figure 4.5a).

Equation 4.1. Figure 4.8 shows computed $\Delta\phi_{opt}$ for $n = 1-4$ as a function of equilibrium parameters (β_N, q_{95}) , both excluding and including the plasma response. In general, $\Delta\phi_{opt}$ is shown to increase with q_{95} , and decrease with β_N . The strongest dependence is on q_{95} , consistent with previous works [93, 95, 120]. Since phase wraps do not represent physics and merely obscure results, phase wraps have been manually removed in this study, which is why in figure 4.8 the range of $\Delta\phi_{opt}$ can exceed 360 degrees. This is done by adding or subtracting integer multiples of 360 to the value of $\Delta\phi_{opt}$ computed with equation 4.1, as sketched in figure 4.7.

However, in contrast with a previous numerical and experimental coil phase scan on DIII-D [91] which found no effect of β_N on $\Delta\phi_{opt}$, this work also finds a weak β_N dependence of $\Delta\phi_{opt}$. The difference in results may be partially explained by different approaches to computational equilibrium scaling. In [91], the core pressure profile is modified to vary β_N , while the pressure pedestal is left approximately constant. In this work however, the entire pressure profile including the pedestal is scaled by a constant factor as shown in Figure 4.6b). Holding the pressure pedestal constant in the β_N scan as in [91], should in principle also keep the equilibrium geometry in the edge region constant, so the modification to alignment with β_N as explained in Figure 4.3 would not occur as it does in this work. In summary, whether or not a dependence of $\Delta\phi_{opt}$ on β_N is detected, may depend on whether the pedestal pressure is allowed to scale with β_N . However, the result reported in [91] is also derived from experimental observations, significantly strengthening the evidence against a β_N dependence of $\Delta\phi_{opt}$. Further experimental studies may shed further light on this issue.




FIGURE 4.7: Since equation 4.1 may only output values between 0 and 360 degrees, the output inevitably contains phase wraps. These would introduce a sharp discontinuity into the scan data, complicating the fitting procedure to derive the parametrisation. Therefore, they are removed manually by adding integer multiples of 360, to smooth the scan result.

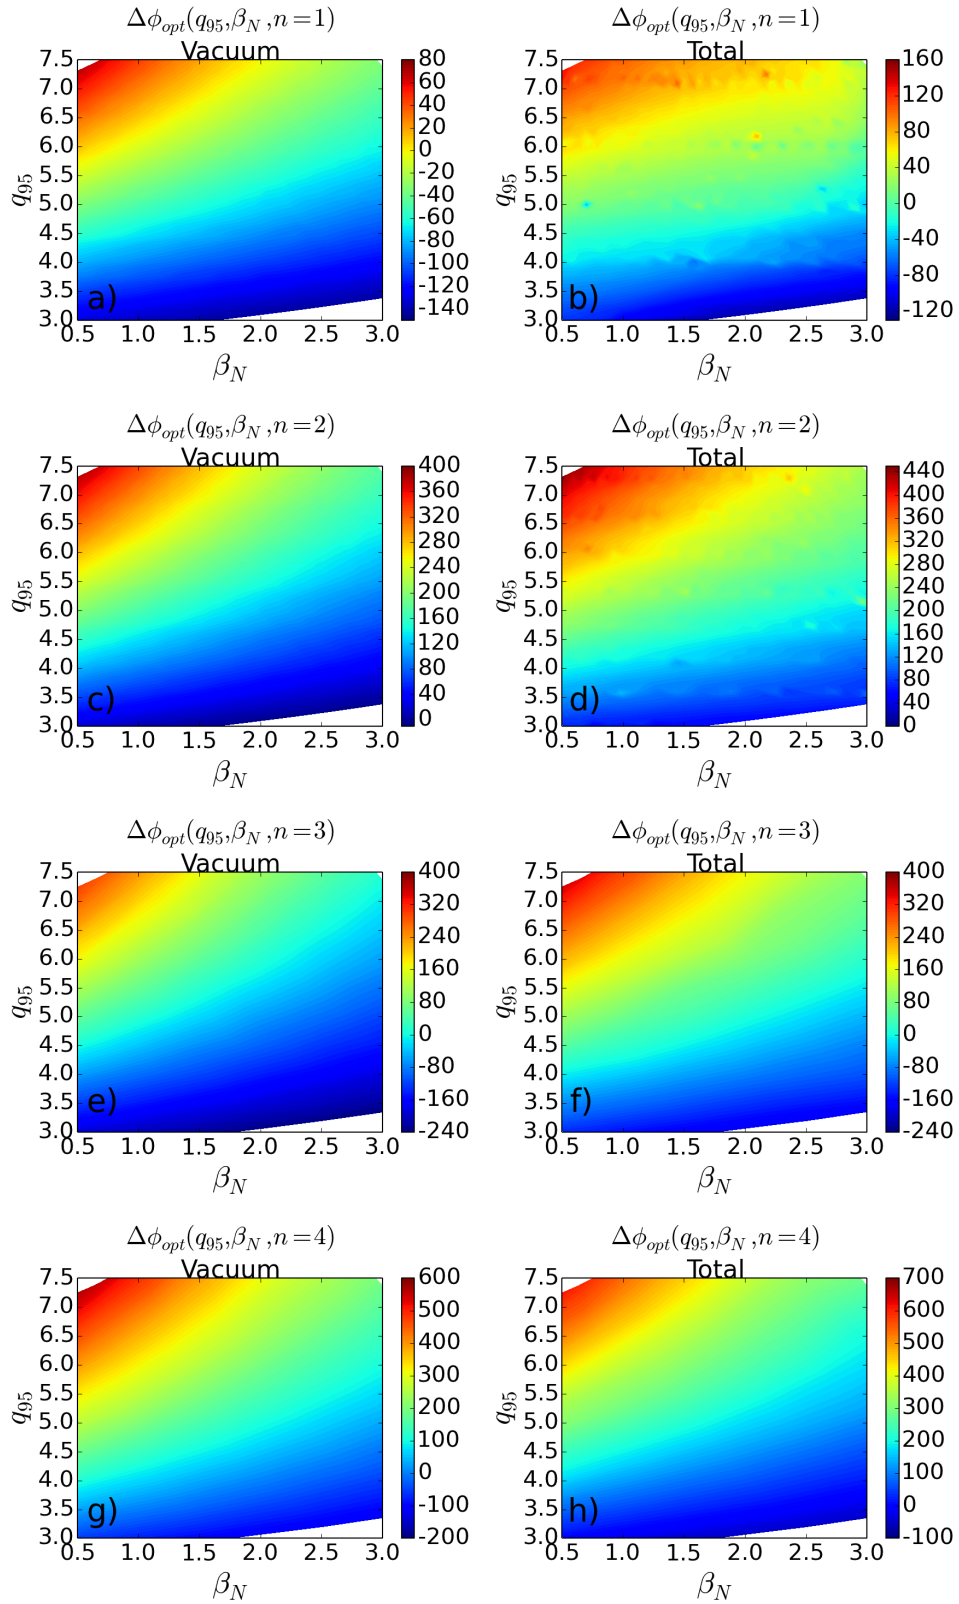


FIGURE 4.8: $\Delta\phi_{opt}$ computed for each equilibrium in the scaled set. MARS-F was used to compute $b_{res}^{1,u}$ and $b_{res}^{1,l}$, and Equation 4.1 for $\Delta\phi_{opt}$. a,c,e,g) Optimal alignment of the vacuum field for $n = 1, 2, 3, 4$ respectively. b,d,f,h) Optimal alignment of the total field for $n = 1, 2, 3, 4$ respectively. In all cases $\Delta\phi_{opt}$ is seen to vary smoothly with β_N and q_{95} , with minor perturbations caused by RFA peaks.

$RMSE_{poly}$	Linear	Quadratic	Cubic
n=1 vacuum	6.65	0.86	0.55
n=1 total	11.80	6.09	5.14
n=2 vacuum	12.80	1.58	0.90
n=2 total	12.99	4.73	4.50
n=3 vacuum	19.74	2.61	1.53
n=3 total	20.41	3.80	2.93
n=4 vacuum	26.25	4.45	3.19
n=4 total	25.73	6.49	4.67

TABLE 4.1: $RMSE_{poly}$ for each scaled equilibrium set, for a linear, quadratic and cubic fit.

coeff	a	b	c	d	e	f	g	h	i
n=1 vacuum	0.13898	0.15842	-1.6741	-0.51686	-7.6899	18.737	-1.2555	65.15	-312.19
n=1 total	0.43305	-5.7	17.097	-2.7405	29.94	-99.267	-0.45866	49.966	-210.18
n=2 vacuum	0.14571	1.7142	-6.3854	-0.2497	-23.719	56.206	-3.1508	127.83	-327.38
n=2 total	0.14047	1.7732	-8.5336	-0.33719	-22.025	63.892	-3.1757	129.07	-286.34
n=3 vacuum	0.28269	1.6533	-6.7085	-0.33683	-34.041	76.084	-4.6137	180.18	-676.44
n=3 total	0.21942	1.8076	-7.1436	-0.56169	-28.501	67.55	-3.9637	171.31	-604.86
n=4 vacuum	0.36048	2.047	-8.1855	-0.50638	-42.161	91.007	-5.1487	219.78	-646.41
n=4 total	0.50969	0.78126	-6.1876	-1.1285	-35.518	85.061	-4.1725	208.44	-572.3

TABLE 4.2: Coefficients of the 2D quadratic parametrisation of $\Delta\phi_{opt}$.

4.2.3 Parametrisation of $\Delta\phi_{opt}$

Figure 4.8 shows that $\Delta\phi_{opt}$ is a smoothly varying function of (β_N, q_{95}) for each n . Therefore the results may be parametrised using a simple analytic function to provide researchers with an estimate $\Delta\phi_{opt}$, which may be used for planning future experiments and interpreting experimental data. A 2D polynomial function in (β_N, q_{95}) is chosen as a compromise between ease of use and providing sufficient degrees of freedom to closely fit the scan results in Figure 4.8. Linear, quadratic and cubic 2D polynomials were fitted to the data in Figure 4.8 by linear regression. An RMSE between the scaled equilibrium data and the polynomial is defined below, to quantify the fit of each polynomial to the scaled equilibrium data.

$$RMSE_{poly} = \left(\sum_i^N (\Delta\phi_{opt,set}^i - \Delta\phi_{opt,poly}^i)^2 / N \right)^{\frac{1}{2}} \quad (4.2)$$

In the above, N is the number of points in the scaled equilibrium set, $\Delta\phi_{opt,poly}^i$ are the optimum coil phases predicted by the 2D polynomials, and $\Delta\phi_{opt,set}^i$ are the MARS-F

computed optimum coil phases for each point in the scaled equilibrium set. For the linear, quadratic and cubic polynomials, the values of $RMSE_{poly}$ for all 8 datasets are listed in table 4.1. The table shows that the linear fit does not closely fit the data, while the cubic fit is not significantly better at representing the data than the quadratic. Therefore the quadratic fit is used as a compromise between accuracy and accessibility. The form of the 2D quadratic parametrisation is described in Equation 4.3, and table 4.2 lists its coefficient values. In the following, let $\Delta\phi_{opt,quad}$ be $\Delta\phi_{opt}$ predicted by the 2D quadratic parametrisation, $x = \beta_N$ and $y = q_{95}$

$$\Delta\phi_{opt,quad} = a(x^2y^2) + b(x^2y) + c(x^2) + d(xy^2) + e(xy) + f(x) + g(y^2) + h(y) + i \quad (4.3)$$

The 2D quadratic parametrisation fits the results of the computed scan to within 6 degrees in all datasets. When the plasma response is included however, small horizontal 'stripes' in the (β_N, q_{95}) domain are apparent. These stripes are caused by sharp peaks in the amplified kink-peeling response as computed by MARS-F, previously detected and discussed in [93, 120], which occur when nq_a is immediately below an integer. These peaks cause a localised upward shift of up to 20 degrees from the underlying trend, but do not significantly affect the 2D quadratic fit, which is dominated by the bulk of the (β_N, q_{95}) domain.

4.3 Parametrisation Benchmarking

4.3.1 Benchmarking Database

In order to quantify the extent to which the 2D quadratic parametrisation reproduced a thorough MARS-F computation, a set of benchmarking points was collected from the ASDEX Upgrade experimental database, consisting of 85 distinct time points from 31 distinct plasma discharges (49 time points from 17 $n = 1$ discharges, 36 timepoints from 14 $n = 2$ discharges). Each benchmarking point consisted of a plasma equilibrium, and set of experimental coil currents and kinetic profiles. Equilibrium reconstructions are routinely performed after each ASDEX Upgrade experiment using the CLISTE [105] code, and these were used as the source of the benchmark equilibria. In order to produce the kinetic profiles for each benchmarking point, the AUGPED tool was used to fit analytic mtanh functions to experimental measurements aggregated over 20ms, of electron temperature T_e , electron density n_e , and ion temperature T_i , and a spline to toroidal bulk plasma rotation v_t . Figure 4.9 shows the resulting kinetic profiles for an example

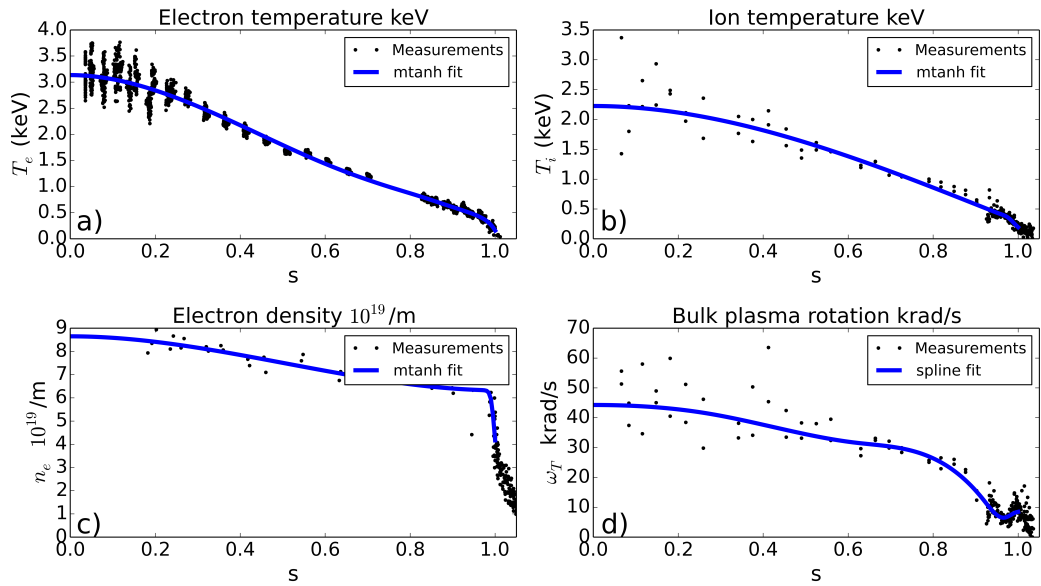


FIGURE 4.9: Curves fitted to kinetic experimental data from ASDEX Upgrade discharge number 30641 at 4.0s, here serving as an example kinetic profile fit for a benchmarking point. The significant scatter is typical for experimental measurements of kinetic profiles. a) Electron temperature data from the Thompson Scattering diagnostic [121] and Electron Cyclotron Emission diagnostic [122], with a fitted mtanh function. b) Ion temperature data from the Charge Exchange Recombination Spectroscopy diagnostic [123], with a fitted mtanh function. c) Electron density data, from the Thompson Scattering diagnostic, the Lithium Beam diagnostic [124] and interferometers, also with a fitted mtanh function. d) Toroidal bulk rotation data from the Charge Exchange Recombination Spectroscopy diagnostics, fitted with a cubic spline.




FIGURE 4.10: Values of (β_N, q_{95}) of the benchmarking points, used to benchmark the 2D quadratic parametrisation. Each benchmarking point consists of a distinct equilibrium, plasma boundary, set of kinetic profiles and RMP coil currents.

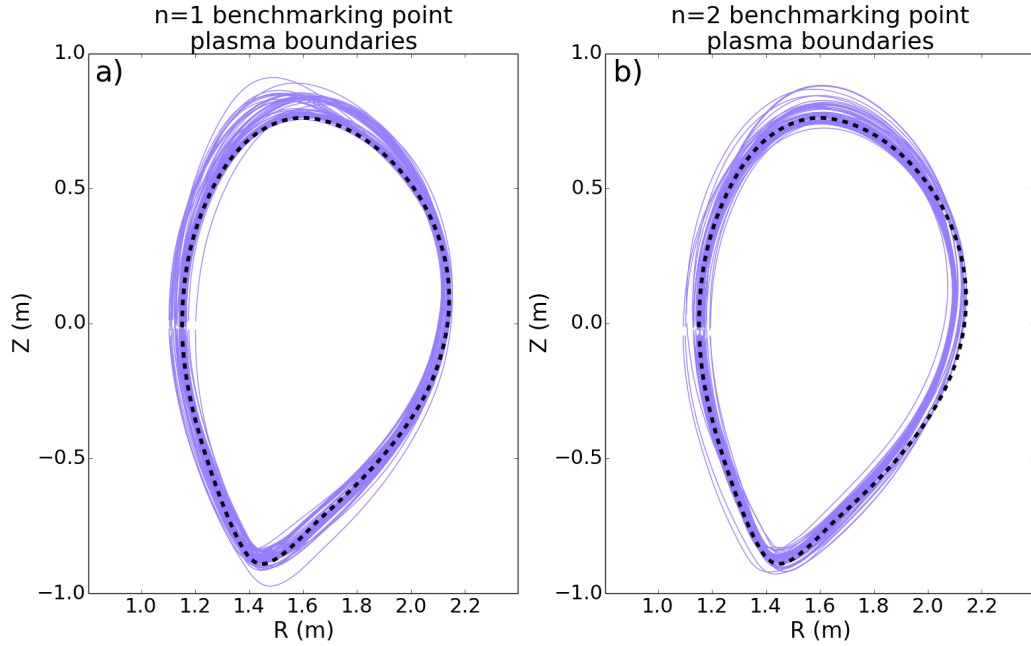


FIGURE 4.11: Plasma boundaries of the reference equilibrium (discharge 30835 at 3.2s) and the benchmarking points. The boundary shapes of the benchmarking points vary significantly.

benchmarking point. Figure 4.10 plots the global plasma parameters (β_N , q_{95}) for the set of benchmarking points, and 4.11 shows the plasma boundary shapes for the reference and benchmarking equilibria. Each benchmarking point consists of kinetic profiles of T_e , n_e , T_i and v_t , a plasma equilibrium and plasma boundary, and the experimentally applied RMP coil currents.

