Optimising energetic particle transport in 3D fields in the ITER tokamak

Samuel Harvey Ward

Doctor of Philosophy

University of York

Physics

January 2022

Abstract

The high-gain ITER baseline tokamak plasma scenario depends upon reliable H-mode operation with edge-localised modes (ELM) suppressed by the application of resonant magnetic perturbations (RMP) using an ELM-control coil (ECC) system. However, these perturbations can lead to significant fast-ion transport and resulting power loads on plasma-facing components (PFC). This is of particular concern to the high-power discharges in ITER, which are extremely challenging to predictively study using conventional experimental devices and the physics codes currently used to study them.

In this work, software was created with which novel, high-performance computing components were assembled to enable, for the first time, routine, high-fidelity simulation of realistic fast-ion transport due to 3D fields in ITER. The primary fast-ion component, `LOCUST`, was verified and validated. The assembled workflow was then deployed to determine methods of operating the ITER ECC system where plasma heating efficiency, PFC power loads and ELM suppression are optimised simultaneously.

The response of fast-ion confinement to ECC operating parameters, such as coil current amplitude and phase, was discovered to correlate with ELM suppression. With this knowledge, the optimal method for operating the ECC system was determined, over a range of plasmas and applied RMP mode spectra, to increase total heating efficiency by 1.7-3.2% points in the baseline scenario. Even in the worst case, PFC power loads were found to be tolerable and rotation of the applied RMP to reduce power loads by up to 0.44MWm^{-2} (64%). In the optimal setting however, rotation may not be required, as minimum power loads and global losses were found to correlate. Methods for experimentally verifying these findings, using low-power plasmas and the diagnostic neutral beam, were also studied. Lowering the ECC coil current was found to be a low-risk approach to predicting fast-ion confinement in the ITER baseline scenario.

Contents

List of Tables

List of Figures

List of publications and conferences

Publications

S.H. Ward et al. *Verification and validation of the high-performance lorentz-orbit code for use in stellarators and tokamaks (LOCUST).* Nuclear Fusion, 61(8): 086029, 2021.

S.H. Ward et al. *LOCUST-GPU predictions of fast-ion transport and power loads due to ELM-control coils in ITER.* Nuclear Fusion, 62(12): 126014, 2022.

S.H. Ward et al. *Optimising fast-ion transport due to ELM-control coils in ITER discharges with LOCUST-GPU (in preparation).*

T.E.G. Nicholas et al. *Re-examining the role of nuclear fusion in a renewables-based energy mix.* Energy Policy 149, 2021.

Conferences

University of Oxford CUDA Programming on NVIDIA GPUs 2017. Oxford, UK. (workshop)
IoP Conference on Computational Plasma Physics 2017. York, UK. (poster)
ITER 2nd Code Camp on the Adaptation of Physics Codes/Workflows to IMAS 2018. St. Paul Lez Durance Cedex, France. (workshop)
General Atomics US TTF Workshop 2018. San Diego, USA. (poster)
Fusenet ITER PhD event 2018. St. Paul Lez Durance Cedex, France. (talk)
ITER 3rd Code Camp on Development, Validation and Demonstration of IMAS Workflows 2019. St. Paul Lez Durance Cedex, France. (workshop)
IAEA 16th Technical Meeting on Energetic Particles 2019. Shizuoka, Japan. (poster)
ITPA 22rd Energetic Particle Physics Topical Group 2019. Rovaniemi, Finland. (talk)
ITPA 23rd Energetic Particle Physics Topical Group 2019. Naka, Japan. (talk)
ITPA 24th Energetic Particle Physics Topical Group 2020. Online. (participant)
ITPA 25th Energetic Particle Physics Topical Group 2021. Online. (participant)

Acknowledgements

The project was partly undertaken on the ITER Organization's Titan GPU server within its Scientific Data & Computing Centre (SDCC).

The project was also partly undertaken on the GPU and CUMULUS clusters at Culham Centre for Fusion Energy.

The project was also partly undertaken on the Viking Cluster, which is a high performance compute facility provided by the University of York. I am grateful for the computational support from the University of York High Performance Computing service, Viking and the Research Computing team.

The work presented in this publication has been carried out under a joint PhD project among the ITER Organization, UKAEA and the University of York and has received financial support from the ITER Organization under Collaboration Agreement LEGAL#6411249v1.

The views and opinions expressed herein do not necessarily reflect those of the European Commission. The views and opinions expressed herein do not necessarily reflect those of the ITER Organization.

Of course, the work presented here would not have been possible without the ITER project and the University of York with its fusion CDT programme - a fantastic institution that produces outstanding scientists. Their commitment to this project is appreciated beyond comprehension. Undoubtedly, a significant portion of this commitment was fulfilled by my internal supervisor, Roddy Vann, who, as part of providing guidance on a whole range of matters, taught me the subtle and intricate dance that is "managing a project with many supervisors". Thanks are also owed to the remainder of this fast-ion triptych, my external supervisors Rob Akers and Simon Pinches. I must also thank Alberto Loarte at ITER, not just for the incredibly entertaining and memorable lunchtime conversations but also for affording this project access to what is, quite frankly, an encyclopaedic knowledge of the ITER tokamak and fusion science in general. Collaboration was at the heart of this work, so I would like to thank those who warmly hosted me at their institutions: Mike Van Zeeland and Bill Heidbrink, as well as Antti Snicker and the fantastic group at Aalto University. I'd also like to thank my other collaborators and co-authors, in particular Asger Jacobsen, Emmi Tholerus, Patrik Ollus, Yueqiang Liu and Todd Evans.

In this case, collaboration also refers to the myriad forms of external support I've been lucky to receive - cheers to you! Thank you to Annabel and Louis Gibb for hosting me in the early days. Cheers to Paddy Smith, Mana and La Cave des Ours for the delicious French food, coffee and wine respectively. Thanks to the friends I made at CCFE, CEA and ITER - probably the best place in the world for pot luck nights. Cheers to those from the original Manchester/London gang (including Archie, Lowell, Ross and Gaby) who came to visit and cause chaos wherever I happened to be. There's also an important group of entities who were an omnipresent lifeline for my mental health at night, abroad and when I was, in a general sense of the word, away: Kurzgesagt, RLM, NTS, in particular the Jamie xx and Magicwire shows, the latter of which I would not have survived without, as well as Lone, Jon Hopkins, Iron Galaxy, Tame Impala, FaltyDL, George FitzGerald's show on Radio 1, Leon Vynehall, Boards of Canada, Skee Mask, Kendrick Lamar, Angel Olsen, Machinedrum, pretty much everyone at Erased Tapes, Night Tapes, the New College Chapel choir, and Ralph Vaughan Williams.

One of the hardest parts of this project was the making, and subsequent stretching, of the incredible ties to members of the fusion CDT that helped hold me together. Thank you to Sam Gibson for all the fun, always making me feel at home and being an all star. Thank you to Jack Leland for indulging me with a great taste in tunes and tactical fleet-based space combat simulators. Cheers to Paul Barron, and our cutest couple award, and Charlie Vincent, for always being lovely. Of course thanks to the other great folk in our year, in particular: Tom Nicholas, Simon Orchard, Omkar Myatra, Will Trickey, Andrew Malcolm-Neale and possibly Bhavin Patel.

Finally I'd like to thank my family - my parents Gail and Richard for always supporting, teaching and loving me. The ability of my late grandparents, Margaret and Peter, to entertain their friends with their newfound - though brief - passion for tokamak science was uplifting. Lastly, my beloved partner Lauran and daughter Wallace: here's to the future.

Declaration

This thesis has not previously been accepted for any degree and is not being concurrently submitted in candidature for any degree other than Doctor of Philosophy of the University of York. This thesis is the result of my own investigations, except where otherwise stated. All other sources are acknowledged by explicit references.

Some of the work presented in this thesis has been published in the entries in the List Of Publications And Conferences section.

Chapter 1

Introduction

1.1 Fusion

Fusion is the nuclear reaction through which two parent nuclei combine to form a daughter nucleus. The daughter nucleus is typically metastable, so it quickly decays into another combination of daughter nuclei and nucleons. For elements lighter than iron, fusion is typically exothermic due to it increasing the binding energy per nucleon. Some examples of common fusion reactions are denoted below and their cross sections are shown in figure 1.1.

$$\begin{aligned}
{}^{2}\mathrm{H} + {}^{2}\mathrm{H} &\longrightarrow {}^{3}\mathrm{He}\,(0.8\text{ MeV}) + n\,(2.5\text{ MeV})\\
{}^{2}\mathrm{H} + {}^{2}\mathrm{H} &\longrightarrow {}^{3}\mathrm{H}\,(1.0\text{ MeV}) + \mathrm{H}\,(3.0\text{ MeV})\\
{}^{2}\mathrm{H} + {}^{3}\mathrm{H} &\longrightarrow {}^{4}\mathrm{He}\,(3.5\text{ MeV}) + n\,(14.1\text{ MeV})\\
{}^{2}\mathrm{H} + {}^{3}\mathrm{He} &\longrightarrow {}^{4}\mathrm{He}\,(3.6\text{ MeV}) + \mathrm{H}\,(14.7\text{ MeV})\\
{}^{3}\mathrm{He} + {}^{3}\mathrm{He} &\longrightarrow {}^{4}\mathrm{He} + 2\,\mathrm{H} + 14.7\,\mathrm{MeV}\\
{}^{3}\mathrm{H} + {}^{3}\mathrm{He} &\longrightarrow {}^{4}\mathrm{He} + \mathrm{H} + \mathrm{n} + 12.1\,\mathrm{MeV}\\
{}^{3}\mathrm{H} + {}^{3}\mathrm{He} &\longrightarrow {}^{2}\mathrm{H}\,(9.5\text{ MeV}) + {}^{4}\mathrm{He}\,(4.8\text{ MeV})\\
{}^{3}\mathrm{H} + {}^{3}\mathrm{H} &\longrightarrow {}^{4}\mathrm{He} + 2\,\mathrm{n} + 11.3\,\mathrm{MeV}\\
{}^{1}\mathrm{H} + {}^{11}\mathrm{B} &\longrightarrow 3\,{}^{4}\mathrm{He} + 8.7\,\mathrm{MeV}
\end{aligned}$$

Nuclear binding energy is extremely large compared to the atomic binding energy released

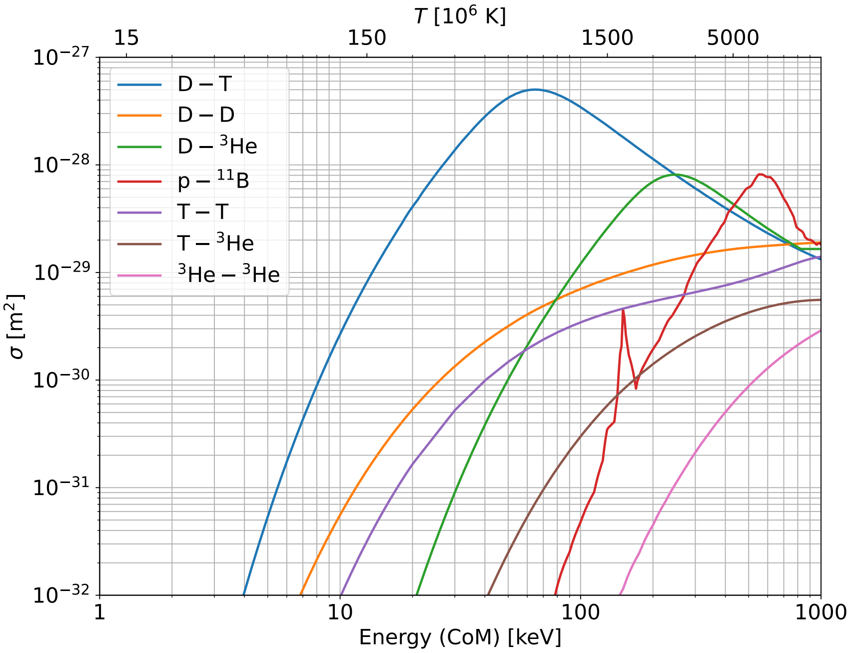

Figure 1.1: Cross sections for fusion reactions as a function of centre-of-mass energy [1, 2]. These are characterised by the existence of a bound state for the intermediate nucleus, which increases the cross section above that which would otherwise be found classically. Deuterium-tritium (DT) fusion exhibits the highest cross section at lower energies, making it attractive for use terrestrially compared with, for example, deuterium-helium-3 fusion, whose fuel would require lunar mining.

by chemical reactions. Out of fission and fusion, fusion typically releases the most energy per nucleon. This is shown by plotting the binding energy per nucleon of the nuclei of different elements in figure 1.2. Hence, moving rightwards via fusion of lighter elements increases the binding energy per nucleon much more quickly than moving leftwards via fission of heavier elements; the gradient essentially represents the energy density of the fuel. This makes fusion one of the densest sources of energy potentially available.

It is understandable then that, since its proposal and initial experimental study in the 1920s-30s, reactors which harness it have been considered a potential solution to many problems characteristic of the decade the research finds itself in. The first of which was a neutron source for the boosting of fission weapons in the late 1940s. In the 1950s, this application extended towards supplying warhead yields directly in thermonuclear weapons. Then, once these weapons had been mutually demonstrated by nations within each of the Eastern and Western Cold War Blocs, attention turned toward sustaining fusion in the laboratory, and the research was declassified in 1958. The hope was for a limitless energy supply, but the

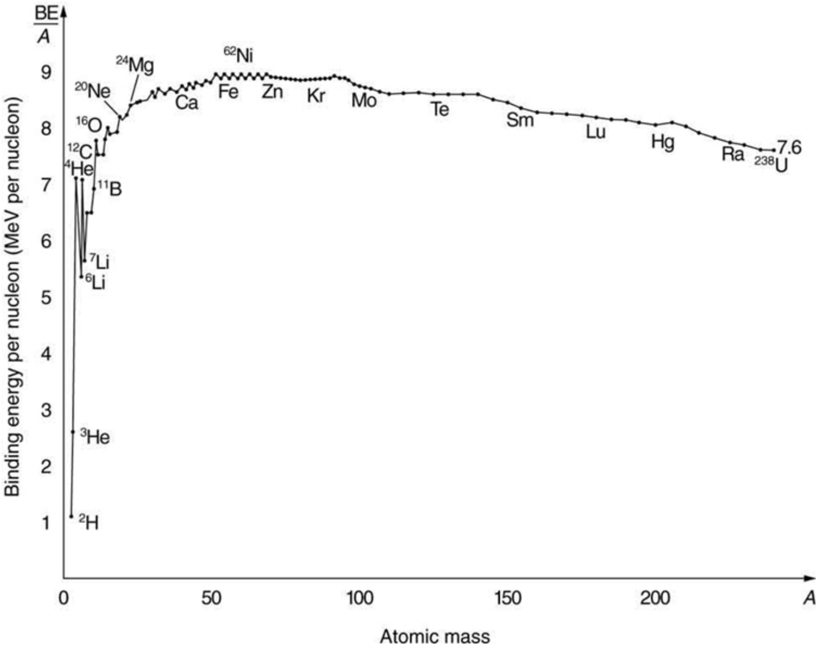

Figure 1.2: The binding energy per nucleon of elements. Lighter elements are shown on the left and, upon fusing, move rightward with increasing nuclear mass - releasing the energy associated with the change in binding energy per nucleon. This process may continue to release energy until iron is synthesised, as is the case in stellar fusion systems. Reproduced from [3].

following period in the 1960s would be marred by the systematic disproval of many reactor concepts. Nevertheless, in 1968 the increasingly globalised and cooperative fusion community observed a breakthrough: electron temperatures approaching 1keV, representing significant improvements in fuel confinement, were observed at the T3 tokamak experiment [4]. This re-instated fusion as a potential energy source. This was further bolstered by the 1973 oil crisis driving nations toward energy independence, a movement embodied by the United States' call for 1000 new nuclear power plants and the development of alternative sources of energy under Project Independence. Though the fission plants were never built, a second oil crisis in 1979 was enough cause for the United States' congress to pass the Magnetic Fusion Engineering Act of 1980, establishing a legal framework for the delivery of fusion energy at a time of peak oil prices. With a clear demand, and the belief that the next generation of experiments would observe more power released from fusion that required to heat the plasma, the 1980s hosted a wave of highly ambitious research experiments - the Joint European Torus (JET), Tokamak Fusion Test Reactor (TFTR), Japanese Torus 60 (JT-60) and Tokamak-15 (T-15). These experiments were conducted at both domestic and international levels, alluding to the willingness of large economies to compete, or small economies to co-operate, over fusion technology. However, net power remained elusive, and their findings made it clearer that significant research was still required. Meanwhile, the leftover oil glut and disaster at Chernobyl made fusion appear increasingly unattractive. Domestic fusion agencies would not consider experiments at this scale again until the mid 2010s. Instead, the next generation net-power experiments were appropriated as part of cold war bipartisanship, shifting the focus to internationally collaborative research which advocated larger, more powerful apparatus for lower costs to individual parties. This was the origin for the ambitious ITER project, which, until the 2010s, remained the only magnetic confinement experiment designed for net-gain to proceed.

It is important to examine the context around fusion, as it allows for more effective strategies for deployment based on an understanding of the historical drivers of research and investment. Evidently, fusion research is tightly coupled to the international sociopolitical climate, its perceived value and its competitiveness versus the technologies available at the time of investment - not necessarily the time of its realisation. In the twentieth century, the uses for fusion shifted from weapons to energy. And for the time it has predominantly been considered an energy source, the origin of its value shifted from abundance of its supply to reliability of its supply, before recently changing once more. It is around these changing values that fusion energy research strategy should adapt, in order to be pulled, as opposed to pushed, onto the grid as a competitive power source once the technology is proven [5]. Otherwise, no amount of physics research can guarantee that the technology will not remain underfunded and underappreciated - despite its clear potential.

Fusion research today is reminiscent of the enthusiastic period of the 1980s. Over half the world's population are participating in the ITER project, and many domestic agencies have recently begun plans for their own net-power experiments, followed closely by a rapidly growing private sector. Whilst the technology has always retained some value due to its inherent safety, high energy returned on energy invested, high energy density and fuel abundance, in recent years fusion's potential to contribute to a decarbonised energy infrastructure has increased this perceived value drastically.

The need for decarbonisation is driven primarily by anthropogenic climate change. Comprehensive assessments of the risks associated with climate change can be found in the reports prepared by the Intergovernmental Panel on Climate Change (IPCC). Recent reports by the

IPCC, such as SR15 [6], have more urgently advocated the need for rapid decarbonisation of the world energy infrastructure; the Earth's atmosphere is already 1.0°C hotter than it was at pre-industrial levels, and permanently capping this rise to 1.5° requires reaching net-zero carbon emissions by 2050. Net zero essentially requires a significant increase in worldwide installed capacity for zero-carbon energy generation, storage and transmission. This is to meet the demands of electrification of sectors such as heating, transport and power generation, whilst also satisfying compounding demands from future growth and development.

As it stands, no single conclusive solution exists to address the resulting supply gap. Due to their extremely low cost [7], it is highly likely that the capacity supplied by modern, utility-scale renewables, such as solar photovoltaic (PV) and wind, will expand to fill a large fraction - if not all - of the energy grid of energy supply and distribution [8]. Nevertheless it remains to be seen whether renewables can provide 100% of energy demands [5]. For this reason, and the advantages associated with renewable grids that combine baseload (sources which can reliably provide a power output) and dispatchable sources (sources that can quickly change their output in response to demand), research to bring fusion energy to the grid remains vital. The urgent challenge is now to facilitate the creation of such a combined grid by ensuring fusion is relevant, scalable and cost-competitive by the time it reaches deployment. It could be argued that the vast majority of fusion energy research today is - and should be - aimed toward this challenge.

1.2 Tokamaks

The ultimate goal of fusion reactors is to sustain a nuclear fusion reaction with minimal input power from external sources. Ignoring the movement of particles via fuel and waste cycles, the total power balance in a steady state fusion reaction can be expressed as

$$P_{\mathrm{heat}} = P_{\mathrm{external}} + P_{\mathrm{recycle}} = P_{\mathrm{loss}}, \tag{1.1}$$

where the power supplied to heat the plasma, P_{heat}, balances the power leaving the plasma P_{loss}, and is supplied by a combination of heating power injected externally from outside the plasma, P_{external}, as well as recycled power, P_{recycle}, generated by fusion and retained within the plasma. Fusion reactors aim to minimise P_{external}. In reality, this external power source is at least required to initially raise the centre-of-mass energy of the reactants to increase the reaction cross section to the maxima in figure 1.1 and thus increase the reaction rate.

Amongst the fusion fuels shown in figure 1.1, DT remains the most accessible due to its high cross-section at lower energies, in addition to fundamental constraints which may limit alternative fuel mixtures [9,10]. Nevertheless, there are significant challenges associated with DT fusion that must be overcome. The first challenge is that the required tritium fuel is radioactive with a half life of just 12 years, meaning there are little to no sizeable sources of extractable tritium outside of fission reactors - which, in some cases, may produce trace quantities as a byproduct. Therefore, fusion reactors operating with DT fuel must breed their own tritium, potentially by recycling the neutrons released by fusion. The second challenge is that these uncharged, high-energy neutrons cannot be magnetically confined like the plasma and instead immediately escape, posing a risk to the materials that comprise the reactor components. These risks stem from the neutron-induced material nuclear transmutations, lattice displacements and sputtering, which may produce activated waste, weaken and embrittle the reactor components, and may alter the transport characteristics of the plasma respectively.

Whilst it is possible to directly accelerate reactants to fuse them, known as beam-on-target fusion, net power is only thought to be practically achievable through the heating of the fuel to a sufficient temperature - typically 10keV (though the Coulomb barrier is ∼1MeV, quantum tunnelling increases the probability of crossing it at lower energies) for a 50:50 mix of deuterium and tritium - to sustain what is termed a "thermonuclear" reaction. In this case the fuel forms a plasma, and a significant source of heating power may come from the reaction products themselves re-depositing some of the energy released from fusion - P_{recycle} in equation 1.1 above. The point at which the total fuel heating power is supplied solely by P_{recycle} is referred to as "ignition", whereupon

$$P_{\text{external}} = 0. \tag{1.2}$$

To estimate the requirements of fulfilling this condition, P_{loss} can be defined as

$$P_{\text{loss}} = \frac{W}{\tau_E}, \tag{1.3}$$

where τ_E defines a characteristic time scale - the confinement time - over which the plasma energy is lost and the total thermal energy stored in a homogeneous plasma, W, can be expressed simply as

$$W = 3nTV, \tag{1.4}$$

where n is density of the plasma and V is the volume of the plasma. In a DT reaction, P_{recycle} is supplied by the re-deposition of the kinetic energy of the alpha particles produced by the fusion reaction, P_α. P_α can be expressed as the product of the reaction rate and the fusion alpha particle energy kinetic energy E_α:

$$P_\alpha = RVE_\alpha. \tag{1.5}$$

Whilst E_α is fixed by kinematics, an expression for the reaction rate per unit volume can be determined by considering the rate of a reaction with cross-section $\sigma(v)$, occurring between two homogeneous species (numbered 1 and 2) of similar density which have isotropic relative velocities v:

$$R = n_1 n_2 \left\langle \sigma v \right\rangle, \tag{1.6}$$

where an integral over the velocity has been performed such that the expectation value of the cross-section, $\langle \sigma v \rangle$, has been calculated. When the plasma is Maxwellian, as is the case for a thermalised plasma, $\langle \sigma v \rangle$ is a function of particle mass and species temperature only. For DT fusion, where the reaction rate is maximised by using a 50:50 mix of deuterium and tritium, equation 1.5 therefore becomes

$$P_\alpha = \frac{n^2}{4} \left\langle \sigma v \right\rangle E_\alpha V. \tag{1.7}$$

Substituting 1.3, 1.4 and 1.7 into 1.1 yields

$$P_{\text{external}} = \left(\frac{3nT}{\tau_E} - \frac{n^2 \left\langle \sigma v \right\rangle E_\alpha}{4} \right) V. \tag{1.8}$$

To achieve the ignition condition,

$$n\tau_E > \frac{12T}{\left\langle \sigma v \right\rangle E_\alpha}. \tag{1.9}$$

This inequality is known as the Lawson criterion. Remembering that $\langle \sigma v \rangle$ is a function of temperature ($\sim T^2$), and that $E_\alpha = 3.5$MeV, for a plasma temperature around 10-20keV, the Lawson criterion translates to

$$nT\tau_E > 3 \times 10^{21}\mathrm{m}^{-3}\mathrm{keVs}. \tag{1.10}$$

The quantity $nT\tau_E$ is known as the fusion triple product, and every fusion reactor aims to maximise it. ITER is expected to surpass the maximum achieved fusion triple product - by a tokamak - by approximately ten times [11]. It is also useful to define the fusion gain factor, Q, as

$$Q = \frac{P_{\mathrm{fusion}}}{P_{\mathrm{external}}}, \tag{1.11}$$

where the expression for P_{fusion} is defined in a similar way to equation 1.5 but for the total energy released by fusion reactions. At the point where internal heating, by α particles in a DT reaction, matches that from external sources, $P_\alpha = P_{\mathrm{external}}$,

$$Q = \frac{P_{\mathrm{fusion}}}{P_\alpha} \approx 5. \tag{1.12}$$

Therefore, $Q = 5$ is often used to label the point at which thermonuclear DT plasmas transition to the "burning" regime, where a significant amount of self-heating occurs. At ignition, $Q \to \infty$.

The confinement of the fuel can be maximised in various ways: using the inertia of the fuel, using magnetic fields or by using some combination of the two. The tokamak reactor concept is one of many proposed designs for a magnetic confinement fusion reactor actively researched today. From the point of view of fusion gain Q, tokamaks have shown the best performance of all magnetic confinement concepts, with JET reaching $Q = 0.67$ in 1997 for a short time and achieving $Q = 0.33$ for a sustained five seconds more recently in 2022.

One way of transitioning to net-gain reactors is widely believed to be through the increase of the reactor size. Whilst the physics behind this reasoning is more complex than simply increasing the surface area to volume ratio, the plasma performance has still been observed to follow scaling laws, which include the radius of the device, that have remained robust across many devices. These scaling laws underpin the physics basis for the design of large-scale tokamaks such as ITER, however it is yet to be seen whether they persist when crossing the boundary to burning plasma experiments. One example of such a scaling law is the ITER98

ELMy H-mode scaling law [12],

$$\tau_E^{\mathrm{IPB98}(y,2)} = 0.0562 I_p^{0.93} B^{0.15} \bar{n_{\mathrm{e}}}^{0.41} P_{\mathrm{heat}}^{-0.69} R^{1.97} \kappa^{0.78} \epsilon^{0.58} M^{0.19}, \tag{1.13}$$

(in s, MA, T, 10^{19}m^{-3}, MW, m) which, for a given plasma shape ϵ (inverse aspect ratio a/R_{maj}) and κ (elongation, whose product with ϵ gives a measure of the plasma height), and ion mass M in atomic mass units, has a strong dependency on the major radius R_{maj} and plasma current I_p.

1.2.1 ITER

The ITER tokamak is a high-performance experimental reactor designed for achieving high-gain ($Q \sim 10$), stationary ($\sim$ 500s) plasma discharges [13]. In demonstrating the novel physics of burning, magnetically confined plasmas, novel technologies must also be integrated, tested and successfully demonstrated. Examples include tritium fuel breeding and processing, handling and maintenance remotely due to activated or toxic materials, plant heat ejection systems for removing plasma exhaust power, and large-scale cryosystems for cooling superconducting magnets. These support systems are required for future demonstration and commercial plants, whose designs will build on ITER. ITER therefore aims to conduct incremental operational stages of increasing plasma power to safely facilitate the integration of these technologies. Likewise, this also provides opportunities for incremental physics studies prior to burning plasma scenarios.

The importance of ITER to the fusion community is therefore paramount, as it represents a key milestone for demonstrating both the physics and technology required to bring fusion to the grid. Nevertheless, there are still physics challenges to address in order to de-risk the project. These risks can be understood by examining various design parameters describing ITER in comparison with MAST-U and JET, one of the largest tokamaks capable of operating with DT at present. These are listed in table 1.1. It can be seen that the ITER design mirrors that of JET, with the same proportions but at a larger scale to lengthen the confinement time defined by equation 1.13. The larger plasma volume, and the increased toroidal magnetic field and plasma current required for confinement of a hotter plasma, typically lead to plasmas with much higher total stored energy. Therefore, the consequences and risks associated with degradation of the plasma confinement, and release of energy stored

Table 1.1: High-level design and performance parameters describing three types of tokamak: a relatively small, spherical device with low aspect ratio, a medium-sized device with conventional aspect ratio, and ITER, a large, near reactor-scale experimental device.

	MAST-U	JET	ITER
Major radius [m]	0.9	3.0	6.2
Minor radius [m]	0.7	1.2	2.0
Plasma volume [m^3]	8	100	840
Plasma current [MA]	2	3-5	15
B on axis [T]	0.8	3.5	5.3
Heating power (NBI) [MW]	5	20	33
Fusion power (max) [MW]	N/A	16	500

in the plasma, are amplified for ITER; reductions in plasma performance due to the loss of plasma confinement, and the resulting power fluxes to plasma-facing components (PFC), both have the potential to greatly impact the ability for ITER to demonstrate its primary goals of efficient and reliable fusion power.

There are many mechanisms which can lead to reductions in the confinement of the plasma and thus limit performance. ELMs are one of the most severe, ejecting significant amounts of energy and particles from the plasma. The primary method of mitigating ELMs in ITER is to suppress them using 3D magnetic field perturbations. At the basic level, the resulting transport and loss of ions at the plasma edge is thought to dampen the drivers of the instability (see section 2.3). However, this also affects the confinement of energetic particles injected externally to supply P_{external} in equation 1.1. Losing these particles before they can transmit power to the plasma degrades the external heating efficiency whilst also posing a risk to PFCs, as the external heating power in ITER, the magnitude of which is shown in table 1.1, is much larger than previous devices; if concentrated over a small surface area, this external power could surpass the $\approx$1-5MWm^{-2} maximum power load that the ITER reactor chamber wall is designed to handle. Therefore, it is extremely important that studies are performed to predict the impact of 3D fields on transport and confinement of energetic particles at the plasma edge, and to calculate the resulting PFC power loads; this would enable the application of perturbations in such a way that mitigates the risk of ELMs whilst also minimising reductions in plasma heating efficiency and risks of large PFC power loads. This is the main goal of this work.

The challenges presented by ITER's scale, complexity and novelty also translate to any attempts to predictively study the device; the performance of PFCs in fusion conditions, the physics surrounding burning plasmas and the feasibility of tritium breeding are all difficult to measure without the generation of a burning plasma, which remains unreproducible with the deuterium plasmas in the inefficient (poor confinement), small-scale experimental devices of today. To elaborate, these devices conduct experiments with deuterium plasmas to avoid the radiological, licensing and financial costs associated with using tritium fuel. However, generating a burning deuterium plasma requires maintaining confinement times and temperatures ($\sim$ 1 MeV, as seen in figure 1.1) much larger than unachievable with tokamaks today.

ITER must therefore be studied computationally. This is a core tenet of modern fusion research, where a small number of costly and temperamental practical experiments are supplemented by computational research in parallel. However, the majority of physics codes used to simulate tokamaks today have been designed around older, smaller devices and are often derived from software dating back decades to the previous millennium. Simulating ITER, on the other hand, involves much more expansive and extreme spatiotemporal scales - an increase that would require either approximately two hundred times more computational resources than are currently used or take Moore's law approximately eight doubling periods (16 years) to address passively. Therefore, ITER represents a computational paradigm shift that favours highly tuned, high-performance computing (HPC) codes. This work exploits one such novel HPC code, designed to simulate energetic particle transport efficiently using specialised algorithms and hardware acceleration, to address the aforementioned problem of energetic particle confinement in 3D fields in ITER. After verifying and validating the code against experiments and other well-trusted but less computationally efficient software tools, a software framework was constructed to enable its deployment at scales that were previously unattainable for ITER. These resulting simulations of energetic particles in ITER are at fidelities higher than achieved with other codes, using computing nodes that are comparatively low-power, cheap and small in size. As a consequence, the impacts of 3D fields on fast-ion confinement have been studied at high precision, over various plasma conditions and applied perturbation spectra, to successfully optimise both the overall plasma heating efficiency and power flux to PFCs.

Chapter 2 provides an overview of how these magnetic perturbations affect the fast-ion population in tokamaks. It first introduces the basic principles of magnetic confinement fusion, including the dynamics of fast-ion, processes which lead to their transport, and the structure of 3D perturbations. Then, chapter 3 describes the computational model used in this work to study fast ions in 3D fields. The fast-ion component of this model is explained in detail, as well as its algorithmic implementation on specific computational hardware and its coupling to other physics components and data sources. This model is then tested in chapter 4 for integration and convergence, with the fast-ion component specifically verified against similar codes and validated against experiments - including those which measure fast-ion transport due to 3D fields. After testing, the model is deployed in chapter 5 to study fast ions in the ITER 3D field system. By estimating fast-ion losses, and the associated power loads to in-vessel components, methods for operating the coil system to optimise fast-ion transport and improve the plasma heating efficiency of the device are qualified. The feasibility of testing these results experimentally using diagnostics available in ITER is also judged. Finally, chapter 6 summarises the work of this thesis, as well as discussing avenues for future research arising from the findings made herein.

Chapter 2

Fast particles in tokamak plasmas

This chapter outlines concepts that are key to magnetic confinement fusion and fast-ion transport, and it provides an overview of the current research aiming to understand the impact of ELM-control coils (ECC) on the heating performance of fast-ion populations. The specifics which underpin the tokamak approach to fusion are introduced in section 2.1. The dynamics of fast ions are described in section 2.2, where the distinction between anomalous and collisional transport is made to differentiate transport arising from ECCs and that which is otherwise naturally expected. The motivation behind and implementation of ECCs is explained in section 2.3, before an overview of their theorised and measured impact on fast-ion confinement is given in section 2.4.

2.1 Tokamak plasmas

Tokamaks are ultimately devices designed to satisfy the Lawson criterion in equation 1.9 by maximising the triple product in equation 1.10. This means first confining a fuel of sufficiently high density for a sufficiently long time. To maximise the DT fusion reaction cross-section, this confined fuel must then be heated to approximately 10-20keV, at which point it exists as a plasma composed of ionised nuclei and their corresponding free electrons. The combination of these electrons and ions is approximately quasineutral, especially at length scales much greater than the distance over which charge imbalances are naturally screened - the Debye length. The tokamak fusion reactor approach is to confine this plasma using electromagnetic fields created by a system of high-strength magnets.

2.1.1 Particle motion and confinement

To understand the tokamak concept, it is important to first discuss the dynamics of ions in an electromagnetic field to understand how a plasma is confined. The equations of motion of these ions [14,15] can be obtained by applying Lagrange's equations to the particle phase-space Lagrangian. For a particle moving in a Lorentz-invariant velocity-dependent potential,

$$U = q\phi - q\mathbf{A}\cdot\mathbf{v}, \tag{2.1}$$

where the potential can be used to define the electric and magnetic fields,

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t} \tag{2.2}$$

$$\mathbf{B} = \nabla\times\mathbf{A}, \tag{2.3}$$

the Lagrangian L, consisting of the particle kinetic and potential energy T and U respectively, becomes

$$L = T - U = \frac{1}{2}mv^2 - q\phi + q\mathbf{A}\cdot\mathbf{v}. \tag{2.4}$$

Applying Lagrange's equations, in the case where the generalised coordinates q_j are the real-space position components of the particle,

$$\frac{d}{dt}\left(\frac{\partial L}{\partial \dot{q_j}}\right) - \frac{\partial L}{\partial q_j} = 0, \tag{2.5}$$

to the expression in equation 2.4 yields the dynamical equation for position $\mathbf{r}$ and velocity $\mathbf{v}$ of an ion of charge q and mass m:

$$\frac{d^2\mathbf{r}}{dt^2} = \frac{d\mathbf{v}}{dt} = \frac{q}{m}(\mathbf{v}\times\mathbf{B}(\mathbf{r}) + \mathbf{E}(\mathbf{r})). \tag{2.6}$$

The main mechanism by which magnetic fields in tokamaks confine ions can now be seen by examining equation 2.6; ions may freely move parallel to magnetic field lines (with velocity $v_{\|}$), whilst perpendicular motion (with velocity $v_{\perp}$) leads to an orbit around the field lines. The centre of this orbit is referred to as the guiding centre, and the frequency at which these orbits are completed is called the cyclotron (or gyro) frequency,

$$\Omega_\mathrm{c} = \frac{q}{m}B. \tag{2.7}$$

The radius of the orbit is

$$\rho = \frac{v_\perp}{\Omega_\mathrm{c}}, \tag{2.8}$$

which is influenced by the particle pitch,

$$\lambda = \frac{v_\parallel}{v}. \tag{2.9}$$

The orbit around the field line also defines a current loop, giving an equivalent magnetic moment which is defined, to zeroth order [16], by

$$\mu = \frac{mv_\perp^2}{2B}. \tag{2.10}$$

These parameters remain constant in a homogeneous magnetic field. In this case, the guiding centre is fixed in the plane normal to the field lines. However, where the magnetic field is changing and these parameters vary, particle motion can become more complex. For example, if a particle orbit crosses a boundary where the magnetic field strength suddenly changes, the resulting change in the orbit width causes the guiding centre to move. In these cases, the particle can move in this plane normal to the magnetic field, and this component of its motion is referred to as cross-field transport.

In the case where the magnetic field strength varies along the field lines, the guiding centre velocity changes:

$$m\frac{dv_\parallel}{dt} = -\mu\nabla_\parallel B. \tag{2.11}$$

In this case, the magnetic field may reflect the ion due to what is called the magnetic mirror effect, in an analogy with particle movement on the surface of a cone [17]. The underlying mechanics may be understood after realising the magnetic moment is an approximate constant-of-motion - an adiabatic invariant [16]. If the magnetic field strength steadily increases, then the point in space at which ions are eventually reflected is dictated by their pitch angle; as a particle moves into a mirror, and its pitch steadily decreases, if the starting pitch of the particle is sufficiently high, then it will pass through the mirror before its pitch has reached zero. Particles with certain pitches may not be reflected at all, for example a particle with a pitch of unity. The region of phase space containing these particles is referred to as the loss cone. Drifts in the particle guiding centre motion arise due to magnetic field curvature,

$$v_{\text{curve}} = -\frac{v_{\parallel}^2}{\Omega}\frac{\mathbf{R}_c \times \mathbf{B}}{R_c^2} \tag{2.12}$$

where $\mathbf{R}_c = R_c^2 \nabla B / B$ points outward from the origin of the radius of curvature [18], leading to

$$v_{\text{curve}} = \frac{v_{\parallel}^2}{\Omega_{\text{c}}}\frac{\mathbf{B} \times \nabla B}{B^2}; \tag{2.13}$$

perpendicular magnetic field gradients,

$$v_{\nabla B} = \frac{v_{\perp}^2}{2\Omega_{\text{c}}}\frac{\mathbf{B} \times \nabla B}{B^2}; \tag{2.14}$$

and external forces,

$$v_{\mathbf{F}\times\mathbf{B}} = \frac{\mathbf{F} \times \mathbf{B}}{qB^2}, \tag{2.15}$$

the latter of which also describes the guiding-centre motion in the presence of an electric field (the $\mathbf{E} \times \mathbf{B}$ drift) or gravitational field, whose small contribution is typically ignored.