The vacuum field and plasma response were computed for each benchmarking point using MARS-F in order to compute the pitch aligned components, and Equation 4.1 was then used to compute $\Delta\phi_{opt}$. For the set of benchmarking points, the values of $\Delta\phi_{opt}$ predicted by the 2D quadratic parametrisation were compared with values computed using a MARS-F plasma response computation in Figure 4.12, for both $n = 1$ and $n = 2$ applied RMPs. Insufficient $n = 3, 4$ RMP experiments could be found to repeat this benchmarking for the $n = 3, 4$ scan. An RMSE between the 2D quadratic parametrisation and the MARS-F computations is defined below, to quantify the agreement between them, where N is now the number of benchmarking points

$$RMSE_{bench} = \left(\sum_i^N (\Delta\phi_{opt,bench}^i - \Delta\phi_{opt,quad}^i)^2 / N \right)^{\frac{1}{2}} \quad (4.4)$$

$RMSE_{bench}$ is found to be 14.1 for $n = 2$ vacuum predictions, 21.3 for $n = 2$ total predictions, 6.9 for $n = 1$ vacuum predictions, and 34.2 for $n = 1$ total predictions. The

predictions of the 2D quadratic parametrisation are plotted against the MARS-F predictions in Figure 4.12. Referring to Figure 4.2, it is apparent that the gradient of $|b_{res}^1|$ is flat in the region of $\Delta\phi_{opt}$, so small misalignments on the order of $RMSE_{bench}$ have a very small effect on $|b_{res}^1|$. Therefore the uncertainty in the 2D quadratic parametrisation is low enough for it to be useful in experimental planning. However, some significant deviations between the 2D quadratic parametrisation and the MARS-F computation are apparent in Figure 4.12, possible causes for which will now be discussed.

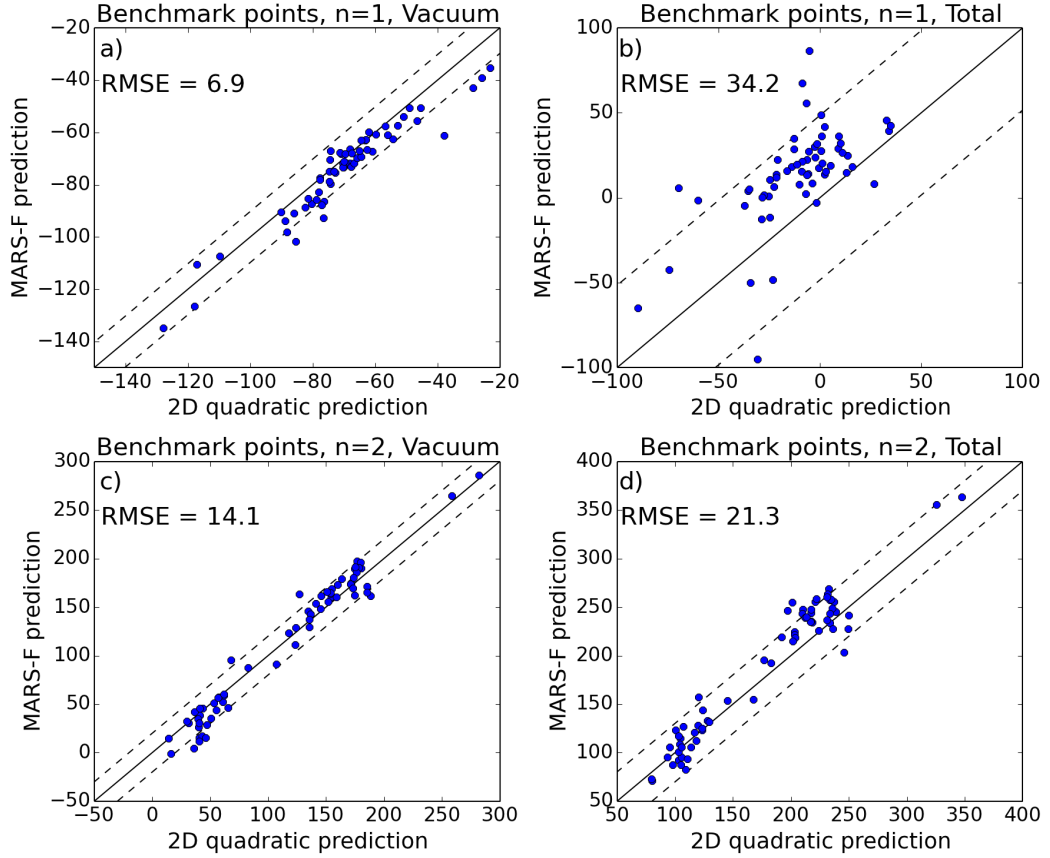


FIGURE 4.12: At each benchmarking point, $\Delta\phi_{opt}$ as predicted by the 2D quadratic parametrisation is plotted against $\Delta\phi_{opt}$ as predicted by a full MARS-F computation. The quadratic parametrisation requires only β_N , q_{95} and n as input, whereas the MARS-F computation requires a plasma equilibrium, plasma boundary shape, coil currents and set of kinetic profiles as input. The $x = y$ line is annotated on in solid black, while the dashed black lines are $RMSE_{bench}$ degrees from the 1:1. While the MARS-F computation is the more rigorous approach, but the 2D quadratic parametrisation requires no specialist software, HPC hardware or expertise and far less input information, and is far simpler and quicker.

4.4 Sources of Uncertainty for $\Delta\phi_{opt,quad}$

Rotation

The set of benchmarking points has significant variation in rotation profiles, but in the (β_N, q_{95}) scan the rotation profile was held constant. To determine the extent to which rotation contributes to the scatter in Figure 4.12, the rotation profile of the reference equilibrium was scaled from 10 times to $1/10^{th}$ of its experimental value while $\Delta\phi_{opt}$ was computed. The computed values of $\Delta\phi_{opt}$ with varying rotation speed are plotted in Figure 4.13. The results indicate that $\Delta\phi_{opt}$ is robust to scaling of the rotation profile to within 10 degrees.




FIGURE 4.13: The rotation profile is scaled while $\Delta\phi_{opt}$ of the reference equilibrium computed by MARS-F. $\Delta\phi_{opt}$ appears to be robust to scaling of the rotation profile.

Kinetic Pedestal

As shown in Figure 4.9, the kinetic profiles used as input for the MARS-F computation are produced by manually fitting curves to aggregated experimental data. It is possible that uncertainty in the profile fitting may propagate to the equilibrium computation or plasma response computations, which may increase the uncertainty in the vacuum or total field values of $\Delta\phi_{opt}$. To investigate this possibility, the kinetic pedestal width is scanned within its uncertainty, and for each pedestal width the equilibrium is recomputed for a self consistent equilibrium, and $\Delta\phi_{opt}$ is recomputed using MARS-F. The pedestal profiles used in the scan are plotted in Figure 4.14. The pedestal width is scanned from -2σ to $+2\sigma$, where σ is the uncertainty in pedestal width of the mtanh fit to the kinetic data. Repeating the equilibrium reconstruction for each pedestal width resulted




FIGURE 4.14: The fitting parameters of the mtanh fits to kinetic profiles, specifically the pedestal width, has an associated uncertainty σ . The pedestal width was scanned from -2σ to 2σ resulting in the above profiles.




FIGURE 4.15: The equilibrium was recomputed using CLISTE for each profile in the uncertainty scan, resulting in self consistent modifications to the equilibrium. Figure a) shows the edge safety factor, b) equilibrium current profile and c) the equilibrium pressure profile. d) shows optimum coil phase $\Delta\phi_{opt}$ computed with MARS-F for each equilibrium in the scan. The results indicate that $\Delta\phi_{opt}$ is also not sensitive to pedestal width within $\pm 2\sigma$ of the fit uncertainty.

in slight changes to the current, pressure and q profiles, as shown in Figure 4.15 a), b) and c). Figure 4.15 d) shows $\Delta\phi_{opt}$ as computed with MARS-F including the plasma response, as the pedestal width is scanned. The plot indicates that $\Delta\phi_{opt}$ is quite robust to uncertainties in the kinetic pedestal width.

Plasma Boundary

The plasma boundaries of the benchmarking points have significant deviation from the reference equilibrium used to derive the 2D quadratic parametrisation, as evidenced in Figure 4.11. A recent work also using MARS-F and the same reference equilibrium, found that a change in the upper triangularity $\Delta\delta_u$ of 0.08 can cause $\Delta\phi_{opt}$ to change by 60 degrees[96]. Therefore it seems plausible that much of the deviations between the 2D quadratic parametrisation and rigorous MARS-F computations observed in Figure 4.12 may be caused by variations in plasma boundary shape. Closer examination of Figure 4.12b) shows an apparent systematic deviation between the 2D quadratic parametrisation and MARS-F predictions for the $n = 1$ discharges, which may be the result of a systematic deviation in the plasma shapes between the reference equilibrium and the $n = 1$ benchmark points. Figure 4.16 plots the plasma horizontal minor radius a_{hor} and outermost major radius extent R_{aus} , for the $n = 1, 2$ benchmarking points and the reference equilibrium. The figure shows that for the $n = 2$ points the shape parameters are randomly scattered around the reference equilibrium, whereas the $n = 1$ points are systematically shifted from it. This further motivates a study of the dependence of $\Delta\phi_{opt}$ on plasma boundary shape, which is left for future work. The uncertainties in the 2D quadratic parametrisation may also be reduced by recomputing the coefficients in table 4.2, using a more representative plasma equilibrium.

4.5 Parametrisation Experimental Validation

Comparing the 2D quadratic parametrisation to MARS-F computations of $\Delta\phi_{opt}$ is useful for benchmarking; however, to increase confidence in the parametrisation, a validation against direct experimental observations is required. In this section, a study is carried out to provide experimental validation of the 2D quadratic parametrisation and MARS-F optimal coil phase predictions. The optimum coil phases for ELM mitigation and density pump out are extracted from coil phase scan experiments on ASDEX Upgrade, and these coil phases are then compared with predictions of full MARS-F modelling, and the 2D quadratic parametrisation.




FIGURE 4.16: Variation in plasma shape of the benchmarking points, shown using horizontal minor radius and maximum major radius. The $n = 2$ points are scattered around the reference equilibrium, while the $n = 1$ points are systematically shifted from it.

4.5.1 Coil Phase Scan Experiments

The ASDEX Upgrade experiments used for this study are summarised in table 4.3. ASDEX Upgrade has two independent power supply units for the ELM coils, which allows the coil phase to be scanned by applying sinusoidally varying currents such that the toroidal waveform of the current in the upper coil set moves relative to the waveform in the lower set. Figure 4.17 shows an example discharge (number 33143) in which the coil phase is scanned through several complete 2π cycles, and f_{ELM} and n_e are modulated by the changing coil phase.

Extracting Experimental $\Delta\phi_{opt}$

Since density pump out is reduction in density, pump out is characterised here using $-1 \times n_e$. In order to extract the optimal coil phase for f_{ELM} and density pump out, the time domains are first divided into complete cycles of $\Delta\phi_{ul}$ from $0 \rightarrow 2\pi$. f_{ELM} and $-1 \times n_e$ are normalised to between $0 \rightarrow 1$ using their maximum and minimum values within each cycle. Using the trace of $\Delta\phi_{ul}(t)$, $f_{ELM}(t)$ and $-1 \times n_e(t)$ are mapped from the time domain to the $\Delta\phi_{ul}$ domain. A function $|(1 + e^{i(\Delta\phi_{ul} + \Delta\phi_{opt})})|/2$ is then fitted to each normalised $f_{ELM}(\Delta\phi_{ul})$ and $\hat{n}_e(\Delta\phi_{ul})$ to determine $\Delta\phi_{opt}$ from f_{ELM} and n_e measurements. This fitting function is chosen because it matches the form of $|b_{res}^1(\Delta\phi_{ul})|$. Figure 4.18 shows an example of this procedure, using shot number 33143. Figure 4.19 demonstrates the same procedure using shot 30680, which contains a larger

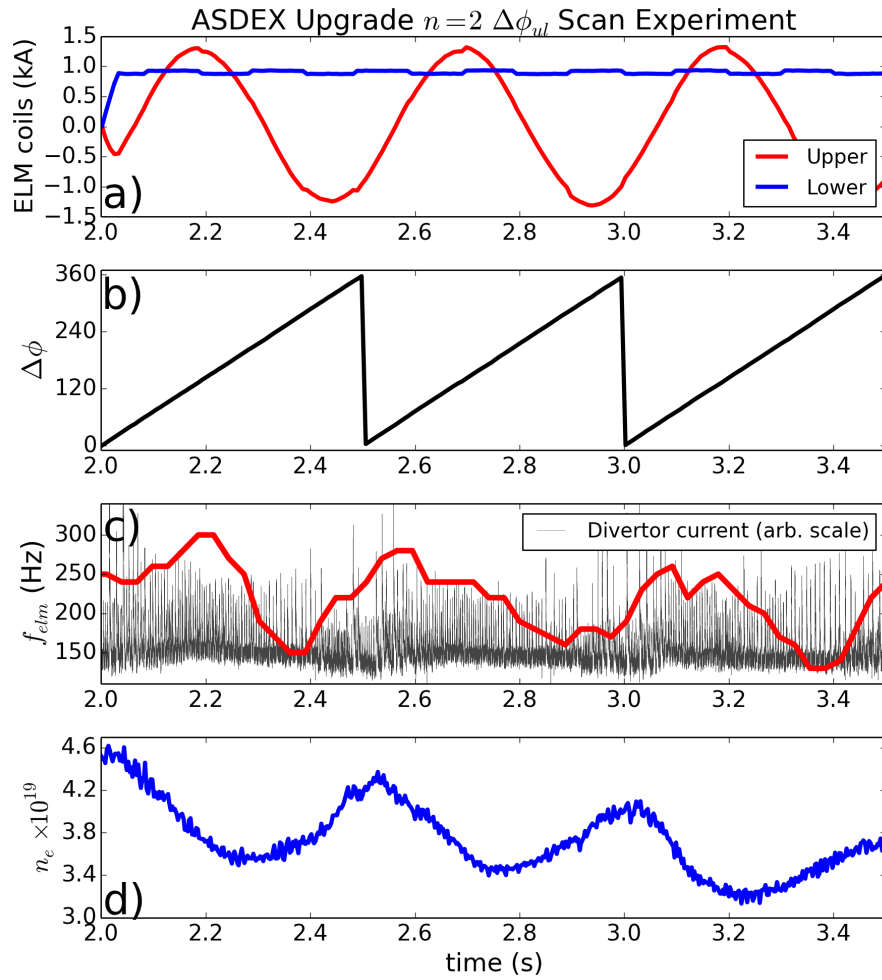


FIGURE 4.17: a) A sinusoidally varying current is used to create a travelling toroidal waveform in the upper coils, while the lower coils are held fixed. b) The motion of the upper toroidal waveform relative to the stationary lower waveform causes a continuously varying phase difference $\Delta\phi_{ul}$ between the upper and lower coils. c) ELMs are diagnosed as spikes in the divertor current trace (grey trace). Counting these spikes yields the ELM frequency (red trace), which varies with $\Delta\phi_{ul}$. d) The line averaged density also varies with $\Delta\phi_{ul}$.

scatter in the measured values of $\Delta\phi_{opt}$. For each shot $\Delta\phi_{opt}$ is approximated as constant, so $\Delta\phi_{ul}$ cycles may be treated as repeated measurements, and used to compute the mean and variance for each shot. Since the measured data is directional (an angle), the mean angular direction and circular standard deviation are used, as explained in [125].

To aid vertical plasma stability, ASDEX Upgrade has two large copper rings installed above and below the midplane inside the vacuum vessel, directly in front of the two rows of ELM coils. These Passive Stabilisation Loops (PSLs) have no effect on static magnetic fields from the ELM coils, but if the ELM coil currents are time varying then eddy currents in the PSLs are induced which act to reduce and delay the field which reaches the plasma. When extracting $\Delta\phi_{opt}$ from the experimental measurement, it

n	Shot	$\Delta\phi_{ul}$ cycles	B_t	I_p	q_{95}	$n_e \times 10^{19} m^{-3}$
1	31023	4	2.58	1.00	4.36	7.12
1	31034	3	2.59	1.00	4.31	7.66
1	32100	4	2.49	0.80	5.27	4.82
1	32116	5	2.48	0.80	5.26	5.11
2	30680	5	2.48	0.80	5.24	8.44
2	30681	6	2.47	0.80	5.43	6.34
2	30682	2	1.77	0.80	3.66	5.20
2	30684	5	2.47	0.80	5.17	6.31
2	30824	2	2.49	0.80	5.41	5.95
2	30826	2	1.79	0.80	3.67	5.65
2	31542	2	2.43	0.80	4.99	5.34
2	31543	2	2.44	0.80	4.99	5.40
2	31545	2	2.43	0.80	4.99	5.51
2	33143	3	1.82	0.80	3.84	5.02

TABLE 4.3: ASDEX Upgrade discharges used to measure the experimental optimal coil phase. $\Delta\phi_{ul}$ cycles refers to the number of complete rotations of coil phase which are used in this study (in general, a small subset of the total number of complete $\Delta\phi_{ul}$ rotations)

is necessary to account for the lag in $\Delta\phi_{ul}$ seen by the plasma, caused by PSL eddy currents. A finite element modelling code has been developed for this purpose (see [102], and references therein), which is used here. For each shot studied here, the PSL induced lag is calculated. Figure 4.20 shows the computed $\Delta\phi_{ul}$ both with and without the effect of the PSLs. It is found that for $\Delta\phi_{ul}$ scans of 2Hz or higher the PSL induced phase lag was 32 degrees, while a scan at 1Hz corresponded to a lag of 25 degrees. The lag is either added to or subtracted from the measured $\Delta\phi_{opt}$, depending on whether $\Delta\phi_{ul}$ is scanned 'upwards' or 'downwards'.