The tokamak is a result of designing a magnetic field topology which accounts for the total motion due to the above drifts so as to confine test particles indefinitely over a complete orbit. A straight, homogeneous magnetic field successfully removes cross-field transport in a system composed of non-interacting test particles, however the number of ions streaming out either end along the field lines is too great - even in the case where the field strength is increased at each end to reflect some ions, as in a magnetic mirror machine. By instead bending this field to create an axisymmetric torus, the ions may freely stream along the toroidal field lines without end. However, this bending introduces a $1/R$ dependency on the magnitude of the magnetic field, which is now stronger on the inboard side. This gradient in the field strength, and resulting magnetic drift, leads to a vertical charge separation and electric field, and the resulting radial $\mathbf{E} \times \mathbf{B}$ drift leads to a rapid loss. Therefore, a poloidal field component is added to enable ions to move between the upper and lower halves of the device. Any vertical charge imbalance is now quickly shorted-circuited by rapid transport along the now-helical field lines to preserve quasineutrality.

An illustration of this geometry is given in figure 2.1. On the left is a plan view of the torus looking down on the plasma, shown in pink. The coordinate system origin is shown in black at the centre of the torus, overlapping the central column which extends vertically. In addition, cylindrical coordinates denote the toroidal angle ϕ and radius R. An elevation view is shown in the middle subplot, where the central column is now shown on the left and two plasmas of different shapes are shown to the right. Another cylindrical coordinate system is adopted, with radius r and poloidal angle θ. The aspect ratio, R_{maj}/a, the ratio of the major and minor radii of the torus, is used to express the geometrical difference between these plasma shapes. Plasmas with conventional aspect ratios (~ 3) are shaped like typical tori whereas low aspect-ratio plasmas (~ 1) are more spherical and shaped like cored apples. Finally, the rightmost subplot shows another elevation view looking at a toroidal section of the plasma from outside the device. Now the electromagnetic coils are shown as well as the magnetic field components they produce. The toroidal field is typically strongest and is produced by the toroidal field coils (red). The poloidal component of the field is produced primarily by a current in the plasma (green dash), whilst some contribution is made by poloidal field coils (blue). These poloidal coils also help to shape the plasma. The plasma current is mostly generated by ramping a current in the central column, which is arranged as a solenoid to act as a transformer. The resulting electromotive force drives the plasma current, generating the poloidal field component. Tokamaks which require a central solenoid for the creation of a poloidal field component are therefore inherently pulsed devices, as they must constantly ramp up the current in the solenoid during a discharge. The impact this has on efficiency and capacity is why steady state plasma scenarios which use non-inductive current drive are attractive.

The resulting motion of ions in this system is plotted in figure 2.2 for two common types of orbits, whose trajectories along the helical field lines are perturbed due to the magnetic field strength gradient and curvature. The left plot illustrates the orbit of a passing particle. Its pitch is sufficiently high so that it is not reflected at any point in space by the mirror force (equation 2.11), meaning its guiding-centre motion is mostly helical around the torus. However, a radial shift is created by the addition of vertical drift velocity components, given by equations 2.13 and 2.14 above, generated by the radial curvature and gradients in the strength of the magnetic field. Particles with a lower pitch are reflected at radial bounce points due to the $1/R$ dependency of the magnetic field strength, leading to trapped "banana" orbits. Whether the banana orbits open inward or outward is dictated by the direction of the

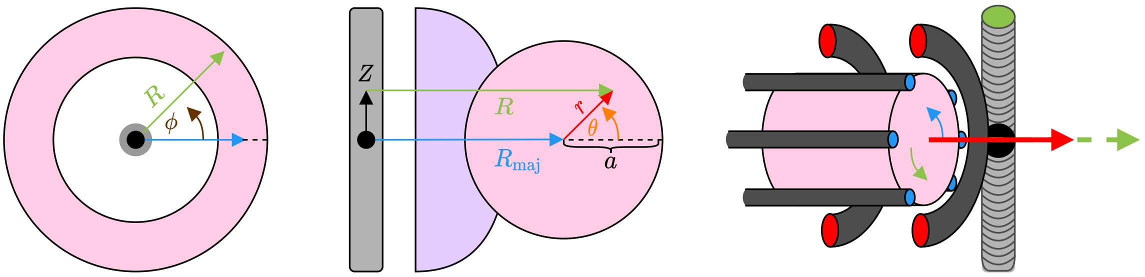

Figure 2.1: Basic coordinate system used to describe the tokamak geometry. Throughout, light grey denotes the central column, pink shows the plasma, dark grey denotes field coils and the black spot denotes the system origin. The left plot shows a plan view of the system, with the cylindrical coordinates R and ϕ describing the radius and toroidal angle respectively. The green and blue arrows, drawn from the centre column, correspond with those shown in the centre plot (not to scale). The centre plot shows the elevation view and the poloidal plane. Two examples for different aspect ratios are shown, with that typical for spherical tokamaks shown beneath and closer to the central column than the conventional, more circular aspect ratio upon which the coordinate system is drawn. The right plot shows a view of a poloidal slice, from outside the tokamak, to show the location of the toroidal (red ends) and poloidal (blue ends) field coils. These coils generate the components of the magnetic field shown by the corresponding coloured arrows. Ramping current in the central column (green top) drives the plasma current toroidally (dashed), which creates the corresponding poloidal magnetic field component (solid).

poloidal component of the magnetic field. For this reason, pitch angle is often defined relative to the plasma current instead. The v^2 dependence of the drift component enables high-energy ions to undergo much larger drift excursions. Where ions are close to topological boundaries, such as the trapped-passing boundary, small changes to the ion pitch can lead to significant transport. Banana - or neoclassical - orbits, which result from toroidal geometry, modify the classical ion transport mechanisms expected for homogeneous magnetic geometries. For example, the prompt loss of trapped ions after one bounce can drive torques much larger than what is possible with only passing ions [19]. Whilst passing and trapped are the two most general and common types of orbit, it is worth mentioning many others are possible too - for example potato, stagnation, pinch, super banana etc., descriptions of which can be found in [20].

The confining magnetic field is calculated by asserting a balance between plasma thermal pressure, defined by the ion momenta, and magnetic pressure, defined by the field line density:

$$\mathbf{J} \times \mathbf{B} = \nabla P \tag{2.16}$$

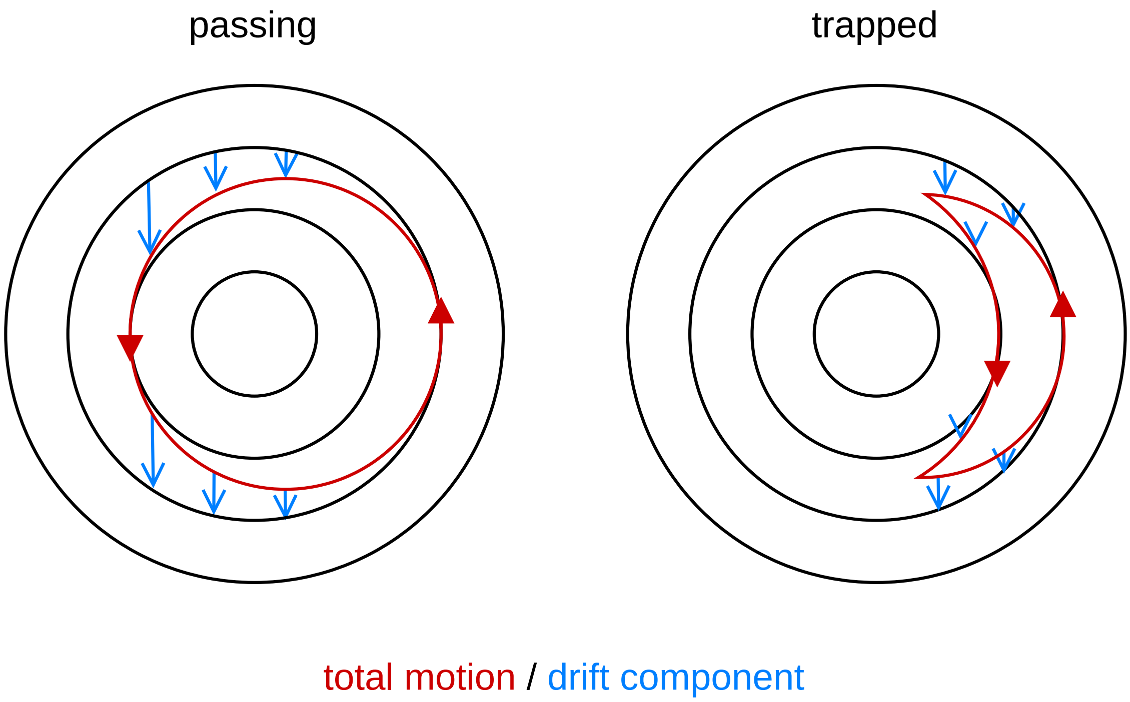

Figure 2.2: Passing (left) and trapped (right) guiding-centre orbits shown in the poloidal plane. The guiding-centre trajectory is shown in red whilst the direction of the magnetic drift component is shown in blue. In this case, the orbits shown correspond to electrons.

Upon inspection, one may notice that

$$\mathbf{B} \cdot \nabla P = \mathbf{J} \cdot \nabla P = 0. \tag{2.17}$$

In other words, the magnetic field lines and plasma current must lie on surfaces of constant pressure. These flux surfaces are analogous to fluid stream function contours. If we wish to express the magnetic field according to some variable ψ, which represents the stream potential of these flux surfaces, then ψ must satisfy

$$\mathbf{B} \cdot \nabla \psi = 0. \tag{2.18}$$

In addition, **B** must still satisfy the divergence-free condition which, in cylindrical tokamak geometry, becomes:

$$\nabla \cdot \mathbf{B} = \frac{1}{R}\frac{\partial}{\partial R}\left(RB_R\right) + \frac{1}{R}\frac{\partial B_\phi}{\partial \phi} + \frac{\partial B_z}{\partial z} = 0. \tag{2.19}$$

For a system that possesses axisymmetry about ϕ, this condition is true if

$$B_R = \frac{1}{R}\frac{\partial\psi}{\partial z} \tag{2.20}$$

and

$$B_z = -\frac{1}{R}\frac{\partial\psi}{\partial R}. \tag{2.21}$$

In this case, equation 2.18 is satisfied, implying ψ indeed defines a series of nested flux surfaces, upon which magnetic field lines flow as streamlines. Repeating this approach, and applying Ampère's law, to define a similar stream potential f for the plasma current density yields

$$\mathbf{B} = \mathbf{B}_\theta + \mathbf{B}_\phi = \nabla\psi \times \nabla\phi + \frac{\mu_0 f}{R}\nabla\phi \tag{2.22}$$

where $f = f(\psi)$ [18]. Upon these flux surfaces, quantities such as plasma density, pressure and temperature are constant, providing the plasma flow is small. In an axisymmetric system, these parameters are therefore one-dimensional. Another of these quantities, the safety factor q, which is defined as

$$q = \frac{1}{2\pi}\oint \frac{1}{R}\frac{B_\phi}{B_\theta}ds \approx \frac{r}{R_{\mathrm{maj}}}\frac{B_\phi}{B_\theta} \tag{2.23}$$

for a large-aspect ratio tokamak, acts as a measure of the field line pitch on a given flux surface - i.e. the number of toroidal transits per poloidal transit. Surfaces at which q is a rational number host field lines which form closed loops.

Figure 2.3 shows an example of a tokamak equilibrium calculated for ITER. Open flux surfaces exist outside of the nested flux surfaces in diverted equilibria such as this one, defined by a magnetic null in the form of an X-point to isolate the core plasma from the area of contact with the first wall. The separatrix marking the boundary between these two topologies of open and closed surfaces is defined by the last-closed flux surface (LCFS). The region containing these open flux surfaces is called the scrape-off-layer (SOL), as particles may freely stream away from the plasma and into the divertor region. The target PFCs surrounding and limiting the inner and outer divertor legs, the region of the separatrix which connects the X-point to PFCs, must therefore be designed to withstand extremely high power loads.

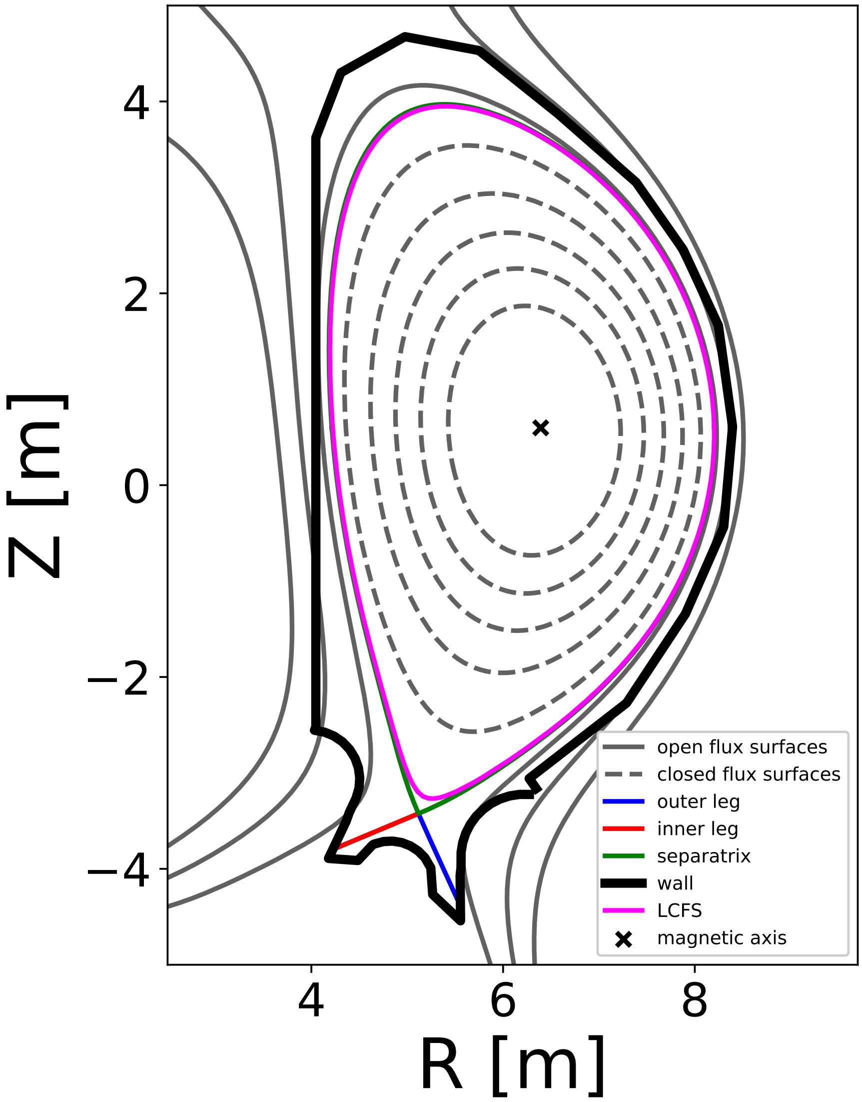

Figure 2.3: Poloidal section showing an axisymmetric tokamak equilibrium. Nested poloidal flux surfaces (dashed) define the isocontours of certain plasma quantities, such as temperature and density. The LCFS defines the boundary between the plasma and SOL, within which heat and particles are rapidly transported along field lines within open flux surfaces towards the divertor section, down the inner and outer legs towards the divertor targets and baffles - the region shown here below ∼3m. Typically, the separatrix is a general term used to describe the boundary between topologies - in this case between the open and closed flux surfaces. Hence it is similar to the LCFS. Any differences shown here unintended and the result of pre-processing the numerical equilibrium. However the difference illustrates the relative accuracy with which the LCFS is calculated by typical fixed-boundary equilibrium codes, compared to the more accurate solution denoted here by the separatrix.

2.1.2 Heating mechanisms

Fusion reactors require a way of heating the plasma to set T in the fusion triple product. With perfect energy confinement (infinite confinement time without any radiated losses), ignition could be achieved with alpha heating alone after an initial injection of energy to spark the reaction. In reality, external sources of power are required to continually replace lost power to sustain a steady-state reaction.

In tokamaks, power is often transferred to the plasma by creating a population of fast, supra-thermal ions, which transfer their energy to the bulk, thermal plasma through the Lorentz force. Ohmic heating, generated by the current propagating through the plasma with finite resistivity, is only efficient until relatively low temperatures ($\sim$3keV) [18], above which the plasma resistivity falls. Therefore, other heating mechanisms are required for efficiently reaching 10keV and higher. Some examples include heating by microwaves or advanced operational procedures such as merging compression.

Another method for creating a fast-ion population is through the injection of a beam of high-energy neutrals - a heating neutral beam (HNB). These particles are unaffected by the electromagnetic fields, and so can freely penetrate the plasma before becoming ionised through charge exchange or short-range Coulomb interactions with ions or electrons, whereby they become confined by the magnetic field. Due to their large size, these neutral beam injectors (NBI) are typically located outside of the tokamak and thus inject into the plasma through outboard entry ports. In short, the beams accelerate a source of ions across a potential difference before neutralising them prior to injection [21]. Upon entry, the beam attenuation is exponential-like, though typically faster due to the increasing plasma density, making it challenging to penetrate high-density plasmas such as those in ITER. An example of the ion deposition from the NBI system in ITER is shown in figure 2.4, for two possible configurations of the beam orientation with respect to the plasma. As the beam propagates from right to left, the ion deposition is rapidly reduced due to the beam attenuation as well as the increasing plasma density. The reduction in attenuation near $R < 5.5$m is due to the remaining ions reaching the other, inboard side of the torus, where the plasma density drops. Some neutrals may pass through the plasma completely. These are called shine-through losses. On the right panel of figure 2.4, the differences between the two beam configurations can be seen more clearly, corresponding to a vertical tilting of the beam. As each beam enters the plasma at a different angle, the resulting pitch of the injected neutrals

- and deposited ions - changes, altering the orbit shape of the deposited ions. This difference in orbit shape, and the resulting effects on the plasma, for example the resulting $\mathbf{J} \times \mathbf{B}$ torque or driven current, is why the flexibility to perform off-axis injection is usually included in the design of many HNB systems - including ITER. The Gaussian-like spreading of the beam in both cases can also be seen.

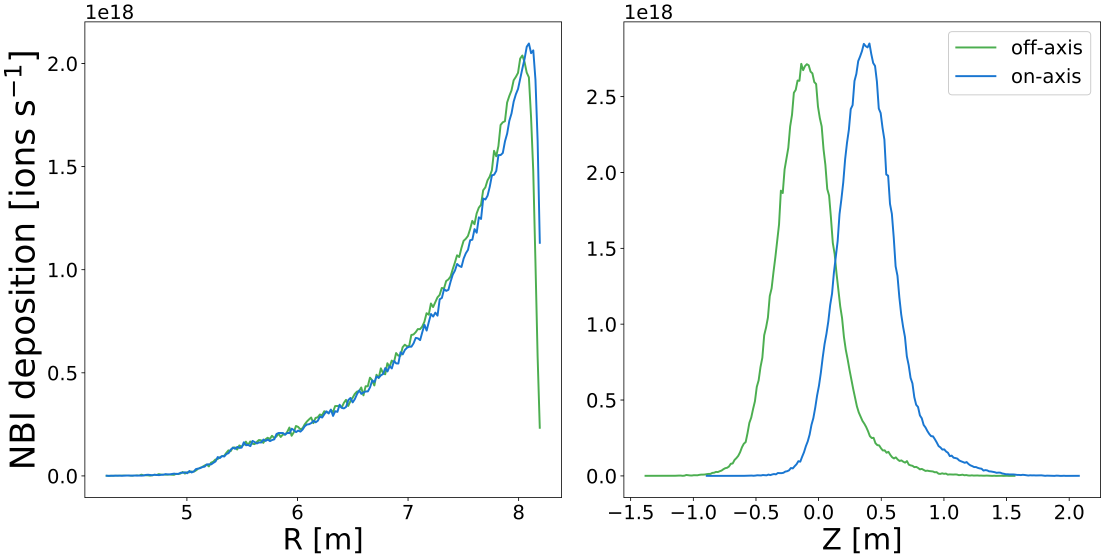

Figure 2.4: Examples of on and off-axis neutral beam deposition of fast ions created by the ITER HNB system. The rapid fall-off of ions as they travel into the plasma can be observed on the left, with an extremely rapid increase in deposition as soon as the beam enters the LCFS. On the right, the difference between on and off-axis deposition can be seen more clearly. The shift in height Z also corresponds to a change in the starting pitch, and therefore resulting orbit shape.

In ITER, ions must be injected by HNBs at energies higher than 300keV to penetrate the edge density barrier [21]. However, above 120keV, the efficiency of the ion neutralisation in NBIs drops considerably for positive ions. Therefore, ITER will use negative ions that are accelerated to up to 1MeV. Importantly, negative ion beams are accelerated to a single possible energy - meaning every neutral entering the plasma in ITER has the same kinetic energy. On the other hand, the possible states of molecular deuterium, and their differences in mass, lead to positive ion NBIs injecting fractions of neutrals at full, half and third energies.

The ITER HNB system itself is comprised of two separate NBI injectors, HNB1 and HNB2, each injecting an equal share of the total ≈33.3MW beam heating power. In future, a third heating beam, HNB3, may be installed, however this is ignored in this work. In addition, ITER will also be equipped with another, non-heating neutral beam system - the diagnostic neutral beam (DNB). This NBI works similarly to HNBs but at much lower injection energy and power to instead diagnose the edge plasma without significantly perturbing it.

2.2 Fast-ion transport

Fast ions must remain confined if they are to transfer their kinetic energy to the bulk plasma through Coulomb interactions. Whilst some ions would already be promptly lost if they are deposited on loss trajectories, there are mechanisms that can transport ions to loss orbits which would otherwise remain confined. The dynamics of particles undergoing this transport can be described at the fundamental level by modelling the changes in the particle distribution function f, defined such that

$$n = \int f(\mathbf{x}, \mathbf{v}, t) dv_x dv_y dv_z. \tag{2.24}$$

In a given phase-space volume moving with the plasma in phase-space, changes in the distribution function can only be caused by particle sources or sinks, or redistribution arising from interactions between the particles in that volume, $C(f)$:

$$\frac{df}{dt} = C(f) + \text{sources}. \tag{2.25}$$

Since $f = f(\mathbf{x}, \mathbf{v}, t)$, we can obtain the kinetic equation, describing the changes in the distribution function at a given point in phase-space and time:

$$\frac{\partial f}{\partial t} + (\mathbf{v} \cdot \nabla) f + \frac{1}{m} (\mathbf{F} \cdot \nabla_v) f = C(f) + \text{sources}. \tag{2.26}$$

Which states that the change in the distribution function with respect to time at a fixed point in phase-space is caused by convection in real space and velocity space as well as sources, sinks and redistribution due to internal interactions. Convection in real space arises from the inertia of the particles whilst convection in velocity space arises from external forces acting on the plasma.

Plasma transport is ultimately driven by electromagnetic fields, however it can be categorised according to the length and time scales over which it occurs. The long-range yet reciprocal-like nature of the cumulative electric field from each ion is often separated into smooth, macroscopic and transient, microscopic components respectively. The large difference in length scales means that physics models often choose to capture either the evolution of the macroscopic field, and its impact on the bulk plasma flow, or alternatively model the microscopic dynamics and short-range "collisions" between individual ions. These two theoretical formulations encompass two common families of physics models: magnetohydrodynamic (MHD) fluids and collisional kinetic particles respectively. The first of these is typically used to calculate the structure and stability of tokamak equilibria whereas the second is used to calculate the dynamics of test ions.

The effects of microscopic collisions are an irremovable yet vital transport source in magnetic confinement fusion systems, as they are ultimately responsible for fusion events themselves. In combination with the fixed magnetic topology, their influence determines the classical and neoclassical transport regimes [22] - the latter being the modification of classical transport due to neoclassical orbits shown in figure 2.2. Historically however, the measured transport in tokamaks is typically much greater than that which is predicted by neoclassical theory alone.

Instead, models which more realistically capture the evolution of the macroscopic field components are increasingly able to explain this "anomalous" transport. Despite the plasma being quasi-neutral overall, local fluctuations in the electromagnetic field, created by the externally imposed magnetic field and internal plasma dynamics, can lead to waves, turbulence and instabilities. The spatiotemporal scales and magnitude of the transport arising from each of these phenomena can vary; in some cases it can initiate a series of events that lead to a disruption - a catastrophic loss of plasma confinement [18]. Therefore, it is vital that models include both collisional and anomalous effects if they are to accurately predict transport in devices such as ITER.

It is possible to capture both microscopic and macroscopic field components simultaneously. One way of achieving this is via reduced models, which rely on assumptions to avoid modelling the system consistently. For example, if MHD modes are assumed to be steady they may be treated as as collision-like "kicks" in kinetic models [23–26]. Alternatively, modelling the behaviour of these systems self-consistently is possible using hybrid approaches and integrated models. Hybrid models typically combine fluid and kinetic representations of the plasma, to model - for example - the drive/damping of plasma waves by the dynamics of individual ions [27, 28]. Integrated models are comprised of separate components that each simulate - potentially iteratively - individual aspects of the total system. For example, one component may track ion dynamics in the 3D plasma equilibrium calculated - and possibly iterated - by another component.

Modelling the N-body dynamics and interactions amongst the $\sim 10^{21}$ ions in a tokamak plasma is computationally infeasible; an exascale computer would take $\sim 10^{16}$ years to calculate the combined electrostatic field of this system at the position of each particle. MHD models solve equation 2.26 by taking high-order moments of the distribution function, where equation 2.24 is the zeroth order, and formulating equations which describe finite volumes of isotropic, Maxwellian, collisional plasma (though collisions themselves are not modelled). These assumptions do not necessarily hold for fast ions. Therefore, kinetic models are typically used when studying fast-ion dynamics. In this case, equation 2.26 is solved by tracking a group of markers that form a sample of the distribution function. ”Weights” are assigned to each marker to normalise the marker sample to regain the real injected particle density. However, weight can also be redistributed such that Monte Carlo techniques, such as splitting or roulette, may be used to increase fidelity around regions of interest. These markers can then be tracked as if they were real particles, and sources, sinks and collisions can be represented as random processes via Monte Carlo methods. This is the approach taken by the physics components used within this work. Anomalous transport processes are then accounted for by prescribing electromagnetic fields calculated by external physics models of the plasma, which are added to form an integrated model.

2.2.1 Collisional transport

Collisional transport is driven by short-range Coulomb interactions between charged particles in the plasma. They are responsible for transporting ions in velocity space and atomic processes such as charge exchange, ionisation and neutralisation. Whilst the fast-ion redistribution by collisions is typically weak and diffusive, significant transport may occur if they cause ions to cross topological boundaries - for example by knocking ions from passing to trapped orbits.

Collisions are typically modelled as discrete Markov processes in that they are assumed to be binary and independent. In this case, the velocity v of a particle may be considered a stochastic variable, which changes state according to a master equation:

$$\frac{dP(v,t)}{dt} = \sum_{v' \neq v} w_{vv'}(t)P(v',t) - \sum_{v' \neq v} w_{v'v}P(v,t). \tag{2.27}$$

This essentially states that the rate of change of probability P of a particle being in some state v is equal to the difference between the rate of transition, w, from state v' to v and from v to v'. These changes in v due to collisions also lead to a change in the particle distribution function f, which is given by the Kramers-Moyal expansion of the master equation,

$$C(f) = -\frac{\partial}{\partial v}\left(D^{(1)}f\right) + \frac{1}{2}\frac{\partial^2}{\partial v^2}\left(D^{(2)}f\right) + ..., \tag{2.28}$$

where D represents the "jump moments" - the probability v undergoes a given transition after a given time:

$$\int d(\Delta x)(\Delta x)^l f(x+\Delta x, t+\Delta t) \approx D^{(l)}(x,t)\Delta t + \mathcal{O}(\Delta t^2) \tag{2.29}$$

where D can be thought of as the transition size per unit time. In the test particle case, the analogy can be made between a transition and a collision event, in which the transition size now becomes the average change in velocity per unit time

$$D \approx \langle \Delta v \rangle . \tag{2.30}$$

By assuming small-angle collisions dominate, the v^3 terms of equation 2.28 can be truncated to yield the Fokker-Planck equation in Cartesian coordinates for a single particle species:

$$C(f) = -\frac{\partial}{\partial v^\mu}(f\langle \Delta v^\mu \rangle) + \frac{1}{2}\frac{\partial^2}{\partial v^\mu \partial v^\nu}(f\langle \Delta v^\mu \Delta v^\nu \rangle) \tag{2.31}$$

where the final term is a dyad. v^μ denotes the μth component of the particle velocity and the Δv^μ jump moment is the change in the μth component of the velocity resulting from a collision [29].

The average particle jump moments are then evaluated by calculating the expectation value of the change in the velocity components of a particle due to a $1/r^2$ force acting upon it [30], the differential cross-section of which is known analytically. For like particle collisions this yields:

$$\frac{1}{\Gamma}C(f) = -\frac{\partial}{\partial v^{\mu}}\left(f\frac{\partial h}{\partial v^{\mu}}\right) + \frac{1}{2}\frac{\partial^2}{\partial v^{\mu} v^{\nu}}\left(f\frac{\partial^2 g}{\partial v^{\mu}\partial v^{\nu}}\right) \tag{2.32}$$

where

$$h(\mathbf{v}) \sim \int f(\mathbf{v}\prime)\left|\mathbf{v}-\mathbf{v}\prime\right|^{-1} d\mathbf{v}\prime \tag{2.33}$$

and

$$g(\mathbf{v}) \sim \int f(\mathbf{v}\prime)\left|\mathbf{v}-\mathbf{v}\prime\right| d\mathbf{v}\prime \tag{2.34}$$

and Γ represents the collision rate. $\mathbf{v}$ denotes the velocity of the test particle with $\mathbf{v}\prime$ the velocity distribution of the species it is colliding with. Now it can be seen that the evolution of the distribution function can be thought to result from the presence of potentials g and h in velocity phase-space - so-called Rosenbluth potentials. In vector form, the Fokker-Planck equation can now be written more succinctly as:

$$\frac{\partial f}{\partial t} = -\frac{\partial}{\partial \mathbf{v}}\cdot\left[\underline{A}f - \frac{1}{2}\frac{\partial}{\partial \mathbf{v}}\cdot\left(\underline{D}f\right)\right] \tag{2.35}$$

where $\underline{A}$ is the dynamical friction and $\underline{D}$ is the velocity diffusion tensor. In general, there is one equation for each plasma species that is not the test particle species - the collisions between which are considered infrequent and ignored. To obtain the velocity space dynamics of a test particle, one may substitute solutions for f into the Fokker-Planck equation above [31,32]. This leads to the test particle Fokker-Planck equation in spherical coordinates:

$$\frac{\partial f}{\partial t} = \frac{D_{\perp}(v)}{v^2}\mathcal{L}f + \frac{1}{v^2}\frac{\partial}{\partial v}\left[v^2\left(Ff + \frac{1}{2}D_{\parallel}\frac{\partial f}{\partial v}\right)\right]. \tag{2.36}$$

In this expression,

$$F = \nu_s v + \frac{D_{\parallel} - D_{\perp}}{v} + \frac{1}{2}\frac{dD_{\parallel}}{dv} \tag{2.37}$$

and the operator

$$2\mathcal{L}f = \frac{\partial}{\partial\lambda}\left(1-\lambda^2\right)\frac{\partial f}{\partial\lambda} + \frac{1}{1-\lambda^2}\frac{\partial^2 f}{\partial\chi^2}. \tag{2.38}$$

Spherical coordinates are used to split the diffusion tensor into scalar perpendicular and parallel parts after assuming that the particle distribution is isotropic. This also simplifies the expression for $\underline{A}$ to $-\nu_s\mathbf{v}(v)$. Here, ν_s represents the slowing-down rate, λ measures the particle pitch angle and χ the gyro-phase. Therefore, the first term in equation 2.36

represents changes in the direction of the test particle velocity (the pitch) and the second term pertains to changes in the particle energy.

At this point, Monte Carlo operators which solve the test particle Fokker-Planck equation for pitch and energy can be derived by finding expressions for the change in their expectation values,

$$\frac{d\left\langle x\right\rangle}{dt} = \int xfdV, \tag{2.39}$$

where x is an arbitrary test particle quantity, such as pitch or energy, and V is the complete phase-space volume over which f is defined. It is typically assumed that fast ions are equally distributed in gyro-phase to remove the χ dependency in the operator $\mathcal{L}$ [33]. These operators are

$$\Delta\lambda = -\lambda\nu_{\mathrm{d}}\Delta t \pm \left[\left(1-\lambda^2\right)\nu_{\mathrm{d}}\Delta t\right]^{1/2} \tag{2.40}$$

and

$$\Delta\epsilon = -2\nu_\epsilon\Delta t\left[\epsilon - \left(\frac{3}{2} + \frac{\epsilon}{\nu_\epsilon}\frac{d\nu_\epsilon}{d\epsilon}\right)k_{\mathrm{B}}T_{\mathrm{th}}\right] \pm 2\left(k_{\mathrm{B}}T_{\mathrm{th}}\epsilon\nu_\epsilon\Delta t\right)^{1/2}, \tag{2.41}$$

where $\pm$ is a random sign with equal probability and the ν_ϵ and ν_{d} terms represent collision frequencies as functions of the Coulomb logarithm $\ln(\Lambda)$ and error function [29, 31, 34]. For a given fast-ion population, these expressions are evaluated for each bulk species.

From these expressions, the separate drag and scattering effects of collisions can be observed. In a typical tokamak plasma, there are three velocity scales:

$$v_{\mathrm{thermal}} << v_{\mathrm{fast}} << v_{\mathrm{electron}}. \tag{2.42}$$

Shortly after deposition, fast ions are typically bombarded by fast-moving electrons, which reduce the fast-ion energy without significantly altering their pitch angle. Only after sufficiently slowing down do these fast ions then begin to interact more with the thermal population. The point at which drag due to electrons and ions is equal is termed the critical energy. Below this energy, significant pitch-angle scattering takes place, and the rate of energy transfer to the ion population increases - thermalising the fast ions and heating the thermal ions. Therefore, fast ions must be confined for most of their slowing-down time, the time taken for them to thermalise, for them to significantly heat the thermal population directly through collisions. As it increases like $\sim T^{3/2}$ [18], the slowing-down time in ITER, where the plasma temperature and beam energy are high, leads to typical slowing-down times of the order ∼1s

- ten to a hundred times larger than conventional devices.

2.2.2 Anomalous transport

There are numerous potential sources and explanations for anomalous transport, but in general it arises due to high-order effects that are neglected from neoclassical transport models that assume static, axisymmetric equilibria. Two categories which describe the primary drivers of anomalous transport are: perturbations to the externally imposed magnetic geometry and the electromagnetic fields resulting from the plasma dynamical response to external or internal perturbations.

MHD predicts a wide range of plasma dynamics in response to various types of perturbations, with experiments, theory and simulations all having discovered various phenomena that mostly enhance fast-ion transport. These responses can be uncorrelated and turbulent, or they can appear in the form of stable or unstable modes of oscillation - in space and often time. As the wide fast-ion orbits traverse MHD structures, they typically receive kicks in phase space, the statistics of which determine the nature of the average transport - convective or diffusive [35]. Diffusive transport from randomly distributed kicks, arising from interactions with turbulence [26, 36] for example, is generally weaker than the convective transport resulting from correlated kicks. Convection typically occurs when fast-ion motion resonates with MHD structures or if there are extreme changes in the plasma topology, such as the introduction of magnetic islands or creation of open field lines. The general resonance condition for ions is

$$\Omega = l\Omega_c + m\Omega_\theta + n\Omega_\phi, \tag{2.43}$$

where l, m and n are integer numbers and Ω is the frequency of the perturbation. Here, the terms marked by the θ and ϕ subscripts represent the poloidal and toroidal transit frequencies respectively.

Many sources of convective fast-ion transport that have been observed in previous devices may be present in ITER, with potentially unknown modifications due to the transition to a burning plasma regime [13, 37–39]. Of potential concern are toroidal Alfvén eigenmodes (TAE), which are weakly-damped fast-ion-driven MHD modes arising due to toroidal geometry [40]. They are driven by nonlinear interactions between shear Alfvén waves and fast ions, and lead to fast-ion redistribution. The large super-Alfvénic population in ITER is potentially a large source of free energy for TAE drive [41]. However, simulations have predicted minimal transport of HNB ions in both the inductive baseline [42] and reverse shear plasma [43] ITER scenarios. Sawtooth oscillations are pressure or current-driven n=1, m=1 internal kink modes which relax the core ion temperature and density profiles in periodic crashes. The large sawtooth periods in ITER may trigger rotating neoclassical tearing modes (NTM) [44]. Crashes themselves greatly redistribute core fast-ions [45, 46] but may not lead to significant losses [47]. This is especially true if sawteeth can be stabilised by energetic ions [48]. Edge-localised modes (ELM) are explosive instabilities driven by steep edge pressure gradients. They expel significant amounts of energy and particles [49] and have the potential to quench a reaction, rapidly depositing significant amounts of energy onto PFCs. ELM-driven fast-ion transport has been studied relatively little, however it is hoped that ELMs will be mitigated by the ITER ECC system.

In addition, there also exists a multitude of perturbations to the imposed magnetic geometry in ITER. The finite number of toroidal field coils creates a toroidally periodic variation in the toroidal field strength, enhancing diffusion of resonant particles, either due to trapping or chaotic orbit motion [50]. This has been observed both in JET [51] and TFTR [52], with numerical studies accurately resembling experimental observations, so long as finite Larmor radius (FLR) effects are included [53] - that is, the difference between the position of the fast-ion and its guiding-centre is accounted for. In ITER, ripple transport [54] could be largely negated by ferritic inserts [55]. Ferritic inserts are first-wall components designed to smooth unwanted perturbations in the equilibrium field. However, they introduce another symmetry-breaking component to the magnetic field in ITER. The inserts are optimised to operate at field strengths above 5T; operating below 2T at low power may lead to increased fast-ion transport during plasma start-up or low-power scenarios. Another in-vessel component containing ferromagnetic elements that may perturb the externally imposed magnetic field are the test tritium breeding blanket modules (TBM). However their impact on fast-ion losses is likely to be small [56].

Compared with the aforementioned examples however, one of the more prominent sources of anomalous fast-ion transport in ITER is thought to be the redistribution and subsequent loss of HNB ions due to the 3D fields created by the ECC system [39]. Unlike the other sources of transport, which may be transient or mitigated in various ways, transport due to ECCs will always be present in high-confinement ITER plasmas, which require consistent

ELM suppression. The ITER Research Plan [13] therefore prioritises studies which aim to predict this transport to develop methods for mitigation.

2.3 ELM-control coils

To achieve high-power discharges, the ITER tokamak aims to operate in high-confinement mode (H-mode) [57], as opposed to low-confinement mode (L-mode). Whilst the edge transport barrier created by H-mode enhances the plasma confinement, the steep edge pressure gradients it creates leads to ELM crashes through the driving of the peeling-ballooning MHD mode. Such is the risk to confinement posed by ELMs that the inclusion of an ELM-control system is considered a necessity to the ITER baseline scenario [58].

The ITER ELM-control system is a series of window-frame coils arranged on the outboard, outward-facing side of the vacuum chamber in three rows of nine coils each. Figure 2.5 shows this geometry from multiple angles. The coils are regularly spaced in the toroidal direction, with centres of the first coils in each row located at $\phi = 30°, 26.7°$ and $30°$ for the upper, middle and lower rows respectively. By the time that the HNB system is commissioned on ITER, each ECC coil will be independently powered. Each coil consists of six windings, carrying a maximum of 15kA each (90kAt total magnetomotive force) [59].

The functional purpose of the ECC system is to impose magnetic perturbations that produce a resonant response in the plasma. To produce resonant magnetic perturbations (RMP) with toroidal mode number n_0, the current in each ECC in a given row is defined by

$$I\left(\phi_{\text{coil}}\right) = I_c \cos\left(n_0\left[\phi_{\text{coil}} + \Phi\right] - \omega t\right), \tag{2.44}$$

where ϕ_{coil} is the toroidal location of the coil centre, ω ($\sim$ 5Hz [58]) is a toroidal rotational frequency and Φ defines a toroidal phase shift - typically applied to each coil row independently to alter the poloidal spectrum of the 3D magnetic field. From an operational point of view, I_c, n_0, $\Phi_{\text{u,m,l}}$ (referring to upper, middle and lower coil rows respectively) and ω are the controllable degrees of freedom, subject to the operational limitations and the required level of ELM suppression - which can be changed depending on the values of these parameters. It is worth noting that, as well as the fundamental mode, n_0, the toroidal spectrum, $I(n)$, of an RMP imposed by coils of finite width and number may include significant contributions from additional, higher-frequency harmonics - n_1, n_2... etc. This can be seen by evaluating

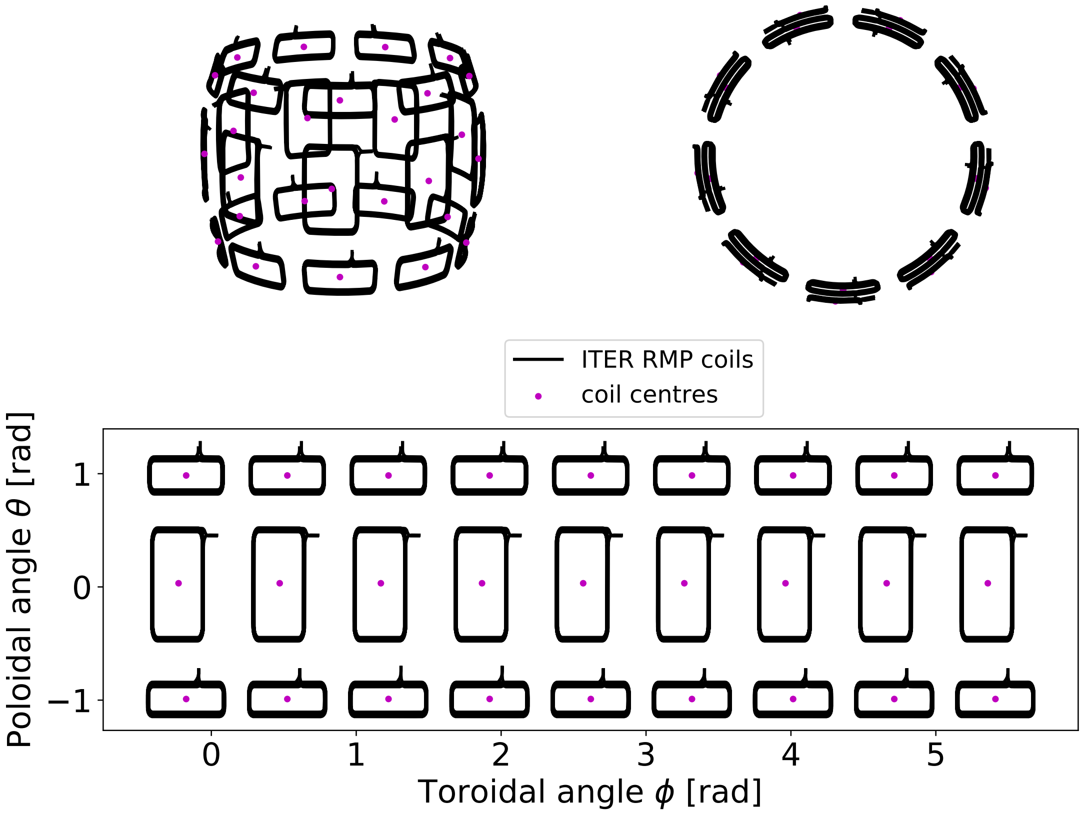

Figure 2.5: Top left: View of the the ITER ECC system (not to scale). The three rows of rectangular coils are responsible for imposing 3D perturbations onto the plasma by each carrying a current defined by equation 2.44. Top right: View looking on the ECC system from above. Coil spacing is toroidally symmetric within a given row, but offsets exist between rows. Importantly, all coils are positioned on the outboard side of the plasma, meaning that the power deposited by heating systems must be transmitted through the 3D field. Bottom: Distribution of ECCs in poloidal (θ) and toroidal (ϕ) angle. Coil centres are marked in magenta in all subplots. The finite poloidal and toroidal width of each coil is responsible for the presence of significant higher harmonics in the spectrum of the imposed field. Figure reproduced from [60].

$I(n)$ for a single row of window frame coils:

$$I(n) = \frac{1}{\pi}\int_{-\infty}^{\infty} I(\phi)e^{-i\phi n}d\phi. \tag{2.45}$$

If the current is constant across the width of each coil $\Delta\phi$, then equation 2.45 becomes

$$I(n) = \sum_{j=1}^{N}\frac{1}{\pi}\int_{\phi_j-\frac{\Delta\phi_j}{2}}^{\phi_j+\frac{\Delta\phi_j}{2}} I_j e^{-i\phi n}d\phi, \tag{2.46}$$

where I_j for each coil j is given by

$$I_j = I_c\cos\left(n_0\left[\phi_j+\Phi\right]\right), \tag{2.47}$$

as defined by equation 2.44 for N total coils in a coil row c. After realising $\Delta\phi_j = \Delta\phi$ is constant across a single row of ITER coils, expression 2.46 becomes

$$I(n) = \frac{I_c}{\pi n}\sin\left(\frac{n\Delta\phi}{2}\right)\sum_{j=1}^{N}\left[e^{in_0(\phi_j+\Phi)-in\phi_j} - e^{-in_0(\phi_j+\Phi)-in\phi_j}\right]. \tag{2.48}$$

Then, upon recognising that

$$\phi_j = \phi_0 + \frac{2\pi\left(j-1\right)}{N}, \tag{2.49}$$

expressions can be found for modes satisfying $n \pm n_0 = kN$, where $k \in \mathbb{Z}$. For the fundamental mode, where $n_0 = n = 3, 4$ is the lowest frequency in the harmonic spectrum, the spectrum takes the form

$$I(n) = I_c C_{c,n} e^{-in_0\Phi}, \tag{2.50}$$

with $C_{c,n}$ capturing the mode amplitude. For the second harmonics, where $n_0 = 3, n = 6$ or $n_0 = 4, n = 5$,

$$I(n) = I_c C_{c,n} e^{in_0\Phi + 9i\phi_0}. \tag{2.51}$$

It can be seen that expression 2.51 introduces a dependency on ϕ_0 - the 3D geometry of the system - as opposed to simply the phase shift Φ between coil rows. Furthermore, the harmonics can be seen to counter rotate, in opposite directions, leading to variations in the poloidal spectrum throughout rotation - potentially altering the fast-ion dynamics and resulting PFC footprint. Typical spectra for ITER coils are illustrated in Figure 2.6. The relative significance of the so-called sidebands (second harmonic) can be seen by calculating the ratio of amplitudes $C_{c,n}$ of the first and second harmonics, yielding approximately 70-85% for $n = 3, 6$ and 88-95% for $n = 4, 5$ in ITER, depending on the coil geometry (for example $\Delta\phi$, which can be seen to differ between rows in figure 2.5).

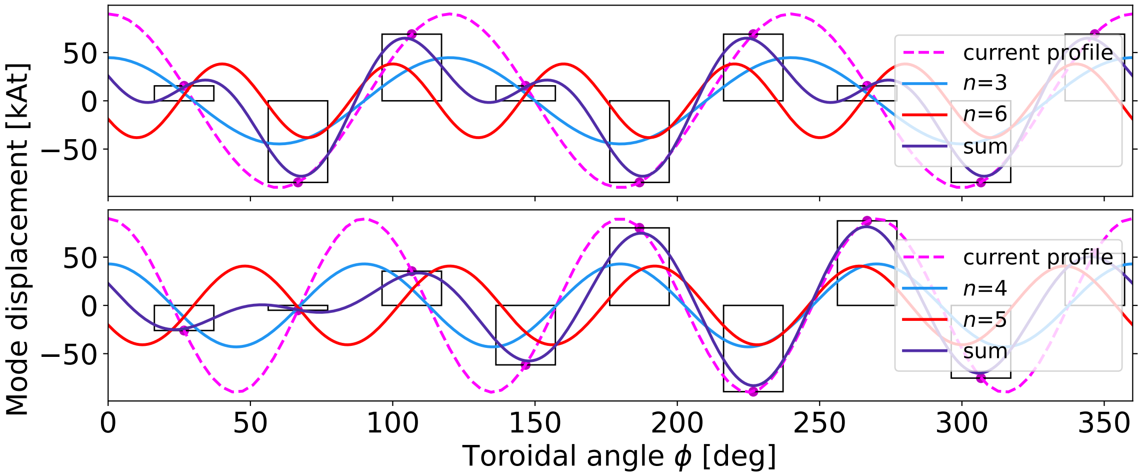

Figure 2.6: Current distribution (magenta) defined by equation 2.44 for $n_0 = 3$ and $n_0 = 4$ RMPs in ITER ECCs (black quadrangles where height represents current carried), as well as the first two corresponding harmonics (blue, red) and their sum (purple), for $\Phi = 0$. The pink scatter points serve as visual guide to confirm that the ECC centres adhere to the dashed current profile. In both cases, the finite number of coils means the higher harmonic significantly contributes to the total waveform. Figure reproduced from [60].

The 3D field created by the ECC system is complex [61, 62] but can be split into two components: the field created by the coils as if in a vacuum with no plasma present and the field generated by the plasma in response to this vacuum field. In similar fashion to convective fast-ion transport, the effect of the vacuum field is to perturb those trajectories of field lines which have orbits that resonate with the imposed perturbation. For field lines, which do not undergo cyclotron motion, the resonance condition in equation 2.43 for static perturbations becomes

$$0 = m\Omega_\theta + n\Omega_\phi \rightarrow \Omega_\theta/\Omega_\phi = m/n. \tag{2.52}$$

Since $\Omega_\theta/\Omega_\phi$ is just the safety factor q defined in equation 2.23, this resonance condition is satisfied by field lines on flux surfaces defined by rational values of $q(\psi) = m/n_0$. The perturbation creates a chain of m magnetic islands on the rational surface, where the island width grows like the square root of the perturbation amplitude [35]. Hence the q-profile plays an important role in determining the structure of the field created by the ECC system. A Hamiltonian description of these field line dynamics reveals an analogy with the Kolmogorov-Arnold-Moser (KAM) theorem of invariant tori. When the perturbation amplitude is increased, the islands in the chain may begin to overlap - dubbed the Chirikov criterion. Whilst flux surfaces within the islands persist, those which lie close to the island separatrices are destroyed, upon which field line motion becomes chaotic. The regions of high q shear, where there is a high concentration of rational surfaces, are sometimes completely destroyed by this island overlap - that is, until a break in the island chain occurs. This "stochastic layer" is therefore created just within the plasma separatrix where the q shear is high and the ECCs are in proximity. Additional topology changes occur near the X-point, where a homoclinic tangle of intersecting, oscillating flux surfaces is created by the unstable dynamics of individual field lines, leading to strike-point splitting [63] - ultimately reducing confinement near the X-point and modifying the divertor power load footprints. Hence the plasma displacement amplitude at the X-point (XPD) is often treated as a metric for ELM suppression [64, 65], as it is in this study. Here, the plasma displacement **zeta** is typically defined by the magnetohydrodynamic fluid velocity:

$$\mathbf{v} = \partial\zeta/\partial t. \tag{2.53}$$

However, in the studies which establish a correlation between XPD and ELM suppression mentioned above, this velocity is modified by toroidal flow shear to account for the change in bulk flow velocity the fluid parcel experiences whilst moving along **zeta** [66, 67].

The mechanisms through which ECCs suppress ELMs are also complex, however they may be explained by the vacuum description of RMPs provided thus far. As well as direct coupling between the peeling-ballooning mode and RMPs [68], ELM suppression is thought to result from the formation of magnetic islands near the pedestal, the radial positions of which are controlled by the q profile, creating q_{95} "windows" of observed ELM suppression [69–72]. Because ions may freely stream along the closed flux surfaces within an island, convective radial ion transport occurs over the island width, relaxing the pressure profile across the island and removing the mode drive. This purposeful decrease in edge confinement of thermal ions is sometimes termed "density pump-out", and it acts as a simple indicator of how fast-ions may be similarly transported by ECC fields.

Lastly, the second component of the ECC field is created by the plasma response to external perturbations. This response is a natural consequence of Lenz's law and the conductivity of the plasma. Similar to a metal, introducing magnetic perturbations induces flows of mobile charge carriers in the plasma, often electrons due to their relatively low inertia, which correspond to the generation of currents that screen out the external field. The length and timescales of this shielding are set by the charge carrier temperature and plasma resistivity. A significant factor to influence RMP penetration and subsequent dynamics, like rotational braking [73], is the plasma rotation. If the plasma resistivity is sufficiently low, MHD models [67, 74, 75] predict:

$$\frac{\delta B_{\mathbf{tot}}}{\delta B_{\mathbf{vac}}} \sim -\omega_{\phi}^{\frac{4}{5}} \tag{2.54}$$

where the δB terms denote the amplitude of the radial components of the perturbation, in the vacuum approximation and including the total plasma response, and ω_{ϕ} denotes the bulk plasma rotation. For ITER therefore, high temperatures act to reduce the plasma resistance, making it efficient at screening out external perturbations, whilst low rotation does the opposite [76]. The effect of this screening repairs the flux surfaces otherwise broken by a pure vacuum field. In reality, including this plasma response [77] can be significantly more realistic than the vacuum modelling approach alone; it has been observed experimentally in MAST that ECCs sometimes have no effect on ELMs [78], implying that, despite the plasma satisfying conditions for island overlap, screening is effective enough to significantly prevent the development of stochastic regions. On the other hand, plasma responses can sometimes act to amplify the perturbation, enhancing transport [79].

2.4 Fast ions in RMP fields

The effects of RMPs on the plasma are complex, and their resulting impact on fast ions in realistic, collisional plasmas with non-circular cross-sections is also complex - especially after the plasma response is taken into account [80]. Nevertheless, there have been many clear experimental observations of ECC-driven fast-ions losses [49, 81–83]. At the highest level, this is due to the coincidence of the edge-localised RMP and the intense, anisotropic source of fast ions deposited by the HNB system [56, 84–86]. For this reason, modelling the impact of RMPs on HNB ions remains a priority for the ITER Research Plan.

The primary features of the RMP field which characterise its effects on fast-ion transport are that it is a nearly-static, magnetic perturbation with both stochastic and island structures. The introduction of rapidly varying, axisymmetry-breaking magnetic field components alters the fast-ion constants-of-motion and adiabatic invariants. As the guiding-centre equations of motion rely upon the conservation of these quantities, tracking the full orbits of fast ions becomes increasingly necessary in such fields. One such otherwise conserved quantity is the canonical angular momentum P_ϕ, given by

$$P_\phi = mv_\phi R + q\psi, \tag{2.55}$$

which is derived as a consequence of toroidal axisymmetry - now broken by the RMP. Therefore, fast ions are now free to undergo significant radial excursions corresponding to large variations in P_ϕ [86, 87]. This effect is thought to be most pronounced at the plasma edge, where the perturbation amplitude is strongest. This edge layer can lead to the inward or outward transport of fast ions, depending on the perturbation phase, and has large implications for a highly localised ion source such as the HNB.

Due to the creation of both stochastic and island regions, there is potential for both convective and diffusive fast-ion transport [88]. For example, the RMP introduces island structures to the fast-ion drift orbits, which widen as the ion canonical angular momentum now varies throughout the orbit [89]. This leads to rapid convection of ions over large poloidal excursions, the size of which is dictated by the island widths. Likewise, the overlap of magnetic and drift-orbit islands may introduce stochastic components to the ion drift orbit [89–95]. Fast ions which more closely follow field lines, for example on strongly passing orbits, may undergo enhanced diffusive transport as they follow field lines which diffuse out of the plasma through the stochastic boundary layer [96]. Collisional transport may also become enhanced, as more topological boundaries are introduced by the RMP.

As to whether screening increases confinement, the plasma response has been observed to both increase [62, 97–99] and decrease [100, 101] fast-ion losses. This makes it highly important to include the response in simulations [89, 102], though the choice of response model has been shown to largely affect results [85]. Due to the lack of a general reduced transport model, ad-hoc simulations of each case are therefore still necessary.

Whilst it can be seen that ECC-driven transport has been studied both computationally and experimentally, and a relatively developed understanding of the underlying transport mechanisms has been formulated and applied to ITER, realistic simulations of ITER transport and the development of robust mitigation strategies have remained challenging. This is mostly due to the computational demands of high-precision ITER simulations. Specifically, a large number of these studies abandon the inclusion, or make no mention, of FLR effects, converged high-precision RMP fields, high-frequency perturbation harmonics or collisional models - despite sometimes acknowledging their importance [103].

For example, few comprehensive analyses of the fast-ion power loads in ITER have been performed. However, experimental observations [63] that ECCs concentrate lost ions in localised, structured footprints make this a priority. This is especially true considering the anisotropy and high power of the ITER HNB ion source. Indeed, the only study [104] which has examined power loads at high resolution suggests some components are at risk from fast-ion losses and that mitigation strategies should be pursued.

Until now, the optimal methods for mitigating power loads whilst maintaining ELM control remain to be determined, again due to the computational demands of expanding the ECC operational space. And in the trade-off between studying ITER discharges that are early and low-power versus high-power and high-priority, the latter has understandably garnered the most attention. However, this detracts from the development of mitigation strategies that may be experimentally verified early on in the ITER campaign, prior to deployment in high-power scenarios where the margin for error is smaller.

This lack of computational flexibility has also meant priority has been placed on identifying the relative prevalence of particular transport mechanisms [85, 86, 105] - at times at the expense of realism. Whilst this approach is sensible in these circumstances, where an understanding of the transport mechanisms may provide information across ITER scenarios that are impossible to study as a whole, it is not a substitute for directly evaluating the real transport.

This study therefore chooses not to attempt to formally identify the dominant transport mechanisms, but instead opts for the highest level of realistic detail. The potential for high-fidelity simulations of fast-ion transport in ITER has only recently become feasible, and this is the first work to routinely apply and exploit that at scale.

Chapter 3

Computational model

The high cost and fusion power of the ITER tokamak demands pre-emptive study for risk mitigation, however the machine's scale and novel plasma physics make this challenging to accomplish using only conventional, experimental devices. Instead, methods for operating ITER safely and optimally must be determined predictively *in silico.* Today, integrated modelling is one of the primary methods of accurately studying tokamak plasmas. These software platforms aim to perform more accurate simulations by coupling multiple, specialised physics codes to create models that are more comprehensive than each of their constituent parts. Some examples of successful integrated modelling platforms are `OMFIT` [106] and `JINTRAC` [107]. Meanwhile, the new Integrated Modelling and Analysis Suite (`IMAS`), designed at ITER, aims to extend this paradigm further. However, the extreme spatiotemporal domain associated with ITER's unique scale and physics makes routine, high-precision simulations computationally challenging for integrated models built from traditional physics codes - typically designed around studying conventional, smaller-scale devices. A way of solving this problem is to design physics codes that can exploit novel computational hardware - for example, general-purpose graphics processing units (GPGPU). `LOCUST` is one such code which uses GPGPU acceleration to make ITER-scale simulations computationally tractable. Its greater speed and efficiency over other fast-ion codes today offers the potential for the discovery of new physics through study of tokamaks at extremely high fidelity. For this reason, the `LOCUST` code was chosen as the primary component of the fast-ion model assembled and used in this work.

This chapter aims to introduce and describe the computational tools used within this study and the motivations behind the selection of each component - including the need for the hardware accelerated `LOCUST` code. GPGPUs are described in section 3.1, along with their advantages/disadvantages over traditional central processing units (CPU) and the resulting implications for physics code design. The `LOCUST` code is described in section 3.2. Finally, section 3.3 describes the additional computational components - `IMAS` [108], physics components and specific data - that are needed for the study of fast ions in 3D fields in ITER and coupled later in chapter 4.

Parts of this chapter have been published in reference [109].

3.1 GPGPUs

CPUs are a type of computational hardware designed to accomplish a wide variety of tasks. Their flexibility has led to most modern computing systems and software, including fusion simulation codes, being designed around their usage. Modern CPUs are usually comprised of a number of processor cores, which can each perform tasks in parallel to increase program execution speed. CPU cores are well-suited for complex workloads, requiring distributed memory or communication between computational threads. Unfortunately, their flexibility is wasted when undertaking parallel workloads consisting of simple, similar types of tasks executing many times. In both cases, execution may be sped up through parallelisation, for example via Message Passing Interface (`MPI`) or Open Multiprocessing (`OpenMP`) application programming interfaces (API). To track many uncoupled fast-ions, where message passing, distributed memory and precise thread control is not required, the small number ($\mathcal{O}(1-10)$) of cores available in CPUs is likely to create an early bottleneck.

Graphics processing units (GPU or GPGPU in the case of a general-purpose GPU, used for computation) on the other hand have a highly specialised architecture, designed around maximising the number of available cores. A simple diagram to illustrate the differences between CPU and GPU architectures is shown in figure 3.1 below. Whilst each CPU core (blue) has dedicated space for control logic (orange) and access to its own memory cache hierarchy (yellow, green, pink), GPU cores reduce that space in exchange for more cores arranged in groups, which are controlled together and share memory. For tasks which require less thread control, such as vector operations like matrix multiplication (the original motivation behind the GPU architecture), the additional cores, and the deep pipelines they form, can easily increase execution speed by orders of magnitude. For example, an Nvidia

Tesla V100 GPGPU card has 5120 cores, compared to the Xeon Phi with typically fewer than 100 cores. Crucially, the large increase in simulation realism and detail made possible by GPGPUs, and the resulting opportunities to discover new physics that was previously unresolvable with conventional CPU fast-ion codes, is the reason why GPGPU codes are adopted for studying ITER in this work.

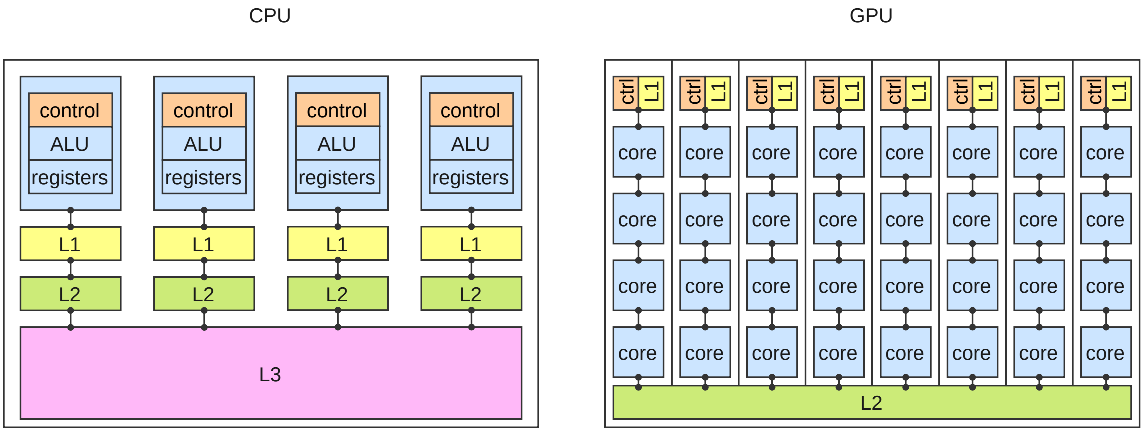

Figure 3.1: Typical CPU architecture (left) shown against GPU architecture (right). L1, L2 and L3 denote memory cache hierarchies of increasing volume, whilst arithmetic logic units (ALU) and registers are responsible for computation. At the cost of thread-level control and memory, the increased space for compute cores on the GPU chip leads to an increase in parallel processing throughput. On the GPU card, cores are arranged in groups (columns) which share L1 memory and are controlled together for vectorised operations.

3.2 LOCUST

Whilst simulations traditionally supplement empirical tokamak science, the large spatiotemporal scales in ITER make such computational studies challenging. For example, compared to DIII-D, fast ions must be tracked for roughly 20 times longer across surface areas roughly 13 times larger, implying a factor 260 increase in simulation time for serial codes. This is compounded by other problems, such as the need for smaller timesteps to minimise the truncation errors of numerical methods and the evaluation of gyro-resolved trajectories to resolve PFC heat loads precisely.

To alleviate these problems, the Lorentz-Orbit Code for Use in Stellarators and Tokamaks (`LOCUST`) is chosen as the basis of the fast-ion simulation workflow assembled in this study. `LOCUST` is a novel high-performance computing (HPC) algorithm that uses multiple levels of single instruction, multiple thread (SIMT) parallelism and GPGPU hardware to enable the routine study of ITER-scale devices impossible with other CPU codes.

3.2.1 Fast-ion model

The primary goal of `LOCUST` is to calculate the steady-state distribution function of fast-ion species as efficiently as possible, whilst also resolving individual ion trajectories in a realistic 3D geometry to calculate PFC power loads. To complement massively parallel memory-bound hardware, `LOCUST` opts for a mathematically simple but computationally intensive approach: solving the Lorentz equation of motion (2.6) for real-space positions $\mathbf{r}_\mathrm{i}(t)$ and velocities $\mathbf{v}_\mathrm{i}(t)$ of individual kinetic markers i, representing Monte Carlo samples of the fast-ion distribution function. To introduce a boundary condition, these samples are typically tracked from some known state - for example, the phase-space ionisation positions of ions injected by an external source such as an NBI.

In `LOCUST`, these trajectories are evaluated using fixed-step numerical schemes to minimise thread divergence - the degree to which computational threads progress through their workloads at different speeds, increasing the time that threads must wait idle to synchronise and, in turn, decreasing the efficiency of the computation. These schemes track either the particle's position in real space throughout a full gyro-orbit (FO) or track its guiding centre (GC) and gyro-phase by solving the modified equations of motion in [16] (see Chapter 7.1). Multiple such numerical integration schemes are included: Strang-Splitting [110]; BGSDC [111]; Runge-Kutta-type integration methods such as McClements-Thyagaraja-Hamilton, Fehlberg [112], Cash-Karp [113], Dormand-Prince [114], and Goeken-Johnson [115]; and Euler methods such as the popular Leapfrog/Boris [116].

`LOCUST` relies on approximations to enhance the computational tractability of the physics model. Firstly, the electromagnetic field experienced by each fast particle is evaluated without the relatively weak contributions from other fast-ion species. If it is assumed that the background plasma is in static equilibrium with the fast-ion distribution, the computation can be simplified in three ways:

1. The influence of fast ions on the thermal plasma is ignored. If and only if this is done then it becomes valid to track independent sub-samples of the fast-ion distribution (see item 2). Without this approximation, these sub-samples of the fast-ion population would be coupled together via their interactions with the bulk plasma. However, the samples can be treated independently in the case where the background equilibrium is

assumed to be in static equilibrium with the final steady-state fast-ion distribution. In practice, equilibria are taken from pre-converged time-dependent transport simulations, which use simplified fast-ion models. In a given `LOCUST` simulation, the equilibrium is held constant by fixing the background plasma temperature, density and rotation profiles as well as the magnetic field.

2. The entire trajectory history of each marker is utilised when calculating the distribution function, such that a marker ensemble of constant size need only be tracked from source to sink once to estimate the steady-state distribution function. By extending the logic from item 1, one can treat individual points along a marker trajectory as independent subsamples of a common distribution function. This effectively parallelises the calculation across time similar to other fast-ion codes such as `ASCOT` [117].

3. Each fast ion is tracked independently in parallel by non-blocking processes - computational processes which may safely proceed independently. This is again enabled by item 1.

Without interactions between fast-ion species, the equilibrium $\mathbf{E}(\mathbf{r}_\mathrm{i})$ and $\mathbf{B}(\mathbf{r}_\mathrm{i})$ terms are dominated by background sources, such as the fields due to external coils and the resulting plasma response, and short-range Coulomb interactions with the thermal species. Whilst the background equilibrium is prescribed numerically, the short-range interactions with bulk species are replicated using a stochastic perturbation to the marker velocity vector term in equation 2.6. This is implemented with a Monte Carlo Fokker-Planck collision operator that includes terms for diffusion and drift of pitch angle (λ) and energy (ϵ). For the computational convenience of evaluating uniformly distributed random numbers, using the Nvdia `CURAND XOR` shift, these collision operators are given by expressions 2.40 and 2.41, based upon the binomial operators derived in [33]. An option to accelerate the effects of these equations is also included - similar to the goosing algorithm in `NUBEAM` [118].

By evaluating these expressions, `LOCUST` calculates and outputs a range of physics results in many formats. 3D power loads are derived from the intersections of orbits and PFCs. Similarly, Hamiltonian field line trajectories can be efficiently evaluated to create various types of Poincaré maps, each designed to illustrate particular magnetic field structures. The distribution function is generated by binning markers typically every 100ns - roughly four gyro-periods for a 1MeV deuteron in ITER - and it can be generated in R-Z-v-λ and constants-of-motion ϵ-P_ϕ-σ-μ spaces. Here P_ϕ is the fast-ion canonical angular momentum and μ the instantaneous magnetic moment, which is expanded to first order [119] in gyroradius to improve accuracy when binning markers in devices with steep field gradients. Another variable σ is required to track the sign of the first-order guiding-centre pitch, however this is not always a constant of motion. `LOCUST` also calculates one-dimensional poloidal flux profiles of fast-ion-driven current, torque ($\mathbf{J} \times \mathbf{B}$ and collisional), pressure and heating to bulk species channels.

3.2.2 Implementation

The overall data flow to and from `LOCUST` is illustrated in figure 3.2, along with preprocessing stages and related external physics solvers. The background equilibrium fields describing the 2D axisymmetric and 3D perturbative components are passed to `LOCUST` as separate numerical representations. These can be in the form of `IMAS` Interface Data Structures (IDS), GEQDSKs, 3D rectilinear grids, or Fourier-decomposed data. The bulk species temperature and density are supplied to the collision operator numerically as interpolated functions of poloidal magnetic flux. The initial fast-ion ensemble is read from individual marker phase-space positions, often generated by external plasma heating codes. Finally, the 3D PFC power flux can be calculated from an axisymmetric limiter outline or an unstructured volumetric tetrahedral mesh, avoiding the need for runtime octrees - the recursive subdivision of regions of 3D space into eight sub-regions. 3D meshes are generated by defeaturing and repairing elements of computer-aided design (CAD) engineering models, typically using CADfix or SpaceClaim, before re-meshing volumetrically in Attila [120]. Equally, this geometry can be represented by the `IMAS` generic grid description (GGD) in the wall IDS.

Once runtime settings and input data are specified, execution proceeds through stages 1-7 in figure 3.2. First, the X-point, magnetic axis and last-closed flux surface (LCFS) are precisely located. Since `LOCUST` is not storage bound (limited by on-board CPU or GPGPU memory volume), the rectangular 2D field is then effectively cached by storing pre-computed bicubic spline coefficients for each knot - either just the required derivatives and cross-derivatives or the entire set of coefficients. Whilst this method is also offered in `LOCUST` for storing 3D fields as tricubic splines on a rectilinear grid, a Fourier-decomposed format is also available. This latter option is preferred if on-board memory begins to limit grid resolution or when resolution is only needed in particular dimensions; the freed space

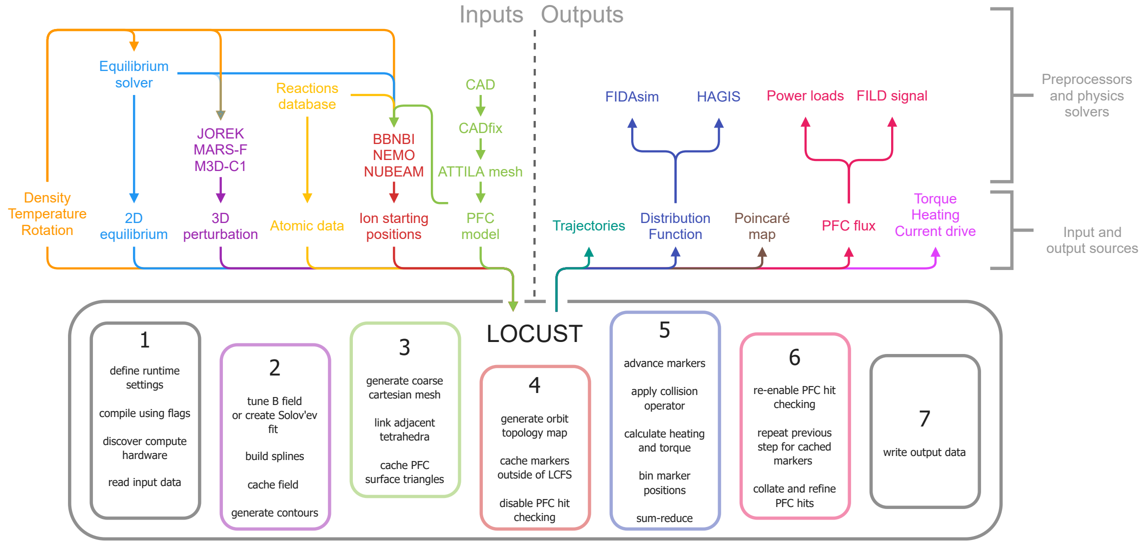

Figure 3.2: Information flow (arrows) to and from `LOCUST`, itself shown in the solid grey box with main preprocessing, runtime and postprocessing stages. These stages execute in ascending order one to seven and are colour-coded to match relevant input/output data. Included are examples of external physics codes which have been used to pre-process or post-process data [118, 120–129]. Figure reproduced from [109].

may be used to redistribute resolution (for example using fewer harmonics at higher spatial grid resolution), reducing magnetic field divergence enough to enable linear interpolation - which is faster. The 3D mesh is then cached by labelling vacuum-facing triangles as PFC surfaces. For rapid and synchronous tracking of tetrahedra traversal, nodes comprising the unstructured wall mesh are also mapped to a coarse cartesian grid and adjacent tetrahedra are identified [130].

After runtime preprocessing, marker tracking is performed in two stages: markers which start and remain inside the LCFS are first tracked *without* PFC interception checks before tracking all remaining markers *with* PFC checks. Considering the constraints outlined in section 3.2.1, optimal performance is achieved if the number of occupied threads is maximised, indicating GPGPUs are currently the most suitable hardware - with each on-board streaming multiprocessor running thousands of threads, each able to track an individual marker in parallel. The Nvidia `CUDA` API is used to interface with this hardware, and the core tracking algorithm is written in PGI `FORTRAN`. `CUDA` also offers libraries for efficient random number generation, including `CURAND` and a Mersenne Twister [131] implementation which are both optional in `LOCUST`.

Implementing SIMT algorithms of this type requires some hardware-specific design considerations. The marker tracking algorithm, which comprises the stages in box 5 of figure 3.2, is illustrated in figure 3.3. The required data are first transferred from the host CPU to the GPGPU device before a non-blocking device kernel is triggered to advance markers in phase space. Simultaneously, fast-ion positions from the previous timestep are sum-reduced on the host - binning their positions on the distribution function grid. Care is taken to overlap these processes; upon their completion, all processes are synchronised by a transfer of the next fast-ion positions to the host. These positions take the integer form of the corresponding distribution function bin indices to lower data throughput. This process is scaled across multiple devices within a node using `OpenMP`.

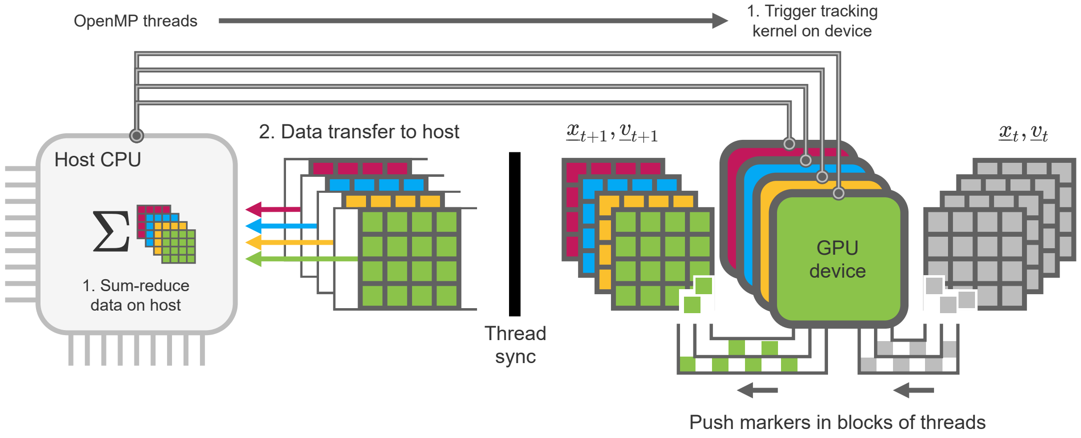

Figure 3.3: Execution model for the kernel-level marker tracking algorithm. Steps 1 and 2 execute in serial, but operations therein are performed asynchronously. In step 1 the GPGPU tracking kernel is triggered whilst the previous marker positions are sum-reduced on the CPU to form the distribution function. These processes are synchronised by a memory transfer of the new marker positions from the device to host in step 2. Figure reproduced from [109].

As marker trajectories are evaluated, the latest positions are cumulatively binned to build up the steady-state distribution function. This process, illustrated in figure 3.4, occurs until markers either strike a PFC, reach a tracking time limit or slow to a prescribed velocity cut-off - typically $3T_{\mathrm{bulk}}/2$. Upon completion, incident PFC power is collated and adaptively refined across the surface mesh by splitting those tet faces (triangles) which record many marker interceptions.

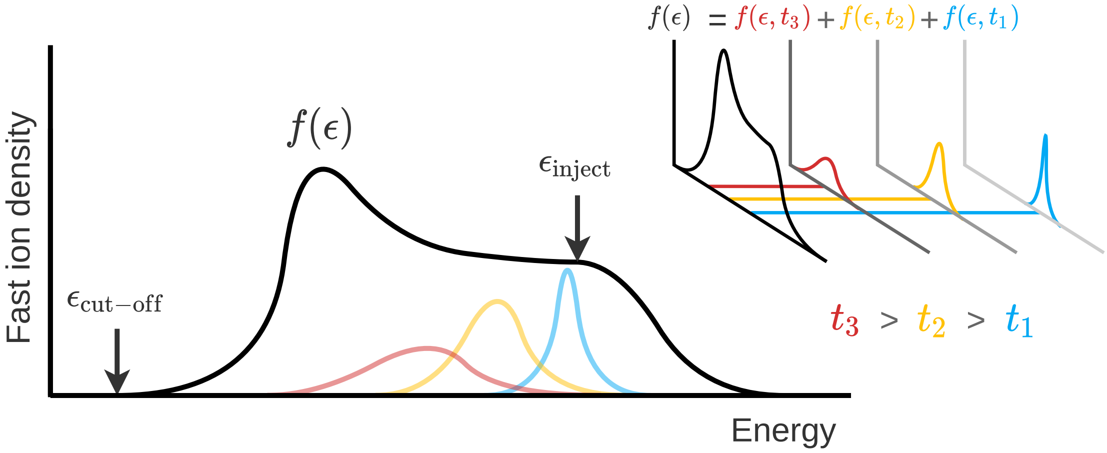

Figure 3.4: Numerical algorithm for generating the steady-state distribution function f, shown here in black as a single function of marker energy $f(\epsilon)$. This is formed by cumulatively summing $f(\epsilon, t_i)$, which are calculated independently at discrete time intervals t_i. Figure reproduced from [109].

3.3 Simulating ITER

To build an accurate integrated fast-ion model of ITER, additional software components to calculate, pre-process, post-process and feed data to and from `LOCUST` are required. Furthermore, to fully exploit `LOCUST`'s speed to perform many simulations, automated control systems are required to manage this flow of data between the software components and `LOCUST`. This section describes the data streams, software components and control systems required to provide the physical quantities outlined in figure 3.2, including the motivation behind using each in the fast-ion workflow, which is assembled later in chapter 4.