For all cycles and shots studied, the measured optimal phases for both ELM mitigation and density pump out, corrected for PSL lag, are shown in figure 4.21. The $n = 1$ discharges studied here did not exhibit a measurable modulation in n_e with the $\Delta\phi_{ul}$ scan, and are therefore not included in the plot. For the $n = 2$ discharges, it is observed that there tends to be a significant offset between the optima for f_{ELM} and n_e . On average the optimum coil phase for ELM mitigation is 59.0 ± 60.6 degrees lower than for maximum density pump out. In the context of the observed correlation between density pump out and ξ_X [87], this result is interesting given that recent computational works [96] have found that b_{res}^1 and ξ_X are synchronised as a function of $\Delta\phi_{ul}$ (ie, no offset between the figures of merit associated with f_{ELM} and pump out). However, it is noted that the size of the shift in the time domain is of the same order as the particle confinement time, and that the density is measured close to the plasma core. Therefore it is quite possible that the shift is due to a time lag between the density and ELM frequency response, caused by the relatively slow transport timescale compared with the much

faster timescale at which the ELM frequency responds. We may therefore not expect to see this shift for static fields, and is consistent with current theoretical expectations. Accounting for this lag is beyond the scope of this work. Since the focus of this work is on ELM mitigation, mitigated ELM frequency f_{ELM} and not density pump out is used to determine the $\Delta\phi_{opt}$ experimental value for comparison with predictions of the 2D quadratic parametrisation.




FIGURE 4.18: a) $\Delta\phi_{ul}$ is divided into $0 \rightarrow 2\pi$ cycles. b) The density is multiplied by -1, so the plot now shows density pump-out. For each cycle the density pump out and mitigation are mapped from the time domain to $\Delta\phi_{ul}$. c) Each cycle of f_{ELM} is normalised to its maximum and minimum value in that cycle, and the optimum extracted by fitting a $|(1 + e^{i(\Delta\phi_{ul} + \Delta\phi_{opt})})|/2$ function. d) Each cycle of density pump out is normalised to its maximum and minimum value in that cycle, and the optimum extracted by fitting a $|(1 + e^{i(\Delta\phi_{ul} + \Delta\phi_{opt})})|/2$ function.

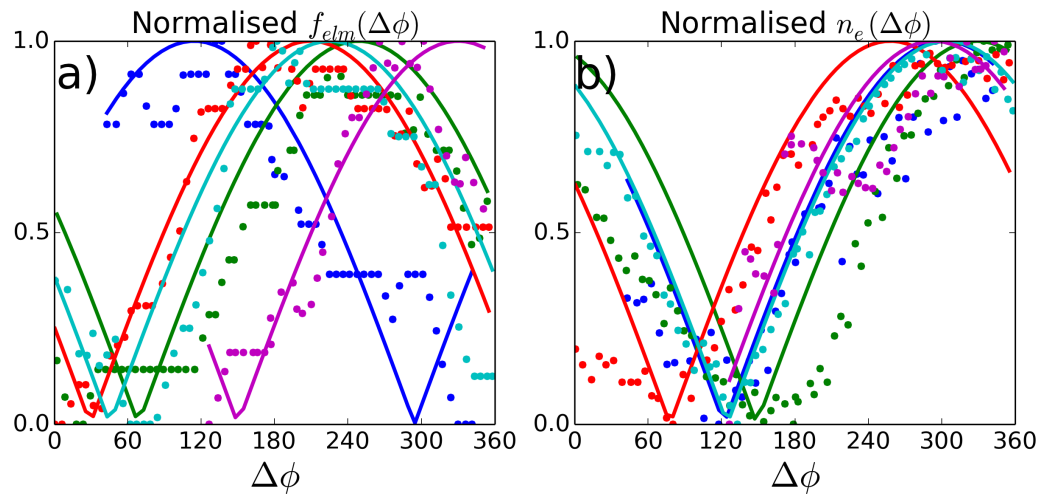


FIGURE 4.19: A repeat of figure 4.18 using shot 30680, which demonstrates scatter in the $\Delta\phi_{opt}$ measurements. To quantify uncertainties, each cycle in a shot is treated as a repeat measurement, to define an average and standard deviation for that shot.




FIGURE 4.20: Eddy currents in the Passive Stabilisation Loops cause the value of $\Delta\phi_{ul}$ as seen by the plasma to lag behind the value at the ELM coils.

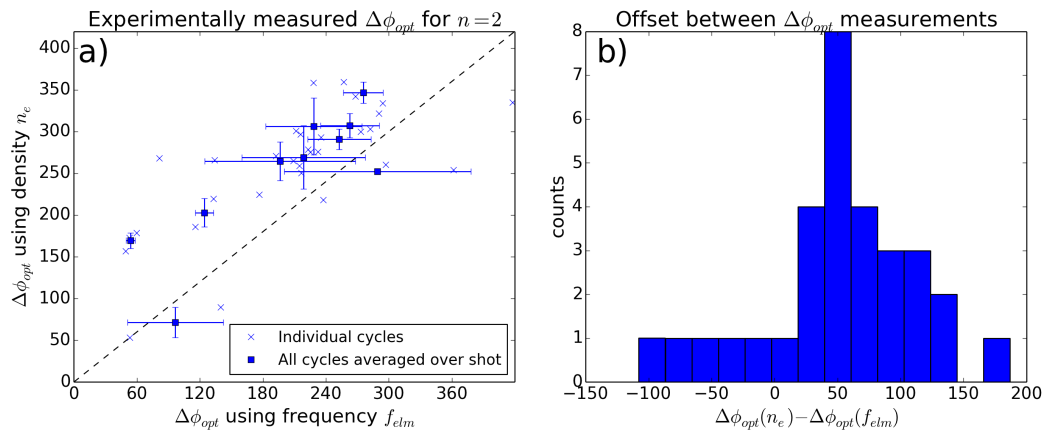


FIGURE 4.21: a) $\Delta\phi_{opt}$ measured using the n_e trace against $\Delta\phi_{opt}$ measured using the f_{ELM} trace. Only $n = 2$ shots included. b) Histogram of shift in $\Delta\phi_{opt}$ as measured with the two methods. Although the scatter is wide, it is clear that the optimal coil phase for density pump out is shifted upwards relative to the optimum for ELM mitigation.

4.5.2 Comparison of measured $\Delta\phi_{opt}$ with predictions

For each cycle studied, the midpoint of the cycle is chosen as a representative point in the cycle. At these representative points, CLISTE based equilibria were downloaded from the AUG database, and profiles are fitted to measurements of T_e , T_i , n_e , v_t . These are used as input for a plasma response computation using the MARS-F code, from which $\Delta\phi_{opt}$ defined using b_{res}^1 is extracted. Also, using the values of β_N and q_{95} at these representative points as input, $\Delta\phi_{opt}$ is computed using the 2D quadratic parametrisation derived in section 4.2.3. Furthermore, using a numerical scan of $\Delta\phi_{ul}$, the dependence of ξ_X on $\Delta\phi_{ul}$ was computed, in order to compare it with the dependence of n_e on $\Delta\phi_{ul}$. It was found, as in previous studies [96], that b_{res}^1 and ξ_X were maximised at the same coil phase. A correlation has previously been observed between density pump out and ξ_X , which implies that we would expect the density pump out to be maximised at maximal ξ_X , which coincides with maximal b_{res}^1 . This expectation is contrary to the observed shift in optimal coil phase for f_{ELM} and density pump out apparent in figure 4.21.

Figure 4.22 shows $\Delta\phi_{opt}$ measured using f_{ELM} and corrected for PSL lag, against $\Delta\phi_{opt}$ predicted by MARS-F computations and the 2D quadratic parametrisation. The figure shows that the predictions of MARS-F and the 2D quadratic parametrisation are both within the uncertainty of the measurement, but in all cases the uncertainty in the experimentally measured $\Delta\phi_{opt}$ is quite high: typically around 40 degrees but as high as 90. The primary cause for the large uncertainty is the low number of cycles used in each shot, since angular standard deviations are rather sensitive to small sample sizes.

For comparison with the previous benchmarking of the 2D quadratic parametrisation, an RMSE between the experimental measurements and MARS-F or 2D quadratic parametrisation predictions is defined as

$$RMSE_{quad} = \left(\sum_i^N (\Delta\phi_{opt,quad}^i - \Delta\phi_{opt,exp}^i)^2 / N \right)^{\frac{1}{2}} \quad (4.5)$$

$$RMSE_{mars} = \left(\sum_i^N (\Delta\phi_{opt,mars}^i - \Delta\phi_{opt,exp}^i)^2 / N \right)^{\frac{1}{2}} \quad (4.6)$$

In the above, $\Delta\phi_{opt,quad}$ is the cycle averaged optimal coil phase for each shot predicted by the 2D quadratic parametrisation, $\Delta\phi_{opt,mars}$ is the cycle averaged optimal coil phase for each shot predicted by MARS-F, and $\Delta\phi_{opt,exp}$ is the cycle averaged experimentally measured optimal coil phase for ELM mitigation for each shot. Note these RMSEs do not represent or take account of the uncertainty in the measurement, and do not replace

a thorough statistical analysis. They are primarily used for a self consistent comparison with the previous result, or the benchmark between the 2D quadratic parametrisation and MARS-F. The values are found to be $RMSE_{mars} = 35.0$ and $RMSE_{quad} = 40.3$.

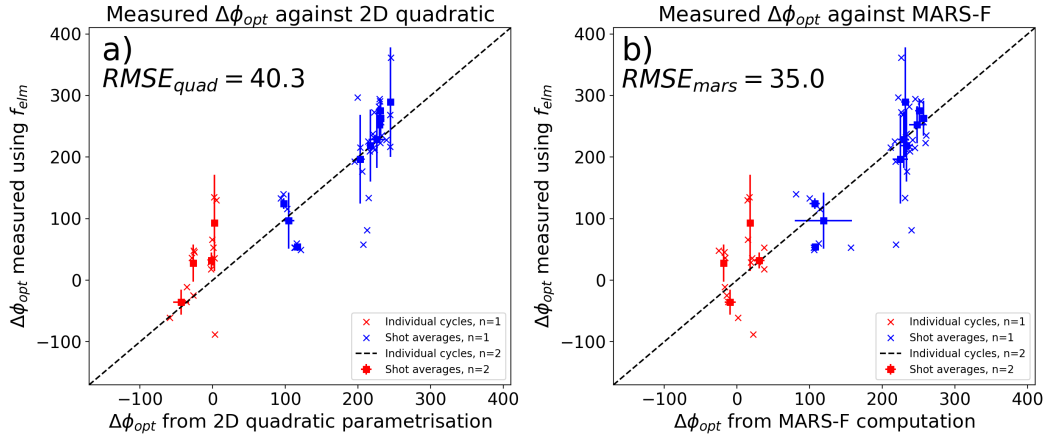


FIGURE 4.22: a) $\Delta\phi_{opt}$ measured from the f_{ELM} trace compared with $\Delta\phi_{opt}$ predicted with the 2D quadratic parametrisation. b) $\Delta\phi_{opt}$ measured from the f_{ELM} trace compared with $\Delta\phi_{opt}$ predicted with a MARS-F plasma response computation.

4.6 ITER Coil Phase Parametrisation

ITER will need active ELM mitigation even during the current and pressure ramp, when the edge safety factor q_{95} and β_N change drastically before settling at their flattop values. Therefore during the ramp, the optimal coil alignment will vary considerably, and it may be necessary to actively track the optimal alignment to sustain sufficient ELM mitigation. This motivates the optimal coil phase for ITER to be parametrised using the previously developed technique, which is benchmarked against MARS-F computations and validated against experimental measurements on ASDEX Upgrade.

In addition to the upper and lower rows of ELM coils, ITER will also have a third middle row which rings the outboard midplane, as sketched in figure 4.23. The principle of superposition is again used, with the minor change that the field is now the linear sum of three sets of coils instead of two. The mid plane coil set is held fixed and the upper and lower sets rotated, such that the field due to all three coils b_{UML} is

$$b_{UML} = b_M + b_U e(-i\Delta\phi_{UM}) + b_L e(-i\Delta\phi_{LM}) \quad (4.7)$$

where b_M , b_U and b_L are the fields due to the midplane, upper and lower coil sets respectively, and $\Delta\phi_{UM}, \Delta\phi_{LM}$ are the coil phase differences between the midplane coil set and the upper and lower coils respectively. The coil phase differences are defined as

$\Delta\phi_{UM} = \Delta\phi_M - \Delta\phi_U$, $\Delta\phi_{LM} = \Delta\phi_M - \Delta\phi_L$, where $\Delta\phi_U$, $\Delta\phi_M$, $\Delta\phi_L$, are the phases of the coil current toroidal waveforms of the upper, midplane and lower coils. We also here define the optimal upper and lower coil phases $\Delta\phi_{opt,UM}$ and $\Delta\phi_{opt,LM}$ as the upper and lower coil phases which maximise b_{res}^1 .

Applying the same reasoning to computing $\Delta\phi_{UM,opt}$ and $\Delta\phi_{LM,opt}$ as previously used for $\Delta\phi_{UL,opt}$ for ASDEX Upgrade, it is clear that the pitch aligned component due to all three coils is optimised when both $b_{res,U}^1$ and $b_{res,L}^1$ are parallel to $b_{res,M}^1$ in the complex plane. Therefore, $\Delta\phi_{UM,opt}$ and $\Delta\phi_{LM,opt}$ both have unique values for a given equilibrium and perturbation, given by

$$\Delta\phi_{UM,opt} = \pm \arccos\left(\frac{b_{res}^{1,m} \cdot b_{res}^{1,u}}{|b_{res}^{1,m}| |b_{res}^{1,u}|}\right) \quad (4.8)$$

$$\Delta\phi_{LM,opt} = \pm \arccos\left(\frac{b_{res}^{1,m} \cdot b_{res}^{1,l}}{|b_{res}^{1,m}| |b_{res}^{1,l}|}\right) \quad (4.9)$$

where $b_{res}^{1,u}, b_{res}^{1,m}$ and $b_{res}^{1,l}$ are the outermost pitch aligned components due to the upper, midplane and lower coils. Figure 4.24 plots b_{res}^1 as both $\Delta\phi_{UM}$ and $\Delta\phi_{LM}$ are scanned, as well as the values of $\Delta\phi_{UM,opt}$ and $\Delta\phi_{LM,opt}$ are predicted by equations 4.9. The figure shows that the formulae match the scan results, which makes 2D scans in $(\Delta\phi_{UM}, \Delta\phi_{LM})$ space unnecessary. It also demonstrates that the optimal phases are independent of the amplitudes of $b_{res}^{1,u}, b_{res}^{1,m}$ and $b_{res}^{1,l}$, and therefore independent of the absolute currents in the coil rows. As before, the sign uncertainty is resolved by simply choosing the sign which results in the larger amplitude of b_{res}^1 . Since the optimal coil phase is independent of the amplitude of the applied coil currents, the toroidal current waveforms applied in the upper, lower and middle coil rows are assumed to be perfect sinusoids with amplitude of 1kAt. As with the ASDEX Upgrade study, four cases of toroidal mode number are considered: $n = 1, 2, 3$ and 4. The phase differences $\Delta\phi_{UL,LM,UM}$ may be converted to the real space toroidal angle offsets by multiplying by n , as explained in figure 3.7.

4.6.1 ITER reference equilibrium and scaled equilibrium set

The reference equilibrium used here is the product of a CORSICA simulation [126], as used in previous works [75, 127, 128], which is intended to be the standard ELMy H-mode scenario for 15MA Q=10 ITER discharges[129]. Plasma profiles of the equilibrium are plotted in figure 4.25. In the original plasma rotation profile, the rotation sharply decreases in the pedestal and reaches zero at the plasma edge. It has been found recently that points in which the plasma rotation crosses zero can drastically alter the plasma




FIGURE 4.23: When completed, ITER will have 3 rows of 9 coils each, an upper, lower and midplane row. Therefore both the phase between the middle and upper coils $\Delta\phi_{UM}$, and between the middle and lower coils $\Delta\phi_{LM}$, may be varied.




FIGURE 4.24: b_{res}^1 plotted as a function of $\Delta\phi_{UM}$ and $\Delta\phi_{LM}$, as well as the location of $(\Delta\phi_{UM,opt}, \Delta\phi_{LM,opt})$ predicted using formulae 4.9.

response locally[130]. Since this zero crossing would coincide with the outermost pitch aligned components, and therefore interfere with this study, the rotation profile was modified to exclude the zero crossing and instead remain finite everywhere.

Using the CHEASE code, this reference equilibrium was scaled in current and pressure to produce a set of self consistent plasma equilibria, which are plotted in figure 4.26. The figure also plots 3 example cases of the scaled equilibria which demonstrates how the scaling was performed.




FIGURE 4.25: Profiles of the synthetic ITER reference equilibrium used for this study. a) Electron and Ion temperatures. b) Electron number density. c) Plasma bulk rotation velocity. d) Safety factor profile. e) Plasma pressure. f) Plasma current.

4.6.2 Parametrisation of ITER coil phase

For each point in the equilibrium set plotted in figure 4.26, the vacuum and total fields due to applied RMPs with the upper, midplane and lower coil sets were computed using MARS-F for $n = 1 - 4$. From these results, the optimal coil phases $\Delta\phi_{UM,opt}$ and $\Delta\phi_{LM,opt}$ are computed using equations 4.9. Figure 4.27 shows the resulting $\Delta\phi_{UM,opt}$ and $\Delta\phi_{LM,opt}$ as a function of (β_N, q_{95}) , for the case of $n = 2$. The figure shows the same behaviour as the ASDEX Upgrade coil phase: $\Delta\phi_{LM,opt}$ and $\Delta\phi_{UL,opt}$ decreasing with q_{95} and increasing with β_N . $\Delta\phi_{UM,opt}$ displays the opposite behaviour, which can be understood by considering that in the case of $\Delta\phi_{UL}$ and $\Delta\phi_{LM}$, the coil which is static is above the coil which is rotated, whereas this is reversed for $\Delta\phi_{UM}$, explaining the opposite behaviour with (β_N, q_{95}) . The behaviour described is general for all n studied here.

For each $n = 1 - 4$ and for the vacuum and total fields, a 2D quadratic parametrisation of the form in equation 4.3 is fitted to the scan for $\Delta\phi_{UM,opt}$, $\Delta\phi_{LM,opt}$ and $\Delta\phi_{UL,opt}$. The

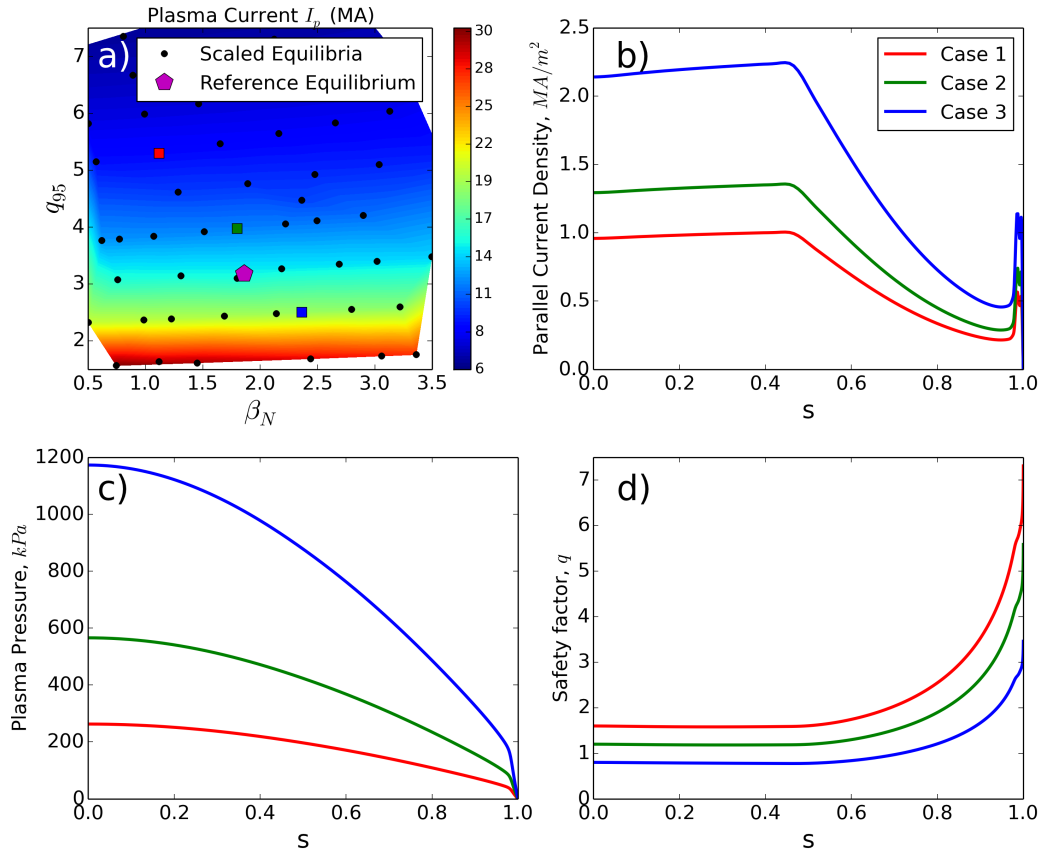


FIGURE 4.26: Equilibrium scaling is performed in the same way as section 4.2.1. a) The set of scaled equilibria, and three example equilibrium points used to demonstrate the variation of the profiles. For these example points the figure plots b) the plasma current density, c) the plasma pressure and d) the safety factor profile.

coefficients of the 2D quadratic parametrisation for $\Delta\phi_{UM,opt}$, $\Delta\phi_{LM,opt}$ and $\Delta\phi_{UL,opt}$ are listed in tables 4.4, 4.5 and 4.6 respectively. The RMSE values between the 2D quadratic parametrisation and the scan, as defined in equation 4.4, are listed in table 4.7, which demonstrates that the 2D quadratic parametrisation closely fits the scan data. The RMSEs between the scan data and fitted 2D quadratic parametrisation are lower than for the ASDEX Upgrade scan data, due to the ITER scan being coarser (fewer points for similar range of q_{95} and β_N) and therefore excluding the sharp RFA peaks evident in the ASDEX Upgrade (q_{95}, β_N) scan. Since there is not yet any experimental ITER data to use, the previous benchmarking and validation will not be possible. However, since the derivation procedure will be the same, the previous benchmarking and validation increases confidence in the ITER 2D quadratic parametrisation.




FIGURE 4.27: $\Delta\phi_{UM,opt}$, $\Delta\phi_{LM,opt}$ and $\Delta\phi_{UL,opt}$ as a function of (β_N, q_{95}) for $n = 2$. As found previously, the optimum coil phase is a smoothly varying function of (β_N, q_{95}) .

coeff	a	b	c	d	e	f	g	h	i
n=1 vacuum	-0.093339	0.65196	-1.7561	0.56293	-2.6838	9.1661	-0.32757	-11.609	-31.172
n=1 total	-0.030495	0.0022716	0.75543	0.42393	-1.6034	-3.3731	-0.23629	-13.996	-41.55
n=2 vacuum	-0.015844	-0.44548	0.70874	0.20367	4.4907	-6.2807	0.77959	-38.327	355.25
n=2 total	0.015959	-0.48245	1.4953	0.23243	2.7549	-11.035	0.53742	-37.057	333.41
n=3 vacuum	-0.099686	-0.16439	0.71514	0.57838	4.7165	-7.6413	0.88612	-53.879	354.39
n=3 total	-0.080711	0.081354	0.89951	0.71276	1.1883	-9.4082	0.90171	-54.174	334.55
n=4 vacuum	0.42932	-4.239	6.8794	-2.0015	25.825	-38.641	4.4783	-93.381	388.34
n=4 total	0.56781	-4.2563	5.9856	-2.1775	21.225	-33.297	4.5186	-91.101	360.36

TABLE 4.4: Coefficients of 2D quadratic parametrisation of $\Delta\phi_{opt,UM}$ for the ITER ELM coils.

coeff	a	b	c	d	e	f	g	h	i
n=1 vacuum	-0.0040945	0.23158	-0.16038	0.041565	-3.228	4.0099	-1.0093	26.254	-1.4536
n=1 total	-0.13202	1.8511	-6.1309	0.21643	-5.7448	24.184	-0.45666	22.314	23.664
n=2 vacuum	0.0073706	0.70125	-1.4941	-0.18402	-5.5805	9.2475	-0.67467	38.524	9.2087
n=2 total	0.13834	-0.67627	0.18032	-0.74705	0.64012	7.4795	-0.082304	35.022	32.787
n=3 vacuum	0.14185	0.076632	-0.82116	-0.81665	-4.3284	7.9093	-0.56298	54.108	9.3993
n=3 total	-0.11585	1.9406	-5.2864	-0.21655	-7.001	22.22	-0.83178	57.457	22.492
n=4 vacuum	0.45059	-1.6385	1.908	-2.3938	3.0733	-4.9654	0.94217	59.418	28.866
n=4 total	0.18459	-1.6625	3.042	-1.9643	7.9322	-9.5015	0.33647	63.183	43.615

TABLE 4.5: Coefficients of 2D quadratic parametrisation of $\Delta\phi_{opt,LM}$ for the ITER ELM coils.

coeff	a	b	c	d	e	f	g	h	i
n=1 vacuum	0.089244	-0.42037	1.5958	-0.52136	-0.54417	-5.1562	-0.68175	37.863	29.719
n=1 total	-0.10153	1.8488	-6.8863	-0.2075	-4.1414	27.557	-0.22037	36.311	-294.79
n=2 vacuum	0.023215	1.1467	-2.2028	-0.38769	-10.071	15.528	-1.4543	76.851	13.963
n=2 total	0.12238	-0.19382	-1.3149	-0.97948	-2.1147	18.514	-0.61972	72.079	-300.62
n=3 vacuum	0.24153	0.24102	-1.5363	-1.395	-9.0449	15.551	-1.4491	107.99	15.012
n=3 total	-0.035142	1.8592	-6.1859	-0.92931	-8.1893	31.628	-1.7335	111.63	-312.06
n=4 vacuum	0.021267	2.6004	-4.9714	-0.39234	-22.751	33.675	-3.5361	152.8	-359.47
n=4 total	-0.38322	2.5938	-2.9437	0.2132	-13.293	23.796	-4.1822	154.28	-316.74

TABLE 4.6: Coefficients of 2D quadratic parametrisation of $\Delta\phi_{opt,UL}$ for the ITER ELM coils.

rmse	n=1 vac	n=1 tot	n=2 vac	n=2 tot	n=3 vac	n=3 tot	n=4 vac	n=4 tot
UM	1.0421	0.50941	1.0409	1.0989	1.1638	1.7153	2.1527	4.8985
LM	1.741	0.86585	0.88217	1.5932	1.288	2.2488	3.4603	4.5708
UL	1.1213	1.1899	1.8665	2.3799	2.4254	3.7072	4.1296	4.3279

TABLE 4.7: RMSEs of 2D quadratic parametrisation of $\Delta\phi_{opt,UL}$ for the ITER ELM coils.

4.7 Database Regression Analysis

In previous studies[70, 87], it was established that b_{res}^1 is associated with ELM mitigation, while ξ_X is associated with density pump out, and it is observed that b_{res}^1 and ξ_X have common optima in computational scans of $\Delta\phi_{ul}$ [96], implying that a coil configuration optimised for ELM mitigation will also maximise density pump out, with negative implications for plasma performance. Furthermore, in chapter 3 it was demonstrated that the pitch aligned components may be driven by the amplified peeling response via poloidal harmonic coupling. Having collected a large number of experimental data points and their plasma response computations affords the opportunity to apply statistical methods to search for dependencies between variables in the dataset.

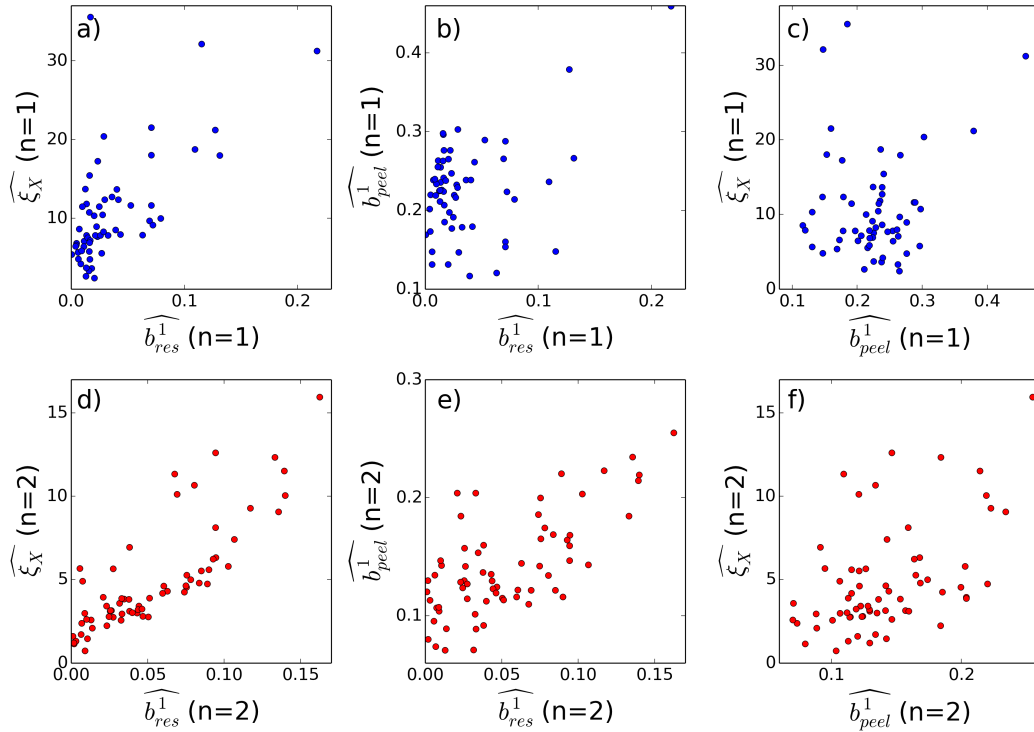


FIGURE 4.28: Normalised X point displacement $\widehat{\xi}_X$ against normalised pitch aligned component \widehat{b}_{res}^1 for a) $n = 1$ and d) $n = 2$ benchmarking points. Normalised peeling response \widehat{b}_{peel}^1 against normalised pitch aligned component for b) $n = 1$ and e) $n = 2$ points. Normalised X point displacement against normalised peeling response for c) $n = 1$ and f) $n = 2$ benchmarking points.

Figure 4.28 plots the outermost pitch aligned component b_{res}^1 against the local maximum plasma displacement at the X point ξ_X , as well as the peeling response measured as the maximum total field in the $m > nq$ region, against b_{res}^1 and ξ_X , for each of the benchmarking points described in section 4.3. Each point is normalised with the $n = 1, 2$ component of the applied coil currents. Since the optimum coil phase $\Delta\phi_{opt}$ varies between the benchmarking points, using a constant value of $\Delta\phi_{ul}$ would cause the result to be convoluted with the previously observed[96] correlation between b_{res}^1 and ξ_X as coil phase is varied. Therefore in order to remove $\Delta\phi_{ul}$ dependence from the study, each point is plotted using its computed value of $\Delta\phi_{opt}$. For the $n = 1$ results, the figure shows a weak correlation between ξ_X and b_{res}^1 , and no discernible correlations between the peeling response and ξ_X or b_{res}^1 . The $n = 2$ results however appear to show weak correlations between the peeling response and ξ_X and b_{res}^1 , and between b_{res}^1 and ξ_X , consistent with theory[88] and previous numerical observations[96].

In section 4.2.2, the question was raised as to whether the equilibrium pressure scaling should scale the pressure pedestal. This was investigated using the benchmarking dataset. For each of the benchmarking points, the pedestal pressure is approximated

by the pressure at the $\psi_N = 0.95$ flux surface, which is then normalised to the magnetic field. The approximate pedestal β values, defined as $2\mu_0 P_{ped}/B_0$ where P_{ped} is the plasma pressure at the 95% flux surface, and B_0 is the magnetic field strength at the magnetic axis, are plotted in figure 4.29 against normalised total pressure β_N . The results support the intuitive expectation that there should be a correlation between β_N and the pedestal pressure, given that the pedestal acts to raise the pressure globally, which supports the approach used to scale the pressure in section 4.2.1.




FIGURE 4.29: Plasma pressure β_N against approximate pedestal beta β_{ped} , defined as $2\mu_0 P_{ped}/B_0$ where P_{ped} is the plasma pressure at the 95% flux surface, and B_0 is the magnetic field strength at the magnetic axis. The strong correlation vindicates the approach to pressure scaling used in section 4.2.1.

To search for relationships between sets of variables, applying multivariate analysis is often a useful preliminary step. Previous studies have shown that the amplification of the plasma response increases with parallel current[89, 116], plasma pressure[116, 131], and plasma triangularity[116]. In this case, we may expect to see in the benchmarking dataset some (unspecified) relationship between equilibrium and pedestal properties, the amplified peeling response, and eventual RMP figures of merit. In this instance, the plasma parameters used as independent variables are $(\beta_N, q_{95}, \text{edge current peak and pedestal pressure})$, while the dependent variables are the peeling response, b_{res}^1 , and ξ_X . The pedestal pressure is approximated as described earlier, while the edge current peak is defined as the maximum parallel current density in the region of $\psi > 0.92$, which is expected to be dominated by the bootstrap current.

Since the relationship is not specified before the study, we know from the start that no conclusions may be drawn from the results, only preliminary indications to guide future studies. This is because it is not valid to test a hypothesis using the same data which is used to form the hypothesis, but rather a separate and distinct dataset must be used. However to provide insight from which some hypotheses may be formed (though not tested), multivariate ('brute force') linear regression was applied to the

benchmarking dataset. Applying this method to the benchmarking dataset detected no correlations between equilibrium quantities, and eventual figures of merit (normalised to the applied field). Since the equilibrium quantities are solutions of the same equation (Grad-Shrafranov), they are not independent. Interdependence of equilibrium quantities may be obscuring the expected correlations by partially compensating them. Further investigation of correlations between equilibrium and pedestal properties and the plasma response, is left for future work.

4.8 Choice of Toroidal Waveform

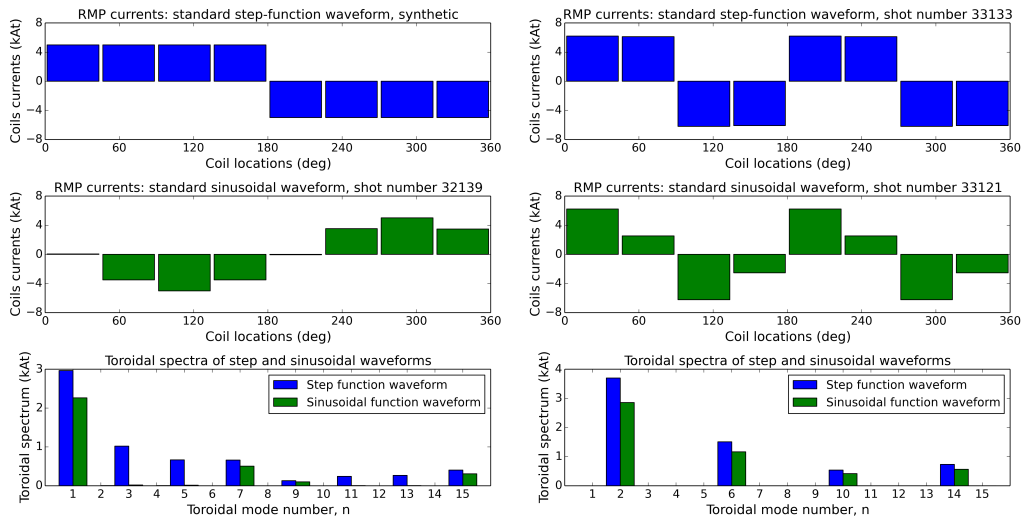


FIGURE 4.30: The choice of toroidal waveform - sinusoidal or rectangular - presents a trade off between flexibility and amplitude.




FIGURE 4.31: \hat{b}^1 represents the maximum penalty which may be incurred to b_{res}^1 due to coil phase misalignment. Since the function is flat around zero, small misalignments do not result in serious penalties, which makes rectangular waveforms viable.