3.3.1 `IMAS`

`IMAS` is the primary simulation platform designed and advocated by ITER to design, build, deploy and run simulations of integrated models of ITER and other tokamak devices. It achieves this by:

- Defining a data model to facilitate the storage and communication of data representing physical quantities between software components. The model consists of a data schema to describe physical quantities; the Interface Data Structure (IDS) storage format; a database environment for storing machine descriptions and IDSs which describe simulation and experimental data; and a multi-language API for enabling community physics codes to access and exchange IDSs.

- Providing integration with software tools, such as Kepler, for defining and automating integrated modelling workflows that consist of physics codes that have been adapted to `IMAS` and exchange IDSs.

There are many benefits to using `IMAS` compared to other available integrated modelling systems. Its data model makes assembling workflows using `IMAS`-compliant codes extremely easy due to each component only requiring IDSs. Once a code is adapted to `IMAS` and compiled as an "actor", it becomes available via the `IMAS` API as a function whose arguments consist purely of IDS instances. This abstraction means that workflows can be created in any scripting language where the API is available through the arbitrary calling of functions in place of physics codes. This is simpler and more scalable than calling standalone physics codes, which each likely require data in different formats and different locations. Therefore, it is generally preferable to use `IMAS` components where possible.

3.3.2 RMP and plasma response model

To accurately calculate the fast-ion transport due to the ITER ECC system, a fast-ion workflow must include a physical model for describing the magnetic field created by the ECC system. As explained previously in section 2.4, it is crucial that the model also include the effects of the plasma response in its calculation of the perturbation field. However, due to the aforementioned computational challenges associated with ITER, there are no `IMAS`-compliant codes able to calculate the ECC plasma response at high fidelities without bottlenecking workflow performance. Therefore, 3D field data were instead taken from a recently pre-calculated database [64, 65] produced outside of `IMAS` by the `MARS-F` code [124, 125] for ITER. The code solves the linearised forms of the full-MHD equations in toroidal geometry. The database contains 3D fields calculated for the scenarios described later in table 5.1, for multiple toroidal mode numbers, harmonics and coil rows of the ECC system. In this work, only fields which included the plasma response were used.

The use of the `MARS-F` database has many positive consequences for the subsequent design of the fast-ion workflow. The linear nature of the physical MHD model means that the individual field components can be rescaled, phase-shifted and re-combined to simulate arbitrary ECC configurations and levels of ELM suppression - without re-running `MARS-F`.

There are also challenges presented by the database. The 3D fields, and supporting data, are stored in a variety of formats. Additionally, the 3D fields must be pre-processed using a separate code prior to reading by `LOCUST`. However, whilst this requires a small amount of time and increases the workflow complexity, the pre-processing allows for the generation of 3D fields on grids of arbitrary precision. As discussed in section 3.2.2, this can increase `LOCUST`'s efficiency through optimisation of the space-time trade-off between grid resolution and the order of 3D field interpolation. Furthermore, the pre-processor is unique in that it uses different orders of splines near the separatrix to better resolve discontinuities in the 3D field due to the plasma response creating fine current sheets. To ensure this approach is valid, a convergence test is performed later in section 4.5.

3.3.3 Equilibria

For consistency, the same equilibrium plasma data - plasma temperature, density, rotation and equilibrium magnetic topology - used to calculate the 3D field database were used as a pre-computed data source in place of a physics component in the workflow. This was previously determined by `ASTRA` [132] for different assumptions of plasma density profile shape, Prandtl number (ratio of toroidal momentum to thermal diffusivity in the core) and ratio of toroidal momentum to thermal confinement times - (τ_ϕ/τ_E). The latter two quantities affect the plasma rotation at the core and edge respectively, in turn influencing the plasma's ability to screen the external perturbation. This effect is shown more clearly in figure 3.5 below. For the fields used in this study, the differences in rotation act only to scale the overall response amplitude without affecting the perturbation structure - though the increase in amplitude is small. As shown later (figure 5.15), this likely amounts to linearly scaling the total fast-ion losses. The studies herein assumed $\tau_\phi/\tau_E = 2$ and a Prandtl number of 0.3 [133], as expected from turbulent transport simulations which correspond to a smaller XPD, due to the low plasma resistivity and relatively high absolute value of the toroidal rotation [134] of 10-30kRad/s, which is low when normalised to the Mach number in ITER.

As it is stored in the same location as the 3D field data, the equilibrium data are highly accessible. Another advantage is that it avoids having to re-run `ASTRA` to re-generate the same data repeatedly within the workflow execution. However, like the 3D field data, the equilibria data are not stored in IDS, and are instead stored in a variety of file formats - some requiring conversion prior to reading by `LOCUST`.

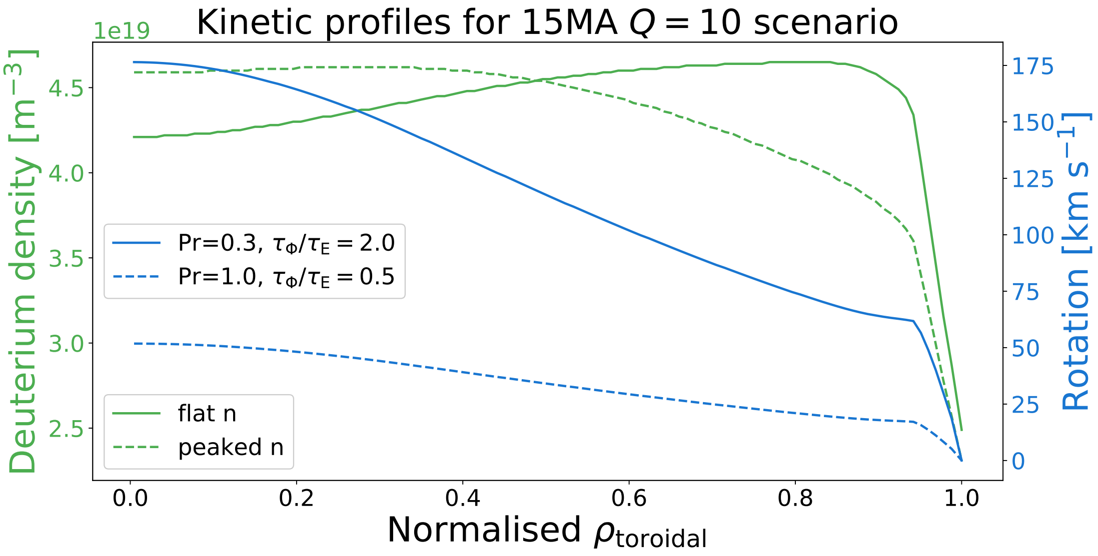

Figure 3.5: Kinetic profiles for the $Q = 10$ scenario under different assumptions of the Prandtl number (Pr) and ratio of toroidal momentum to thermal confinement times (τ_ϕ/τ_E). Profiles are plotted against normalised toroidal flux. These parameters greatly affect the rotation profile, which can greatly affect the plasma's ability to screen external perturbations. $\tau_\phi/\tau_E = 2$ and a Prandtl number of 0.3 were chosen in this study, whilst the "flat n" profile is used when studying the $Q = 10$ scenario. This was shown to be the worse option for fast-ion confinement due to the high edge density.

3.3.4 Neutral beam model

No pre-calculated NBI deposition is available in the 3D field database, hence a physics component is needed to dynamically calculate it when required. Multiple NBI codes are available in `IMAS`, and two were tested: `NEMO` [128] and `BBNBI` [127]. However, unlike `BBNBI`, it was found that `NEMO` cannot correctly write depositions in 3D geometry, which is crucial for this study. Therefore, `BBNBI` was chosen as the NBI component of the workflow.

As an actor available through `IMAS`, the `BBNBI` actor reads and writes IDSs - one which describes the NBI geometry, such as injecting on or off axis, and one which describes the axisymmetric plasma quantities, such as temperature, density and rotation. Because of this, some file conversion between the 3D database, `BBNBI` and `LOCUST` is required. In the case where the neutral beam geometry is missing, for example when using the diagnostic neutral beam as is done later in section 5.3, the IDS must first be created.

3.3.5 Plasma-facing component model

`IMAS` contains functionality for the definition, creation, manipulation and re-meshing of PFC geometries. This is mostly encapsulated by the combination of the General Grid Description (GGD), part of the `IMAS` data schema, and the `SMITER` [135] field line tracing code, which contains a suite of tools used to manipulate GGD meshes. However, because of the precision required to resolve the more vulnerable PFCs, such as the under-dome cooling pipes shown in figures 3.6 and 3.7 below, inefficiencies in memory allocation during wall IDS generation made the `IMAS` and the GGD impractical.

Instead, a pre-existing volumetric mesh (57.6M tetrahedra), derived from a defeatured computer-aided design (CAD) model of ITER, was used to represent the first-wall geometry. This mesh is detailed enough to resolve the divertor geometry and cooling pipes, as well as the subtler geometrical features of larger components, such as the gaps between divertor cassettes and the shaped surfaces of the first-wall tiles. Gaps exist for heating and diagnostic ports. This is illustrated in Figure 3.6, which shows the PFC geometry and its individual components highlighted.

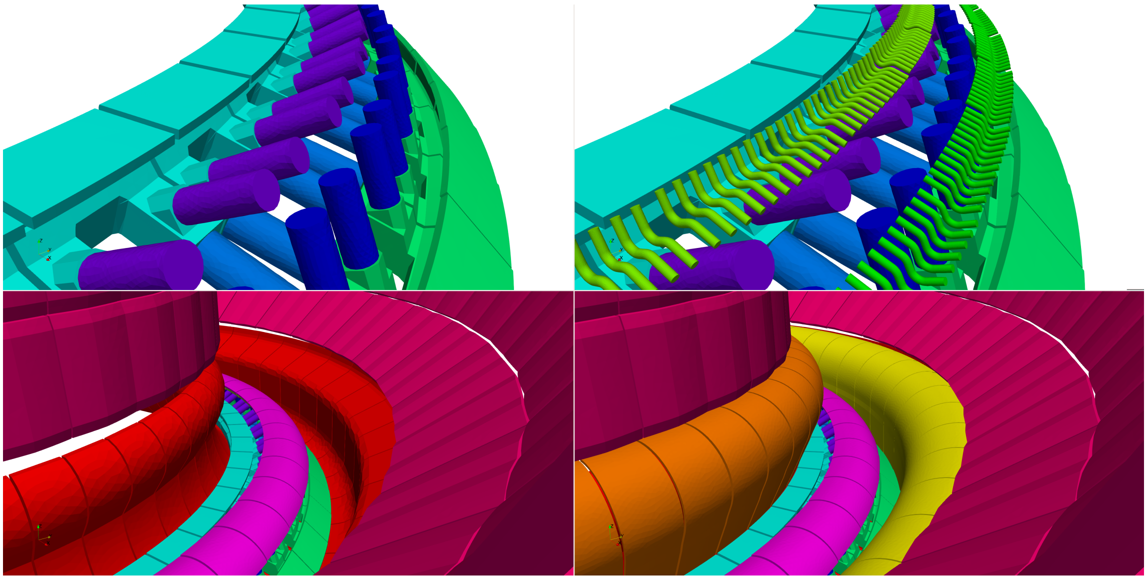

Figure 3.6: ITER mesh used by `LOCUST`, shown in four separate panels as groups of components are incrementally added. Clockwise from top left: Close-up of divertor inner and outer reflector plates with inner, outer and horizontal supports; adding inner and outer under-dome pipes; zoomed out with divertor base, first-wall and dome added; inner and outer divertor target and baffle added. The underdome pipes, whilst shielded by the dome from the private flux region above, may be reached by fast ions from the exposed sides. Fast ions in this region may also strike the rounded vertical supports too. Colours match those used later in figure 5.3. Figure reproduced from [109].

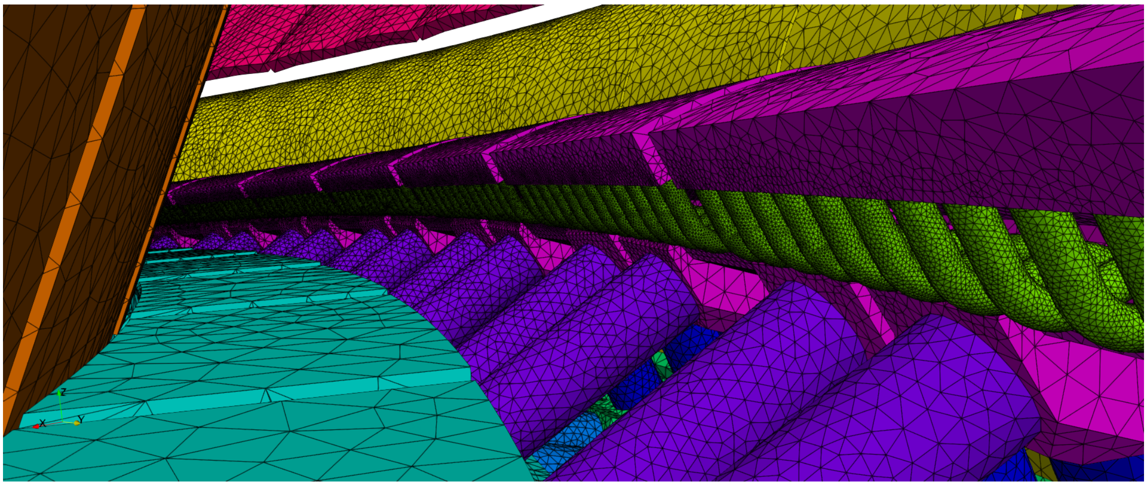

Figure 3.7: Close-up of the divertor region of the mesh model used by `LOCUST`. Tet outlines are highlighted to show the required model precision to resolve power loads on the vulnerable under-dome pipes (shown in green). Colours match those used later in figure 5.3.

Chapter 4

Model testing and integration

Due to the lack of any prior testing in the literature, the `LOCUST` code was first tested to ensure its accuracy when applied to predictively studying ITER and enable confident future use outside of this work. To verify that the code is implemented correctly, `LOCUST` was compared against multiple, well-established codes that perform equivalent calculations. To validate that the simulated physical models realistically represent the real world, results from `LOCUST` were compared to experimental results. In order to isolate and examine individual features of `LOCUST`, a range of increasingly complex scenarios were studied. These tests are outlined in sections 4.1, 4.2 and 4.3.

A software system was then designed and built to supply `LOCUST` with input data from the software components described in section 3.3. This system was also designed to automate `LOCUST` preparation and execution, so that the speed of the code could be fully utilised to perform larger parameter scans more quickly than previously possible. This system is described in section 4.4 and tested in section 4.5, and it was used to generate the results presented in chapter 5.

Parts of this chapter have been published in references [109] and [60]. All codes were operated by the author except `ORB_GC`, `MPI_ORBF`, `FIDASIM`, `M3D-C1` and `probe_g`.

4.1 Verification of axisymmetric simulations

The most fundamental test possible was performed to check that `LOCUST` correctly solves the equations of motion (equation 2.6) for individual markers and, for the test case used in this chapter, orients the system and magnetic field geometry correctly to recreate the main orbit topologies. Previous studies have shown orbits to match well with other codes and conserve constants of motion satisfactorily when expected [136]. In figure 4.1, collisionless marker trajectories generated by `LOCUST` using the Boris integrator are compared with equivalent from `ORB_GC` (see appendix of [137]), a guiding-centre fast-ion code. Each marker was initialised at a fixed location with the full DIII-D injection energy of 80keV. The trajectories were calculated in a realistic tokamak magnetic equilibrium with different values of pitch angle, corresponding to subplots (a) to (d), to check that the different orbit types described in section 2.1.1 are reproduced. The magnetic equilibrium was a reconstruction of DIII-D discharge #157418 [81], an experiment that studied fast-ion confinement in a plasma with an ITER-similar shape and externally applied RMP fields. Because of this ITER-similar shape, the discharge is well-suited for verifying `LOCUST`'s performance in ITER-relevant conditions, and so it will be used throughout chapter 4. A static plasma was simulated by using only a single time slice at 3s to remove variations in the background data.

The full-orbit trajectories calculated by `LOCUST` in figure 4.1 closely resemble the motion produced by `ORB_GC` solving the guiding-centre equations of motion, showing that `LOCUST` correctly simulates both the orbit shape and transitions to different topologies - trapped, co-passing and counter-passing - as the marker pitch varies.

To then test the implementation of 2D and 3D models of the tokamak wall in `LOCUST`, the previous study was repeated with checks for wall interceptions in the presence of a 2D and 3D wall model of DIII-D. Markers were also spaced over R-λ space along the outer horizontal mid-plane and at an injection energy of 80keV to produce an ensemble that contained a wide range of orbit shapes. To make the comparison more accurate, `LOCUST` was compared with the full-orbit code `MPI_ORBF` [138], which used a 3D wall model.

Figure 4.2 shows that `LOCUST` and `MPI_ORBF` measure the same prompt loss boundary to within ($\mathrm{d}\lambda = 0.04$, $\mathrm{d}R = 1.0$cm). Variations of this magnitude amount to absolute changes of ∼0.1-0.4 percentage points in total loss fraction: 6.8% (`MPI_ORBF`), 6.9% (`LOCUST` with 2D wall) and 6.5% (`LOCUST` with 3D wall). Therefore, the differences between the trajectories calculated by `LOCUST` and `MPI_ORBF` are within the variations caused by the wall model. These variations are expected in `LOCUST`, as the 2D model, being a perfectly smooth, axisymmetric contour, did not include the mid-plane port box cut-outs; this can be seen in the real-space loss patterns in figure 4.2b., where marker interceptions gather at the box edge after they pass outside the axisymmetric contour boundary. Nevertheless, all three runs still predicted

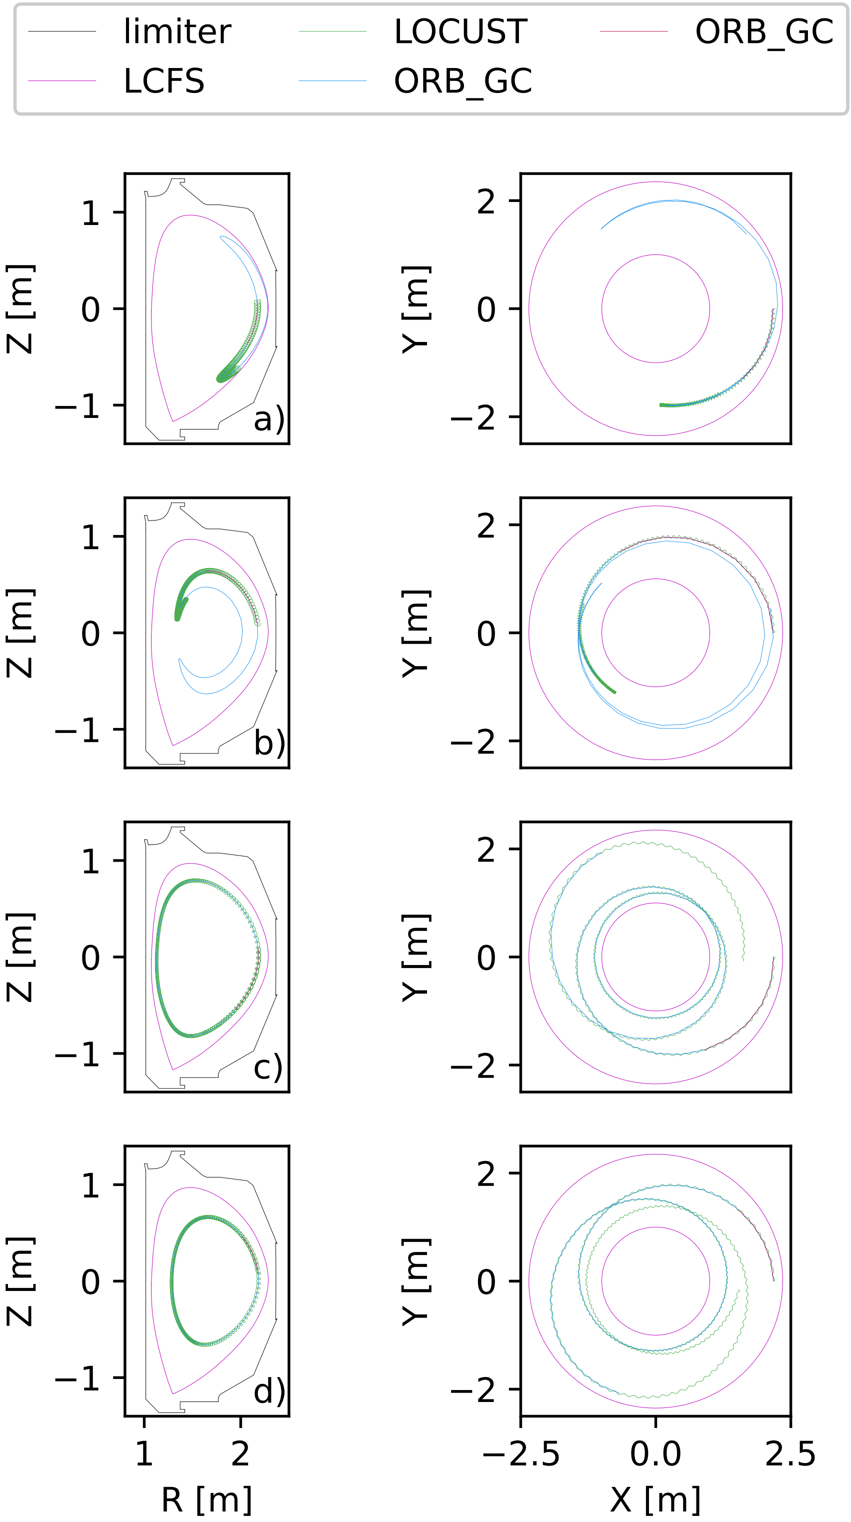

Figure 4.1: Collisionless marker trajectories calculated by `LOCUST` and `ORB_GC` fast-ion codes. An initial portion of the `ORB_GC` trajectory is shown in red to better illustrate the direction of travel of markers. Subplot rows (a)-(d) correspond to pitch values of -0.4, 0.6, -0.8 and 0.8 respectively. The markers were started at $R = 2.2$m, $Z = 0.1$m, $\phi = 0$.

most losses to be located at the outboard mid-plane ($\theta = 0°$) and divertor ($\theta = -100°$) regions.

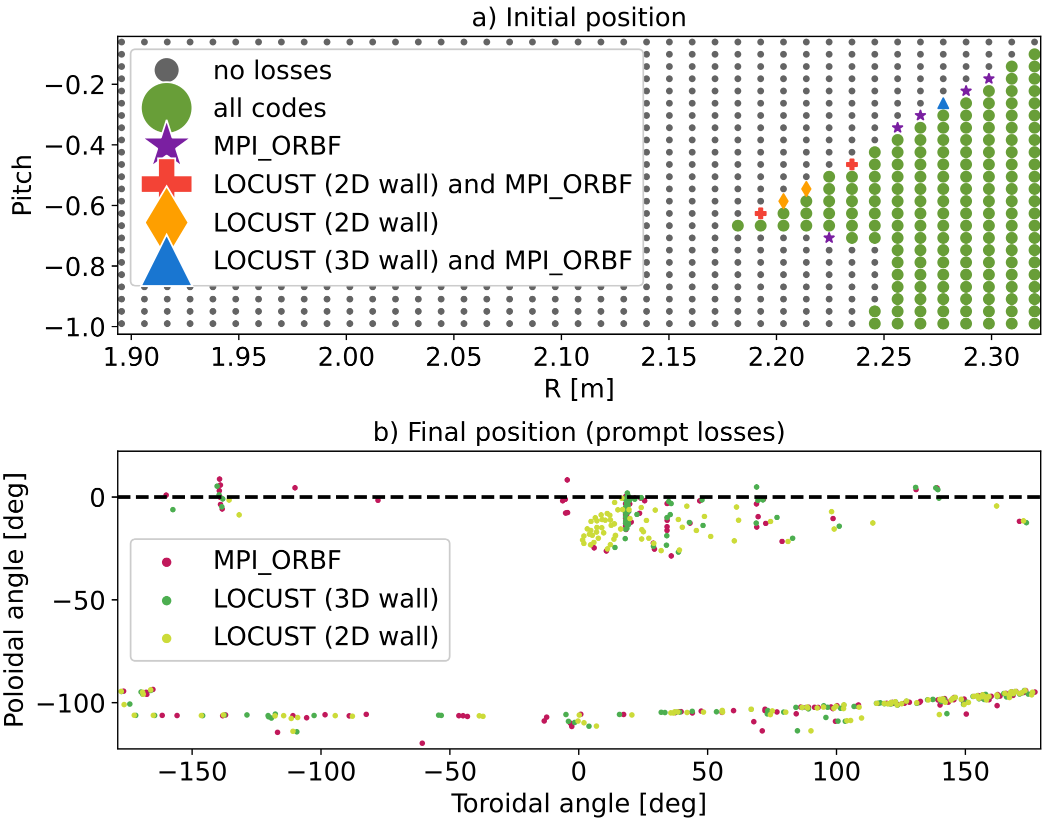

Figure 4.2:
a) Phase-space locations for markers initially distributed on a uniform pitch-radius grid along the outer horizontal mid-plane. Pitch is measured relative to the plasma current. Marker type denotes which simulations measure losses. Discrepancies are localised to the loss boundary within $\mathrm{d}\lambda \mathrm{d}R$.
b) Final real-space positions of all promptly lost markers. All three runs used a slightly different wall model but predict most losses at the outboard mid-plane ($\theta_{\mathrm{pol}} = 0°$) and divertor ($\theta_{\mathrm{pol}} = -100°$). The structure of the mid-plane port box cut-outs can be observed in the loss patterns produced by 3D simulations. Figure reproduced from [109].

To validate the collision operator implementation in `LOCUST`, simulations that included a toroidally symmetric background plasma were compared with `ASCOT4` [117] (version 5 has since been shown to produce similar results [139]) and the `NUBEAM` [118] module of `TRANSP` [140] via the `OMFIT` [106] interface. A realistic deposition of a single, counter-current-injected 80keV deuterium neutral beam was calculated by `NUBEAM` and tracked by each code. To maximise the range of energies explored by the markers, whilst keeping the deposition simple to make errors more noticeable, only the full-energy component was retained.

Notably, it was found that fast-ion distributions generated by co-current beams are much less sensitive to plasma conditions, meaning it is much more difficult to clearly expose incorrect simulations. This is because the orbits of fast-ions injected counter to the plasma current open outwards (net increase in average ψ over one orbit), and so they can explore a larger volume of phase space, meaning they are also more likely to be lost from the plasma. Therefore, an important conclusion of this work is the proposal that future fast-ion benchmarking exercises, such as those proposed by the ITPA-EP, do not rely solely on pulses with co-injection neutral beams - which are often studied due to their superior confinement characteristics in experiment. Ultimately, verification and validation exercises should aim to demonstrate agreement in as wide a phase-space volume as possible, with separate focus on regions detectable by real-world experiments.

Because all codes were forced to use an identical starting marker list from `NUBEAM`, which only provides the weight, real-space position, pitch and energy, the study was limited to guiding-centre tracking. Consequently, when calculating collisional effects and wall interceptions, there is an unavoidable systematic spatial error $\sim \rho$ in the marker position introduced by differences in the finite Larmor radius (FLR) model implemented by each code:

- `NUBEAM` assigns markers a random gyrophase, assuming a circular orbit.
- `LOCUST` tracks gyrophase from birth.
- `ASCOT` ignores FLR corrections unless near the PFC wall, where a random gyrophase is chosen.

Guiding centres were followed for 100 milliseconds and collisions disregarded for markers outside the LCFS to mimic the behaviour of `NUBEAM`. A static, pure-deuterium plasma was used to again avoid complications and smearing out errors by high-order physical models; impurities, neutral species, bulk rotation, electric fields and neutral particle interactions were removed. To effectively remove any beam-beam interaction in `NUBEAM`, the neutral beam density was artificially lowered by reducing the injected NBI power to 1 watt in `NUBEAM` using `OMFIT`, before scaling the deposited marker ensemble power to 1 watt across all simulations. This is valid because there is no interaction between the fast ions in these simulations, meaning their weights can be scaled arbitrarily at any point in time to match a desired power deposition. All codes employed a lower energy cut-off at $3T_{\text{bulk}}/2$, with ions below this recorded as thermalised and removed from the simulation. For the following studies, `LOCUST` used the Cash-Karp integrator scheme to track marker guiding-centres and the Strang-Splitting scheme to track full-orbit positions. It is worth noting that, to achieve the required fidelity without resorting to Monte Carlo smoothing, which may mask possible discrepancies, `NUBEAM` simulations for DIII-D were conducted with the untypical settings shown in table 1 in appendix .2.

At this point, it would be ideal to remove the wall model to isolate fast-ion dynamics within the plasma, however this is impossible in `NUBEAM`. Instead, an artificial, axisymmetric PFC surface was created, concentric to the LCFS but 5% further from the magnetic axis ($R_{\mathrm{fac}} \equiv r_{\mathrm{limiter}}/r_{\mathrm{LCFS}} = 1.05$) - the closest permitted by `NUBEAM`. This was also to ensure that fast ions between the plasma and first wall, which are usually removed by interactions with neutral species, do not build up in the absence of such a neutral model (which is not present here for sake of simplicity). This limiter, as well as others at larger values of R_{fac}, is shown in figure 4.3. Figures 4.4 and 4.5 show the R-Z and ϵ-λ projections respectively of the distribution functions calculated by each code and collected at the marker guiding centres. Subfigures (a) show the same density contours as produced by each code whilst subfigures (b) and (c) both show the absolute element-wise differences, $\delta f(R, Z)$ and $\delta f\prime(R, Z)$, between the `ASCOT` and `LOCUST` distribution functions. A similar element-wise comparison against `NUBEAM` could not be performed reliably, since `NUBEAM` collates the fast-ion density onto a flux-aligned grid. Because `NUBEAM` and `ASCOT` use seemingly different collision operators, two forms of collision operator were used by `LOCUST` to generate the data in subfigures (a) and (b), where `LOCUST` aims to match the `ASCOT` collision operator, and in figures (c), where `LOCUST` aims to match `NUBEAM`. In the former case, the collision operator was fully expanded and the `ASCOT` definition of the Coulomb logarithm was used (see appendix .3). In the latter, the `NUBEAM` definition of the Coulomb logarithm was used, and collisions against some plasma species p were truncated to $\sim \mathcal{O}([v_{\mathrm{i}}/v_{\mathrm{p}}]^5)$ whilst thermal ion accumulation was disabled.

In most regions, $\delta f \sim 3\%$ between both codes, which is within the fundamental uncertainties in the theoretical formulation of the Coulomb logarithm - of the order 5% [141, 142]. Switching to a `NUBEAM`-like collision operator, as in $\delta f\prime$, creates a lower density in the core and outboard edge but higher density towards the X-point. Within the LCFS, this leads to the average δf increasing from 3.6% to 4.7%, a change still in line with theoretical variations. Nevertheless, when the collision operator was matched, δf shows only regions near

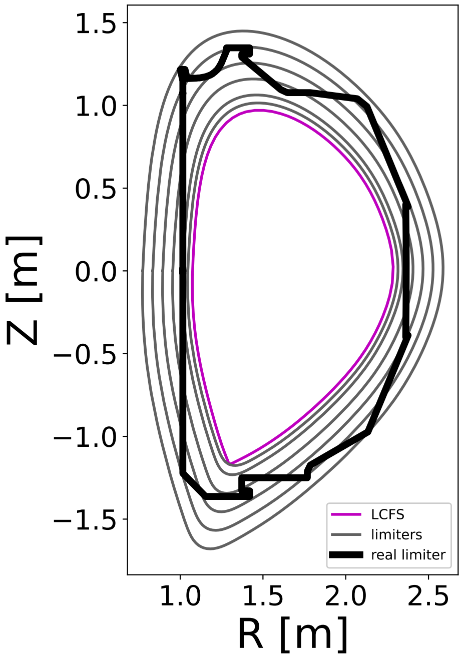

Figure 4.3: A realistic 2D representation of the DIII-D first-wall outline (black) used in discharge #157418, along with the reconstructed LCFS (pink). In grey are the artificial limiters created from the shape of the LCFS. The varying distances, between $1.05 < R_{\mathrm{fac}} < 1.50$, were used to control how many fast ions were able to populate regions outside the plasma.

the X-point and wall retain any distinguishable difference - likely caused by the influence of FLR model on wall interceptions. The remaining noise in the core plasma is likely caused by differences in any tuning applied to the equilibrium field by each code, which, by recalculating the position of the magnetic axis, private flux region and LCFS, may perturb the position of the flux surfaces and cause a flux-aligned noise pattern.

The calculated flux profiles in figure 4.6 also show good agreement, but they highlight the importance of the plasma boundary. The uptick in edge electron heating is caused by the FLR corrections in `NUBEAM` spreading deposited power over a gyro-orbit width; the orbits of markers with guiding centres located just outside the LCFS concentrate their deposited power into a thin shell where the orbit overlaps the plasma. In `LOCUST`, this power is instead collected at the guiding centre - outside the plasma for these markers - and is thus ignored. This effect is artificial, and simulations could avoid this by imposing ion sinks outside the plasma boundary, such as a neutral density for charge-exchange or an extrapolated plasma density - both of which are present in reality and typically implemented in fast-ion simulations.

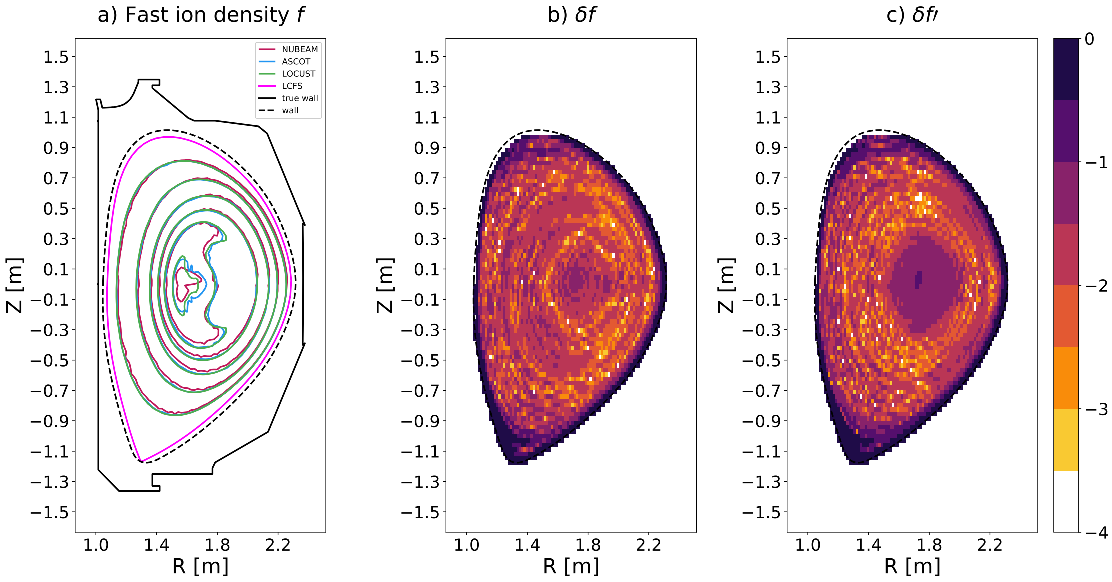

Figure 4.4:
a) Contours of fast-ion density integrated over velocity space. Contours are at equal levels across codes. The real and artificial limiter profiles are shown in black solid and dashed respectively. `NUBEAM` bins according to a flux-aligned helical grid whereas `ASCOT` and `LOCUST` use rectilinear, hence some variation near the magnetic axis is expected due to finite bin widths.
b) Absolute element-wise difference between `LOCUST` and `ASCOT` distribution functions $\delta f \equiv \log_{10}(|f_{\text{LOCUST}} - f_{\text{ASCOT}}|/\max(f_{\text{LOCUST}}, f_{\text{ASCOT}}))$ where max() is the local, element-wise maximum. A high-order collision operator was used in `LOCUST`, as well as the `ASCOT` definition of $\ln(\Lambda)$.
c) The same as plot b) except with `LOCUST` using a truncated collision operator and the `NUBEAM` definition of $\ln(\Lambda)$. This choice primarily affects the core region, though some differences on the outboard side are noted. The loss region near the X-point remains. Figure reproduced from [109].

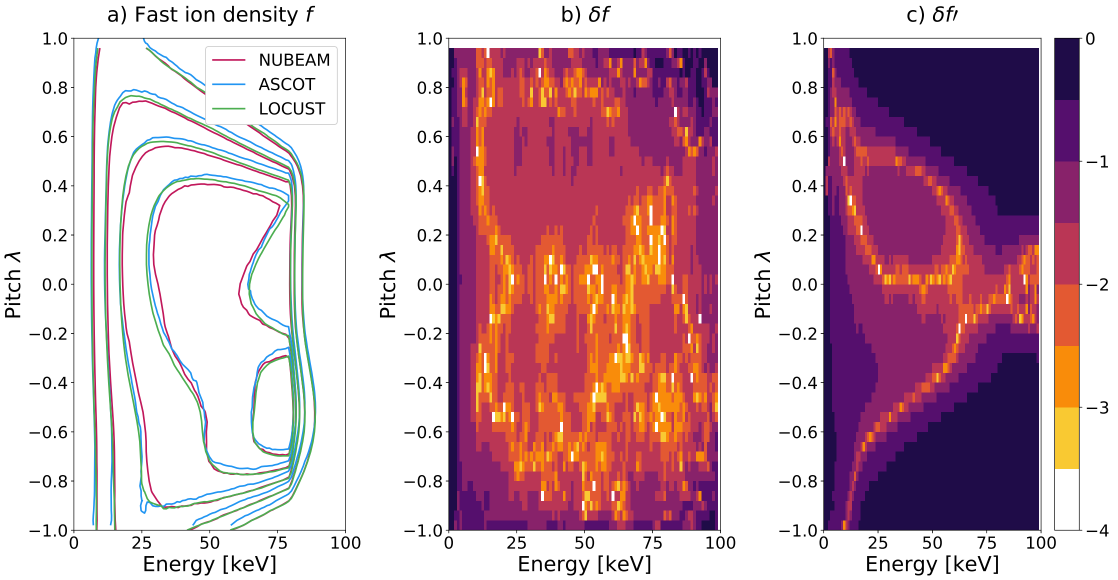

Figure 4.5:
a) Contours in pitch-energy space of fast-ion density integrated over real space, where λ is defined with respect to the direction of toroidal current flow as is convention in `NUBEAM`. Contours are at equal levels across codes.
b) The absolute element-wise difference between the `LOCUST` and `ASCOT` distribution functions $\delta f \equiv \log_{10}(|f_{\text{LOCUST}} - f_{\text{ASCOT}}|/\max(f_{\text{LOCUST}}, f_{\text{ASCOT}}))$ where max() is the local, element-wise maximum. A high-order collision operator was used in `LOCUST`, as well as the `ASCOT` definition of $\ln(\Lambda)$.
c) The same as plot b) except with `LOCUST` using a truncated collision operator and the `NUBEAM` definition of $\ln(\Lambda)$. The injection energy is 80keV, so random variations due to collisional diffusion can be expected above this energy level. Figure reproduced from [109].

The simulated $\mathbf{J} \times \mathbf{B}$ torque suggests some discrepancy in orbit topology [19], but this information cannot be directly extracted from `NUBEAM`. To explore this further, additional simulations were performed with similar artificial limiters up to $R_{\text{fac}} = 1.5$, the largest limiter shown in figure 4.3. This allowed more orbits to populate the vacuum region between the plasma and first wall. Importantly, it was found that it was necessary to increase the `TRANSP` setting `XBMBND` from 1.3 to 1.5; it was found that `TRANSP` uses this hidden setting to register any fast ions located at $\sqrt{\psi_{\text{toroidal}}} >$ `XBMBND` as hitting a PFC (as well as including the original limiter), meaning all prior DIII-D simulations performed by the community will contain erroneous losses caused by this condition. Figure 4.7 shows how `XBMBND` influences the location of these losses, with markers clearly being deposited far away from the realistic PFC wall present in the simulations, which is only correctly illuminated when `XBMBND` is increased to 1.4.