The previous discussion focussed on optimisation of b_{res}^1 by tuning the coil phase difference $\Delta\phi_{ul}$. Precise alignment of the coil phase requires sinusoidal waveform which can take any arbitrary phase. However, when producing a sinusoidal waveform, most of the RMP coils are at less than maximum capacity. Figure 4.30 demonstrates this, comparing the toroidal waveform of a rectangular RMP toroidal waveform with an equivalent sinusoidal waveform, for $n = 1$ and $n = 2$ RMPs. Since the sinusoidal toroidal waveform more closely adheres to a true $n = 1, 2$ sinusoid, it has much lower sidebands than the rectangular waveform. However, because its coils are not all at maximum capacity, it also has a lower dominant toroidal component n_{dom} than the rectangular waveform. In this case, the dominant n component of the sinusoidal waveform $n_{dom,sin}$ is approximately 77% of the dominant n component of the rectangular waveform, $n_{dom,rec}$. This is also approximately true for $n = 2$ perturbations by symmetry, as figure 4.30 demonstrates. In the linear MARS model the n_{dom} component of the total field increases linearly with the n_{dom} component of the applied current, and therefore a reduction in n_{dom} causes a corresponding reduction in b_{res}^1 . When a square waveform is used, all the coils may be run at full capacity, but can only be tuned to certain discrete values of coil phase depending on number of coils toroidally and the toroidal mode number n . Namely, $\Delta\phi_{ul}$ may take values $j(360n/N)$ where $j = 0, 1, 2, (N/n) - 1$. For example with 8 coils, an $n = 2$ RMP field may have $\Delta\phi_{ul} = 0, 90, 180, \dots$, and an $n = 1$ may have $\Delta\phi_{ul} = 0, 45, 90, \dots$. Therefore, assuming the rectangular waveform is tuned to the nearest available phase to $\Delta\phi_{opt}$, the maximum misalignment due to the use of a rectangular instead of sinusoidal waveform will be $(360 * n/N)/2$, ie, 22.5 degrees for $N=8, n = 1$, or 45 degrees for $N=8, n = 2$.

In order to quantify the maximum penalty to b_{res}^1 which may be caused by misalignment, a normalised \hat{b}^1 is defined such that $\hat{b}^1(\Delta\phi_{opt}) = 1$ and $\hat{b}^1(\Delta\phi_{opt} + 180) = 0$ (ie, maximum misalignment). Figure 4.31 plots $\hat{b}^1(\Delta\phi_{ul})$. The figure shows that the worst case penalty (factor by which misalignment reduces b_{res}^1) of a 22.5 and 45 degree misalignment are 0.97 and 0.91 respectively. However in the general case the penalty will be less severe, as the minimum of $|b_{res}^1(\Delta\phi_{ul})|$ is generally not zero. Taking the case of $n = 2$, we now let $b_{opt,sin}$ be b_{res}^1 for an optimally aligned sinusoidal waveform, $b_{opt,rect}$ be b_{res}^1 for an optimally aligned rectangular waveform, and $b_{mis,rect}$ be b_{res}^1 for a 45 degree misaligned rectangular waveform. The previous discussion established that $b_{opt,sin} = 0.77b_{opt,rect}$, and from figure 4.31 we see that $b_{mis,rect} = 0.91b_{opt,rect}$. Therefore, $b_{opt,sin} = 0.85b_{mis,rect}$. That is to say, a 45 degree misaligned rectangular waveform leads to a larger b_{res}^1 than an optimally aligned sinusoidal waveform. The advantage of sinusoidal waveforms is that they may be precisely aligned so $\Delta\phi_{ul} = \Delta\phi_{opt}$, but this is more than compensated for by the required reduction in total coil current compared with a 'minimally misaligned' rectangular waveform.

4.9 Chapter Results Summary

It is demonstrated that β_N , q_{95} and n are key parameters for determining RMP coil alignment and hence $\Delta\phi_{opt}$. The variation of optimal RMP coil phase $\Delta\phi_{opt}$ with plasma equilibrium parameters β_N and q_{95} is quantified for $n = 1 - 4$ RMP perturbations and ASDEX Upgrade plasmas. A reference ASDEX Upgrade equilibrium was scaled in plasma pressure and current to produce a set of equilibria spanning the typical (β_N, q_{95}) parameter space of ASDEX Upgrade. For each point in the set, $\Delta\phi_{opt}$ was computed using MARS-F in the vacuum approximation and including the plasma response, for $n = 1 - 4$, which found that for a given n , $\Delta\phi_{opt}$ varies smoothly with β_N and q_{95} . A 2D quadratic function is used to parametrize $\Delta\phi_{opt}(\beta_N, q_{95}, n)$, for experimentalists to use in the planning of future experiments.

In order to assess the uncertainty in the 2D quadratic parametrization, a set of benchmarking points was assembled, each consisting of a free boundary equilibrium reconstructed from magnetic measurements using CLISTE, a set of kinetic profiles fitted to experimental data, and the experimentally applied RMP coils currents. For each benchmarking point $\Delta\phi_{opt}$ was computed both with the 2D quadratic parametrisation and a rigorous MARS-F calculation, which showed that the 2D quadratic parametrisation is accurate relative to a MARS-F computation to within 6.9 degrees for $n = 1$ vacuum predictions, 34.2 degrees for $n = 1$ total predictions, 14.1 degrees for $n = 2$ vacuum predictions, and 21.3 degrees for $n = 2$ total predictions. Using previously conducted coil phase scans from the ASDEX Upgrade database, the coil phase for optimal ELM mitigation was experimentally measured for a range of discharges with varying experimental conditions. The optimum coil phase was also predicted using MARS-F and the previously derived 2D parametrisation. It was found that the MARS-F predictions matched the experimental predictions to within 35 degrees, while the 2D quadratic parametrisation matched the experimental measurements to within 40 degrees. Since the 2D quadratic parametrisation is lightweight and simple, it may be used to compute $\Delta\phi_{opt}$ rapidly for large datasets, or as part of a feedback control mechanism for automatic coil phase optimisation, such that the coil phase may be optimised throughout a plasma discharge. This may be useful in ITER, which will require reliable RMP mitigation even during the current ramp up phase when q_{95} and β_N are varying. Further work is required to refine this approach however. The benchmarking procedure must be repeated for $n = 3, 4$ RMPs, and the quadratic parametrisation should be refined by recomputing it using a more representative plasma discharge, as well as for other tokamaks. Further study is also required to ascertain whether the primary source of uncertainty is the boundary shape, and if so, to incorporate shaping parameters into the parametrisation.

Using a predicted equilibrium of the ITER tokamak, the above procedure was repeated to parametrise the optimum upper and lower coil phase of the ITER ELM coils, and the coefficients of the resulting parametrisation are provided as a basis for a coil phase optimisation system for ITER.

With the benchmarking points collected for benchmarking the 2D parametrisation, a regression analysis was performed to test for correlations between equilibrium and plasma parameters and the eventual figures of merit relevant for ELM control. No significant relationships were found. However, a weak correlation was found between the amplitude of the $n = 2$ peeling response and the outermost pitch aligned component. This is consistent with the premise that the pitch aligned components may be driven by the amplified peeling response via poloidal harmonic coupling. A correlation between the pitch aligned component and the X point displacement was also found. This is consistent with previous theoretical works[88], which explain that the peeling response also drives X point displacement, which explains the correlation between the X point displacement and pitch aligned component.

By comparing the penalty (in terms of reduction in maximum b_{res}^1) of using a sinusoidal waveform and precisely aligning the RMP coils, to the penalty of using a rectangular waveform and imprecisely aligning the coils, it was shown that using a rectangular toroidal current waveform which may only be approximately aligned, produces a greater pitch aligned component than a precisely aligned sinusoidal waveform. This result relaxes the condition on the accuracy to which the optimal coil phase must be known, to the phase precision of a rectangular waveform. For a set of 8 toroidal RMP coils, this precision is 22.5 and 45 degrees respectively for an $n = 1$ and $n = 2$ RMP.

Chapter 5

Conclusion

5.1 Summary and Discussion

Chapter 1

Controlled nuclear fusion offers a clean, safe and permanent solution to energy scarcity, and the most promising route to a viable fusion reactor so far developed is the tokamak. The ITER tokamak, currently under construction, is designed to produce and sustain the world's first burning plasma. However, to ensure the success of the ITER experiment, robust ELM control techniques must be developed to avoid rapid erosion of the plasma facing components. ELMs may be mitigated (reduced in size) or suppressed entirely by the application of resonant magnetic perturbation fields at the plasma edge, and a set of ELM coils will be installed on ITER for this purpose. However, a robust predictive theory of ELM control with which an ELM control strategy may be developed is currently lacking.

The broad aim of this work was to study the plasma response to applied RMPs, in order to gain insight into the ELM suppression and mitigation mechanisms, as well as to develop and optimise the ITER ELM mitigation strategy.

Chapter 2

In chapter 2, the theory of Magnetohydrodynamics (MHD) is introduced, which combines Maxwells equations of electromagnetism with the equations of fluid dynamics, to model the plasma as an electrically conducting fluid, moving in response to external and self generated electric and magnetic fields. Models of plasma equilibrium and plasma

stability derived from the MHD model, which are central to tokamak science, are explained. It is explained that ELMs are driven by steep pressure gradients at the plasma edge and the resulting high plasma currents, and are triggered when the plasma exceeds the Peeling-Ballooning stability boundary.

Current working theories of RMP ELM mitigation and suppression are outlined. The original theory of ELM suppression on which the ITER ELM coil design is based[75], proposes that RMPs drive islands to form at rational surfaces, which may overlap to create a region of stochastic field at the plasma edge. This region has enhanced transport relative to unperturbed flux surfaces, and so prevents the pressure pedestal from building too high and crossing the P-B stability threshold, suppressing ELMs[64]. Another theory proposes that ELM suppression requires a zero crossing of the electron rotation to be co-located with a rational surface at the top of the pedestal. MHD simulations predict that this condition would cause a large island to form at the pedestal top, which would limit the pedestal growth such that it does not cross the P-B stability threshold[80–84]. A distinct theory proposes that ELM mitigation by RMPs may be explained by considering the stability properties of the 3D equilibria which result from RMP induced plasma boundary distortions. These 3D equilibria have been shown to be more unstable to P-B modes and therefore ELMs are triggered earlier in the ELM cycle, causing them to be more frequent and smaller[86].

Even in the absence of a robust theoretical understanding of ELM mitigation or suppression, ELM control strategies can be developed using figures of merit which are correlated with ELM effects. Three figures of merit derived from the plasma response are of particular interest. The pitch aligned component b_{res}^1 which drives island formation at the plasma edge, is correlated with mitigated ELM frequency[70, 72]. The X point displacement ξ_X is the maximum plasma displacement around the X point, and has previously been correlated with mitigated ELM frequency[70], and RMP induced density pump out[87]. The edge peeling response, referring to an amplified MHD mode localised at the plasma edge, is correlated with ELM suppression[92, 93], density pump-out[87], and ELM mitigation[70, 94]. Computing these figures of merit and the plasma response more generally, necessitates the use of the MARS-F single fluid spectral code, which computes the resistive linear plasma response to a magnetic perturbation applied to a given tokamak equilibrium, in realistic geometry and including toroidal plasma rotation.

Chapter 3

In chapter 3, the MARS-F code is used to compute the plasma response to an $n = 2$ RMP with a substantial $n = 6$ sideband, applied to an ASDEX Upgrade plasma discharge in

which significant ELM mitigation was observed. Using the Biot-Savart code ERGOS, the vacuum magnetic field was carefully benchmarked. The benchmark demonstrated that although the n_{dom} component of the field matched closely, using only the single largest toroidal harmonic results in a poor representation of the field in real space. However, when the largest sideband was also included, the ERGOS vacuum field was well recovered globally by the fourier based MARS-F vacuum field. This indicates that if an accurate representation of the global field is required (such as for particle tracing for example), the sidebands must be included. This may be accomplished by computing the fields for different toroidal harmonics separately, and summing them in post process. However, for many applications such as coil phase optimisation, the single n approach is sufficient.

The computed plasma response is dominated by two distinct mode structures; a core kink type mode which is apparent in the $n = 2$ response, and an edge localised peeling type mode which is apparent in both the $n = 2$ and $n = 6$ responses. The peeling mode has previously been predicted on DIII-D[93] where it is correlated with observed ELM suppression[92], and on MAST where it is correlated with density pump out[87]. In this instance - an ASDEX Upgrade discharge with significant ELM mitigation - the peeling response is predicted to be the dominant response.

At the X point of a divertor plasma is a null in the poloidal field, which in MARS-F would cause a numerical singularity. To avoid this, the X point is smoothed in MARS simulations to resemble a limiter plasma with an otherwise identical shape. In order to test the robustness of the predicted peeling response to the truncation of the X point geometry, a scan of the 'smoothness' of the X point is carried out and the resulting peeling response recomputed. The scan showed that the peeling response was robust in amplitude to the edge truncation, and the edge poloidal structure was radially distorted following the movement of the resonant curve $m = nq$. This indicated that the predicted peeling response was not merely a consequence of the X point truncation, as has previously been suggested[93].

It is shown that although the pitch aligned components are strongly screened by the plasma response, when finite resistivity is included they are not screened to zero in the edge region, where resistivity is larger than the plasma bulk. This implies that the strong screening of the pitch aligned components does not entirely preclude the formation of islands, but it does reduce their size and make them far less likely in regions of very low resistivity. Whereas ideal MHD would predict exact screening of all pitch aligned components, the resistive results indicate a stochastic region at the plasma edge, although a much thinner one than would be predicted excluding the plasma response.

By scanning the phase difference between the upper and lower toroidal waveforms $\Delta\phi_{ul}$, it was shown that both the peeling response and the pitch aligned components are strongly dependant on the poloidal spectrum of the RMP, which may be tuned using $\Delta\phi_{ul}$. It was also shown that the plasma response shifts the $\Delta\phi_{ul}$ at which b_{res}^1 is maximised by 60 degrees from its vacuum value. This indicates the importance of including the plasma response in interpreting RMP experiments. As well as greatly overestimating the size of induced islands and the extent of stochasticity, a vacuum prediction would fail to predict the optimal coil phase for outermost island width. Current best estimates for the efficacy of the ITER RMP system and the ELM control strategy rely on vacuum computations of stochasticity[75], based on a correlation observed on DIII-D between vacuum stochasticity and ELM suppression. Although the correlation is not in dispute, the underlying physical assumptions are clearly flawed, and therefore extrapolating this result to ITER may not be sound. A plasma response based assessment of the ITER ELM coils is required. It was also shown that the $n = 6$ pitch aligned components were strongly dependent on the $n = 6$ coil phase, which in the case of this $n_{dom} = 2$ applied field was shown to be the inverse of the n_{dom} coil phase, ie, $\Delta\phi_{n=6} = -\Delta\phi_{ul}$. When this was accounted for, it became apparent that the $n = 2$ and $n = 6$ pitch aligned components were maximised at similar values of $\Delta\phi_{ul}$. A previous work showed that the vacuum island overlap (ie, stochasticity neglecting the plasma response) may be enhanced significantly by the pitched aligned components of side bands, which fill in the gaps between islands of the dominant toroidal mode number[109]. It may be that the $n = 2$ and $n = 6$ components being maximised at the same $\Delta\phi_{ul}$, explains why single n plasma response studies are sufficient for finding correlations between the ELM frequency and pitch aligned component, at least for $n_{dom} = 2$ fields, since the $n = 2$ and $n = 6$ share the same $\Delta\phi_{ul}$ dependence. The alternative explanation is simply that the $n = 6$ component has a negligible effect on the ELM frequency. Examining the coil phase dependence of the sidebands of $n_{dom} = 1, 3$ fields may be a useful step to determine whether this behaviour is general, and whether it has any bearing on the ITER ELM control strategy.

By applying single m perturbations as a boundary condition at the plasma edge, it was shown that the pitch aligned components may be driven by poloidal harmonic coupling. This result, suggesting that a strongly amplified peeling response can also drive the pitch aligned components, has some interesting implications for ELM control theories. Some works suggest[64] that RMP induced stochasticity or a single large island[80, 81] is the cause of ELM suppression, while other theories propose that the edge perturbation due to the peeling response is responsible for ELM mitigation[86, 132]. The result here suggests that these two may be very difficult to decouple in experiments, since it implies a strong correlation between the peeling response, pitch aligned components and plasma

displacement. It may be then, that works which report a connection between the peeling response and ELM suppression or mitigation, also imply a connection between the pitch aligned components and these effects. It may also imply that the mechanisms of ELM suppression and mitigation mechanisms may interfere with and inhibit each other. If it were the case that the peeling response drove both the pitch aligned components and edge displacement, then the effect of the peeling response would be both stabilising (pitch aligned components driving islands to keep the plasma below the P-B stability boundary, causing ELM suppression), and simultaneously destabilising (edge displacement lowering the P-B stability boundary making ELMs more frequent, causing ELM mitigation). Therefore the ELM suppression and mitigation mechanisms may inhibit each other to some extent. It may also explain why suppression and mitigation are often seen to have overlapping parameter spaces[63], if both were driven by the peeling response. Although understanding is still obscured, from a practical perspective the correlations between numerous figures of merit associated with ELM suppression and mitigation may be useful for ELM control. For example, if a strategy aimed at inducing ELM suppression is not entirely successful, it may at least achieve significant mitigation by merit of the two effects occupying similar parameter spaces. That is, a strategy designed to achieve one, if unsuccessful may at least achieve the other.

In order to provide some experimental validation of the MARS-F predictions, the MARS-F code was used to compute magnetic and displacement perturbations for $n = 1$ and $n = 2$ applied RMPs, which were compared with direct measurements of the magnetic and displacement perturbations provided by collaborators. The magnetic measurements were made with arrays of B_p probes arranged around the machine, while the displacement was measured with ECE, ECE-I, CXRS and SXR diagnostics which measured displacement within the plasma. The MARS-F simulations predict correctly the magnetic and displacement poloidal structure of the edge, as well as the amplitude and phase of the displacement perturbation at the edge and for most of the plasma bulk. However, the measurements show a significant core kink response, which is not recovered in the MARS predictions. This result compliments validation of MARS performed on DIII-D[131], and of particular importance to this work is the validation of the $\Delta\phi_{ul}$ dependence, which improves confidence in the other $\Delta\phi_{ul}$ scans performed herein. However, the lack of a robust match in the amplitude, particularly in the core, remains a concern. Further investigation implicated the uncertainty in the equilibrium reconstruction as a possible cause. The safety factor profile and core plasma pressure are crucial parameters for the plasma response, and are highly uncertain in the core. This raises the issue of equilibrium uncertainty, not just for plasma response computations, but also for stability computations and global mode growth rates, which are of crucial importance for predicting the stability and performance of tokamak experiments[133]. In

order for uncertainties in plasma response and stability to be rigorously quantified, the uncertainties in the equilibrium reconstruction must be known. This is left for future work.

Chapter 4

Chapter 4 was focussed on developing and testing an optimisation scheme for ELM mitigation, based on tuning the coil phase $\Delta\phi_{ul}$ to maximise b_{res}^1 . It is explained that varying q_{95} or n modifies the alignment of the RMP (and therefore $\Delta\phi_{opt}$), by shifting the resonance curve relative to the applied spectrum. Varying β_N also modifies alignment, by distorting the magnetic geometry by Shafranov shift and thereby distorting the representation of the applied field in SFL coordinates. A previous result conversely found no dependence of $\Delta\phi_{opt}$ on β_N . It is suggested that this is due to differences in the approach to pressure scaling. This work scales the pressure profile self similarly including the pedestal, whereas in [91] the pressure pedestal is left constant.

An ASDEX Upgrade reference equilibrium is scaled in β_N and q_{95} to span the experimentally accessible range of (β_N, q_{95}) , and the optimal coil phase $\Delta\phi_{opt}$ was computed for each point using MARS-F. Experimentalists often need to know the optimal coil phase in advance of experiments, but using a plasma response computation this is often impractical due to the detailed input data required. The optimal coil phase is therefore parametrised with a simple 2D quadratic function of (β_N, q_{95}, n) , for use as a guide for future experiments. This 2D quadratic parametrisation is also envisaged to have a use in a feedback control system, which may control the RMP coil phase for optimum ELM mitigation throughout a discharge. This would be an especially useful feature for ITER, which requires ELMs to be well controlled even during the initial current ramp up, during which β_N and q_{95} and hence $\Delta\phi_{opt}$ vary considerably.

In order to benchmark the 2D quadratic parametrisation against rigorous MARS-F computations, a set of benchmarking points is gathered from the ASDEX Upgrade experimental database. Each point consists of an experimental ASDEX Upgrade reconstructed equilibria, experimentally measured kinetic profiles, and experimentally applied coil currents. For each benchmarking point, $\Delta\phi_{opt}$ is computed both using the 2D quadratic parametrisation (requiring only β_N, q_{95} and n as inputs), and using a rigorous MARS-F computation (which requires a equilibrium, set of kinetic profiles and coil currents as input). The parametrisation is found to be accurate relative to MARS-F to within 6.9 degrees for $n = 1$ vacuum predictions, 34.2 degrees for $n = 1$ total predictions, 14.1 degrees for $n = 2$ vacuum predictions, and 21.3 degrees for $n = 2$ total predictions. Sources of uncertainty are assessed and discussed. Having ruled out the likely sources,

the plasma shape is suggested as the primary source of error. This is further supported by an apparent systematic error in the $n = 1$ $\Delta\phi_{ul}$ predictions, and a corresponding systematic deviation of the $n = 1$ shapes from the reference equilibrium, while the $n = 2$ shapes show only random scatter around the reference equilibrium.

In order to experimentally validate the 2D quadratic parametrisation, the optimum coil phase is measured experimentally by measuring the ELM frequency and density during experimental scans of $\Delta\phi_{ul}$. These measurements are then compared with optimal coil phase as predicted by with MARS-F and the 2D parametrisation. The 2D parametrisation agrees with the experimental measurement to within 40.3 degrees, while the MARS-F computation agrees to within 35.0 degrees. It is interesting to note that the modelling used to achieve this agreement with experiment was purely single n . From this we may deduce that the effect of sidebands on ELM frequency is either small, or acts in synchrony with the dominant n component of the field. For each of the experimental scans, the dependence of ξ_X on $\Delta\phi_{ul}$ was also computed using numerical scans, and it was found that the optimum coil phase for ξ_X coincided with the optimum coil phase for b_{res}^1 , in agreement with a previous study[96]. Since ξ_X has previously been correlated with density pump out, and b_{res}^1 with ELM mitigation, this suggests that an RMP optimised for ELM mitigation will also necessarily cause the maximum reduction in density and associated confinement degradation. However, when the optimum coil phases for density pump out and ELM mitigation were measured experimentally, it was observed that the coil phase at which ELM mitigation was maximised was on average 59.0 ± 60.6 degrees lower than the coil phase at which density pump out was maximised, which suggests that mitigation and pump out may be decoupled, and that it may be possible to optimise ELM mitigation without also maximising confinement degradation. However, this may also be due to a time lag between the ELM frequency and density response, caused by the relatively slow particle transport timescale. It may be possible to distinguish between these possibilities using either a very slow coil phase scan, or a scan in which the coil phase was ramped in long flat steps.

In order that this coil phase parametrisation approach may be applied to ITER, the derivation process is repeated using a synthetic ITER equilibrium as the reference equilibrium and the ITER ELM coils. The coefficients of the 2D quadratic parametrisations for the upper and lower coil phases (relative to the middle row) are given as a starting point for future research, and as the basis for an automatic ELM mitigation optimisation system, which may be used to maintain good coil alignment even while plasma parameters vary.

To test a proposal that the amplified peeling response may be sensitive to edge pressure and current[116], the database of validation points is mined for correlations between

figures of merit and plasma equilibrium parameters. No significant correlations are discovered. This may be due to oversimplified approximations for the edge current and pedestal pressure. However, the database does show correlations between the peeling response and pitch aligned components, the peeling response and the X point displacement, and the pitch aligned components and X point displacement. Previous studies have reported these correlations for single discharges during which $\Delta\phi_{ul}$ was scanned, however correlations reported here are found in distinct discharges all at their respective values of $\Delta\phi_{opt}$. This further supports the theory that the amplified peeling response may drive the pitch aligned components via poloidal harmonic coupling.

The trade off between coil amplitude and coil phase flexibility was examined, by comparing the penalty to the pitch aligned components due to the minimum misalignment of a rectangular waveform, to the penalty due to using a sinusoidal waveform which does not use the RMP coils to full capacity. It was shown that the penalty was less severe for a minimally misaligned rectangular toroidal waveform than a precisely aligned sinusoidal waveform, and therefore the optimal strategy for maximising b_{res}^1 would be to apply a rectangular waveform, with a coil phase as close to $\Delta\phi_{opt}$ as the n and N numbers allow. However, it has been shown that strike point splitting caused by applied RMPs may cause localised regions of high heat flux at the divertor, which may be mitigated by rotating the RMP in order to distribute the heat flux over a larger area[134]. If a rigid rotation of the RMP is required to avoid material damage, rectangular waveforms will not be possible and therefore the precise alignment of a rotating sinusoidal RMP would be the best remaining strategy for maximising ELM mitigation.

5.2 Outlook

Preparing for successful ITER operation is a core objective of the UK Tokamak Science Programme and EFDA Fusion Roadmap. The motivation and objective of this and similar works, is to provide a robust prediction of the expected efficacy of the ITER ELM control system, and optimise the ITER ELM control strategy. Despite recent advances by other authors and the incremental contribution of this work, much more research in this field is required to achieve this.

Currently several ELM control theories exist, with varying amounts of experimental evidential support. However the evidence bases of these theories are mostly disparate, and are solely drawn from experiments on a single machine. In order to distinguish which theory is supported by a preponderance of the evidence, a multi-machine database approach is required, in which any theory may be tested against a comprehensive database from many tokamaks.

While MARS-F was shown in this work to be quite successful at recovering experimental measurements, further validation of the MARS-F code against experiment is required, particularly at high plasma pressures.

In order to expand the database of experimental observations and plasma response computations, against which theories of ELM control may be tested, the process of computing the plasma response for a given plasma equilibrium and set of kinetic profiles should be automated. This would allow plasma response computations to be performed for entire discharges, in a similar manner as is currently performed for equilibrium reconstruction. An automated plasma response code would be able to generate databases of plasma response computations of statistically significant size, spanning wide parameter space, which would be useful for finding correlations which may lead to further theory developments. Furthermore, since many RMPs contain significant sidebands, MARS-F computations should be multi- n as standard.

The 2D parametrisation of the optimal coil phase developed in this work, should be further refined using a more representative ASDEX Upgrade equilibrium and incorporating variations in plasma boundary shape. It may be used to form the basis of an automated feedback control of the ITER coil phase, to maintain optimal ELM mitigation during ramp up, and during systematic scans of plasma shape and safety factor. If a figure of merit is identified for ELM suppression which is also controllable by RMP coil phase, then this method could be applied to it also. The method should also be extended and tested thoroughly using $n = 3$, since ITER will use primarily $n = 3$ perturbations. DIII-D habitually uses $n = 3$ RMPs, making it an ideal test bed for this work.

The uncertainties inherent to equilibrium reconstruction should be thoroughly investigated, so that these errors may be propagated forward into the plasma response computations and the plasma response and stability uncertainty may therefore be assessed rigorously.

Currently the ITER ELM control strategies and assessments of the ELM coils are based only on vacuum predictions, which while a useful first step, have some serious deficiencies. It was shown in this work that the plasma response both drastically reduces the extent of the expected stochastic region at the edge, and also shifts the coil phase at which maximum ELM mitigation is expected by 60° . In light of this, assessment of the ITER ELM coils should be enhanced by considering the plasma response to the applied RMPs.

It is vitally important that progress in this field continues with urgency, in order to settle the question of how to prevent material damage by ELMs before the commencement of ITER operations.

Appendix A

Derivation of displacement resonance condition

We begin with the linearised ideal induction equation

$$\underline{b} = \underline{\nabla} \times (\underline{\xi} \times \underline{B}) \quad (\text{A.1})$$

We then choose a cylindrical coordinate system (s, χ, ϕ) , and consider only the component in the \hat{s} direction, by taking $J\underline{\nabla} \cdot$ of the above equation, where J is the (arbitrary) Jacobian.

$$\begin{aligned} J\underline{\nabla} \cdot \underline{b} &= J\underline{\nabla} \cdot [\underline{\nabla} \times (\underline{\xi} \times \underline{B})] \\ &= -J\underline{\nabla} \cdot [\underline{\nabla} \times (\underline{\xi} \times \underline{B})] \\ &= -J\underline{\nabla} \cdot [(\underline{\nabla} \cdot \underline{B}) - (\underline{\nabla} \cdot \underline{\xi})\underline{B}] \end{aligned} \quad (\text{A.2})$$

In the above the vector identities $\underline{A} \cdot (\underline{B} \times \underline{C}) = -\underline{B} \cdot (\underline{A} \times \underline{C})$, and $\underline{A} \times (\underline{B} \times \underline{C}) = (\underline{A} \cdot \underline{C})\underline{B} - (\underline{A} \cdot \underline{B})\underline{C}$ are used. We now define $b_1 = J\underline{\nabla} \cdot \underline{b}$, and in the above we use that $(\underline{\nabla} \cdot \underline{B}) = 0$ since \underline{B} has no radial component. The above then becomes

$$\begin{aligned} b_1 &= J\underline{\nabla} \cdot [(\underline{\nabla} \cdot \underline{\xi})\underline{B}] \\ &= J[(\underline{\nabla} \cdot \underline{\xi})(\underline{\nabla} \cdot \underline{B}) + \underline{B} \cdot \underline{\nabla}(\underline{\nabla} \cdot \underline{\xi})] \end{aligned} \quad (\text{A.3})$$

In the above, the identity $\underline{\nabla} \cdot (f\underline{A}) = f\underline{\nabla} \cdot \underline{A} + \underline{A} \cdot \underline{\nabla} f$ is used (where \underline{A} is an arbitrary vector and f an arbitrary scalar function). We now apply the non-existence of magnetic monopoles $\underline{\nabla} \cdot \underline{B} = 0$ to the above, and define $\xi_1 = \underline{\nabla} s \cdot \underline{\xi}$, to yield

$$\begin{aligned} b_1 &= J\underline{B} \cdot \underline{\nabla}(\underline{\nabla} s \cdot \underline{\xi}) \\ b_1 &= J\underline{B} \cdot \underline{\nabla} \xi_1 \end{aligned} \quad (\text{A.4})$$

Now $\xi_1(s, \chi, \phi)$, so we may expand out the dot product applying the chain rule to arrive at

$$\begin{aligned} b_1 &= J \frac{\partial \xi_1}{\partial \chi} \underline{B} \cdot \underline{\nabla} \chi + J \frac{\partial \xi_1}{\partial \phi} \underline{B} \cdot \underline{\nabla} \phi + J \frac{\partial \xi_1}{\partial s} \underline{B} \cdot \underline{\nabla} s \\ b_1 &= J \frac{\partial \xi_1}{\partial \chi} \underline{B} \cdot \underline{\nabla} \chi + J \frac{\partial \xi_1}{\partial \phi} \underline{B} \cdot \underline{\nabla} \phi \end{aligned} \quad (\text{A.5})$$

where we have used that $\underline{B} \cdot \underline{\nabla} s = 0$ since \underline{B} has no radial component. We now choose to represent the equilibrium field \underline{B} as

$$\begin{aligned} \underline{B} &= \underline{\nabla} \phi \times \underline{\nabla} \psi + F \underline{\nabla} \phi \\ &= \frac{d\psi}{ds} \underline{\nabla} \phi \times \underline{\nabla} s + \frac{JF}{R^2} \underline{\nabla} s \times \underline{\nabla} \chi \end{aligned} \quad (\text{A.6})$$

Using this representation also gives us $\underline{B} \cdot \underline{\nabla} \chi = \frac{d\psi}{ds} \frac{1}{J} = \frac{\psi'}{J}$ and $\underline{B} \cdot \underline{\nabla} \phi = \frac{F}{R^2}$, and from these we find also $\frac{\underline{B} \cdot \underline{\nabla} \phi}{\underline{B} \cdot \underline{\nabla} \chi} = q = \frac{JF}{\psi' R^2}$. Applying these the above becomes

$$b_1 = \left(\psi' \frac{\partial}{\partial \chi} + \frac{JF}{R^2} \frac{\partial}{\partial \phi} \right) \xi_1 \quad (\text{A.7})$$

Now we choose a straight field line coordinate system so all quantities may be fourier decomposed into poloidal and toroidal harmonics, such that $\frac{\partial}{\partial \chi} \rightarrow im$ and $\frac{\partial}{\partial \phi} \rightarrow -in$. Applying these to the above we arrive at

$$\begin{aligned} b_1 &= (\psi' im - \frac{JF}{\psi' R^2} in \psi') \xi_1 \\ b_1 &= i\psi'(m - qn) \xi_1 \end{aligned} \quad (\text{A.8})$$

rearranging the above we find the displacement resonance condition, which explains that ξ_1 will always be maximised around $q = m/n$

$$\xi_1 = \frac{b_1}{i\psi'(m - nq)} \quad (\text{A.9})$$

Bibliography

- [1] M Roser. Our World In Data. ourworldindata.org, 2016. Accessed: 02-2017.
- [2] The World Bank. World Development Indicators. data.worldbank.org, 2017. Accessed: 02-2017.
- [3] The World Bank. World Development Indicators, Life Expectancy at Birth. data.worldbank.org/indicator/SP.DYN.LE00.IN, 2017. Accessed: 02-2017.
- [4] The World Bank. World Development Indicators, Primary Completion Rate. data.worldbank.org/indicator/SE.PRM.CMPT.ZS, 2017. Accessed: 02-2017.
- [5] The World Bank. World Development Indicators, Poverty headcount ratio at \$1.90 a day (2011 PPP). data.worldbank.org/indicator/SI.POV.DDAY, 2017. Accessed: 02-2017.
- [6] M Roser. War and Peace. Published online at OurWorldInData.org. ourworldindata.org/war-and-peace, 2016. Accessed: 02-2017.
- [7] M Roser and E Ortiz-Ospina. World Population Growth. Published online at OurWorldInData.org. ourworldindata.org/world-population-growth, 2017. Accessed: 02-2017.
- [8] S Niu, Y Jia, W Wang, R He, L Hu, and Y Liu. Electricity consumption and human development level: A comparative analysis based on panel data for 50 countries. *International Journal of Electrical Power & Energy Systems*, 53(1):338–347, Dec 2013.
- [9] US Energy Information Administration of the US Department of Energy. International Energy Outlook 2016. Published online at eia.gov. eia.gov/outlooks/ieo/, 2016. Accessed: 02-2017.
- [10] IPCC Core Writing Team, R K Pachauri, and L A Meyer (eds.). Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. *Climate Change 2014: Synthesis Report*, page 151, 2014.

-
- [11] International Energy Agency. Key World Energy Statistics. Published online at iea.org. iea.org/publications/freepublications/publication/key-world-energy-statistics.html, 2016. Accessed: 02-2017.
- [12] British Petroleum. BP Statistical Review of World Energy June 2016. Published online at bp.com. bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy.html, 2016. Accessed: 02-2017.
- [13] D F Morehouse. The Intricate Puzzle of Oil and Gas Reserves Growth. *Natural Gas Monthly*, Jul 1997. Energy Information Administration, Office of Oil and Gas, US Department of Energy.
- [14] World Health Organisation. Ambient air pollution: A global assessment of exposure and burden of disease. Published online at who.int. who.int/phe/publications/air-pollution-global-assessment/en/, 2016. Accessed: 02-2017.
- [15] World Health Organisation. Preventing disease through healthy environments: a global assessment of the burden of disease from environmental risks. Published online at who.int. who.int/quantifying_ehimpacts/publications/preventing-disease/en/, 2015. Accessed: 02-2017.
- [16] M Wang and J E Overland. A sea ice free summer Arctic within 30 years? *Geophysical Research Letters*, 36(7), Apr 2009.
- [17] S Hales, N de Wet, J Maindonald, and A Woodward. Potential effect of population and climate changes on global distribution of dengue fever: an empirical model. *The Lancet*, 360(9336):830834, Sept 2002.
- [18] S R Loarie, P B Duffy, H Hamilton, G P Asner, C B Field, and D D Ackerly. The velocity of climate change. *Nature*, 462:1052–1055, Dec 2009.
- [19] O Hoegh-Guldberg and J F Bruno. The impact of climate change on the worlds marine ecosystems. *Science*, 328:1523–1528, Jun 2010.
- [20] J Sheffield and E F Wood. Projected changes in drought occurrence under future global warming from multi-model, multi-scenario, IPCC AR4 simulations. *Climate Dynamics*, 31(1):79 – 105, 2008.
- [21] T R Knutson, J L McBride, J Chan, K Emanuel, G Holland, C Landsea, I Held, J P Kossin, A K Srivastava, and M Sugi. Tropical cyclones and climate change. *Nature Geoscience*, 3:157–163, Feb 2010.
- [22] A L Westerling and B P Bryant. Climate change and wildfire in california. *Climate Change*, 87:231–249, Mar 2008.