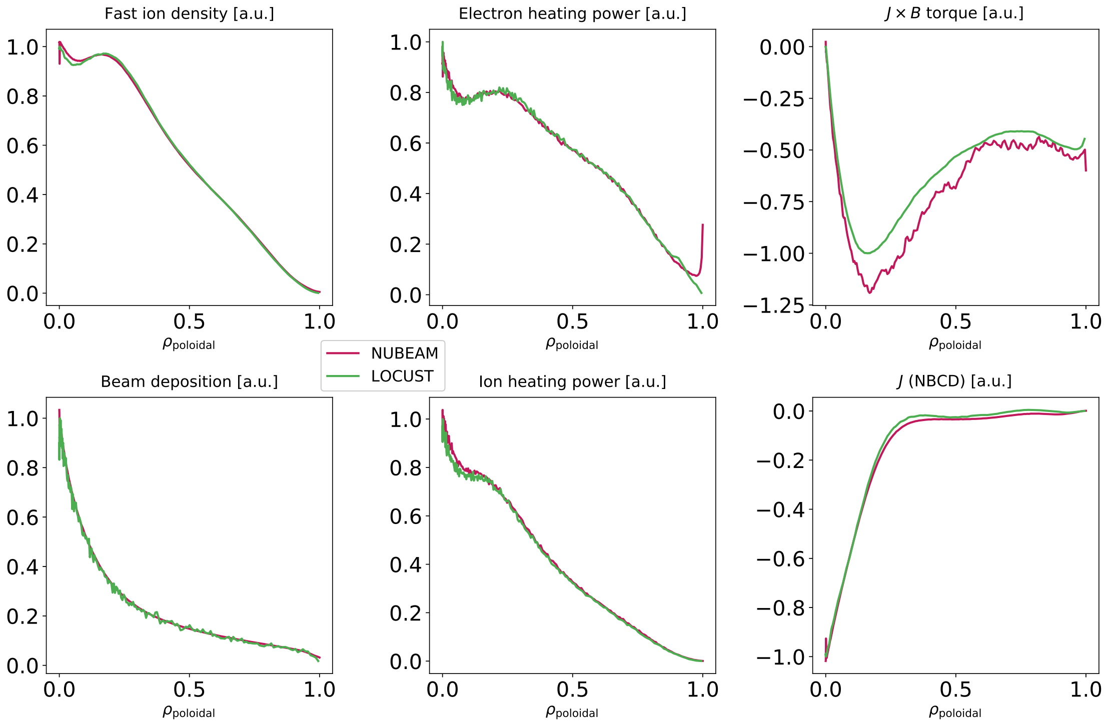

Figure 4.6: Normalised quantities measured against normalised poloidal flux ρ_{poloidal}, including the neutral beam current drive (NBCD). Some discrepancy around $\rho = 0$ is expected due to finite bin widths. Most important is the discrepancy in the $\mathbf{J} \times \mathbf{B}$ torque, which suggests a difference in the measured orbit topology. Figure reproduced from [109].

Figure 4.8 shows the measured steady-state PFC power flux as the limiter distance is increased. As first-wall losses are sensitive to the wall model, some disagreement between the codes is expected, especially at high R_{fac} when orbits between the plasma and outboard PFCs may have a significant ρ. Hence the agreement between `ASCOT` and `LOCUST` is satisfactory, as it is mostly within the variations expected from differences in wall model. However, it is unclear why the `NUBEAM` power fluxes diverge so quickly, though it is encouraging that the resulting discrepancies in the distribution function remain solely at the plasma edge.

The previous methodology was repeated for a spherical tokamak topology. Such devices tend to have steeper gradients [143], meaning any inaccuracies in fast-ion models will be exacerbated; in MAST for example, the Larmor radius ρ can approach orders of $\sim a$, and the impact on the validity of the guiding-centre approximation has long been questioned [144, 145]. MAST shot #29034 was selected to allow for comparison with measured fast-ion D-α (FIDA) emission [146]. The NBI deposition code `BBNBI` was included to enable comparison of full-orbit simulations, as the code is capable of calculating the full 3D space and

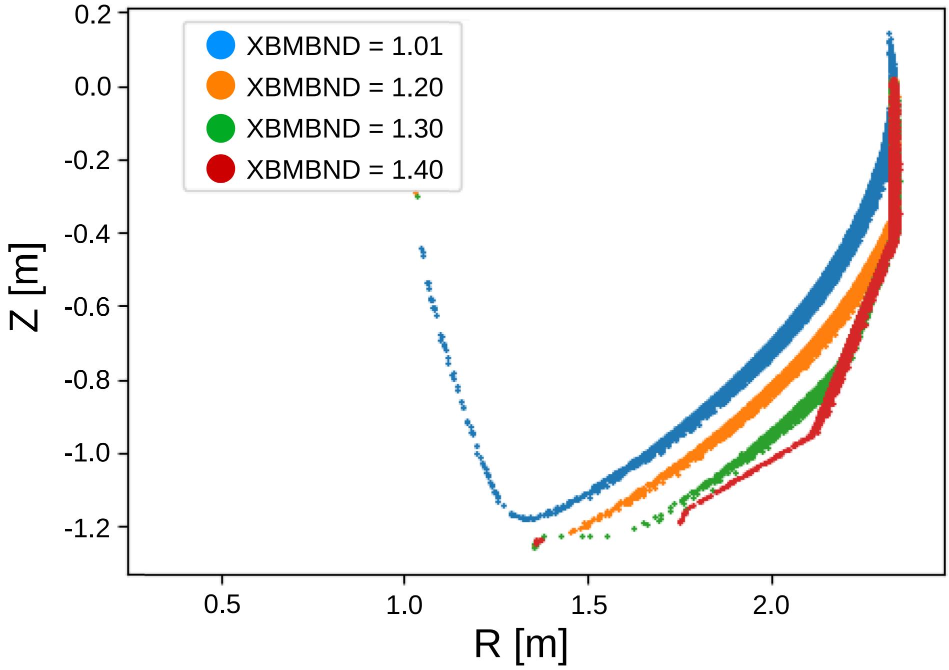

Figure 4.7: Positions of lost fast ions for various values of `XBMBND` in `TRANSP`. The artificial boundary created by `XBMBND` is concentric to the LCFS, existing alongside the real 2D first wall and dominating the observed losses for values of `XBMBND` < 1.4.

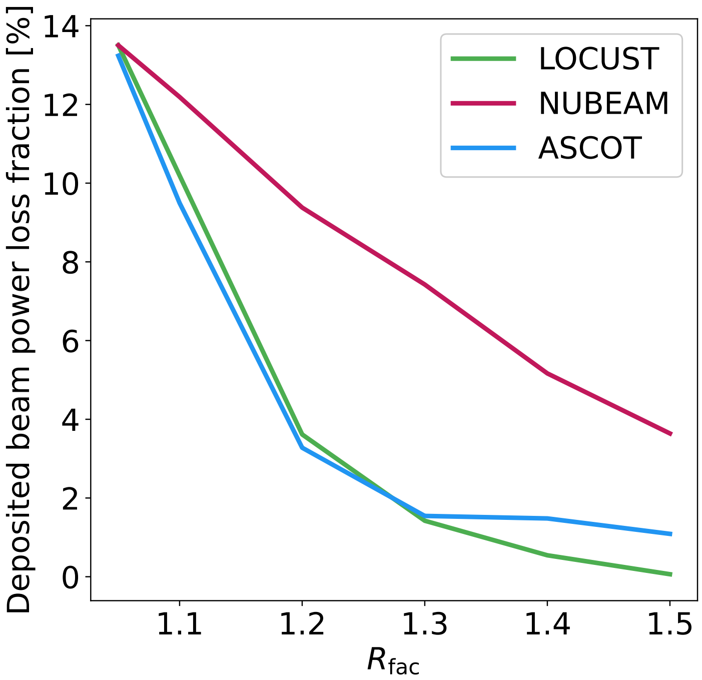

Figure 4.8: Fraction of deposited power lost as a function of limiter wall radius where $R_{\mathrm{fac}} \equiv r_{\mathrm{limiter}}/r_{\mathrm{LCFS}}$ with r representing minor radius. Figure reproduced from [109].

velocity components of deposited markers. Like before, a single time slice of data describing the background plasma was extracted at 360ms. A quiescent period was chosen during the flat-top phase when core electron temperature and density were relatively constant. To create a realistic deposition, time-resolved NBI data from `OMFIT` were used to generate the `NUBEAM` deposition for the south-south neutral beam. To create a similar deposition using `BBNBI`, a code which could not be used with multiple time slices, beam settings were calculated by averaging values in the `NUBEAM` input files across time slices: a 62keV co-current beam with 62% full, 27% half and 11% third energy fractions. Hence it is technically inappropriate, and out of the scope here, to cross-compare the results attached to each beam code in this case, as the input data are not exactly equivalent. GC and FO trajectories were then calculated by each code over 100ms - enough to reach steady state.

The co-current NBI confines the fast ions to the plasma core, where discrepancies are hard to distinguish and there is a systematic shift in spatial density due to the FLR displacement, so instead $f(\epsilon)$ is examined, the distribution function integrated over all dimensions except energy ϵ. This still encodes some real-space information through the effects of the steep temperature and density gradients on the fast-ion diffusion rate. The only discrepancy in $[R,Z]$ space is caused by noise in the `NUBEAM` distribution function, since this case decreased `NPTCLS` to 10^5.

The average $f(\epsilon)$ across all simulations is shown normalised in black in figures 4.9 and 4.10. The residuals - defined as the difference between each simulation and the un-normalised average - are also shown, normalised against the average. Most simulations are within ±3% of the mean. The near-symmetrical splitting of predictions at lower energies is caused by the collision operator truncation whilst systematic differences can be mainly attributed to the $\ln(\Lambda)$ used.

As a figure of merit, for simulations using the deposition calculated by `BBNBI` in figure 4.9, the maximum difference in total fast ions between any two simulations is 5%. This falls to $\leq 2\%$ for pairs of simulations which follow similar assumptions, even including the `ASCOT` GC simulation, which differs from the `LOCUST` equivalent by 1.7% (the FO equivalents differ by 0.9%). The reason for the increased density at high energies measured by `ASCOT` GC is possibly due to the characteristics of the adaptive time step near the magnetic axis. `NUBEAM`, which also uses orbit acceleration, shows a similar feature in figure 4.10. For comparison, in comparing both FO and GC simulations of a homogeneous plasma by both `LOCUST` and

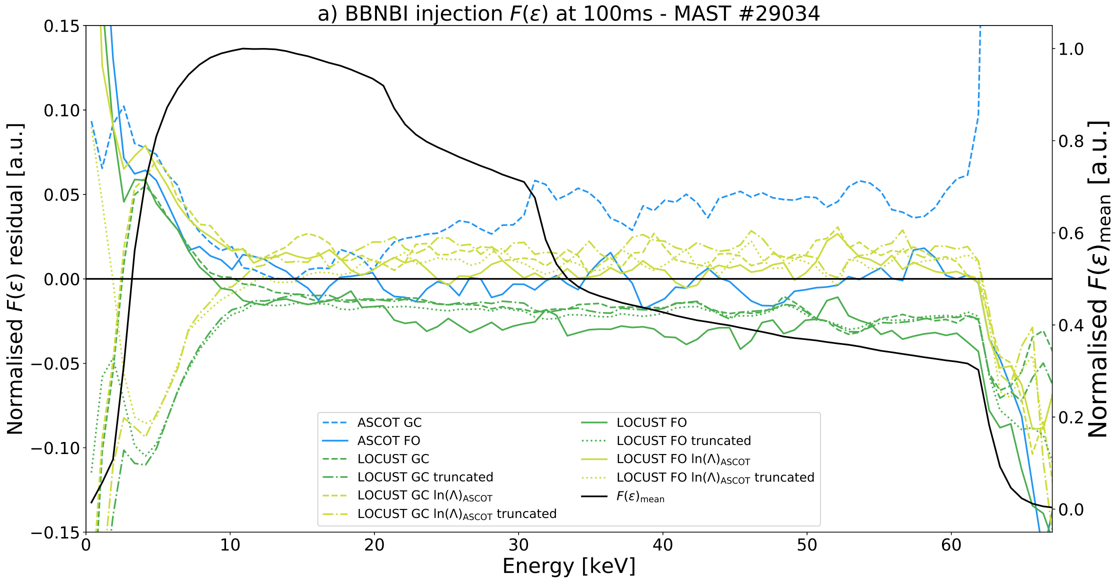

Figure 4.9: The fast-ion density $f(\epsilon)$ after 100 milliseconds, integrated over all dimensions except energy, ϵ, is averaged across all simulations, normalised and shown in black as $f(\epsilon)_{\text{mean}}$. The residual differences $(f(\epsilon) - f_{\text{mean}})/f_{\text{mean}}$ for each simulation are plotted against the density residual axis. Simulations using GC and FO tracking (solid and dashed respectively); truncated and high-order collision operators (truncated and non-truncated labels respectively); and `LOCUST` $\ln(\Lambda)$ and `ASCOT` $\ln(\Lambda)$ (dark green and light green respectively), are all shown, with combinations of these linestyles representing the corresponding combinations of options. Figure reproduced from [109].

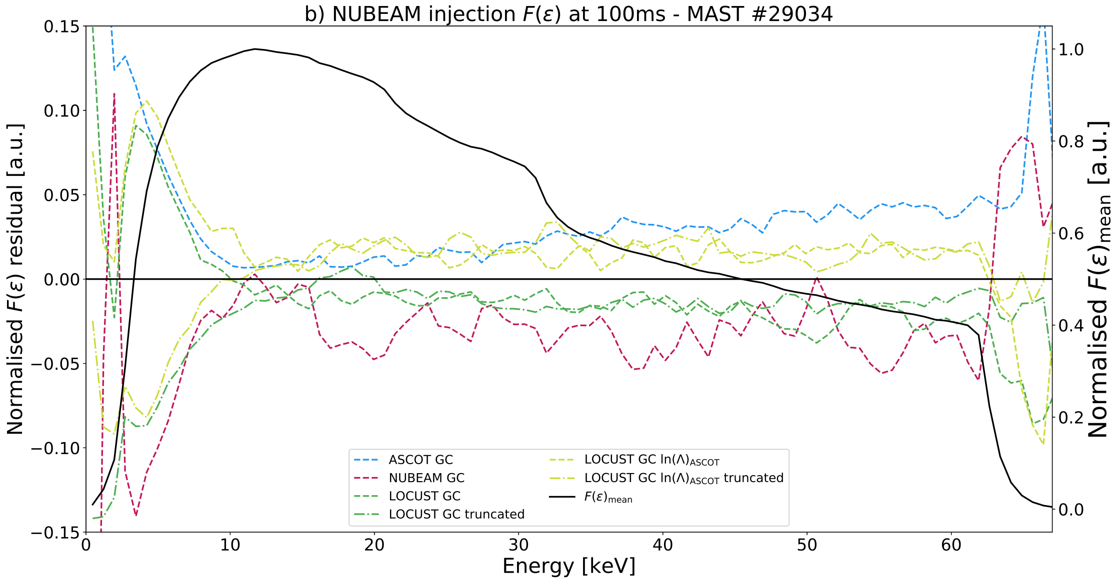

Figure 4.10: Equivalent to figure 4.9 but using deposition from `NUBEAM`, meaning only GC simulations can be performed. The overall trend is similar to figure 4.9 but the `ASCOT` tail is reduced. Figure reproduced from [109].

`ASCOT`, the only discrepancy is a higher fast-ion density of $\sim 20\%$ in the high-energy diffusive tail measured by `ASCOT` GC. Whilst this effect is much more pronounced in MAST than DIII-D, the extra $\sim 20\%$ in the homogeneous case still only leads to a $\sim 0.2\%$ difference in total fast ions - small enough to be affected by slight variations in beam deposition, as shown in figure 4.10 where the effect is lessened when the beam deposition is varied.

To compare these predictions more quantitatively, Kolmogorov-Smirnov (KS) statistics D [147], and their corresponding probabilities $P(D_{\text{measured}} < D)$, were calculated. These are shown in figure 4.11 for matching pairs of $f(\epsilon)$. Typically, the null hypothesis of the KS test, that the two empirical distribution functions to be compared are drawn from the same distribution function, is *rejected* if the measured KS statistic D_{measured} satisfies $P(D_{\text{measured}} < D) \leq \alpha$, with α typically chosen to be ≤ 0.05. Here, α was increased to 0.1 for added rigour. For calibration, figure 4.11a first shows the KS statistic for the GC and FO `LOCUST` simulations. In this case $P \approx 1$ ($D_{\text{measured}} \approx 0$), meaning these distributions comfortably agree as expected. This is in contrast to the next KS statistic, where the collision operator has been truncated, and the test clearly fails. Next, it can be seen that `LOCUST` FO simulation easily agrees with the `ASCOT` equivalent, with $P \approx 1$ after using the same Coulomb logarithm. The equivalent GC comparison fails solely due to the high-energy tail effect, since the equivalent measurement for the `NUBEAM` deposition in figure 4.11b passes. This result is very similar to that reported in [145], where the same difference in high-energy anomalous diffusion between these codes was used to calculate the corresponding synthetic neutron emission rates, as detected by the MAST neutron camera system - though this difference is too small to be resolved this way (see figure 11a of [145]). Hence it is unlikely that any of the differences in anomalous diffusion measured here would be resolvable by the MAST neutron camera system. Finally, it is observed that `NUBEAM` decisively matches `LOCUST`.

4.2 Validation with FIDA measurements

To validate that these comparisons for MAST are realistic, the distribution functions calculated from the `NUBEAM` deposition by `LOCUST` (GC) and `NUBEAM` were fed into `FIDASIM` to generate synthetic FIDA measurements. The predicted and total measured signals for this spectral range are shown in figure 4.12a. Signals within a 660.5-661.5nm gate filter are shown radially resolved in figure 4.12b, along with each predicted signal from `FIDASIM`.

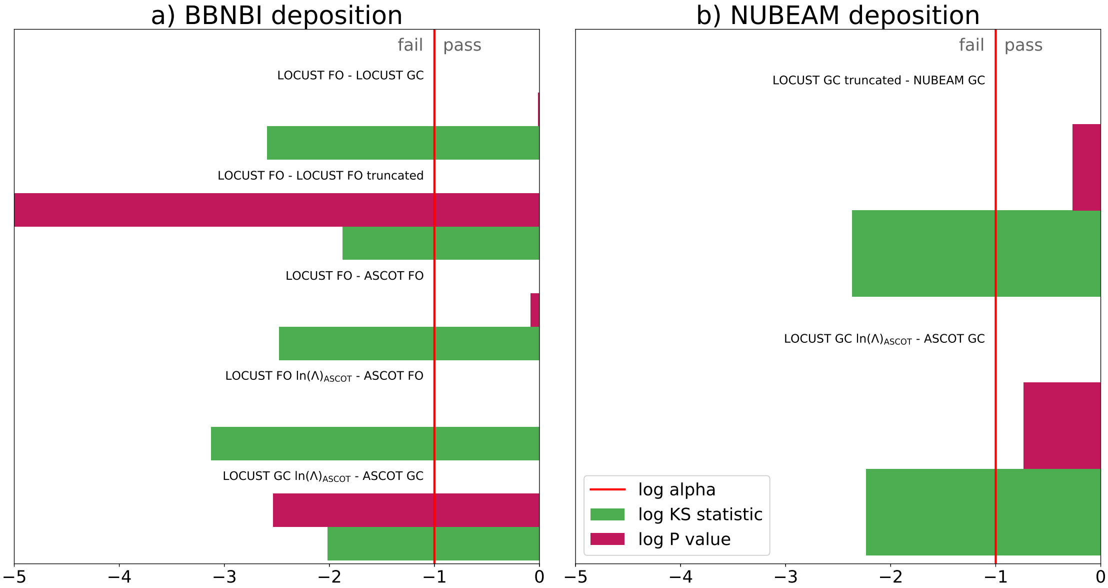

Figure 4.11: Kolmogorov-Smirnov test statistics (in green) for combinations of distribution functions shown in figures 4.9 and 4.10. The associated P=$P(D_{\text{measured}} < D)$ values are shown in red, with the arbitrary test pass-fail boundary α shown vertically in bright red. To pass the test, P must be lower than α, as demonstrated by comparing GC and FO `LOCUST` simulations. Figure reproduced from [109].

Within the core plasma, signals from `LOCUST` and `NUBEAM` are within the smallest error bars of each other so as to be indistinguishable by the FIDA diagnostic. Despite the gate, the presence of background bremsstrahlung emission is still observed outside of $R_{\text{major}} = 1.25$m and the LCFS in figure 4.12b. Current error estimates do not take background light into account, or even the lack of a time-dependent background plasma, but if the FIDA measurement errors are increased globally by $\sim 70\%$ then the reduced χ^2 approaches unity for data points within $R_{\text{major}} < 1.25$m.

4.3 Verification of non-axisymmetric simulations

To test the implementation of 3D fields in `LOCUST`, the DIII-D discharge #157418 was modelled again using the same input data from [81]. The corresponding results from `LOCUST` were compared to the simulations performed in [81] by the `SPIRAL` code [148].

The fundamental harmonic of the perturbed field was calculated by the 3D nonlinear MHD code `M3D-C1` [126] at 3.7mm, which is well within the precision required for convergence as shown later in figure 4.19. Because `LOCUST` uses the same deposition as `SPIRAL`, which

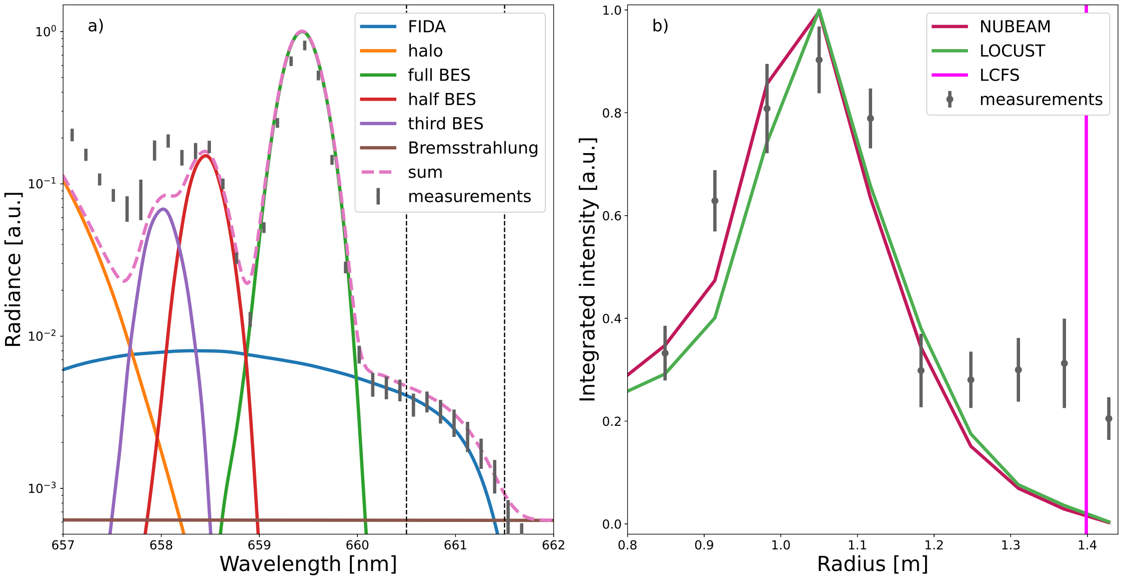

Figure 4.12:
a) Theoretical and measured signal intensities in MAST shot #29034 at 300ms for one channel. The wavelength filter denoted by vertical dashed lines encompasses 660.5-661.5nm and is used to integrate all signals across all channels to produce plot b). Some signals are still left out near 660nm to avoid integration of beam emission in other channels.
b) Radially resolved FIDA signal measured in MAST and as produced by `FIDASIM` for `NUBEAM` and `LOCUST`. An FLR correction was applied post-simulation. Figure reproduced from [109].

aims to be realistic, it includes both co and counter-injected beams and all three beam energy components. Furthermore, the deposition also accounts for the effects of the plasma displacement on the temperature and density profiles due to the external RMPs.

Shown in figures 4.13 and 4.14 are the differences in the steady-state fast-ion distribution f caused by the addition of the 3D field to the background equilibrium at the orbit-following stage (the neutral beam deposition was not altered). Only the distribution function at the plasma edge, where $\kappa > 0.77$, is shown, where κ in this context represents the square root of the normalised toroidal flux. This window is different to that used in [81], where $\kappa > 0.7$. However, as the `LOCUST` distribution function is mapped onto finite, rectilinear bins, any value of κ chosen within one `LOCUST` bin width of $\kappa = 0.7$, including $\kappa > 0.77$, will be sensible, as the comparison is limited by grid resolution.

Figure 4.13 shows this difference, $f_{3\mathrm{D}}(\epsilon, \lambda) - f_{2\mathrm{D}}(\epsilon, \lambda)$. There is good qualitative agreement between the codes, with both showing that the 3D field mostly affects edge-localised, co-passing particles (where particle pitch is typically between 0.5 and unity). Of these particles, those with energies above 40keV, half the injection energy, do behave somewhat differently

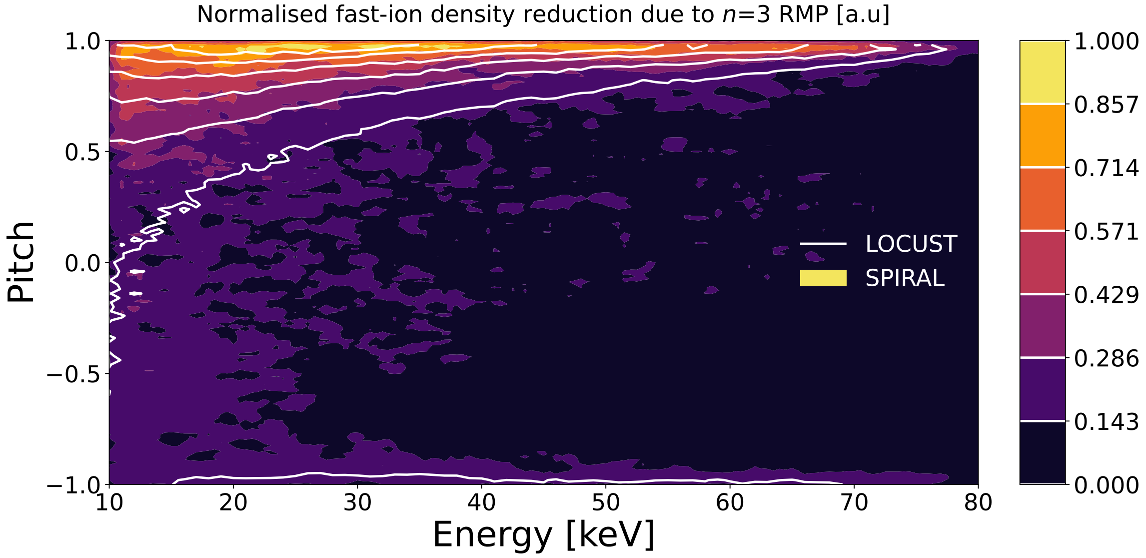

Figure 4.13: The normalised, element-wise difference between $f(\epsilon, \lambda)$ calculated with and without the RMP field, where $f(\epsilon, \lambda)$ is the distribution function integrated over all dimensions except energy and pitch. Filled contours are calculated by `SPIRAL` whilst the intermediate white contours by `LOCUST`. Figure reproduced from [60].

between the simulations. Figure 4.14 illustrates this more clearly, showing the RMP-induced transport as a fraction of the axisymmetric distribution function, $(f_{3\mathrm{D}}(\lambda) - f_{2\mathrm{D}}(\lambda))/f_{2\mathrm{D}}(\lambda)$. The effect of these high-energy markers is a 10% point difference in the fractional transport within $0.8 < \lambda < 1.0$, 40keV $< \epsilon$. This could be because `SPIRAL`, unlike `LOCUST`, uses an electron density profile that extends beyond the last-closed flux surface, causing additional transport of high-energy markers. Likewise, `LOCUST` did not include effects due to plasma rotation, impurities or electric fields.

There are a few reasons as to why this difference may appear. Firstly, the gap between $\kappa = 0.7 - 0.77$ is a small volume in real-space, comprising roughly one bin width in `LOCUST`, which increases the effects of Monte Carlo noise. Meanwhile, the implementation of the κ boundary in `LOCUST` is 0D, with rejection of rectilinear real-space distribution function bins based only on their centres, which introduces another source of error of the order one bin width. Another source for the additional transport in `SPIRAL` could be its use of plasma profiles that extend past the last-closed flux surface, which were not used in `LOCUST`. Additionally, `LOCUST` did not account for any impurities or plasma rotation. Nevertheless, these results provide significant evidence that the jump from 2D to 3D simulations in `LOCUST` has not introduced errors, and that the 3D field model is implemented correctly.

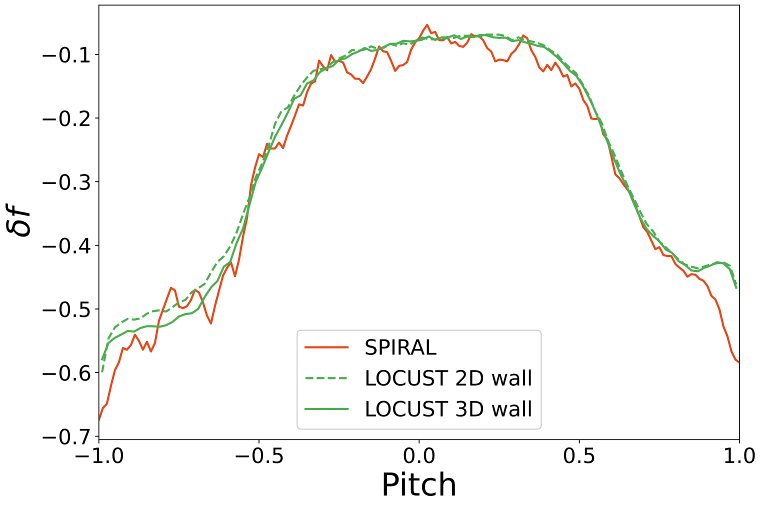

Figure 4.14: Local, fractional reduction of NBI fast-ion density due to RMPs in DIII-D #29034, defined by $\delta f \equiv (f_{3D}(\lambda) - f_{2D}(\lambda))/f_{2D}(\lambda)$. For comparison, the losses when `LOCUST` used a 2D PFC outline and fully 3D mesh as tokamak wall models are shown. The difference between `LOCUST` and `SPIRAL` above 0.8 is solely due to particles with energies above half the NBI injection energy. Figure reproduced from [60].

4.4 `LOCUST-GPU` in a workflow

After studying and establishing the accuracy of `LOCUST`, a software framework, `LOCUST_IO`, was constructed to enable the study of ITER scenarios at scale. This framework unifies the tools described in section 3.3 and solves a number of problems, including:

- Coupling `LOCUST` to the input data sources described in chapter 3.3. These sources comprise physics codes, databases and other calculated data, and they transfer data amongst each other and `LOCUST` via many un-standardised formats. As a result, the framework was initially built upon code made to transfer data between the physics components studied in the verification validation chapters above. However, functionality was extended to allow for conversion between an arbitrary number of data formats.

- Running `LOCUST` at scale. As mentioned in chapter 3.2, a typical run cycle of `LOCUST` involves many steps, often requiring users to manually remove out-dated caches or edit and re-compile source code by hand. This latter requirement for dynamic re-compilation is incompatible with the `IMAS` ecosystem, which relies on physics codes being installed as system-wide modules. Whilst the time taken for this overhead is less of a problem for extremely lengthy runs, it becomes a limiting factor in situations requiring many shorter runs - such as parameter scans. This process is also prone to human error and requires wasted effort to reproduce. Therefore, to be able to fully leverage `LOCUST`'s speed in all situations, functionality was required to abstract users away from the `LOCUST` run cycle by automating and combining the steps involved, allowing the user to checkout, edit, compile and run `LOCUST` with a single command. The ability to leverage batch systems to automate the execution of this process at scale was also required - something which is not native within `IMAS`. Additional steps to retain run settings, copies of edited source code files and the Git commit hash were also desirable to improve reproducibility and provenance.

- Dynamically running physics and pre-processing codes to generate new data. In some cases, it is beneficial to regenerate input data for `LOCUST` by re-running the original source code that produced it. For example, when performing convergence tests, data may need to be regenerated at various levels of fidelity. Combined with the need to automate `LOCUST`, this meant the framework was extended to execute arbitrary sets of ordered, user-defined tasks which comprise a 'workflow'.
- Running ITER simulations on different computing platforms. To leverage additional computing power, and to increase the portability of the framework to institutions lacking certain requirements, such as access to `IMAS`, the framework also abstracts away system environments and queue/batch systems.

Considering these requirements and features, `LOCUST_IO` was built in pure Python for flexibility, deployability, ease of use, accessibility as an open-source language, and its large, active community. This also enables `LOCUST_IO` to interface with the `IMAS` platform, leveraging its data model and physics codes, without being constrained by the limitations associated with `IMAS` - for example the lack of re-compilation or inefficiencies in the wall model (mentioned in section 3.3.5). Nevertheless, as part of this work, `LOCUST` was still adapted to `IMAS` to optionally read/write data directly to/from `IMAS` IDSs, allowing for future use of `LOCUST` in `IMAS` workflows. In this case, `LOCUST` would likely be provided as a pre-compiled actor.

At a fundamental level, `LOCUST_IO` is a set of Python classes that describe the data that are input to and output from `LOCUST`, each with methods for plotting, and reading and writing to various formats. Alongside these exist additional classes that encapsulate aspects of a simulation. The `Build` class encapsulates methods for editing and compiling pieces of external, non-Python source code from within Python, whilst the `Environment` class providing similar functionality for system environments, such as modules to load or environment variables to export. Instances of these classes are often manipulated by the `Workflow` class, which encapsulates an arbitrary set of user-defined tasks, for example executing shell commands, performing operations on `LOCUST_IO` objects within Python or even executing other `Workflows`. As an example, as one of its tasks, a `Workflow` may instantiate a `Build` object to describe a piece of `FORTRAN` code, providing it with an `Environment` instance that describes the compile-time system environment, before calling the `Build.make()` method to compile the code. Finally, the `Batch` class enables the execution of many `Workflow` objects by providing methods for creating and submitting batch scripts in parallel or running each `Workflow` interactively in serial. The hierarchical interaction and possession amongst these classes for a typical group of simulations is illustrated in figure 4.15, with classes described by a box possessing anything within the box perimeter. In brief, the `Batch` instance possesses all of the sets of arguments required for all simulations, and supplies each set of arguments to individual `Workflow` instances describing a single simulation. These `Workflows` may then possess additional, arbitrarily nested `Workflows`, each optionally possessing run-time `Environments` and `Build` objects with compile-time `Environments`.

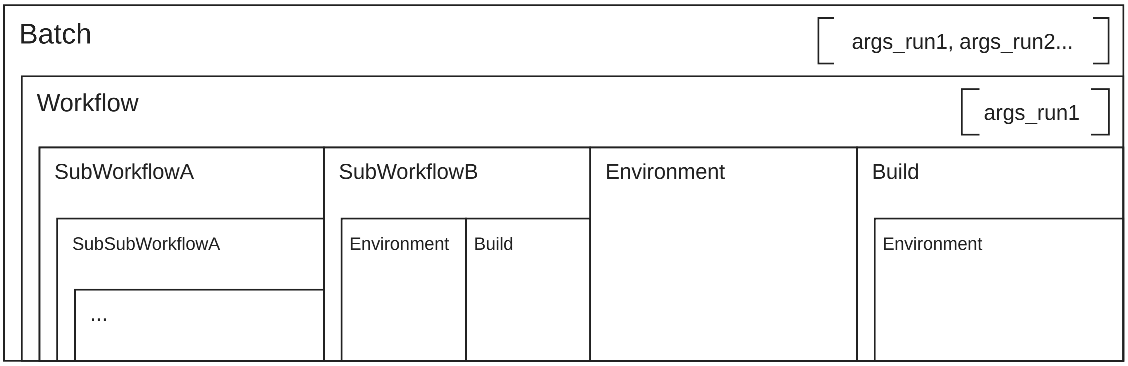

Figure 4.15: Scope of simulation classes within `LOCUST_IO`. Box boundaries represent possession, for example an instance of the `Batch` class possesses an instance of the `Workflow` class. Also shown are the list of sets of arguments [`args_run1`,...] possessed by `Batch`, each element of which is used to instantiate a `Workflow` that possesses set of `args_run` as member data. Many `Workflows`, one for each `args_run`, are executed this way. A `Workflow` itself can also contain other workflows, and `Builds` and `Environments`.

The specific `Workflow` created for this study is shown in figure 4.16, with some components comprising other, nested `Workflows` - for example the `LOCUST` stage comprises its own `Workflow` which moves input data, checks-out, edits, compiles and runs the code before retrieving the output data and cleaning up. Prior to this is a stage which pre-processes data from the `MARS-F` database. Whilst this stage is necessary to obtain perturbation data on the fourier-decomposed R-Z grid expected by `LOCUST`, or could be avoided by reading from the same pre-computed R-Z grid, it adds extra functionality, for example control over the perturbation grid precision to establish a trade-off between accuracy and performance. As mentioned in chapter 3.2, using a perturbation grid that is precise enough to allow for linear interpolation drastically increases the speed of `LOCUST`. It has also been shown [104], and will be shown later, that convergence tests are necessary to avoid erroneous losses when tracking markers in 3D fields. However, these tests are rarely observed in the literature, with some [86] using grid spacings larger than the acceptable levels found later in chapter 4.5. In addition, many data conversion and preparation stages are included in the `Workflow`, for example converting various different formats to that accepted by `LOCUST` as well as generating IDSs for `BBNBI` which are composed of combinations of other IDSs: the NBI geometry from the ITER machine description database (or generated by `LOCUST_IO` in the case of the diagnostic system) and the plasma data from the `IMAS`-stored `ASTRA` database.

To deploy, prepare and execute the workflow shown in figure 4.16 to study fast-ion transport in ITER, a user must do the following:

1. Clone `LOCUST_IO` from the `IMAS` Git repository.

2. Checkout the `ITER_RMP_study_MARSF` branch. This contains templates for executing ITER RMP studies, as well as the launch and plot scripts used in this work.

3. Verify that the correct environment is enabled in the settings - otherwise adding the environment to the `Batch` and `Environment` classes if it does not exist.

4. Copy any required input data to the *input_files/* directory, for example the `MARS-F` 3D field database and the ITER PFC mesh.

5. Create a (or use the included template) run script which defines a `Workflow` sub-class. The class should define individual functions which execute the individual tasks required to perform the desired simulation - for example retrieving input data, executing `LOCUST` or automatically created plots of input or output data.