- [23] International Energy Agency. Global EV Outlook 2016. Published online at iea.org. iea.org/publications/freepublications/publication/global-ev-outlook-2016.html, 2016. Accessed: 02-2017.
- [24] R Perez. The Intricate Puzzle of Oil and Gas Reserves Growth. *IEA SHCP Newsletter*, 62, Nov 1997. Energy Information Administration, Office of Oil and Gas, US Department of Energy.
- [25] T B Cochran, H A Feiveson, W Patterson, G Pshakin, M V Ramana, M Schneider, T Suzuki, and F von Hippel. Fast Breeder Reactor Programs: History and Status. *International Panel on Fissile Materials*, Feb 2010.
- [26] NASA. Sun Fact Sheet. Published online at nssdc.gsfc.nasa.gov. nssdc.gsfc.nasa.gov/planetary/factsheet/sunfact.html, 2016. Accessed: 02-2017.
- [27] International Atomic Energy Agency. *Fusion Physics*. International Atomic Energy Agency, Vienna, 2012.
- [28] Steven C Cowley. The quest for fusion power. *Nature Physics*, 12(5):384–386, May 2016.
- [29] A M Bradshaw, T Hamacher, and U Fischer. Is nuclear fusion a sustainable energy form? *Fusion Engineering and Design*, 86(9-11):2770–2773, Oct 2011.
- [30] P E Stott. The feasibility of using D3He and DD fusion fuels. *Plasma Physics and Controlled Fusion*, 47(8):1305–1338, Aug 2005.
- [31] E Lellouch, B Bezard, T Fouchet, H Feuchtgruber, T Encrenaz, and T de Graauw. The deuterium abundance in jupiter and saturn from iso-sws observations. *Astronomy & Astrophysics*, 370(2):610–622, 2001.
- [32] J G Delene. Updated comparison of economics of fusion reactors with advanced fission reactors. *Fusion Science and Technology*, 19(3P2A):807–812, May 1991.
- [33] A Nechaev, V Onufriev, and K T Thomas. Long-term storage and disposal of spent fuel. *IAEA Bulletin*, 1986.
- [34] D Maisonnier, D Campbell, I Cook, L Di Pace, L Giancarli, J Hayward, A Li Puma, M Medrano, P Norajitra, M Roccella, P Sardain, M Q Tran, and D Ward. Power plant conceptual studies in Europe. *Nuclear Fusion*, 47(11):1524, 2007.
- [35] A Glaser and R J Goldston. Proliferation risks of magnetic fusion energy: clandestine production, covert production and breakout. *Nuclear Fusion*, 52(4):043004, Apr 2012.

- [36] A P Meshkin. The Workings of an Ancient Nuclear Reactor. Published online at scientificamerican.com/article/ancient-nuclear-reactor/, 2009. Accessed: 02-2017.
- [37] John Wesson. *Tokamaks*. Oxford University Press, 2011.
- [38] Francis F. Chen. *Introduction to Plasma Physics and Controlled Fusion: Vol. 1*. Springer, 2006.
- [39] EUROfusion. Tokamak Principle. Published online at euro-fusion.org/euro-fusion.org/2011/09/tokamak-principle-2/, 2016. Accessed: 02-2017.
- [40] A J Webster. Fusion: Power for the future. *Physics Education*, 38(2):135–142, Mar 2003.
- [41] G McCracken and P Stott. *Fusion: The Energy of the Universe*. Academic Press, Vienna, 2012.
- [42] N J Peacock, D C Robinson, M J Forrest, P D Wilcock, and V V Sannikov. Measurement of the electron temperature by thomson scattering in tokamak T3. *Nature*, 224:488–490, 1969.
- [43] K Ikeda. ITER on the road to fusion energy. *Nuclear Fusion*, 50(1):014002, Jan 2010.
- [44] K Ushigusa and the JT-60 team. Steady state operation research in JT-60U. *Proc. 16th IAEA Conf. on Fusion Energy, (Montreal, Canada)*, 1996.
- [45] M Keilhacker, A Gibson, C Gormezano, P J Lomas, P R Thomas, M L Watkins, P Andrew, B Balet, D Borba, C D Challis, I Coffey, G A Cottrell, H P L de Esch, N Deliyankis, A Fasoli, C W Gowers, H Y Guo, G T A Huysmans, T T C Jones, W Kerner, R W T König, M J Loughlin, A Maas, F B Marcus, M F F Nave, F G Rimini, G J Sadler, S E Sharapov, G Sips, P Smeulders, F X Söldner, A Taroni, B J D Tubbing, M G Von Hellermann, D J Ward, and the Jet Team. High fusion performance from deuterium-tritium plasmas in JET. *Nuclear Fusion*, 39(2):209–234, Feb 1999.
- [46] ITER Physics Basis Editors, ITER Physics Expert Group Chairs, Co-Chairs, ITER Joint Central Team, and Physics Integration Unit. The ITER physics basis. chapter 1: Overview and summary. *Nuclear Fusion*, 39(12):2137, 1999.
- [47] M Shimada, D.J Campbell, V Mukhovatov, M Fujiwara, N Kirneva, K Lackner, M Nagami, V.D Pustovitov, N Uckan, J Wesley, N Asakura, A.E Costley, A.J.H Donné, E.J Doyle, A Fasoli, C Gormezano, Y Gribov, O Gruber, T.C Hender,

- W Houlberg, S Ide, Y Kamada, A Leonard, B Lipschultz, A Loarte, K Miyamoto, V Mukhovatov, T.H Osborne, A Polevoi, and A.C.C Sips. Chapter 1: Overview and summary. *Nuclear Fusion*, 47(6):S1–S17, Jun 2007.
- [48] 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, G G Lister, H M Mayer, D Meisel, E R Müller, H Murmann, H Niedermeyer, W Poschenrieder, H Rapp, H Röhr, F Schneider, G Siller, E Speth, A Stäbler, K H 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. *Physical Review Letters*, 49(19):1408–1412, Nov 1982.
- [49] P H Diamond, S-I Itoh, K Itoh, and T S Hahm. Zonal flows in plasma, a review. *Plasma Physics and Controlled Fusion*, 47(5):R35–R161, May 2005.
- [50] J W Connor, A Kirk, H R Wilson, and S Benkadda. Edge Localised Modes (ELMs): Experiments and Theory. In *AIP Conference Proceedings*, volume 1013, pages 174–190. AIP, 2008.
- [51] P T Lang, A Loarte, G Saibene, L R Baylor, M Becoulet, M Cavinato, S Clement-Lorenzo, E Daly, T E Evans, M E Fenstermacher, Y Gribov, L D Horton, C Lowry, Y Martin, O Neubauer, N Oyama, M J Schaffer, D Stork, W Suttrop, P Thomas, M Tran, H R Wilson, A Kavin, and O Schmitz. ELM control strategies and tools: status and potential for ITER. *Nuclear Fusion*, 53(4):043004, Apr 2013.
- [52] A W Leonard, A Herrmann, K Itami, J Lingertat, A Loarte, T H Osborne, W Suttrop, the Database Expert Group, and the ITER Divertor Modelling and Physics Expert Group. The impact of ELMs on the ITER divertor. *Journal of Nuclear Materials*, 266-269:109–117, Mar 1999.
- [53] A Loarte, G Saibene, R Sartori, D Campbell, M Becoulet, L Horton, T Eich, A Herrmann, G Matthews, N Asakura, A Chankin, A Leonard, G Porter, G Federici, G Janeschitz, M Shimada, and M Sugihara. Characteristics of type I ELM energy and particle losses in existing devices and their extrapolation to ITER. *Plasma Physics and Controlled Fusion*, 45(9):1549–1569, Sep 2003.
- [54] D G Whyte, A E Hubbard, J W Hughes, B Lipschultz, J E Rice, E S Marmor, M Greenwald, I Cziegler, A Dominguez, T Golfopoulos, N Howard, L Lin, R M McDermott, M Porkolab, M L Reinke, J Terry, N Tsujii, S Wolfe, S Wukitch, and Y Lin. I-mode: an H-mode energy confinement regime with L-mode particle transport in Alcator C-Mod. *Nuclear Fusion*, 50(10):105005, Oct 2010.

- [55] K H Burrell, M E Austin, D P Brennan, J C DeBoo, E J Doyle, P Gohil, C M Greenfield, R J Groebner, L L Lao, T C Luce, M A Makowski, G R McKee, R A Moyer, T H Osborne, M Porkolab, T L Rhodes, J C Rost, M J Schaffer, B W Stallard, E J Strait, M R Wade, G Wang, J G Watkins, W P West, and L Zeng. Quiescent H-mode plasmas in the DIII-D tokamak. *Plasma Physics and Controlled Fusion*, 44(5A):A253, 2002.
- [56] W Suttrop, M Maraschek, G D Conway, H U Fahrbach, G Haas, L D Horton, T Kurki-Suonio, C J Lasnier, A W Leonard, C F Maggi, H Meister, A Muck, R Neu, I Nunes, T Putterich, M Reich, A C C Sips, and the ASDEX Upgrade Team. ELM-free stationary H-mode plasmas in the ASDEX Upgrade tokamak. *Plasma Physics and Controlled Fusion*, 45(8):1399, 2003.
- [57] E J Doyle, W A Houlberg, Y Kamada, V Mukhovatov, T H Osborne, A Polevoi, G Bateman, J W Connor, J G Cordey, T Fujita, X Garbet, T S Hahm, L D Horton, A E Hubbard, F Imbeaux, F Jenko, J E Kinsey, Y Kishimoto, J Li, T C Luce, Y Martin, M Ossipenko, V Parail, A Peeters, T L Rhodes, J E Rice, C M Roach, V Rozhansky, F Ryter, G Saibene, R Sartori, A C C Sips, J A Snipes, M Sugihara, E J Synakowski, H Takenaga, T Takizuka, K Thomsen, M R Wade, H R Wilson, ITPA Transport Physics Topical Group, ITPA Confinement Database and Modelling Topical Group, and ITPA Pedestal and Edge Topical Group. Chapter 2: Plasma confinement and transport. *Nuclear Fusion*, 47(6):S18–S127, Jun 2007.
- [58] A Herrmann. Overview on stationary and transient divertor heat loads. *Plasma Physics and Controlled Fusion*, 44(6):883, 2002.
- [59] T E Evans. ELM mitigation techniques. *Journal of Nuclear Materials*, 438(SUPPL):S11–S18, Jul 2013.
- [60] P T Lang, G D Conway, and T Eich. ELM pace making and mitigation by pellet injection in ASDEX Upgrade. *Nuclear Fusion*, 44(5):665–677, 2004.
- [61] P T Lang, H Meyer, G Birkenmeier, A Burckhart, I S Carvalho, E Delabie, L Frassinetti, G Huijsmans, G Kocsis, A Loarte, C F Maggi, M Maraschek, B Ploekl, F Rimini, F Ryter, S Saarelma, T Szepesi, and E Wolfrum. ELM control at the L-H transition by means of pellet pacing in the ASDEX Upgrade and JET all-metal-wall tokamaks. *Plasma Physics and Controlled Fusion*, 57(4):045011, Apr 2015.
- [62] L R Baylor, N Commaux, and T C Jernigan. Experimental Demonstration of High Frequency ELM Pacing by Pellet Injection on DIII-D and Extrapolation to ITER. *Proc. 24th Int. Conf. on Fusion Energy 2012*, 2012.

- [63] A Kirk, I T Chapman, T E Evans, C Ham, J R Harrison, G Huijsmans, Y Liang, Y Q Liu, A Loarte, W Suttrop, and a J Thornton. Understanding the effect resonant magnetic perturbations have on ELMs. *Plasma Physics and Controlled Fusion*, 55(12):124003, Dec 2013.
- [64] T E Evans, R A Moyer, P R Thomas, J G Watkins, T H Osborne, J A Boedo, E J Doyle, M E Fenstermacher, K H Finken, R J Groebner, M Groth, J H Harris, R J La Haye, C J Lasnier, S Masuzaki, N Ohyabu, D G Pretty, T L Rhodes, H Reimerdes, D L Rudakov, M J Schaffer, G Wang, and L Zeng. Suppression of Large Edge-Localized Modes in High-Confinement DIII-D Plasmas with a Stochastic Magnetic Boundary. *Physical Review Letters*, 92(23):235003, Jun 2004.
- [65] Y M Jeon, J K Park, S W Yoon, W H Ko, S G Lee, K D Lee, G S Yun, Y U Nam, W C Kim, J G Kwak, K S Lee, H K Kim, and H L Yang. Suppression of Edge Localized Modes in High-Confinement KSTAR Plasmas by Nonaxisymmetric Magnetic Perturbations. *Physical Review Letters*, 109(3):035004, Jul 2012.
- [66] T E Evans, M E Fenstermacher, R A Moyer, T H Osborne, J G Watkins, P Gohil, I Joseph, M J Schaffer, L R Baylor, M Bécoulet, J a Boedo, K H Burrell, J S DeGrassie, K H Finken, T Jernigan, M W Jakubowski, C J Lasnier, M Lehnen, A W Leonard, J Lonroth, E Nardon, V Parail, O Schmitz, B Unterberg, and W P West. RMP ELM suppression in DIII-D plasmas with ITER similar shapes and collisionalities. *Nuclear Fusion*, 48(2):024002, Feb 2008.
- [67] Y Sun, M Jia, Q Zang, L Wang, Y Liang, Y Q Liu, X Yang, W Guo, S Gu, Y Li, B Lyu, H Zhao, Y Liu, T Zhang, G Li, J Qian, L Xu, N Chu, H H Wang, T Shi, K He, D Chen, B Shen, X Gong, X Ji, S Wang, M Qi, Q Yuan, Z Sheng, G Gao, Y Song, P Fu, and B Wan. Edge localized mode control using $n = 1$ resonant magnetic perturbation in the EAST tokamak. *Nuclear Fusion*, 57(3):036007, Mar 2017.
- [68] W Suttrop, T Eich, J C Fuchs, S Günter, A Janzer, A Herrmann, A Kallenbach, P T Lang, T Lunt, M Maraschek, R M McDermott, A Mlynek, T Pütterich, M Rott, T Vierle, E Wolfrum, Q Yu, I Zammuto, and H Zohm. First Observation of Edge Localized Modes Mitigation with Resonant and Nonresonant Magnetic Perturbations in ASDEX Upgrade. *Physical Review Letters*, 106(22):225004, Jun 2011.
- [69] A Kirk, Y Liu, E Nardon, P Tamain, P Cahyna, I Chapman, P Denner, H Meyer, S Mordijck, and D Temple. Magnetic perturbation experiments on MAST L- and H-mode plasmas using internal coils. *Plasma Physics and Controlled Fusion*, 53(6):065011, Jun 2011.

- [70] A Kirk, W Suttrop, I T Chapman, Y Liu, R Scannell, A J Thornton, L Barrera Orte, P Cahyna, T Eich, R Fischer, C Fuchs, C Ham, J R Harrison, M W Jakubowski, B Kurzan, S Pamela, M Peterka, D Ryan, S Saarelma, B Sieglin, M Valovic, and M Willensdorfer. Effect of resonant magnetic perturbations on low collisionality discharges in MAST and a comparison with ASDEX Upgrade. *Nuclear Fusion*, 55(4):043011, Apr 2015.
- [71] I T Chapman, W A Cooper, A Kirk, C J Ham, J R Harrison, A Patel, S D Pinches, R Scannell, and A J Thornton. Three-dimensional corrugation of the plasma edge when magnetic perturbations are applied for edge-localized mode control in MAST. *Plasma Physics and Controlled Fusion*, 54(10):105013, Oct 2012.
- [72] A Kirk, I T Chapman, J Harrison, Y Liu, E Nardon, S Saarelma, R Scannell, and A J Thornton. Effect of resonant magnetic perturbations with toroidal mode numbers of 4 and 6 on edge-localized modes in single null H-mode plasmas in MAST. *Plasma Physics and Controlled Fusion*, 55(1):015006, Jan 2013.
- [73] A Kirk, J Harrison, Y Liu, E Nardon, I T Chapman, and P Denner. Observation of Lobes near the X Point in Resonant Magnetic Perturbation Experiments on MAST. *Physical Review Letters*, 108(25):255003, Jun 2012.
- [74] M E Fenstermacher, T E Evans, T H Osborne, M J Schaffer, M P Aldan, J S DeGrassie, P Gohil, I Joseph, R A Moyer, P B Snyder, R J Groebner, M Jakubowski, A W Leonard, O Schmitz, and the DIII-D Team. Effect of island overlap on edge localized mode suppression by resonant magnetic perturbations in DIII-D. *Physics of Plasmas*, 15(5):056122, 2008.
- [75] T E Evans, D M Orlov, A Wingen, W Wu, A Loarte, T A Casper, O Schmitz, G Saibene, M J Schaffer, and E Daly. 3D vacuum magnetic field modelling of the ITER ELM control coil during standard operating scenarios. *Nuclear Fusion*, 53(9):093029, Sep 2013.
- [76] Y Liu, A Kirk, and E Nardon. Full toroidal plasma response to externally applied nonaxisymmetric magnetic fields. *Physics of Plasmas*, 17(12):122502, 2010.
- [77] R Fitzpatrick. A simple model of the resistive wall mode in tokamaks. *Physics of Plasmas*, 9(8):3459, 2002.
- [78] Y Liu, M S S Chu, Y In, and M Okabayashi. Resonant field amplification with feedback-stabilized regime in current driven resistive wall mode. *Physics of Plasmas*, 17(7):072510, 2010.