6. Create a (or use the included template) launch script that defines the arguments passed to the `Workflow` before launching it in interactive or batch mode. These arguments can be used to control the run stages that the `Workflow` executes - for example whether to

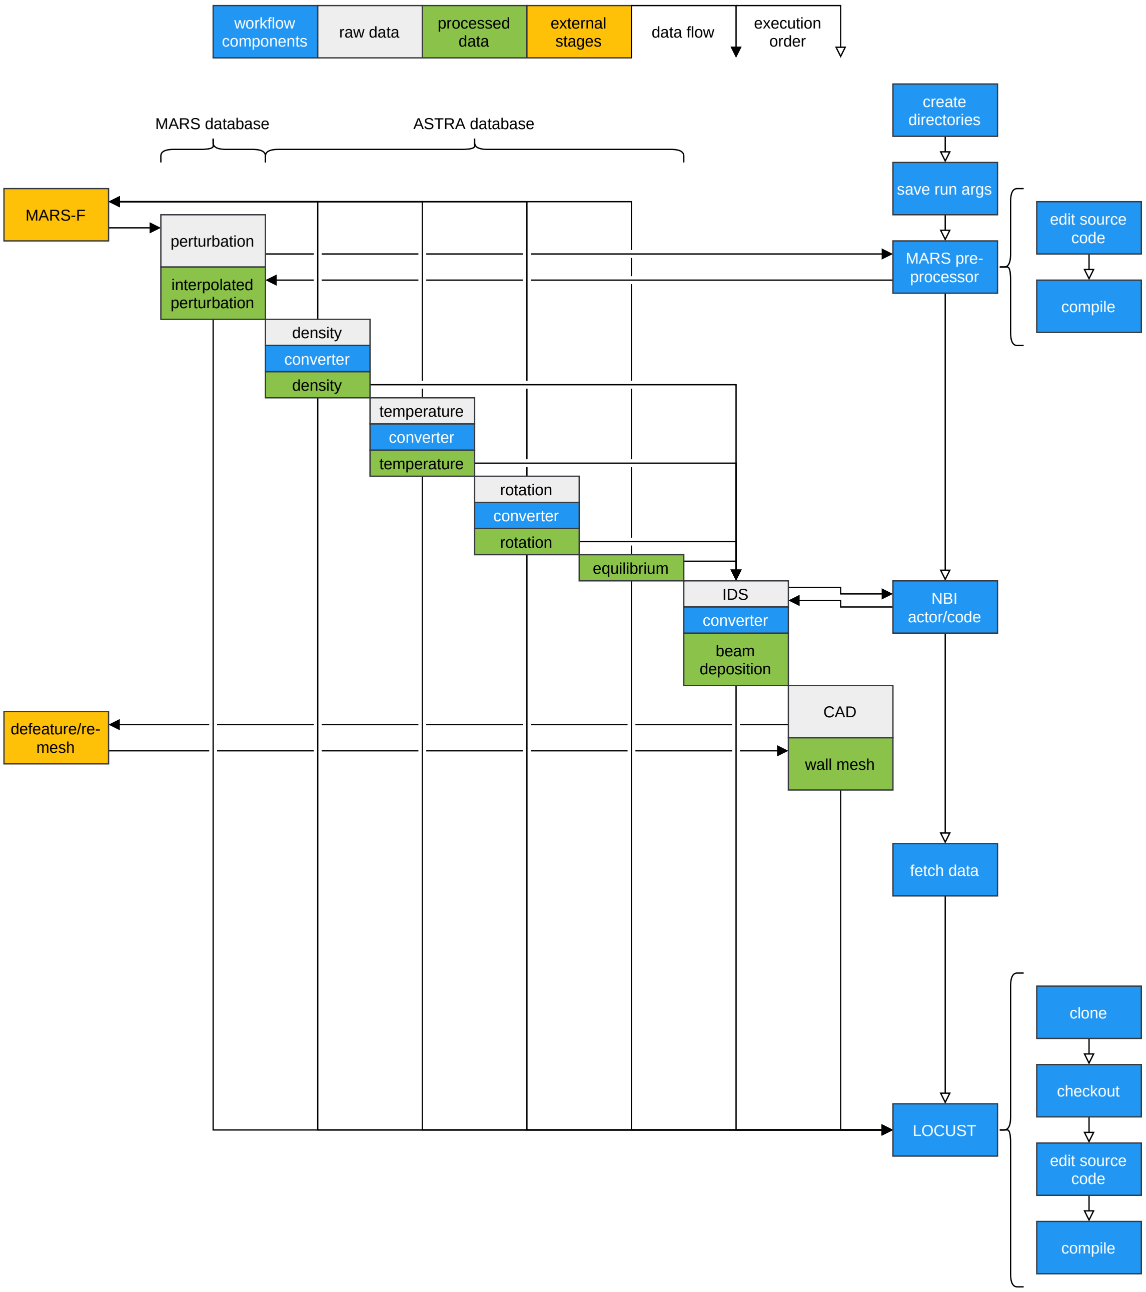

Figure 4.16: Illustration of top-level RMP workflow components (see key at top), along with the data that they operate on and the external components that produce and process that data. Arrows represent data flow, and the workflow components execute from top to bottom. Not shown here are workflow sub-components and every physics code used to produce the raw data.

skip 3D field generation or even running `LOCUST` at all. To perform many simulations, multiple sets of arguments should be created and passed to the `Workflow`. Therefore, parameter scans can be over any arbitrary argument required by the `Workflow`, and can be physics-related, for example the phase of the perturbation generated by the 3D field pre-processor, or arbitrary, such as the resolution of the 3D field grid or the filename of a given output.

7. Execute the launch script with Python. In batch mode, this will submit a job to the batch queue for every parameter set passed to the `Workflow` instance in the launch script. A user may wish to set timers or stagger launches within Python.

8. Optionally define plot scripts. A template module exists containing helper functions for the parallelised reading and plotting of input/output files. These can be called by simply importing the launch script as a Python module and passing that module to a given helper function. All the inputs or outputs of a given type, for example the distribution function, generated by the group of simulations are returned as a list of `LOCUST_IO` objects, which can be manipulated using the `LOCUST_IO` library.

4.5 Workflow tests

Convergence tests were performed to both ensure the workflow is functioning correctly and determine optimal runtime settings for optimising accuracy and performance. First, the 3D field functionality was tested to ensure that the data were manipulated and combined correctly by the pre-processor and read, interpreted and evaluated correctly by `LOCUST`. Figure 4.17 below shows the vacuum perturbation components created by an $n = 3$ RMP in $Q = 10$ as evaluated by `LOCUST` within the workflow after pre-processing and reading from the `MARS-F` database. These components are compared to equivalents calculated by the `probe_g` Biot-Savart code and a Python script which reads the post-processed files before they are read by `LOCUST`. By choosing a single point in the poloidal plane and evaluating the perturbation toroidally, the field is shown to be correctly oriented and to compare well with the field produced by `probe_g`. To show that the necessary phase shifts are implemented correctly across toroidal harmonics, which often counter-rotate, as is the case here, the test is repeated for the second harmonic, $n_1 = 6$, in figure 4.18.

The 3D field model is now tested for convergence. As magnetic fields obey the divergence

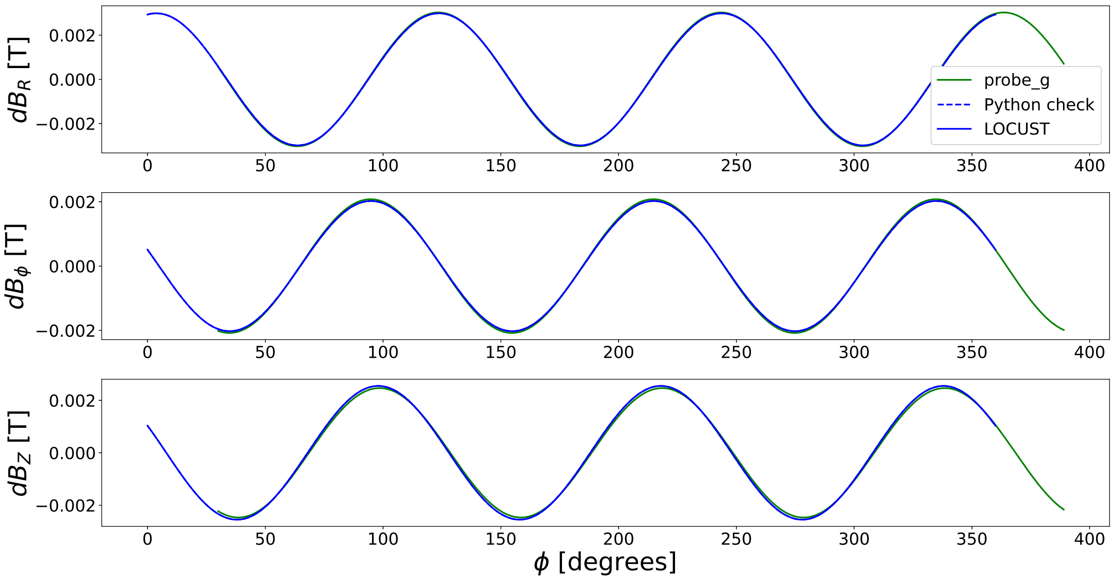

Figure 4.17: Components of the 3D vacuum field in $Q = 10$ evaluated at a single point in the poloidal plane ($R = 6.22$m, $Z = 0.06$m) over the toroidal angle ϕ. ECC current amplitude was set to 90kAt and the upper, middle and lower coil row phases to 86°, 0° and 34° respectively. The components evaluated by `LOCUST` within the workflow, `probe_g` and a Python script operating on the post-processed 3D fields are shown. The results from the Python script and `LOCUST` are indistinguishable here. A 30° shift is required to map the `probe_g` coordinate system onto that used by the other components. The overlapping of each of the components show that the field used by `LOCUST` is scaled, phase-shifted and oriented correctly.

free condition, $\nabla \cdot \mathbf{B} = 0$, its calculation can be used to gauge errors inherent within a numerical model of a magnetic field - in this case the 3D field. Figure 4.19 below plots the average $\nabla \cdot \mathbf{B}$ over all points within the plasma for $Q = 10$ when the first two harmonics are present, as the perturbation grid spacing is changed, before being scaled by one of two lengths: B on-axis and either the Boris step length at the injection energy or the perturbation grid spacing, the latter of which should give an estimate of the global truncation error. As the grid size is reduced, the global error decreases like $\mathcal{O}(\Delta x_{\tilde{\mathbf{B}}})$.

Whilst minimising the grid spacing is shown to maximise accuracy, the minimum acceptable precision is desired to maximise the available on-board GPGPU memory. Therefore, the corresponding NBI power loss was also calculated and plotted in figure 4.19. Convergence begins below 10cm, however it takes until ≈3cm for fast-ion losses to stabilise. Therefore, a grid size of 1cm was chosen as a suitable spacing for all proceeding simulations performed in chapter 5.

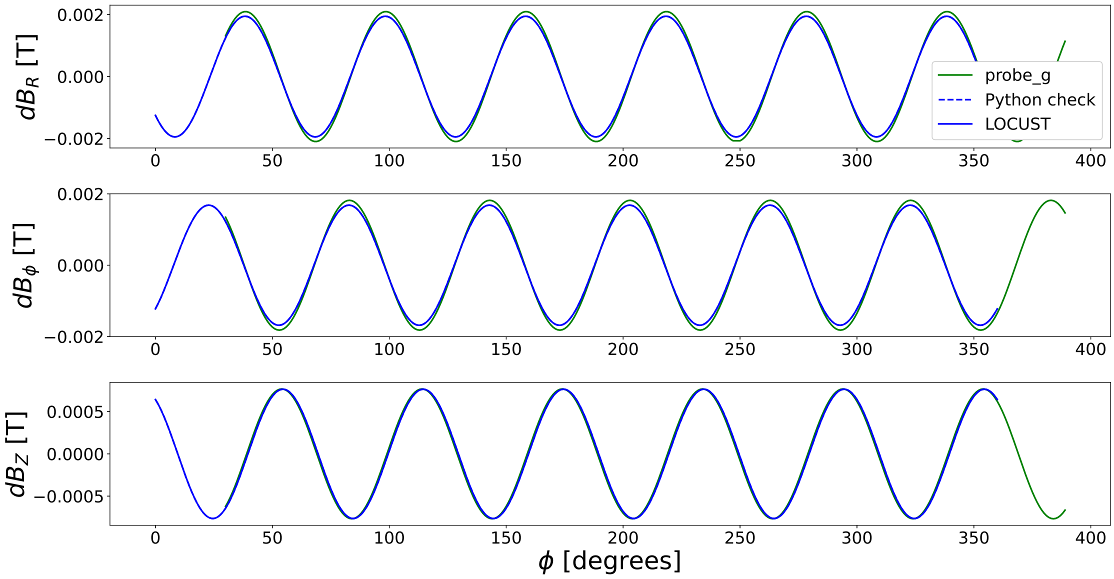

Figure 4.18: The second harmonic components corresponding to the field evaluated in figure 4.17 above. Again, the results from the Python script and `LOCUST` are indistinguishable here. The agreement here also shows that the phase-shifting of counter-rotating harmonics is calculated correctly.

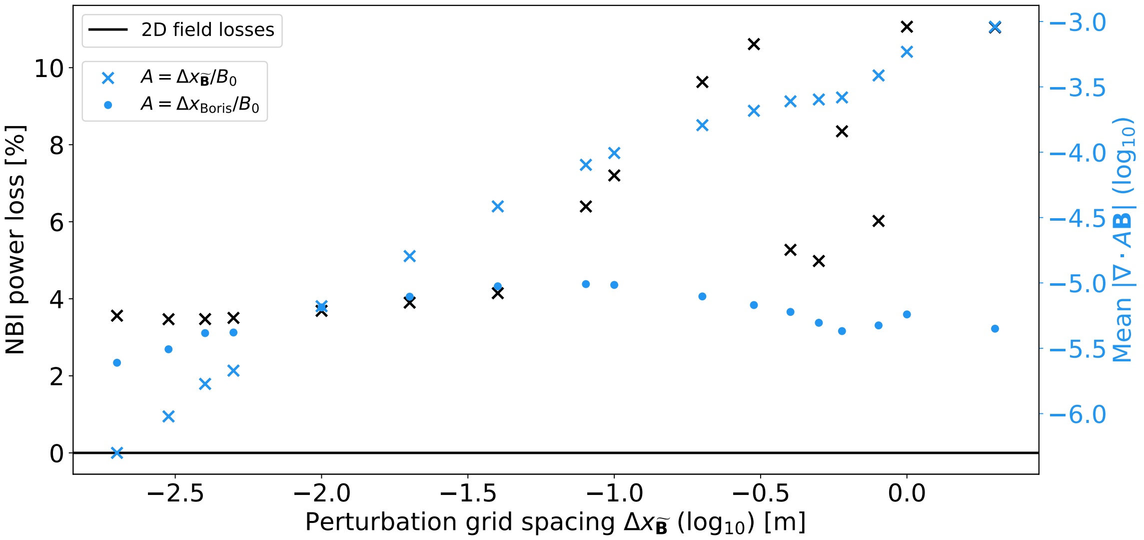

Figure 4.19: Perturbation grid spacing convergence test measuring global fast-ion losses, as a percentage of deposited beam power, along with the average magnetic divergence, $\nabla \cdot \mathbf{B}$, scaled according to two length scales: the perturbation grid spacing and a typical marker step length. The fundamental and second harmonic are included. To show the underlying 2D equilibrium is sufficiently resolved, the global losses in the axisymmetric field, which are expected to be negligible, are also shown. The beginnings of convergence in global truncation error and particle losses coincide at approximately 10cm, with both decreasing approximately linearly. Losses saturate at roughly 3cm; to give some margin for error, a grid spacing of 1cm subsequently then chosen for all simulations. Figure reproduced from [60].

Convergence tests were then performed to maximise the workflow speed and accuracy for general runtime parameters, such as fast-ion movement timestep and total number of required Monte Carlo markers. Figure 4.20 shows the measured fast-ion loss power (in arbitrary units) as a function of loss time and fast-ion energy, for different values of timestep; markers took $10 \leq$ `UNBOR` ≤ 200 Boris timesteps between the application of a particle collision every 0.1μs. `UNBOR`=100, which corresponds to a Boris timestep of 1ns, is used for the rest of the study. This is approximately 25 steps of 1cm per gyroperiod for a 1MeV deuteron in the ITER baseline scenario. The minimum number of Monte Carlo markers was then determined by measuring the global and time-resolved loss power. Figure 4.21 shows this as the number of `CUDA` threads is varied. It was concluded that 2^{15} markers led to the absolute shortest simulation time whilst still allowing for accurate estimation of global fast-ion losses and resolving of prompt losses. Nevertheless, values greater than 2^{13} were found to be sufficient for estimating global losses.

The impact of simplifying the 3D field and plasma models was also explored. Because the 3D field and plasma collisions must be evaluated for every marker periodically, a significant amount of execution time is dedicated to them. In the case of the 3D field, the required grid precision also demands a significant amount of storage.

The effects of limiting the number of toroidal harmonics to the fundamental only are shown in figure 4.22 below for two types of toroidal mode spectra. Shown is a Poincare plot of the intersections made by field lines as they pass through the $\phi = 60°$ poloidal plane, before and after the second harmonic is added. In the case where only one harmonic is present, $n = n_0$, the harmonic amplitude is scaled up to counteract the removal of additional harmonics so as to represent a system capable of generating pure toroidal fields. Whilst the penetration of the RMP does not seem to increase after the addition of the second harmonic, changes appear in the arrangement and sparsity of field lines toward the plasma edge, at $q < 2.5$. For example, in the $n_0 = 3$ case, the second harmonic can be seen to introduce a new island chain near $q \approx 8/3$, outside of which the field line density changes dramatically - signalling a change in the field line diffusion rate. Conversely, the second harmonic in the $n_0 = 4$ case leads to a stochastisation of the flux surfaces near $10/4 < q < 11/4$, especially near the mid-plane at $0° < \theta < 50°$ (possibly due to resonance on the 13/5 surface in between the 10/4 and 11/4 surfaces, causing some overlap). To see whether it has any effect on the fast-ion dynamics, the probability of losing a marker born at the plasma edge was calculated

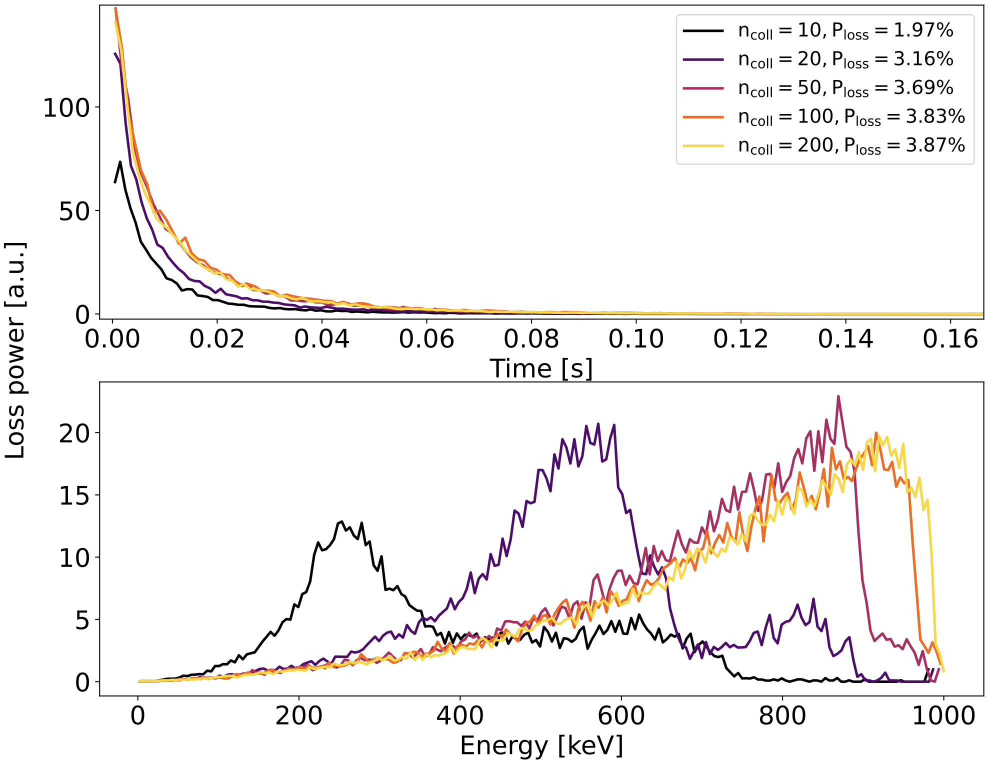

Figure 4.20: Fast-ion loss power measured against marker energy for various values of the `UNBOR` timestep setting in `LOCUST`. `UNBOR`=100 represents 100 Boris pushes between collisions ($n_{\text{coll}} = 100$ in the figure legend) each with a timestep of $\tau_{\text{master}}/$`UNBOR` $= 1$ns, where τ_{master} is the collisional timestep and is usually chosen to be 0.1μs in `LOCUST`. `UNBOR`=100 was deemed satisfactory for use in all simulations proceeding this test.

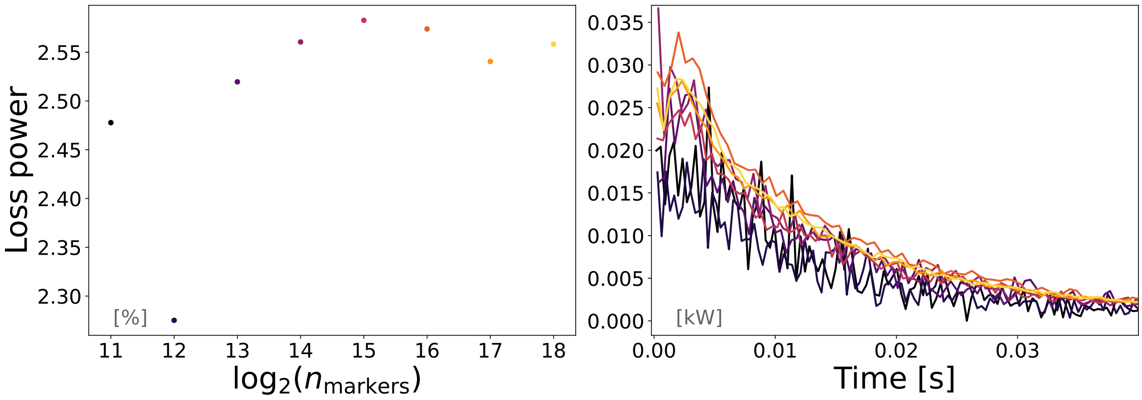

Figure 4.21: Measured total (left) and time-resolved (right) fast-ion loss power to PFCs using various marker population sizes. Computational speed-up ceases below 2^{15} markers, hence this value is chosen when running the workflow to estimate global losses.

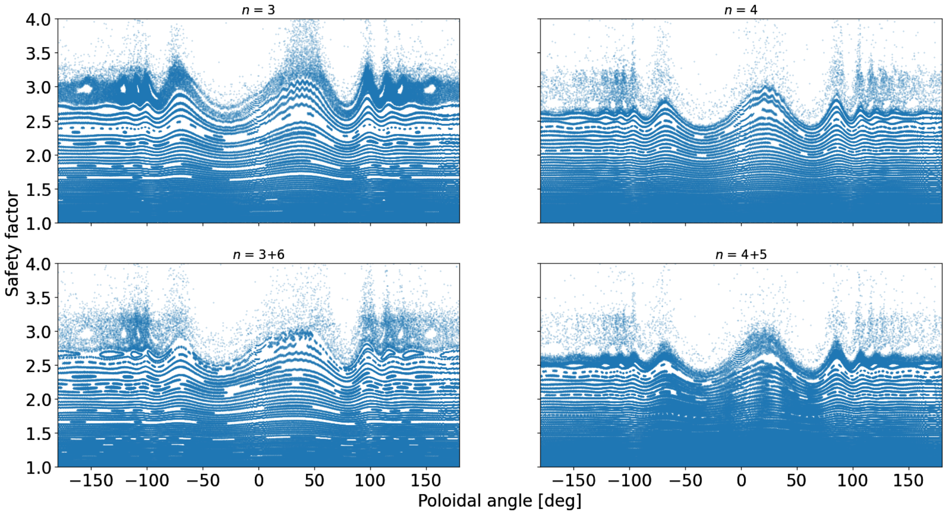

Figure 4.22: Poincare maps for fields at the same Φ_{m} but different toroidal mode numbers and spectra, plotted against safety factor q and poloidal angle θ. The maps are evaluated at $\phi = 60°$. The patterns near $\theta = 0°$ are due to the transform from rectilinear coordinates. Fields with one harmonic have been artificially scaled up to correspond to an ECC system capable of creating pure-n_0 fields. In this case, the RMP penetration is not increased by the secondary harmonic, but it is possible to see that it still changes the field structure at the edge. Figure reproduced from [60].

with and without the second harmonic. Plasma collisions were also toggled. Loss probability was calculated over a number of perturbation phases to examine whether the effects of the changing RMP field were also reflected in the fast-ion dynamics. The upper and lower limits of this loss probability are plotted in 4.23 against marker safety factor, q, and pitch, λ, at the point of deposition ($t = 0$). The peak in deposition occurs at $\lambda \approx 0.6$. The RMP can be clearly seen to have an effect on edge confinement - as it should. One can also see that the addition of the first toroidal harmonic is extremely important, as it acts to decrease the loss probability throughout, regardless of plasma collisions or toroidal mode number. Meanwhile, collisions are also shown to change the average loss probability - especially at lower pitches. In addition, the large changes in loss probability gradient observed in the collisionless case, likely due to topology changes in the 3D field, are smoothed out by the addition of collisions. Therefore, it was concluded that both the second harmonic and plasma collisions were too important to neglect in the model.

Because the plasma model was shown to be vital, tests were then performed to see if it could be simplified. As the plasma background model in `LOCUST` is basic, and cannot

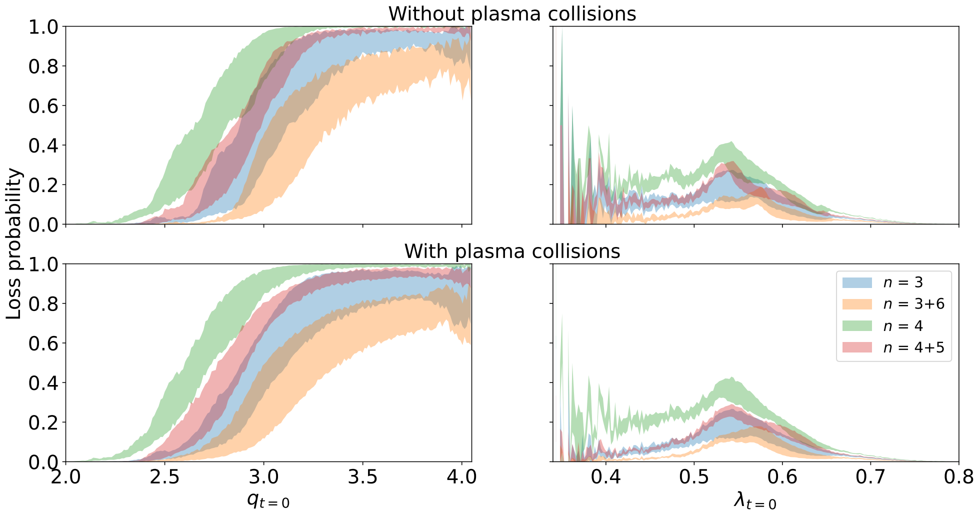

Figure 4.23: Fast-ion marker loss probability plotted against starting safety factor, q, and pitch, λ, (left and right respectively) for $n_0 = 3$ and $n_0 = 4$ RMP fields with and without a second harmonic. Collisions are toggled between the subfigure rows. The influence of including both the second harmonic and plasma collisions was deemed too important to leave out of the final model.

prescribe density profiles for multiple species, tests were performed to determine whether minority species within the plasma affected the measured fast-ion losses. Figure 4.24 below shows the expected rate of energy loss for fast ions moving in a homogeneous plasma with species densities set to those in the $Q = 10$ database. As can be seen, high-Z impurities in low concentrations, such as tungsten, have little effect on the fast-ion slowing down rate. Lighter elements, such as helium, have a much more pronounced effect - however appear similar to hydrogen isotopes at most energies. That is, if scaled so each species possessed the same number of protons in total, the drag rates due to hydrogen and helium can only be distinguished near to the thermalisation energy, making them essentially indistinguishable from the point of view of fast ions. To show that this is the case for lost fast-ions, the fast-ion loss power to the PFCs is plotted in figure 4.25 against fast-ion energy after impurities are included and removed from the $Q = 10$ plasma. Whilst impurities lead to an overall reduction in fast-ion density, by approximately 4%, no effect on fast-ion losses can be seen. Hence they are ignored in this work. Instead, the plasma in `LOCUST` is assumed to consist of a hydrogen isotopes, whose density profile is the same as the electron density profile but linearly re-scaled to satisfy charge neutrality.

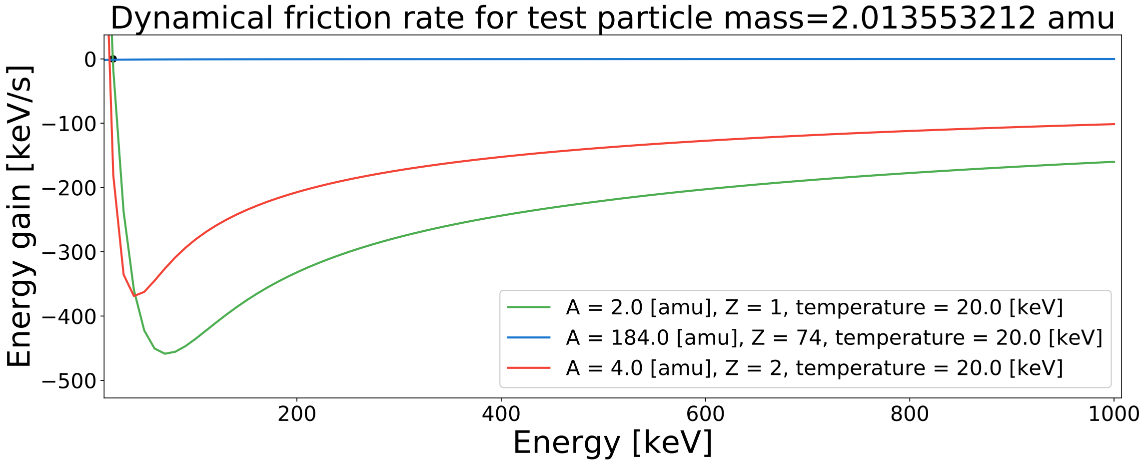

Figure 4.24: Dynamical friction rate is shown here for a deuteron moving in a homogeneous plasma with densities corresponding to the $Q = 10$ ITER baseline scenario. If the densities of each minority species are scaled to contain the same total number of protons, then at most energies the friction due to helium is indistinguishable from that due to deuterium, showing that they can be interchanged from the point of view of fast ions.

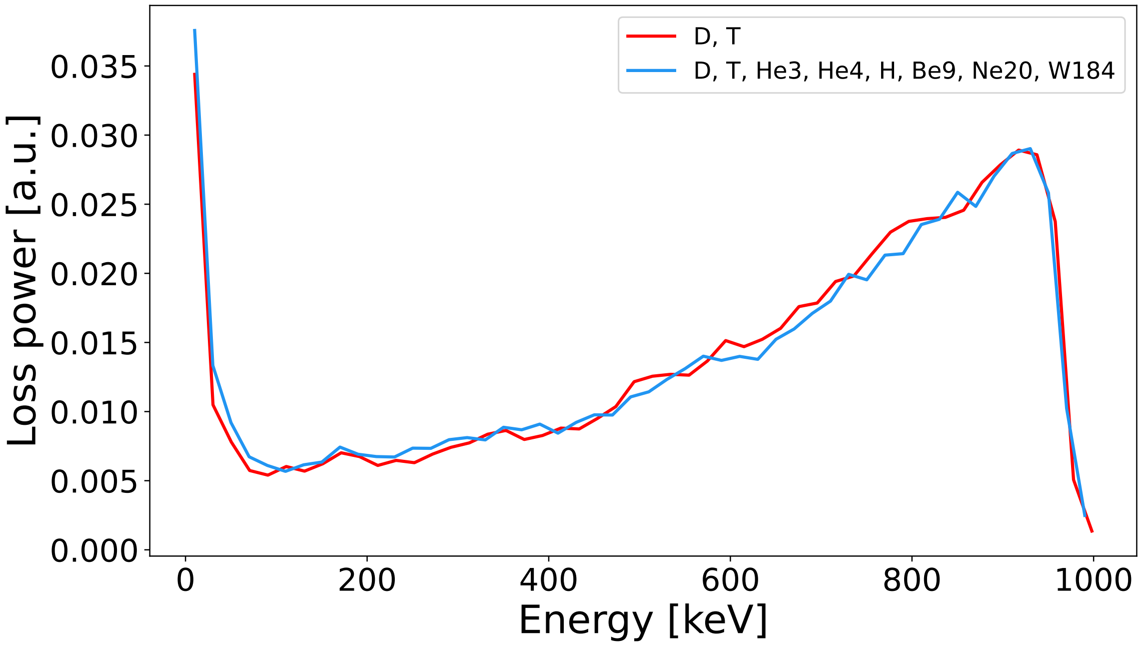

Figure 4.25: Fast-ion loss power in the $Q = 10$ scenario measured as a function of fast-ion energy with and without impurities. The isotopes of each impurity species are shown in the legend.

4.6 Summary

The implementation of `LOCUST` was validated against other fast-ion codes, and its physical models verified against experiment, before a software workflow system was constructed around the code to enable large-scale ITER simulations.

It was shown that, for a range of scenarios, the differences between `LOCUST` and well-established fast-ion codes can be reduced to within expected Monte Carlo fluctuations. This instills confidence that the underlying physical models are correctly implemented in `LOCUST`. Moving forward, it can therefore be expected that the accuracy of `LOCUST` is fundamentally limited by theoretical uncertainties in the physical models - the best any code can hope to achieve, given the constraints of experimental diagnostics and theoretical understanding at the time. Indeed, in comparing synthetic FIDA signals from `LOCUST` and `NUBEAM` to those measured on MAST, it was shown that differences between the codes are smaller than the uncertainties of the diagnostic.

Nevertheless, by isolating and sufficiently resolving the foundational models within each fast-ion code, it was possible to uncover more significant differences that dominate local uncertainties due to the interplay between numerical techniques, approximations, and other assumptions implemented at the fundamental level. Though they are to be expected, it is important to note that the magnitude of the differences observed here should call into question the validity of certain codes in particular regimes. As an important example, the above analyses with co-injected beams provided evidence that the GC and FO formulations are statistically equivalent in spherical devices - something previously debated [144]. But as figure 4.6 shows, the addition of a plasma boundary significantly reduces the performance of GC codes at the edge for specific FLR models - even in conventional tokamaks. Though the lack of coupling between edge and core fast-ions means errors remain localised within $\sim \rho$, it can still span significant plasma length scales - especially at the edge of spherical devices. Indeed, recently it was independently shown [145] that the interplay between edge neutrals and FLR effects at the plasma boundary in a spherical tokamak was a possible explanation for the neutron deficit measured in MAST across some years [149]. For these reasons, extreme caution should be used when applying GC models to spherical tokamaks, calculating PFC heat loads or modelling dynamics at the edge.

Whilst this example demonstrates the general importance of understanding the interference between fundamental and higher-order physical models, this work also generated guidance on how to create tests rigorous enough to resolve this interference in fast-ion models. Firstly, the required level of fidelity was found for each code, with typical `TRANSP` settings proving inadequate for providing the necessary statistics . The settings used are shown in table 1. It was also found that discharges with counter-current NBIs create optimal conditions for validating wiith experiment, as they naturally lead to fast-ions occupying a wider range of phase-space, and are more likely to test the implementation of the plasma boundary and wall models.

Chapter 5

Fast-ion transport in ITER

The ITER tokamak aims to perform discharges of increasing toroidal field, current and, ultimately, fusion power [13]. This staged approach consists of two main phases: Pre-Fusion Power Operation (PFPO), in which operational protocols are developed and initial, low-power experiments are undertaken, and Fusion Power Operation (FPO), during which high-power DT experiments take place.

PFPO is further split into PFPO1 and PFPO2. PFPO1 aims to develop the basic stages and protocols involved in ITER plasma scenarios (start-up, current ramp-up, flat-top and ramp-down). Plasmas during PFPO1 consist of hydrogen to limit neutron production and component activation, and the maximum plasma current reaches up to 5-7.5MA for H and L-mode respectively. ITER will operate in H-mode this way at reduced plasma current and toroidal field in order to lower the plasma stored energy to reduce the risks to the device PFCs from ELMs, VDEs and disruptions in general, including any subsequently generated runaway electrons [13]. But as various methods and systems for plasma control and disruption mitigation are implemented and commissioned, the field and current will gradually be increased. PFPO1 is ignored in this work as it operates without a functional HNB system. This system is commissioned and introduced in PFPO2, which includes HNB1, HNB2 and the DNB sub-systems operating at nominal power for periods of up to 100s. PFPO2 is designed to demonstrate that ITER systems can reliably operate in high-power scenarios. Experiments will be performed with hydrogen and/or helium plasmas, producing information on the confinement of fast ions and other impacts of fast particles on the performance in ITER-scale plasmas. The plasma current and field will reach 15MA and 5.3T respectively in L-mode whilst H-mode sees 7.5MA and 2.65T respectively. At this point, the ECC system

Table 5.1:
Parameters describing various ITER operating scenarios used in `LOCUST` simulations. Shown are the magnetic field strength on-axis (B), the plasma current (I_p), ECC coil current amplitude (I_c) and the safety factor evaluated at the magnetic axis (q_0) and the magnetic flux surface at 95% the flux at the LCFS (q_{95}). Where no species is indicated, a deuterium neutral beam and plasma can be assumed. PFPO1 cases have not been modelled here due to their lack of an operational HNB system. In all cases, the NBI injected power is the full rated 33MW.

Case	2	3	5	7
Stage	PFPO2 or FPO	FPO	FPO	FPO
B [T]	2.65	5.3	5.3	4.5
I_p [MA]	7.5	7.5	15	7.5
q_0	1.3	1.0	1.0	1.4
q_{95}	3.2	6.2	3.1	5.3
Plasma/NBI species	H/H or D/D	D/D	DT/D	D/D
NBI energy [MeV]	0.87 or 1	1	1	1
I_c [kAt]	30-45	30-45	60-90	30-45

will operate with independent power supplies for each coil. Therefore, it is during PFPO2 that techniques to reduce ECC-driven fast-ion transport may be applied and verified.

FPO is similarly split. FPO1 acts to reproduce PFPO2 experiments using DD before DT fuels. H-mode operation will be gradually expanded up to 12.5-15MA of current and fields of 5.3T, whilst pulses reach 50s at $Q \geq 5$. As ITER moves toward FPO2 and FPO3, the discharges increase in duration and performance.

A selection of discharges from these stages has been previously studied with MHD codes and 3D fields calculated. These discharges are simulated in this chapter assuming the parameters given in table 5.1. The corresponding plasma profiles for each of these cases are shown in figure 5.1. Case 5 represents the $Q = 10$ ITER baseline scenario. Whilst the ECC system has a maximum achievable current of 90kAt, the range defined by the two values of I_c in table 5.1 is likely to be used. This is studied in section 5.2 as the basis for two methods of optimisation. In parts of this work, an alternative version of case 2 is studied in which a low-energy hydrogen neutral beam injects into a pure hydrogen plasma. Furthermore, the $Q = 5$ variation of the ITER baseline discharge is neglected due to the similarity of the plasma response to $Q = 10$ variation.

This chapter describes and analyses the primary results of fast-ion simulations of the aforementioned ITER operating scenarios, performed with the workflow assembled and tested in chapter 4. First, section 5.1 provides upper estimates for fast-ion transport and PFC power loads in the high priority $Q = 10$ ITER baseline scenario (case 5) where ELM suppression is

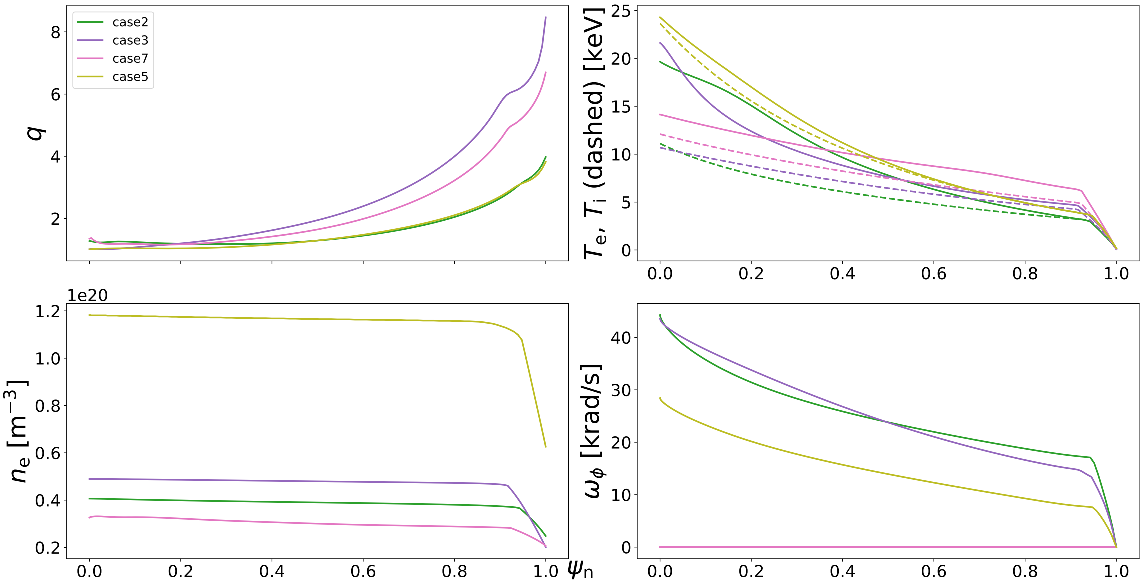

Figure 5.1: Safety factor (q), ion and electron temperature (T_i and T_e respectively), electron density (n_e) and rotation frequency (ω_ϕ) plotting against normalised poloidal flux (ψ_n) for the cases described in table 5.1 as calculated by `ASTRA` and used in this study by `LOCUST`. The high q-profile gradients near the plasma edge ($\psi_\mathrm{n} \approx 1$) will likely increase the stochastisation of the plasma edge by enabling islands, created by the RMP on adjacent rational flux surfaces, to overlap.

maximised. The impact of RMP rotation on the first wall power loads is also assessed. To avoid coupling to core MHD modes whilst ensuring sufficient edge penetration, ITER will likely operate $n_0 = 3$ and $n_0 = 4$ RMP waveforms. Experiments in DIII-D plasmas with ITER-similar shapes have shown these to successfully suppress ELMs in deuterium [150] and helium [151] plasmas. Hence both these mode numbers are studied and compared here, as well as the inclusion of the second harmonic. In section 5.2, fast-ion transport is studied in the more realistic scenario where the ECC system operates at reduced ELM suppression and density pump-out. Two methods for optimising fast-ion confinement in this situation are explored and compared across the cases listed in table 5.1, some of which are low-power compared to the ITER baseline scenario. The potential for using these discharges as low-risk test beds for optimising ECC operation prior to high-power ITER baseline discharges is assessed. Finally, section 5.3 determines whether diagnostic neutral beam confinement [152] may act as a suitable proxy for the equivalent HNB confinement, in order to determine the feasibility of estimating the impact of ECCs on HNB confinement via the sensing of DNB fast-ion losses. The DNB is first implemented in `IMAS` before simulations are performed for the optimum ECC operating strategies determined in the previous sections.

Parts of this chapter have been published in references [60] and [153]. All codes were operated by the author except `ASTRA` and `MARS-F`.

5.1 Fast-ion transport in the ITER baseline scenario

The aim of this section is to evaluate, from the point of view of fast-ion confinement, proposed methods for operating the ECC system whilst maintaining ELM suppression. In this situation, XPD is assumed to remain constant, as XPD is believed to be a reliable metric for ELM suppression (see section 2.3). To estimate the likely worst case scenario for fast-ion confinement, when ELM suppression is maximal, studies where XPD is maximised are performed here. This involves optimising and fixing the relative current phase between the upper and lower ECC rows relative to the middle row, $\Phi_{\mathrm{u,m}}$ and $\Phi_{\mathrm{l,m}}$, and maximising the current amplitude. The coil row current phases which maximise XPD are stated in table 5.2 for each case outlined in table 5.1. The remaining degree of freedom is the absolute phase of the current in each row - i.e. the rotation of the complete 3D field, Φ_{m}. Because XPD, and thus ELM suppression, is independent of Φ_{m}, in the limit where the perturbation spectrum consists of a single mode, continuously rotating the field is currently the proposed method of power load mitigation in ITER. Experiments suggest that rotation of RMP fields modulates the intensity of the fast-ion power flux on a given PFC - for instance, a fast-ion loss detector (FILD) [81, 138]. Whilst this suggests RMP phase may be used to control where the power flux lands, it may also modulate the global loss rate. Therefore, one may instead choose to optimise the RMP phase according to some combination of maximising total NBI heating efficiency and minimising peak PFC power fluxes to avoid localised heating. Either way, an optimal RMP phase could be determined and fixed. However, if large PFC power fluxes are unavoidable, then it may be beneficial to rotate the RMP in time to reduce the root mean square (RMS) power loads. The downside of this is the potential thermal cycling of components and reduction in RMS NBI efficiency. Ultimately, the best scenario is if both criteria can be optimised simultaneously - that is, if peak power flux correlates with global losses. Determining this is of key importance.

As the plasma response to a given ECC perturbation is predicted to be similar in both the $Q = 5$ and $Q = 10$ variations of the $I_{\mathrm{p}} = 15$MA DT baseline FPO scenario, focus is placed on studying case 5 due to its importance. The upper and lower row phases relative

Table 5.2:

Optimal upper and lower coil row current phases relative to middle current phase to maximise XPD as predicted by the `MARS-F` code. Current phase is as defined in equation 2.44.

Case	2		3		5		7	
Toroidal mode number n	3	4	3	4	3	4	3	4
$n\Phi^{\mathrm{max}}_{\mathrm{u,m}}$ [deg]	-190	-230	-330	-215	-200	-250	-290	0
$n\Phi^{\mathrm{max}}_{\mathrm{l,m}}$ [deg]	-145	-75	-20	-100	-140	-95	-55	10

to the middle coil row which maximise XPD were taken from [64, 65]. To generate a general yet realistic distribution of deposited neutral beam ions, both heating neutral beams (HNB1 and HNB2 [152]) were used, with HNB1 injecting off-axis and HNB2 injecting on-axis. In this scenario, the calculated shine-through losses were extremely small and thus assumed to be negligible throughout.

5.1.1 Impact on heating efficiency

To measure the impact of RMP rotation on the global fast-ion loss rate, i.e. the overall heating efficiency of the NBI system, the total power lost to the PFCs was measured for multiple individual RMP phases. Because losses were observed to saturate after ≈30ms, which was subsequently used as a simulation time limit, it was assumed that the time-dependent fast-ion losses due to an ECC system oscillating at 0-5Hz (period as short as 200ms) could be well-approximated by separate simulations using static RMP fields at discrete phase intervals; i.e. it was assumed that $\omega = 0$ in equation 2.44. As the importance of including the second harmonic was demonstrated in section 4.5, fast ions were tracked with and without the second harmonic for comparison. Where the second harmonic was removed, the remaining fundamental mode was artificially scaled up proportionally to mimic an ECC system capable of generating a pure $n = n_0$ RMP field. It is worth pointing out that, in reality, it is impossible for the ITER ECC system to generate a field with a fundamental mode of this amplitude, as the required current exceeds the current carrying capabilities of the ELM control coils.

Figure 5.2 shows the total measured power lost to PFCs from both beams as the RMP is rotated. 2^{15} markers were tracked over different values of absolute perturbation phase, represented in figure 5.2 by middle row phase, Φ_{m}, whilst relative toroidal phase between coil rows ($\Delta\Phi_{\mathrm{u}} - \Delta\Phi_{\mathrm{m}}$ and $\Delta\Phi_{\mathrm{l}} - \Delta\Phi_{\mathrm{m}}$) was maintained.

Over a rotation cycle, the total measured losses vary by approximately $\pm 2.8 - 3.2\%$

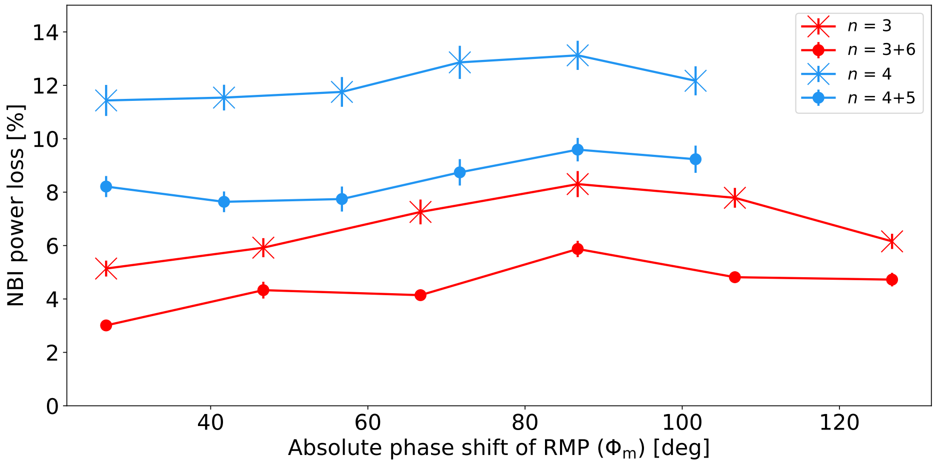

Figure 5.2: Global fast-ion losses to PFCs as a percentage of deposited beam power (≈33MW) for different middle coil row phases, Φ_m, whilst maintaining relative phases between all coil rows. This reflects the situation in which the RMP is rotated to spread the fast-ion power loads. Shown are results from fields with ($n = n_0 + n_1$) and without ($n = n_0$) the second harmonic. The scan begins at the point where the first middle coil is carrying maximum current ($\Phi_\mathrm{m} = 26.7°$). To relate the RMP field phase to the 3D system geometry, such as the toroidal localisation of the NBI deposition with respect to perturbation phase, Φ remains as defined in equation 2.44 - that is, in the ITER machine coordinate system. Hence, absolute phase is varied through $360°/n_0$. Figure reproduced from [60].

points for $n_0 = 3$ and $\pm 1.7 - 2.1\%$ points for $n_0 = 4$ - showing similar room for optimisation for both mode numbers. Whilst other work observes roughly double this variation for an individual beam in an $n = 3$ field [86], the toroidal separation of the beams, which can be seen in figure 5.31, means each is impacted by the perturbation in turn as it moves past, with losses from one beam lagging the other. Therefore the discrepancy is likely smaller. Furthermore, the magnitude of global losses reflects those calculated in [104], and the absolute phases corresponding to minimal and maximal losses align well with those predicted in [86] ($\Phi \approx 22°$ and $82°$ respectively [154]). Interestingly, these phases also align across toroidal mode numbers and spectra, possibly because the upper and middle coil phases are similar in both the $n_0 = 3$ and $n_0 = 4$ fields when XPD is maximised. This alignment, as well as the difference between individual beam losses, suggests that the phase difference between peaks in ECC current and NBI deposition is a critical 3D parameter to tune if the field is to be fixed in place. However, it must be noted that the maximum losses do not occur when these peaks overlap (maximal HNB1 and HNB2 deposition occurs at $\phi = 58°$ and $74°$ respectively). It is clear that $n_0 = 4$ fields are consistently worse for NBI heating efficiency, with minimum losses 6.4% and 4.9% points higher than their respective single and multi-harmonic $n_0 = 3$ equivalents. This could be due to the increased penetration of the stochastic layer which is observed in the Poincare plots (the same as those plotted in figure 4.22). Including the second harmonic is also crucial, as for both $n_0 = 3$ and $n_0 = 4$ fields it acts to significantly lower global losses - by approximately 2.5% and 3.6% points respectively (by $\approx$ 1MW). Because of this, and the fact that pure $n = n_0$ fields are unrealistic, from here on only fields including the second harmonic are studied when measuring PFC power loads.

5.1.2 Fast-ion power loads

It is now important to determine whether it is possible to optimise overall NBI heating efficiency whilst minimising localised PFC power fluxes. Figure 5.3 shows the component-resolved power loads for different toroidal mode spectra as a function of absolute RMP phase. Foremost, it can be seen that the relationship between power load and absolute phase is dependent upon the component and mode spectrum in question, and is, at times, not monotonic; some components do not experience peak power loads when global fast-ion losses are maximised - even in pure $n = n_0$ fields, where rotation does not affect the field structure. Hence the rotation and spectrum of an RMP each act not only to scale but also redistribute fast-ion losses.

It is therefore important to examine the spatial distribution of lost particles, which is plotted in figure 5.4 for a single phase over the perturbation amplitude evaluated at the plasma edge. Whilst the aim is not to quantify the contributions of each loss or redistribution mechanism, fast-ion losses are ultimately caused by the enhancement of radial diffusion within the stochastic layer at the plasma edge. This predominantly affects passing particles, which are typically lost along the inboard divertor leg and strike the inboard divertor targets. As shown in figure 5.31, the fast-ion deposition in ITER is relatively localised in pitch, and it lies near the trapped-passing boundary at pitch $\lambda \approx 0.6$. Hence, many of these particles have also crossed a topological boundary to or from barely trapped orbits. Some of these markers which travel down the inner divertor leg are then reflected and drift onto the inboard under-dome divertor cooling pipes. Trapped particles also experience enhanced radial diffusion at the banana tips, whereupon their orbits may open out to intercept the outboard wall on the outer co-current leg. Whilst these particles typically have banana tips near to the X-point, barely trapped particles may also bounce into the inner divertor leg before bouncing again - similar to those which strike the under-dome pipes - and passing through the private

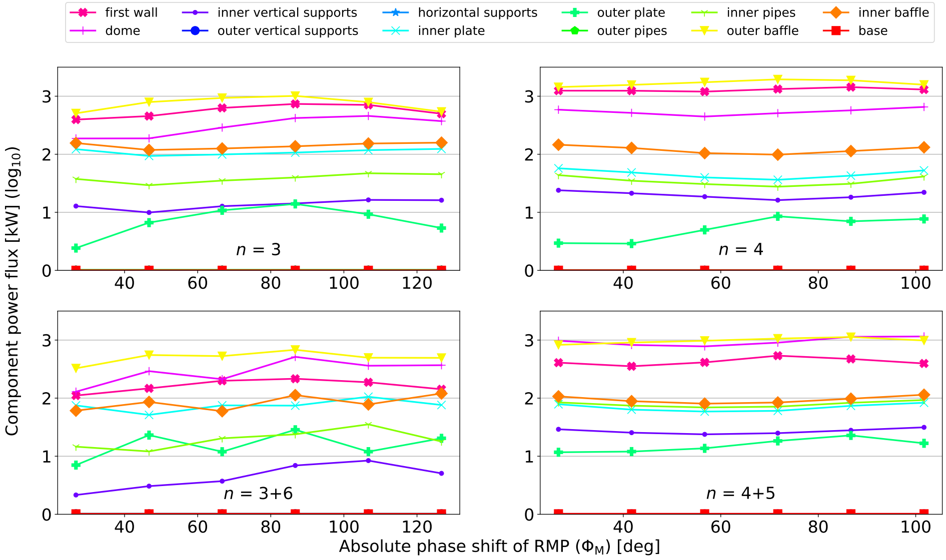

Figure 5.3: Power loads to various tokamak first-wall components as a function of absolute toroidal perturbation phase for different RMP mode spectra at 90kAt coil current amplitude. Components which received negligible power flux are denoted at the bottom (base, outer vertical supports, horizontal supports and outer pipes). Colours and components correspond to those labelled in figure 3.6. Where traces do not follow the same pattern exhibited by global losses in figure 5.2, the field is implied to redistribute fast ions amongst components. Figure reproduced from [60].

flux region to strike outboard divertor targets. This influence of particle orbit topology can be seen in figure 5.4 as a strong correlation between loss location and particle pitch angle ($v_{\|}/v$, as measured against plasma current direction, which is typically clockwise in ITER). This persists as the perturbation phase changes. The loss pattern, which remains field-aligned and n_0-fold toroidally symmetric, rotates with the perturbation without significantly changing shape; any redistribution is observed to occur within the length scales of the loss footprint. This is likely why FILDs detect oscillations in fast-ion losses as RMPs are rotated in experiments [81, 138].

To identify hotspots within the footprint, and determine whether they persist, redistribute, or fluctuate, the number of markers was increased to 2^{21}, and 3D power loads were resolved at the sub-component level. These are shown for $n = 3+6$ fields in a top-down view in figure 5.5 over the same phases used for figure 5.2. For $n = 3 + 6$ fields, the power load footprint is largely contained within the divertor and first-wall regions, specifically panels near the outboard mid-plane - as can be seen in figure 5.4. The same can be said for the $n = 4 + 5$ field; Figure 5.6 shows similar patterns, except a new footprint is introduced on

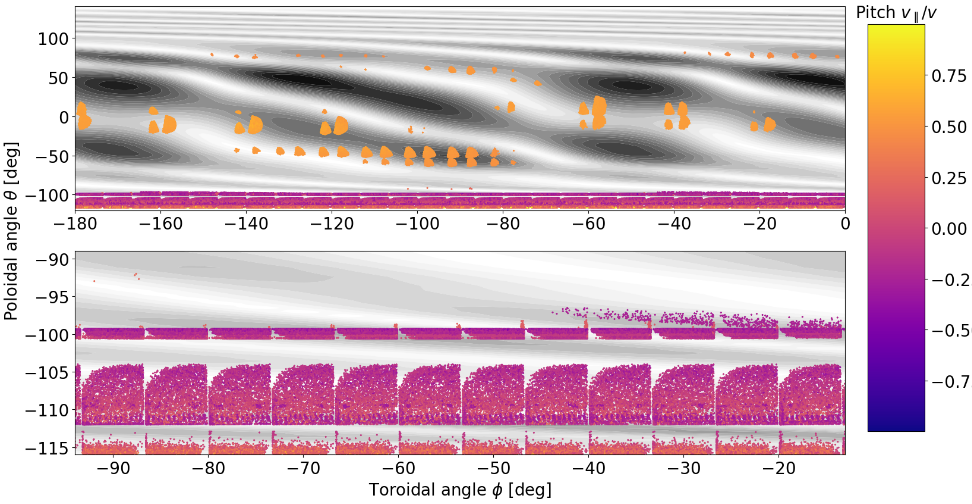

Figure 5.4: Two views of the amplitude of an $n = 3+6$ perturbation, evaluated at $\psi_\mathrm{n} = 0.99$ and shown in grey, with lost markers plotted at their final poloidal angle θ and toroidal angle ϕ locations, coloured according to final pitch λ. The bottom subfigure is a zoomed view of a section of the top subfigure. The particle loss pattern adheres to the perturbation as it is rotated toroidally. Figure reproduced from [60].

the outboard ceiling first-wall panels.

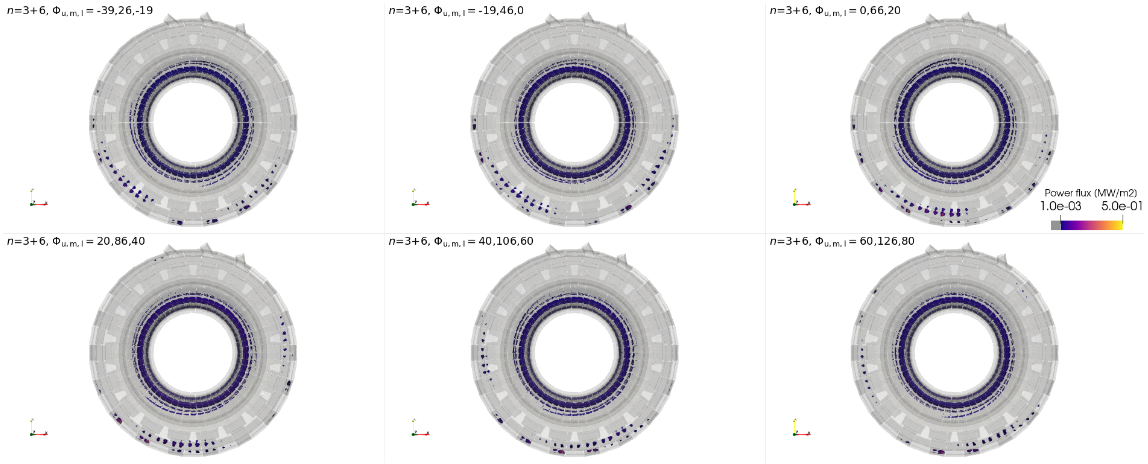

Figure 5.5: Top-down view of the first-wall with power loads rendered in colour. Ports for the three HNBs protrude from the vessel at the top, near $\phi = 90°$, whilst first-wall panels opposite attract a relatively high power load. From this view, the n_0-fold symmetric shape of the power load can be seen throughout the rotation cycle. Untouched surfaces are rendered semi-transparent. Figure reproduced from [60].

Like the loss pattern, for both mode numbers the power load adheres to the perturbation as it rotates. However, first-wall power loads are toroidally asymmetric - mostly limited

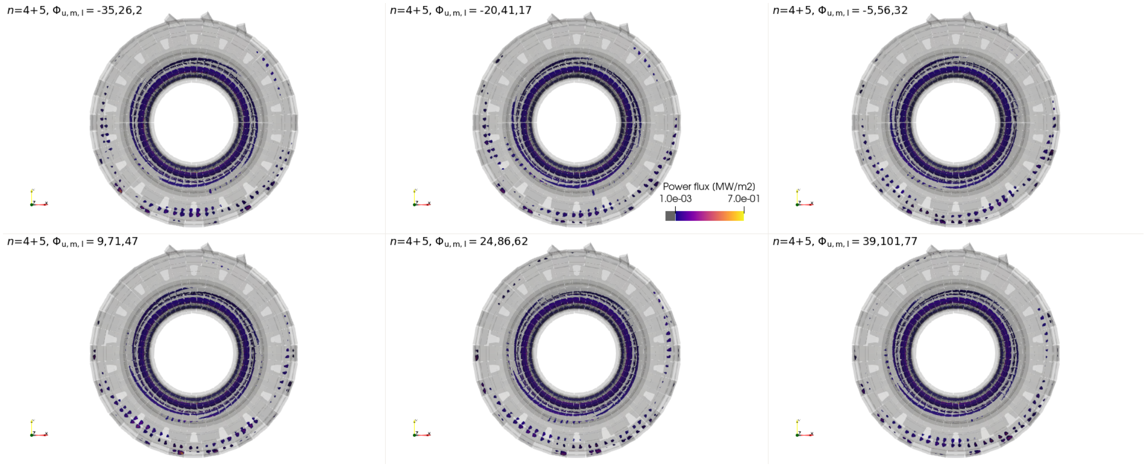

Figure 5.6: Same view used in figure 5.5 but for an $n_0 = 4$ RMP field. A similar behaviour to the $n_0 = 3$ field can be seen. Figure reproduced from [60].

to sections further from the HNB ports. These are not shine-through losses (`LOCUST` only tracks deposited ions, and the regions are blocked by the central column). Figures 5.7 and 5.8 show these particular power loads from outside the machine. It can now be seen more clearly that, as the footprint rotates and first-wall panels are struck in-turn, specific first-wall segments attract far higher power loads. Crucially, power loads are not always present here throughout the rotation, meaning that, for both mode numbers, first-wall RMS power loads could be reduced by RMP rotation. The only loads that do persist are limited to specific divertor regions shown below.

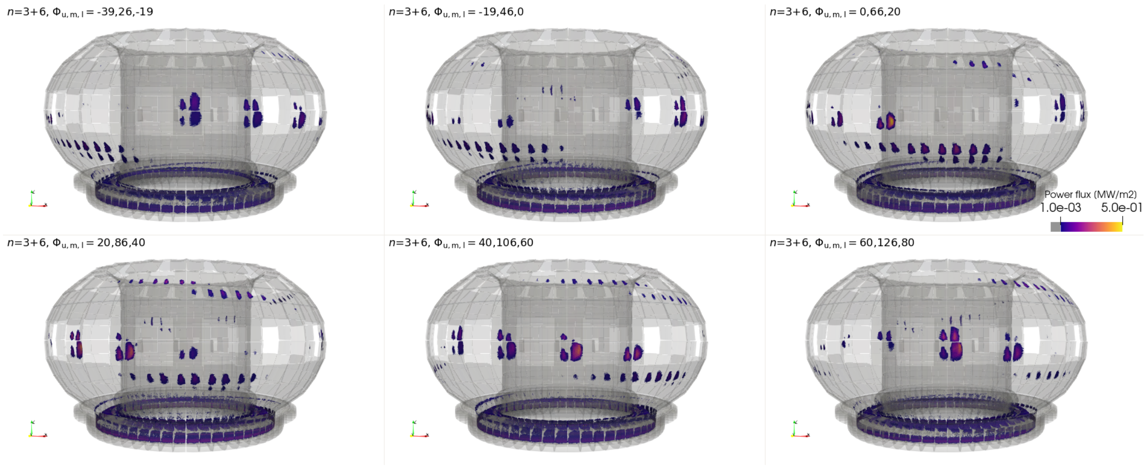

Figure 5.7: Fast-ion power loads due to $n_0 = 3$ RMP field as viewed outside the vessel. The wall panels closest to the camera are located opposite to the HNB ports. Gaps between panels exist in places where diagnostic or entry ports are located. Untouched surfaces are rendered semi-transparent. Figure reproduced from [60].

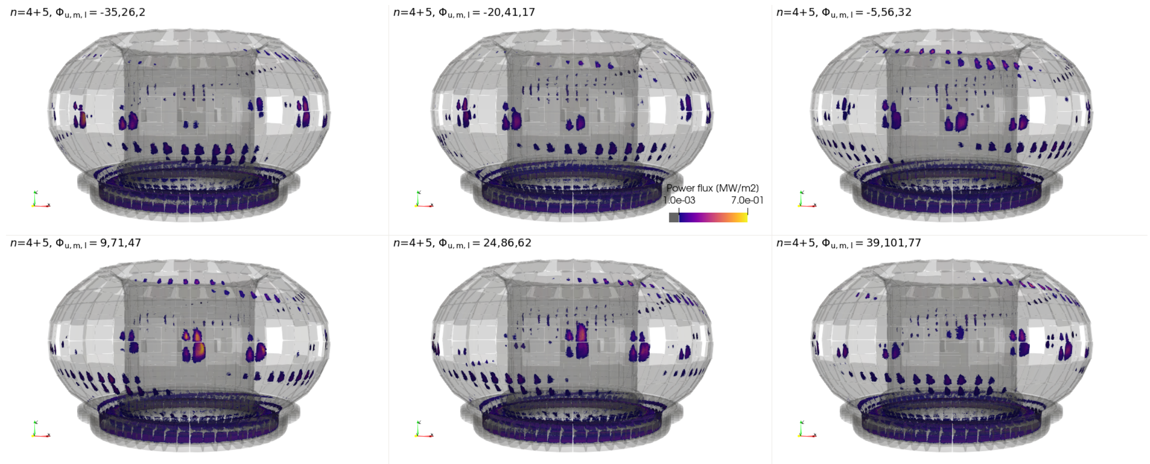

Figure 5.8: Same view as figure 5.7 but for an $n_0 = 4$ RMP field. Not only are the power loads larger in area and more intense at their peak, but additional footprints are created in the upper regions of the first wall. Figure reproduced from [60].

Further close-ups of these first-wall loads are displayed in figures 5.9 and 5.10, which highlight panels diametrically opposite to the HNB ports. Here it can be seen that the peak power loads strike specific sides of the bevelled panels - typically the side facing the HNB ports. Whilst not negligible, the maximum power loads of approximately 0.5-0.7MWm^{-2} are tolerable in these locations. The wall panels in question correspond to blanket modules 14 and 15 [155] which are designated ”enhanced heat flux” panels [156], designed to accommodate up to ≈5MWm^{-2} (compared to ”normal heat flux” panels rated for 1-2MWm^{-2}). As can be seen, the peak power loads can be reduced - and re-positioned - by switching toroidal mode number from $n_0 = 4$ to $n_0 = 3$.

To quantify the effect of rotation of the perturbation on RMS power loads, the peak power load to the first-wall panels with the highest load in Figures 5.9 and 5.10 is plotted at various RMP phases in figure 5.11. For both mode numbers, rotation of the RMP can reduce the RMS peak first-wall power load, by 0.37MW^{-2} and 0.44MW^{-2} for $n_0 = 3$ and $n_0 = 4$ fields respectively. For $n_0 = 3$ fields, peak power flux correlates well with the global fast-ion losses shown in figure 5.2, meaning that choosing the optimum phase that minimises global NBI losses automatically ensures the lowest peak power load, potentially making RMP rotation unnecessary. However, the same cannot be confidently said for $n_0 = 4$ fields. This is corroborated when considering the divertor region, shown in figures 5.12 and 5.13. Like the first-wall, peak power loads to the divertor dome and outer baffle oscillate with global losses in the $n = 3 + 6$ field. Whereas in the $n = 4 + 5$ field, for example on the dome, peak power

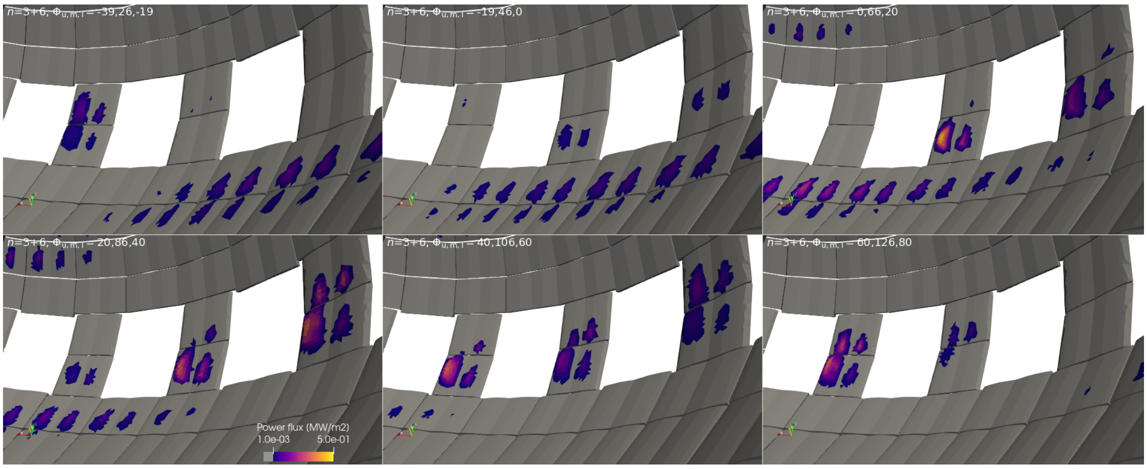

Figure 5.9: Fast-ion power loads shown from the divertor looking upward toward the wall tiles located opposite to the HNB ports for the $n_0 = 3$ RMP field. The power load on a given wall panel is heavily influenced by the panel geometry, with peak heat fluxes typically located on the side facing the neutral beam ports. Figure reproduced from [60].

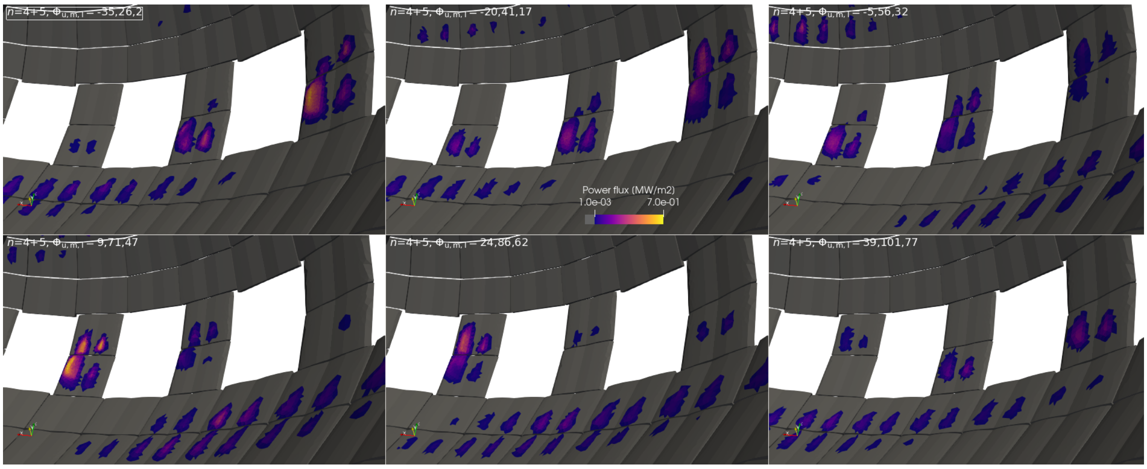

Figure 5.10: Same view and wall model used above in figure 5.9 but for an $n_0 = 4$ RMP field. The panels which receive low power loads for $n_0 = 4$ here instead receive high peak loads in Figure 5.9 above for $n_0 = 3$ fields. The additional power loads in the upper region of the first wall can also be seen. Figure reproduced from [60].

loads only redistribute and can even persist; when global losses are minimal ($\Phi_{\mathrm{m}} = 41°$), higher fluxes continue to reach the inboard side of the dome structure and outer baffle.

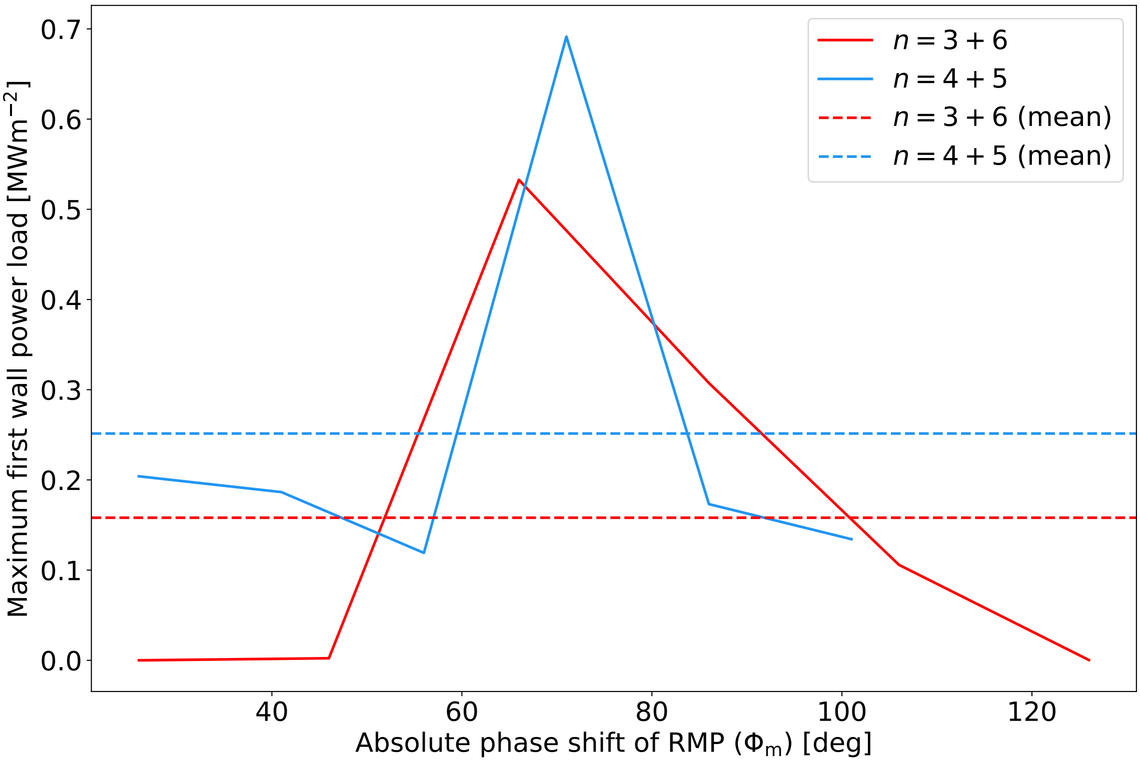

Figure 5.11: Power load reaching the first-wall at the point of maximum power flux (for all RMP phases) versus RMP phase for $n_0 = 3$ and $n_0 = 4$ toroidal mode numbers. The average value is also displayed, which is lower than the peak power flux over the total rotation cycle - demonstrating that the time-averaged power flux to hotspots caused by NBI losses can be reduced when the RMP field is rotated. Figure reproduced from [60].

Overall, $n_0 = 4$ fields consistently lead to higher divertor power loads across all components. And like the bevelled first-wall tiles, the orientation and subtle geometry of PFCs can greatly affect the power received. This is apparent specifically on the dome and outer baffle, where breaks in shadowing and regions with discontinuities, such as gaps between divertor cassettes, is subject to particularly high power loads due to the low angle of incidence. Rotation is unlikely to reduce these power loads between the cassette domes in the $n_0 = 4$ field, again making optimisation less straightforward. However, it should be noted that the dome and outer baffle are designed to handle power loads of 5MWm^{-2}, whilst the highest power loads from NBI losses on these components predicted above are only 0.1MWm^{-2} and thus not of concern from the point of view of power handling.

The same is true for other divertor components which are not designed for direct inter-

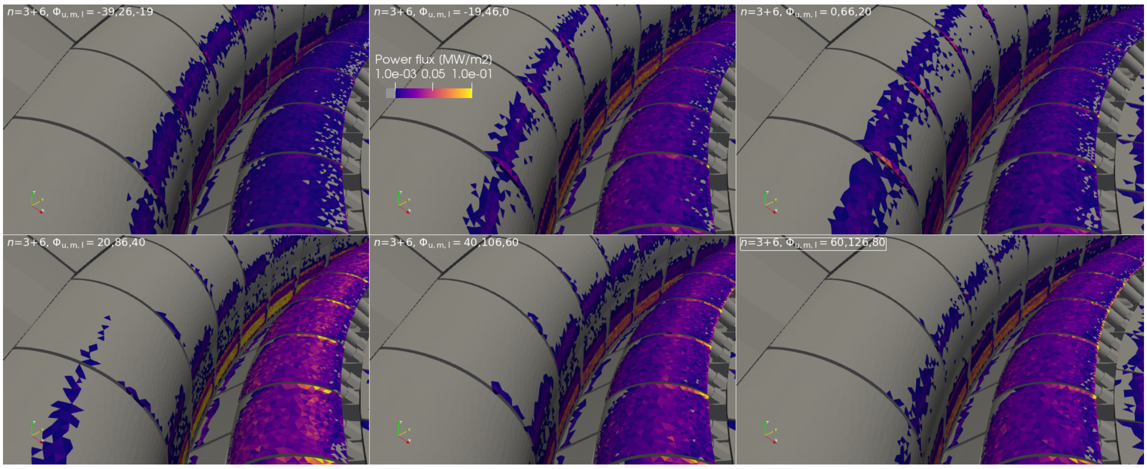

Figure 5.12: Fast-ion power loads looking down onto the divertor dome for an $n_0 = 3$ field. Here it can be seen that maximal power loads correlate well to maximum global losses shown in figure 5.2. The cassette gaps on the outer baffle and divertor dome can attract power loads that approach the peak across the wall. Figure reproduced from [60].