- [79] D A Ryan, Y Q Liu, A Kirk, W Suttrop, B Dudson, M Dunne, R Fischer, J C Fuchs, M Garcia-Munoz, B Kurzan, P Piovesan, M Reinke, and M Willensdorfer. Toroidal modelling of resonant magnetic perturbations response in ASDEX-Upgrade: coupling between field pitch aligned response and kink amplification. *Plasma Physics and Controlled Fusion*, 57(9):095008, Sep 2015.
- [80] M R Wade, R Nazikian, J S DeGrassie, T E Evans, N M Ferraro, R A Moyer, D M Orlov, R J Buttery, M E Fenstermacher, A M Garofalo, M A Lanctot, G R McKee, T H Osborne, M A Shafer, W M Solomon, P B Snyder, W Suttrop, A Wingen, E A Unterberg, and L Zeng. Advances in the physics understanding of ELM suppression using resonant magnetic perturbations in DIII-D. *Nuclear Fusion*, 55(2):023002, Feb 2015.
- [81] R Nazikian, C Paz-Soldan, J D Callen, J S DeGrassie, D Eldon, T E Evans, N M Ferraro, B A Grierson, R J Groebner, S R Haskey, C C Hegna, J D King, N C Logan, G R McKee, R A Moyer, M Okabayashi, D M Orlov, T H Osborne, J-k Park, T L Rhodes, M W Shafer, P B Snyder, W M Solomon, E J Strait, and M R Wade. Pedestal Bifurcation and Resonant Field Penetration at the Threshold of Edge-Localized Mode Suppression in the DIII-D Tokamak. *Physical Review Letters*, 114(10):105002, Mar 2015.
- [82] P B Snyder, T H Osborne, K H Burrell, R J Groebner, A W Leonard, R Nazikian, D M Orlov, O Schmitz, M R Wade, and H R Wilson. The EPED pedestal model and edge localized mode-suppressed regimes: Studies of quiescent H-mode and development of a model for edge localized mode suppression via resonant magnetic perturbations. *Physics of Plasmas*, 19(5):056115, 2012.
- [83] M Becoulet, F Orain, P Maget, N Mellet, X Garbet, E Nardon, G T A Huysmans, T Casper, A Loarte, P Cahyna, A Smolyakov, F L Waelbroeck, M Schaffer, T Evans, Y Liang, O Schmitz, M Beurskens, V Rozhansky, and E Kaveeva. Screening of resonant magnetic perturbations by flows in tokamaks. *Nuclear Fusion*, 52(5):054003, May 2012.
- [84] E Nardon, P Tamain, M Bécoulet, G Huysmans, and F L Waelbroeck. Quasi-linear MHD modelling of H-mode plasma response to resonant magnetic perturbations. *Nuclear Fusion*, 50(3):034002, Mar 2010.
- [85] M F Heyn, Ivan B Ivanov, S V Kasilov, W Kernbichler, I Joseph, R A Moyer, and A M Runov. Kinetic estimate of the shielding of resonant magnetic field perturbations by the plasma in DIII-D. *Nuclear Fusion*, 48(2):024005, Feb 2008.
- [86] I T Chapman, A Kirk, C J Ham, J R Harrison, Y Q Liu, S Saarelma, R Scannell, A J Thornton, M Becoulet, F Orain, W A Cooper, S Pamela, and MAST Team.

- Towards understanding edge localised mode mitigation by resonant magnetic perturbations in MAST. *Physics of Plasmas*, 20(5):056101, 2013.
- [87] Y Liu, A Kirk, Y Gribov, M P Gryaznevich, T C Hender, and E Nardon. Modelling of plasma response to resonant magnetic perturbation fields in MAST and ITER. *Nuclear Fusion*, 51(8):083002, Aug 2011.
- [88] Y Liu, C J Ham, A Kirk, L Li, A Loarte, D A Ryan, Y Sun, W Suttrop, X Yang, and L Zhou. ELM control with RMP: plasma response models and the role of edge peeling response. *Plasma Physics and Controlled Fusion*, 58(11):114005, Nov 2016.
- [89] A Wingen, N M Ferraro, M W Shafer, E a Unterberg, J M Canik, T E Evans, D L Hillis, S P Hirshman, S K Seal, P B Snyder, and a C Sontag. Connection between plasma response and resonant magnetic perturbation (RMP) edge localized mode (ELM) suppression in DIII-D. *Plasma Physics and Controlled Fusion*, 57(10):104006, Oct 2015.
- [90] C Paz-Soldan, R Nazikian, S R Haskey, N C Logan, E J Strait, N M Ferraro, J M Hanson, J D King, M J Lanctot, R A Moyer, M Okabayashi, J K Park, M W Shafer, and B J Tobias. Observation of a Multimode Plasma Response and its Relationship to Density Pumpout and Edge-Localized Mode Suppression. *Physical Review Letters*, 114(10):105001, Mar 2015.
- [91] C Paz-Soldan, N C Logan, S R Haskey, R Nazikian, E J Strait, X Chen, N M Ferraro, J D King, B C Lyons, and J K Park. Equilibrium drives of the low and high field side $n = 2$ plasma response and impact on global confinement. *Nuclear Fusion*, 56(5):056001, May 2016.
- [92] M J Lanctot, R J Buttery, J S de Grassie, T E Evans, N M Ferraro, J M Hanson, S R Haskey, R A Moyer, R Nazikian, T H Osborne, D M Orlov, P B Snyder, and M R Wade. Sustained suppression of type-I edge-localized modes with dominantly $n = 2$ magnetic fields in DIII-D. *Nuclear Fusion*, 53(8):083019, Aug 2013.
- [93] S R Haskey, M J Lanctot, Y Q Liu, J M Hanson, B D Blackwell, and R Nazikian. Linear ideal MHD predictions for $n = 2$ non-axisymmetric magnetic perturbations on DIII-D. *Plasma Physics and Controlled Fusion*, 56(3):035005, Mar 2014.
- [94] F Orain, M Hölzl, E Viezzer, M Dunne, M Bécoulet, P Cahyna, G T A Huijsmans, J Morales, M Willensdorfer, W Suttrop, A Kirk, S Pamela, S Günter, K Lackner, E Strumberger, and A Lessig. Non-linear modeling of the plasma response to RMPs in ASDEX Upgrade. *Nuclear Fusion*, 57(2):022013, Feb 2017.

- [95] Y Liu, D Ryan, A Kirk, L Li, W Suttrop, M Dunne, R Fischer, J C Fuchs, B Kurzan, P Piovesan, and M Willensdorfer. Toroidal modelling of RMP response in ASDEX Upgrade: coil phase scan, q95 dependence, and toroidal torques. *Nuclear Fusion*, 56(5):056015, May 2016.
- [96] L Li, Y Q Liu, A Kirk, N Wang, Y Liang, D Ryan, W Suttrop, M Dunne, R Fischer, J C Fuchs, B Kurzan, P Piovesan, M Willensdorfer, and F C Zhong. Modelling plasma response to RMP fields in ASDEX Upgrade with varying edge safety factor and triangularity. *Nuclear Fusion*, 56(12):126007, Dec 2016.
- [97] A D Turnbull. Plasma response models for non-axisymmetric perturbations. *Nuclear Fusion*, 52(5):054016, May 2012.
- [98] J P Hans Goedbloed and S Poedts. *Principles of Magnetohydrodynamics: With Applications to Laboratory and Astrophysical Plasmas*. Cambridge University Press, 2004.
- [99] E Frieman and M Rotenberg. On Hydromagnetic Stability of Stationary Equilibria. *Reviews of Modern Physics*, 32(4):898–902, Oct 1960.
- [100] M S Chu, J M Greene, T H Jensen, R L Miller, A Bondeson, R W Johnson, and M. E. Mauel. Effect of toroidal plasma flow and flow shear on global magnetohydrodynamic MHD modes. *Physics of Plasmas*, 2(6):2236, 1995.
- [101] Y Q Liu, A Bondeson, C M Fransson, B Lennartson, and C Breitholtz. Feedback stabilization of nonaxisymmetric resistive wall modes in tokamaks. I. Electromagnetic model. *Physics of Plasmas*, 7(9):3681, 2000.
- [102] M Willensdorfer, S S Denk, E Strumberger, W Suttrop, B Vanovac, D Brida, M Cavedon, I Classen, M Dunne, S Fietz, R Fischer, A Kirk, F M Laggner, Y Q Liu, T Odrščil, D A Ryan, E Viezzer, H Zohm, and I C Luhmann. Plasma response measurements of external magnetic perturbations using electron cyclotron emission and comparisons to 3D ideal MHD equilibrium. *Plasma Physics and Controlled Fusion*, 58(11):114004, Nov 2016.
- [103] P Piovesan, V Igochine, Y Q Liu, M Maraschek, L Marrelli, D A Ryan, W Suttrop, M Weiland, D Yadykin, M Gobbin, D Terranova, N M Ferraro, M J Lanctot, A Bogomolov, I Classen, M Dunne, A Gude, R McDermott, M Reich, and E Wolfrum. High-resolution internal measurements of 3D plasma response for model validation in high- β plasmas. In *42nd EPS Conference on Plasma Physics*, pages 2–5, 2015.
- [104] P Piovesan, V Igochine, F Turco, D A Ryan, M R Cianciosa, Y Q Liu, L Marrelli, D Terranova, R S Wilcox, A Wingen, C Angioni, A Bock, C Chrystal, I Classen,

- M Dunne, N M Ferraro, R Fischer, A Gude, C T Holcomb, A Lebschy, T C Luce, M Maraschek, R McDermott, T Odstrčil, C Paz-Soldan, M Reich, M Sertoli, W Suttrop, N Z Taylor, M Weiland, and M Willensdorfer. Impact of ideal MHD stability limits on high-beta hybrid operation. *Plasma Physics and Controlled Fusion*, 59(1):014027, Jan 2017.
- [105] P J 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, Jan 2012.
- [106] H Lütjens, A Bondeson, and O Sauter. The CHEASE code for toroidal MHD equilibria. *Computer Physics Communications*, 97(3):219–260, Sep 1996.
- [107] O Sauter, C Angioni, and Y R Lin-Liu. Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime. *Physics of Plasmas*, 6(7):2834, 1999.
- [108] E Nardon, M Bécoulet, G Huysmans, O Czarny, P R Thomas, M Lipa, R A Moyer, T E Evans, G Federici, Y Gribov, A Polevoi, G Saibene, A Portone, and A Loarte. Edge localized modes control by resonant magnetic perturbations. *Journal of Nuclear Materials*, 363-365:1071–1075, Jun 2007.
- [109] J Kim, G Park, C Bae, S Yoon, H Han, M G Yoo, Y S Park, W H Ko, J W Juhn, and Y S Na. Suppression of edge localized mode crashes by multi-spectral non-axisymmetric fields in KSTAR. *Nuclear Fusion*, 57(2):022001, Feb 2017.
- [110] G T A Huysmans. External kink (peeling) modes in x-point geometry. *Plasma Physics and Controlled Fusion*, 47(12):2107–2121, Dec 2005.
- [111] R Fitzpatrick. Stability of coupled tearing and twisting modes in tokamaks. *Physics of Plasmas*, 1(10):3308, 1994.
- [112] M J Lanctot, H Reimerdes, A M Garofalo, M S Chu, Y Q Liu, E J Strait, G L Jackson, R J La Haye, M Okayabashi, T H Osborne, and M J Schaffer. Validation of the linear ideal magnetohydrodynamic model of three-dimensional tokamak equilibria. *Physics of Plasmas*, 17(3):030701, 2010.
- [113] M. J. Lanctot, H. Reimerdes, A. M. Garofalo, M. S. Chu, J. M. Hanson, Y. Q. Liu, G. A. Navratil, I. N. Bogatu, Y. In, G. L. Jackson, R. J. La Haye, M. Okayabashi, J.-K. Park, M. J. Schaffer, O. Schmitz, E. J. Strait, and A. D. Turnbull. Measurement and modeling of three-dimensional equilibria in DIII-D. *Physics of Plasmas*, 18(5):056121, 2011.

- [114] S P Hirshman. Steepest-descent moment method for three-dimensional magneto-hydrodynamic equilibria. *Physics of Fluids*, 26(12):3553, 1983.
- [115] N M Ferraro. Calculations of two-fluid linear response to non-axisymmetric fields in tokamaks. *Physics of Plasmas*, 19(5):056105, 2012.
- [116] Y Liu, A. Kirk, L Li, Y. In, R Nazikian, Y Sun, W Suttrop, B Lyons, D A Ryan, S Wang, X Yang, and L Zhou. Comparative investigation of ELM control based on toroidal modelling of plasma response to RMP fields. *Physics of Plasmas*, 24(5):056111, May 2017.
- [117] J D King, E J Strait, R Nazikian, C Paz-Soldan, D Eldon, M E Fenstermacher, N M Ferraro, J M Hanson, S R Haskey, R J La Haye, M J Lanctot, S A Lazerson, N C Logan, Y Q Liu, M Okabayashi, J K Park, D Shiraki, and A D Turnbull. Three-dimensional equilibria and island energy transport due to resonant magnetic perturbation edge localized mode suppression on DIII-D. *Physics of Plasmas*, 22(11):112502, Nov 2015.
- [118] W Schneider, P J McCarthy, K Lackner, O Gruber, K Behler, P Martin, and R Merkel. ASDEX upgrade MHD equilibria reconstruction on distributed workstations. *Fusion Engineering and Design*, 48(1-2):127–134, Aug 2000.
- [119] Y Liu, S Saarelma, M P Gryaznevich, T C Hender, and D F Howell. Modelling resonant field amplification due to low- n peeling modes in JET. *Plasma Physics and Controlled Fusion*, 52(4):045011, Apr 2010.
- [120] L Li, Y Q Liu, A Kirk, N Wang, Y Liang, D Ryan, W Suttrop, M Dunne, R Fischer, J C Fuchs, B Kurzan, P Piovesan, M Willensdorfer, and F C Zhong. Modelling plasma response to RMP fields in ASDEX Upgrade with varying edge safety factor and triangularity. *Nuclear Fusion*, 56(12):126007, Dec 2016.
- [121] H Murmann, S Gotsch, H Rohr, H Salzmann, and K H Steuer. The Thomson scattering systems of the ASDEX upgrade tokamak. *Review of Scientific Instruments*, 63(10):4941, 1992.
- [122] W Suttrop, A G Peeters, the ASDEX Upgrade Team, and NBI Group. Practical limitations to plasma edge electron temperature measurements by radiometry of electron cyclotron emission. (*Technical Report IPP 1/306*) Garching (DE): Max-Planck-Institut fr Plasmaphysik, 1996.
- [123] E Viezzer, T Putterich, R Dux, R M McDermott, and the ASDEX Upgrade Team. High-resolution charge exchange measurements at ASDEX Upgrade. *Review of Scientific Instruments*, 83(10):103501, 2012.

- [124] M Reich, E Wolfrum, J Schweinzer, H Ehmler, L D Horton, J Neuhauser, and Asdex Upgrade Team. Lithium beam charge exchange diagnostic for edge ion temperature measurements at the ASDEX Upgrade tokamak. *Plasma Physics and Controlled Fusion*, 46(5):797–808, May 2004.
- [125] P Berens. CircStat : A MATLAB Toolbox for Circular Statistics. *Journal of Statistical Software*, 31(10), 2009.
- [126] J A Crotinger. Corsica; a comprehensive simulation of toroidal magnetic-fusion devices. *Lawrence Livermore National Laboratory Technical Report UCRL-ID-126284*, 1997.
- [127] L Li, Y Q Liu, N Wang, A Kirk, H R Koslowski, Y Liang, A Loarte, D Ryan, and F C Zhong. Toroidal modeling of plasma response to RMP fields in ITER. *Plasma Physics and Controlled Fusion*, 59(4):044005, Apr 2017.
- [128] Y Liu, R Akers, I T Chapman, Y Gribov, G Z Hao, G T A Huijsmans, A Kirk, A Loarte, S D Pinches, M Reinke, D Ryan, Y Sun, and Z R Wang. Modelling toroidal rotation damping in ITER due to external 3D fields. *Nuclear Fusion*, 55(6):063027, Jun 2015.
- [129] T Casper, Y Gribov, A Kavin, V Lukash, R Khayrutdinov, H Fujieda, and C Kessel. Development of the ITER baseline inductive scenario. *Nuclear Fusion*, 54(1):013005, Jan 2014.
- [130] B C Lyons, N M Ferraro, C Paz-Soldan, R Nazikian, and A Wingen. Effect of rotation zero-crossing on single-fluid plasma response to three-dimensional magnetic perturbations. *Plasma Physics and Controlled Fusion*, 59(4):044001, Apr 2017.
- [131] Z R Wang, M J Lanctot, Y Q Liu, J K Park, and J E Menard. Three-Dimensional Drift Kinetic Response of High- β Plasmas in the DIII-D Tokamak. *Physical Review Letters*, 114(14):145005, Apr 2015.
- [132] F Orain, M Hölzl, E Viezzer, M Dunne, M Bécoulet, P Cahyna, G T A Huijsmans, J Morales, M Willensdorfer, W Suttrop, A Kirk, S Pamela, S Günter, K Lackner, E Strumberger, and A Lessig. Non-linear modeling of the plasma response to RMPs in ASDEX Upgrade. *Nuclear Fusion*, 57(2):022013, Feb 2017.
- [133] T C Hender, J C Wesley, J Bialek, A Bondeson, A H Boozer, R J Buttery, A Garofalo, T P Goodman, R S Granetz, Y Gribov, O Gruber, M Gryaznevich, G Giruzzi, S Günter, N Hayashi, P Helander, C C Hegna, D F Howell, D a Humphreys, G T a Huysmans, A W Hyatt, A Isayama, S C Jardin, Y Kawano, A Kellman, C Kessel, H R Koslowski, R J La Haye, E Lazzaro, Y Q Liu, V Lukash, J Manickam,

- S Medvedev, V Mertens, S V Mirnov, Y Nakamura, G Navratil, M Okabayashi, T Ozeki, R Paccagnella, G Pautasso, F Porcelli, V D Pustovitov, V Riccardo, M Sato, O Sauter, M J Schaffer, M Shimada, P Sonato, E J Strait, M Sugihara, M Takechi, A D Turnbull, E Westerhof, D G Whyte, R Yoshino, H Zohm, and the ITPA MHD, Disruption and Magnetic Control Topical Group. Chapter 3: MHD stability, operational limits and disruptions. *Nuclear Fusion*, 47(6):S128–S202, Jun 2007.
- [134] A J Thornton, A Kirk, P Cahyna, I T Chapman, G Fishpool, J R Harrison, Y Q Liu, L Kripner, and M Peterka. ELM mitigation via rotating resonant magnetic perturbations on MAST. *Journal of Nuclear Materials*, 463:723–726, Aug 2015.