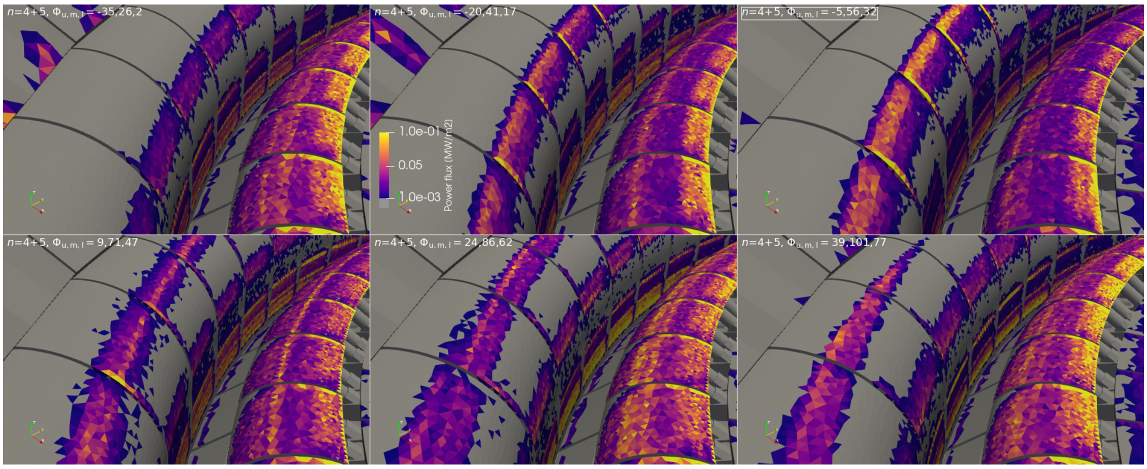

Figure 5.13: Same view as figure 5.12 but for an $n_0 = 4$ RMP field. Unlike $n_0 = 3$ fields, peak power loads are located in multiple regions, such as cassette gaps, the outer baffle and inboard side of the dome. In some of regions, peak power loads persist throughout the RMP cycle. Figure reproduced from [60].

action with the plasma; figure 5.14 shows these for the two phases which exhibit the highest power loads, which rarely exceed 0.25MWm^{-2} - comparable to the expected radiative load on the dome supports [157]. Again, these are sensitive to the local geometry - specifically the curved edges of the inner support legs and gaps between divertor cassette gaps which highlight endmost cooling pipes. The high power loads on the individual surface triangles of the underdome pipes are subject to uncertainties due to a lack of marker statistics, however the presence of more marker hits in the $n = 4 + 5$ field signifies an increased power load compared to $n_0 = 3$.

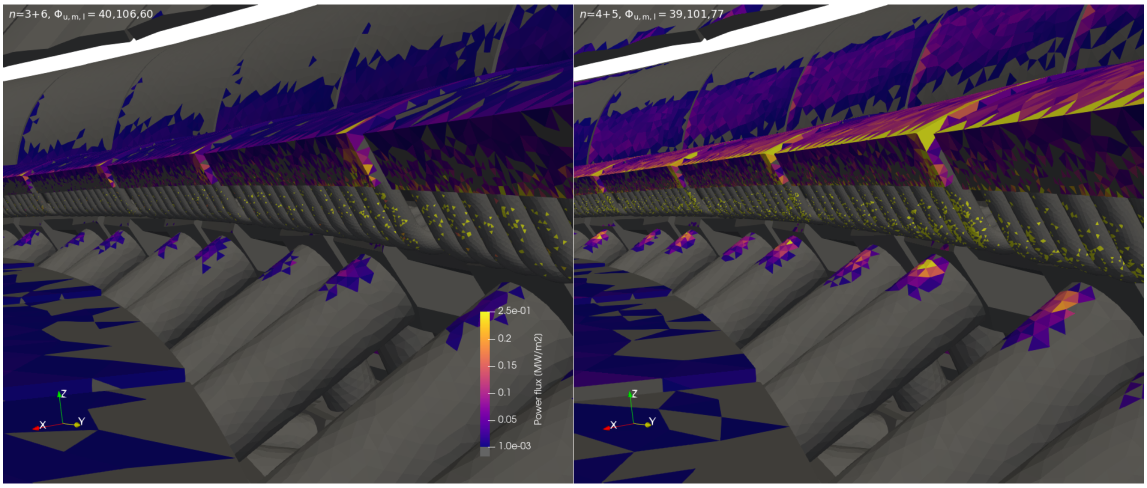

Figure 5.14: View of divertor power loads for $n_0 = 3$ and $n_0 = 4$ fields when power loads are maximal, looking towards the outboard side from the inner divertor plate. Shown are the inner support legs, the inboard edges of the divertor dome structure and its cooling pipes, and a portion of the upper region of the outer baffle. Whilst the number of markers striking the pipes is small, the edge-most pipes exhibit an increased number of hot surface triangles due to the gaps in the divertor dome. Figure reproduced from [60].

5.2 Optimising fast-ion transport across ITER experiments

The previous section showed that, from the point of view of fast-ion power loads, ITER may safely operate with maximal ELM suppression. Nevertheless, the reality is that ITER will conduct experiments nearer to the minimum viable level of ELM suppression so as to minimise the impact of density pump-out on the thermal ion confinement. This could have large ramifications for fast-ion confinement too, as it introduces new possibilities for operating the ECCs; on top of absolute phase, relative phase and current amplitude now become parameters that may be varied whilst maintaining ELM suppression. These alternate methods may be necessary if, in fact, a reduction in ELM suppression alters the fast-ion transport - potentially redistributing power loads. Therefore, the aim of this section is to compare two different ECC operating scenarios in which the coil currents are cycled periodically and XPD maintained.

The first method, the fixed relative phase scan, studied in section 5.2.3, is similar to that studied in section 5.1; relative coil phases are set to those which maximise XPD before the absolute RMP phase is steadily increased. The XPD is then linearly scaled by 2/3 by lowering the ECC current amplitude accordingly. This method is studied using the smaller coil currents in table 5.1.

The second method, the variable relative phase scan, studied in section 5.2.4, is to increase the ECC current amplitude to the larger values in table 5.1, before proportionally lowering the XPD through adjustment of the relative upper and lower coil row phases. Because the XPD is not required to be maximal, there are now multiple possible values of upper and lower coil row phase which yield the same XPD - defining a contour around which the ECCs may be freely cycled. Examples of these contours can be seen later in the plots of XPD in figure 5.17.

To get a general sense of whether fast ions respond to changes in the relative phase and ECC amplitude, even as XPD is maintained, parameter scans of these quantities are first performed in sections 5.2.1 and 5.2.2. Because these types of simulations have remained computationally infeasible at this scale for conventional fast-ion codes, scans like this have not been performed before.

These simulations are performed for each discharge in table 5.1. This is firstly to produce initial estimates of the fast-ion transport in each case, as these cases are scarcely studied - if present at all - in the literature. In addition, expanding the study in this way to multiple, low-power scenarios is used to assess the suitability of using earlier discharges as low-risk testbeds for calibrating ECC operation before the ITER baseline scenarios.

5.2.1 Fast-ion response to ECC current

To better understand how fast ions respond to changes in ECC current and corresponding changes in XPD, as is the case when transitioning between the two proposed ECC operating scenarios, the fast-ion losses to the PFC wall were measured over increasing ECC current amplitude for each case and toroidal mode number. The total losses are shown in figure

5.15. Relative phases which maximise XPD and the absolute phase which minimises global fast-ion losses at 90kAt were used to configure the RMP field.

Foremost, figure 5.15 shows that the fast-ion response is relatively steady in all cases; there are no large oscillations in the total losses, and in some cases the response is linear. This implies that scaling XPD similarly scales global fast-ion losses, though scenario and toroidal mode number can influence this relationship. This could be due to many reasons: changes in the equilibrium plasma, the resulting beam deposition or changes in the perturbation itself, for example the difference in q profile between the cases affecting the amount of stochastisation at the plasma edge. In some cases, the latter reason is extremely influential, as is shown by the variations between mode numbers in a given scenario; at only 10kAt, transitioning from $n_0 = 4$ to $n_0 = 3$ in case 3, for which the q profile changes the most at the edge, leads to losses increasing from 0.07MW to 0.98MW, and in case 5 it changes the response from mostly linear to exponential, before a linear saturation at roughly 50kAt.

As a first estimate, figure 5.15 shows that risks to NBI heating efficiency are not limited to FPO scenarios, for example the $n_0 = 3$ variation of case 3 carries most risk at maximal ELM suppression. Care should therefore be taken to ensure that the ECC current is always minimised where possible. Lowering the currents to those planned in table 5.1 mitigates this risk, at which point the confinement in the worst case is comparable to that in the $Q = 10$ baseline scenario with $n_0 = 4$ fields at maximum ELM suppression.

Though total losses can be reduced by decreasing the ECC current amplitude, their possible redistribution may still lead to specific components receiving disproportional power loads. The total power load reaching each component was therefore evaluated and plotted in 5.16. One can see that the relationship between power load and current is not always monotonic, implying that redistribution is occurring. This means that decreasing the ECC current may indeed increase the power reaching some components. Nevertheless, all components which receive over 100kW - the first wall, outer baffle and dome - are designed for high power loads. This is except for the $n_0 = 4$ variation of case 5, during which the inner pipes may receive just below 100kW at currents over 30kAt. Interestingly, the power loads to some components exhibit threshold and saturation behaviour; power loads to some components only appear once a certain threshold value of coil current is exceeded, whilst for other components the power load saturates.

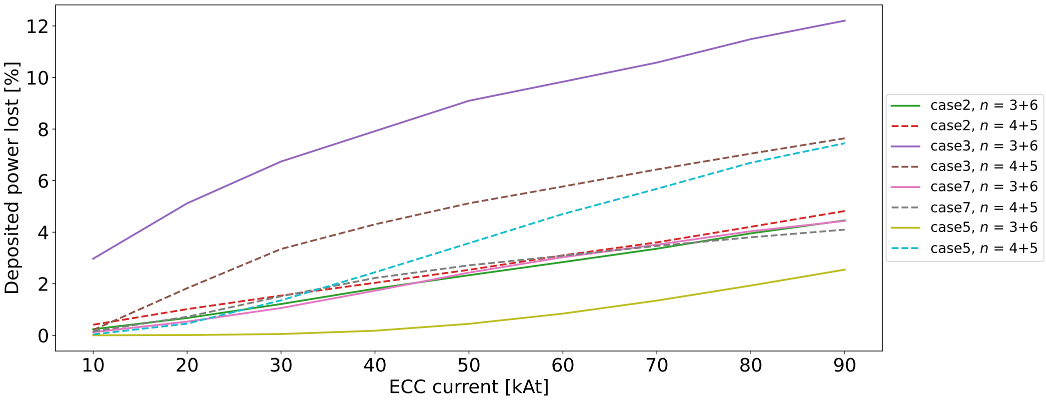

Figure 5.15: The total HNB power loss to PFCs over the cases listed in table 5.1 as ECC current is increased from 10kAt to the maximum possible value of 90kAt. Relative phases between the coil rows are chosen to produce to maximum XPD, whilst the absolute phase of the perturbation is set to that which minimises global losses. At these current increments, the increase in losses is relatively steady - nearly linear - in most cases. Figure reproduced from [153].

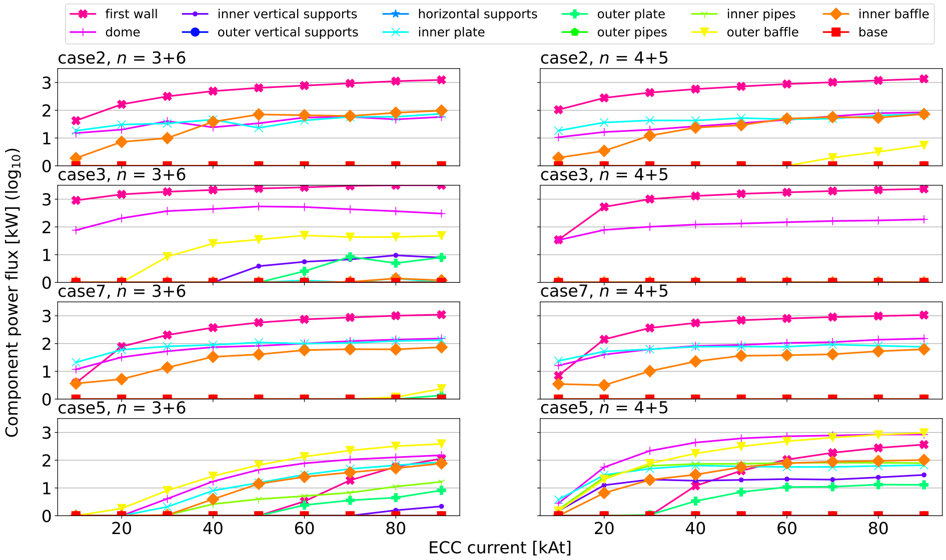

Figure 5.16: The component-resolved power loads corresponding to those losses calculated in figure 5.15. Some redistribution of losses is implied where losses decrease as ECC current is increased. Figure reproduced from [153].

5.2.2 Fast-ion response to relative ECC coil row phase

XPD is also dependent upon on the relative phase between the current flowing in different coil rows. It has been shown previously [104] that the field line leak rate (the rate that markers travelling along magnetic field lines are lost over 200 toroidal transits) responds similarly to XPD when relative phase is varied. However, the impact on fast-ion confinement has not yet been studied. Hence to answer whether fast-ion transport is similarly scaled by relative phase as it is by ECC current, the confinement was calculated over the possible relative phases between upper and lower coil rows relative to the middle row. This confinement is then compared to the XPD due to a single harmonic field calculated by `MARS-F`. For comparison, losses are calculated with and without the inclusion of a second harmonic.

The XPD and total losses for each case and toroidal mode number are shown in figure 5.17. To observe how relative changes in the XPD affect fast-ion losses, each loss grid is normalised according to the values within that grid only. The XPD responds similarly in every case; there is a general symmetry which is translated and then perturbed by 3D features of the system. Case 3 is unique due to its high q_{95}, making its response to an $n_0 = 4$ waveform in the middle coil row almost purely core-kink-like, resulting in little X-point displacement and a sinusoidal dependence on the upper and lower row only. It is noticeable that the global fast-ion losses correlate well with XPD in each case. This means that, within a particular scenario, XPD (and therefore ELM suppression) may be used as a robust metric for optimising total fast-ion confinement and therefore overall NBI heating efficiency.

To examine this proposal more quantitatively, the previous data are replotted in figure 5.18 - which shows the fast-ion losses against XPD. To isolate the losses caused purely by the perturbation, each line has been translated to intercept the origin. Fits are calculated in each case, without and with sidebands in the bottom left and right subfigures respectively. Within each case, the correlation is extremely high. This is particularly true for the first wall and outer baffle, where the Pearson correlation coefficients are 0.84 and 0.69 on average respectively (the mean across components is 0.52). At the relative phase where XPD is maximal, the large difference in XPD between the variations of case 3 could explain the differences in power loads seen in figure 5.16.

Nevertheless, the correlation cannot be claimed to be perfect; fast-ion losses were previously shown in section 5.1 to modulate by approximately 2-3% even whilst using a single harmonic to maintain XPD. Hence it is still important to examine the methods of ECC

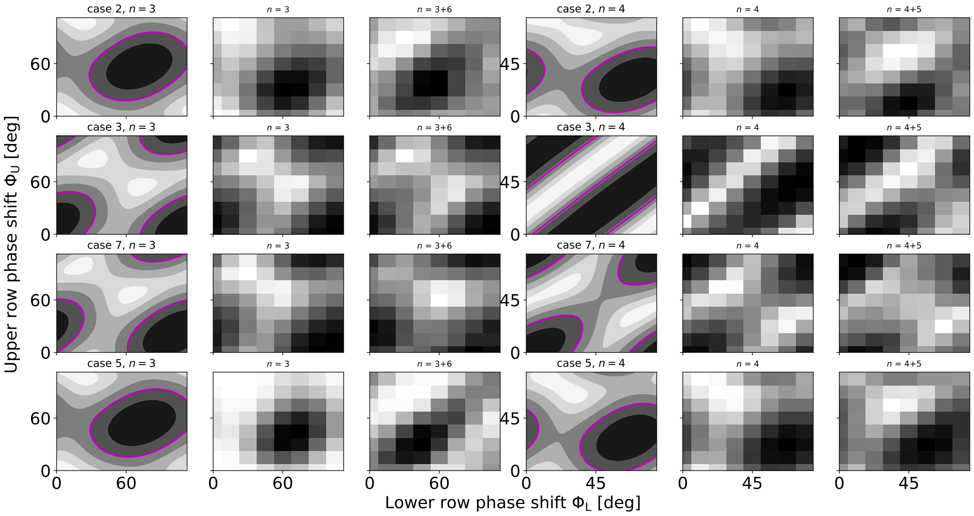

Figure 5.17: XPD amplitude as calculated by `MARS-F` is shown in the first column of each set of three adjacent columns. The following two columns are the corresponding global fast-ion losses calculated by `LOCUST` for single and multi-harmonic RMP fields respectively. The left set of three columns uses $n_0 = 3$ fields whereas the right $n_0 = 4$. The $n_0 = 4$ field in case 3 is unique in that the high edge gradient of the q profile greatly modifies the MHD response to the imposed perturbation. Each grid showing losses are normalised to values within that grid to show how changes in the mode spectrum modify the shape of the loss distribution, not amplitude. `MARS-F` assumes an axisymmetric plasma whilst the `LOCUST` losses were calculated at an absolute RMP phase of 0°. Figure reproduced from [153].

operation proposed earlier by evaluating the transport directly.

5.2.3 Fast-ion transport at fixed relative phase

The first ECC operating method to be studied here is similar to that studied in section 5.1 but at reduced ELM suppression; phases between the coil rows are chosen to maximise XPD whilst ECC current is lowered to the smaller value in table 5.1. Oscillation of the perturbation is achieved by cycling the absolute phase to rotate the field toroidally. This is the most likely method to be used to operate the ITER ECC system.

Again, to mimic the effects of continuously rotating the RMP, fast-ion dynamics were simulated at discrete phase intervals, but this time over all cases in table 5.1. A variation of case 2 is included in which hydrogen is injected by the HNB system into a pure hydrogen plasma, which is one possible variation performed during the PFPO2 phase.

Figure 5.19 shows the calculated loss rate as the absolute phase is varied. Compared to figure 5.2, the response to RMP phase is flattened - varying by $\sim \pm 1\%$ at most. The response

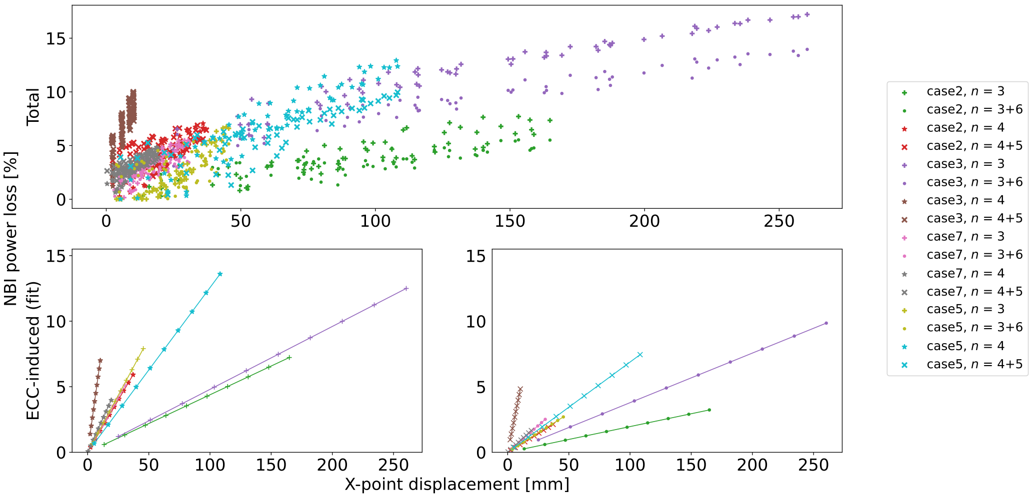

Figure 5.18: Correlation between XPD and total fast-ion loss rate to all components. The upper panel shows the fast-ion losses evaluated at the grid points in figure 5.17. Due to it impacting the XPD, the effects of the second harmonic are shown by plotting linear fits of the single and multi-harmonic fields in the lower two panels on the left and right respectively. In these plots, the intercept is set to zero to isolate and compare the relative transport due to the 3D field as opposed to axisymmetric and prompt losses. Figure reproduced from [153].

is flattest for low-power cases, making them poorly suited for optimising $Q = 10$ operation in the situation where the relative phase is fixed. However, with RMP rotation causing such little variation in the confinement in case 5 at this ECC current, optimisation may not be as necessary from the point of view of HNB heating efficiency. This is especially true given that confinement can be worse in these low-power discharges. An alternative option for optimising the ECC system in the baseline scenario may simply be to reduce the ECC current; for case 5, the relationship between absolute phase and global losses is maintained at both 60kAt and 90kAt ECC current (see figure 5.2). Therefore, if H-mode in the ITER baseline can be accessed at lower XPD, an optimal way of tuning the HNB heating efficiency is to lower the ECC current whilst somehow measuring global losses. On the other hand, if H-mode does require maximal ECC current, then additional studies of fast-ion losses in L-mode discharges are required to confirm that the losses do not significantly diverge from those measured here. In any case, for this strategy to be considered lower risk, the power loads must also be reduced at lower ECC current. Hence these are studied next.

For comparison with figure 5.3, the component-resolved power loads are shown for case 5 in figure 5.20. The relative power distribution amongst components is mostly retained

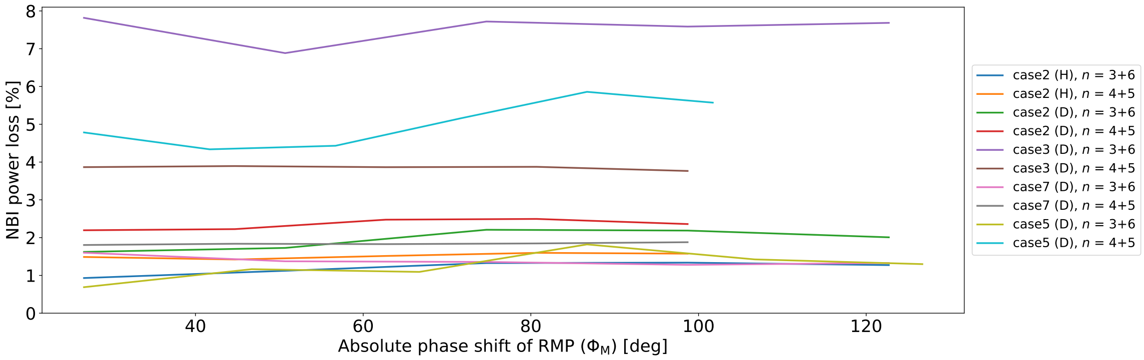

Figure 5.19: Global fast-ion losses to PFCs as a percentage of deposited beam power for different middle coil row phases, Φ_{m}, whilst maintaining relative phases between all coil rows. The scan begins at the point where the first middle coil is carrying maximum current ($\Phi_{\mathrm{m}} = 26.7°$). ϕ remains plotted in the ITER machine coordinates for comparison with 3D geometry. Brackets denote the NBI species. Figure reproduced from [153].

but reduced overall, with the re-scaling reflecting the drop in total losses in figure 5.16. Furthermore, the relationship with absolute phase is also similar - more so for the $n_0 = 4$ variation. Only components designed to tolerate high power fluxes experience power loads exceeding 0.1MW.

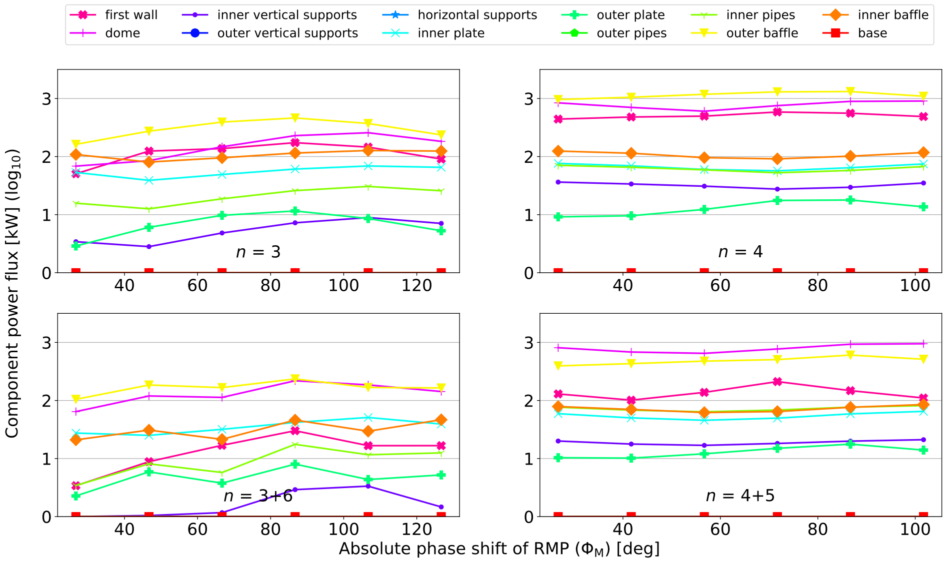

Figure 5.20: Total power reaching each component in the $Q = 10$ ITER baseline scenario at 60kAt ECC current amplitude. Figure reproduced from [153].

The spatially resolved power loads are now evaluated to look for hot spots, as it was previously understood that reducing the ECC current may cause redistribution and an increase in power loads on certain components. For both toroidal mode numbers, the overall shape and behaviour of the power load footprint and hot spots is similar to that at maximum ECC current; wall panels diametrically opposite to the HNB port attract the largest power loads. Examining these wall panels in figures 5.21 and 5.22 shows more closely this similarity. For both toroidal mode numbers, the maximum peak power load lands on the same group of wall panels. However, it can be seen that the power density distribution on a given panel changes. The maximum power load reaching these tiles falls by a factor 5 in the $n_0 = 3$ field and by a factor ~ 2 in the $n_0 = 4$ field compared to 90kAt. This confirms that a low-risk method for optimising the absolute phase of RMPs in the ITER baseline scenario is to measure the losses at lower ECC current.

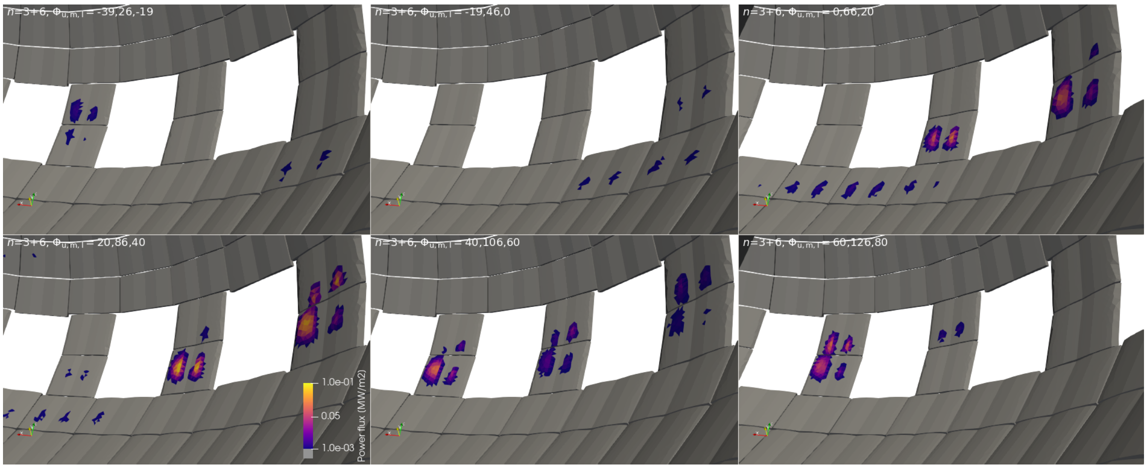

Figure 5.21: Same view as used in figure 5.9 but for 60kAt ECC current. The distribution of power loads is extremely similar to 90kAt, however the peak is now reached at $\Phi_\mathrm{m} = 86°$ instead and the maximum power load is reduced by a factor 5. Figure reproduced from [153].

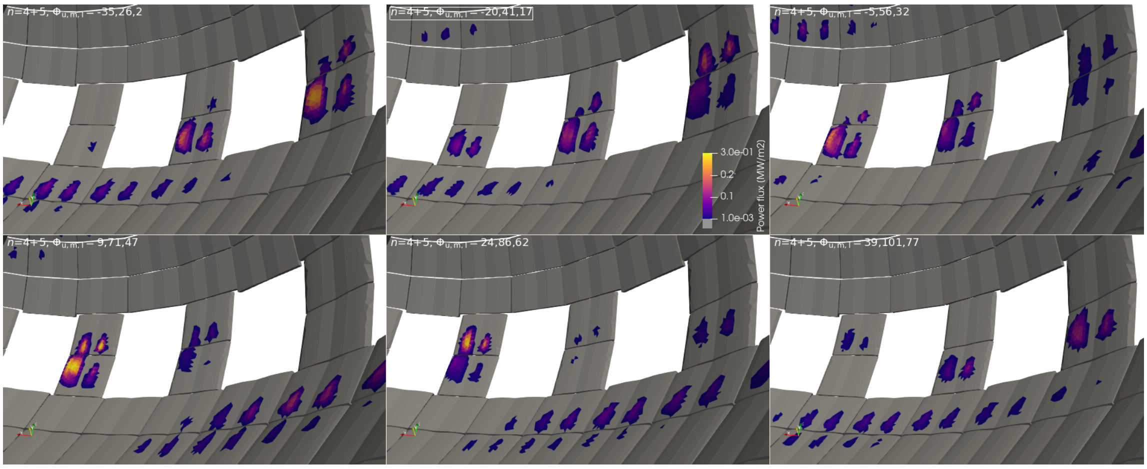

Figure 5.22: Same view as used in figure 5.10 but for 60kAt ECC current. Figure reproduced from [153].

The power density in the divertor is also shown in figures 5.23 and 5.24. Again, the peak power flux is reduced in the $n_0 = 3$ field - by a factor ~2. The power is also reduced in many areas in the $n_0 = 4$ field, however interestingly the power load on the inboard side of the divertor dome remains the same as at 90kAt. This again shows how the impact of scaling ECC current on fast-ion power loads is highly dependent on the PFC in question. Here the threshold-like behaviour observed previously in figure 5.16 corresponds to a persistence of power flux on some components. Whilst this may seem to jeopardise the previous proposal that lower ECC currents can be used to help optimise the fast-ion transport, it should be mentioned that the power loads observed here, in both fields, are still far below the design limits of the PFC components that are affected.

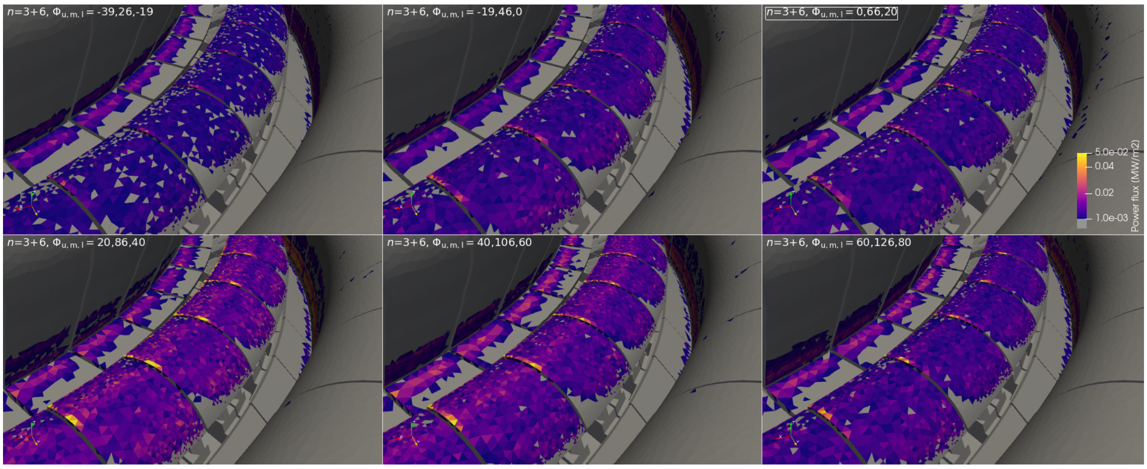

Figure 5.23: Same view as figure 5.12 but for 60kAt ECC current. The power density is reduced by a factor two by the decrease in ECC current from 90kAt. Figure reproduced from [153].

5.2.4 Fast-ion transport at variable relative phase

In this second alternative method to that studied in section 5.2.3, the XPD is reduced by tailoring the relative phase of the upper and lower coil rows with respect to the middle, before the XPD is restored by raising the ECC current. For a given value of absolute phase Φ_{m}, this defines a contour in the $\Phi_{\mathrm{u}} - \Phi_{\mathrm{l}}$ domain along which the ECCs may navigate to approximately maintain ELM suppression. Examples of these contours can be seen in the subplots of XPD in figure 5.17. Like the previous method, the absolute phase may also be cycled at the same time. However, for simplicity, it is fixed at the absolute phase which was observed to produce the minimum global losses measured previously in section 5.2.3. Fast-

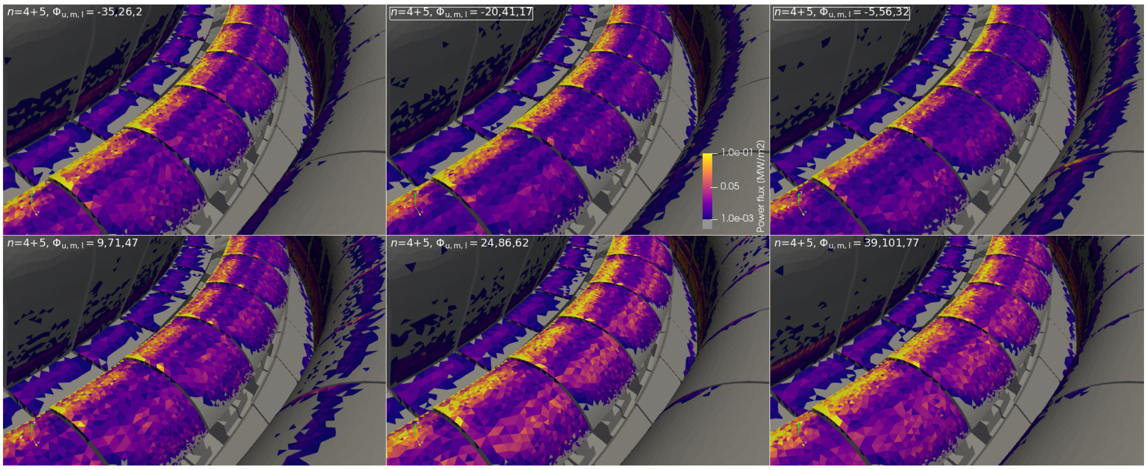

Figure 5.24: Same view as figure 5.13 but for 60kAt ECC current. Whilst the power density is reduced in most areas compared with the field produced by coils carrying 90kAt, the power loads on the inboard side of the divertor dome still persist. Figure reproduced from [153].

ion transport was then measured at eight equidistant points along the isocontour defined by 2/3 the maximum XPD and at the higher ECC current listed in table 5.1. Because the perturbation field due to each coil row is poloidally localised, adjusting each coil row independently this way changes the field structure more dramatically than simply rotating the entire field - as in section 5.2.3. Therefore, the aim of this section is to determine whether this method reduces ion losses and power loads, compared to solid rotation, for a given level of XPD.

The same analysis as performed in section 5.2.3 is repeated here for comparison, starting with the measurement of total losses shown in figure 5.25. Compared to the losses exhibited during rotation of absolute phase shown in figure 5.19, the losses here are higher on average, and there is a much larger degree of variation for all cases and mode numbers. At best, and apart from $n_0 = 4$ fields in case 3, the XPD contour rotation can only reduce power losses to those measured during absolute rotation, making this method much less preferred from the point of view of time-averaged HNB heating efficiency. That is unless the field is to be fixed in place. In this situation, the relationship between beam losses and position along the XPD contour is maintained well across all cases and mode numbers, making it possible to optimise the coil configuration using low-power discharges.

The corresponding component-resolved power loads for the ITER baseline scenario are shown in figure 5.26. In comparison to the previous method, varying relative phase retains the distribution of lost power amongst components but can, in the $n_0 = 3$ field, lead to much

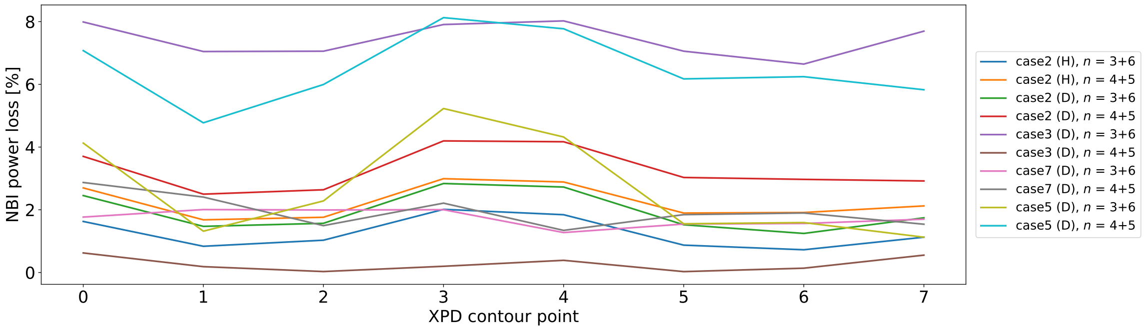

Figure 5.25: Global fast-ion losses to PFCs as a percentage of deposited beam power for different equidistant points along an XPD isocontour in upper-lower relative coil phase space. The middle coil phase is fixed to that which corresponds to minimum global fast-ion losses during RMP rotation for each case and mode number at maximum XPD. Figure reproduced from [153].

larger variations in the loss power to specific components. Again, at some phases the power to individual components approaches that recorded during variation of the absolute phase. The added variation in XPD introduced by the second harmonic could be contributing more strongly for this method, as the difference between the results obtained with single and multi-harmonic fields is much larger here than shown in 5.20.

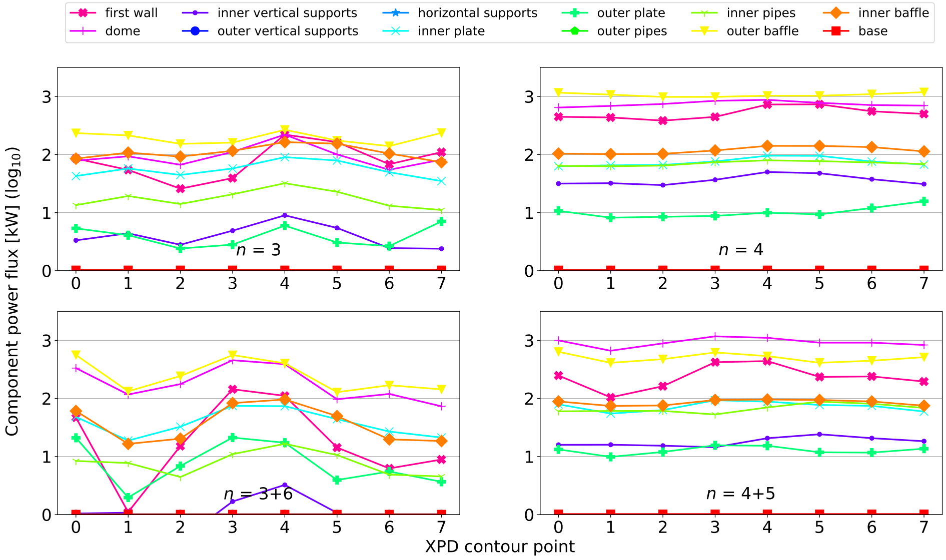

Figure 5.26: Total power reaching each component in the $Q = 10$ ITER baseline scenario at 90kAt ECC current amplitude. Each measurement is taken at a point along an XPD contour. The XPD is the same as in figure 5.20. Figure reproduced from [153].

The first-wall power loads for these contour points in the ITER baseline scenario are also shown in figures 5.27 and 5.28. Compared with the previous method, the peak power load during XPD contour rotation is much higher; it increases from ~ 0.1MWm^{-2} to ~ 0.5MWm^{-2} in $n_0 = 3$ fields and from ~ 0.3MWm^{-2} to ~ 0.8MWm^{-2} in $n_0 = 4$ fields. For both mode numbers, the variation in the RMP poloidal spectrum during rotation can be seen to change the wall panel upon which the maximum power load impinges compared to the previous method. The average and peak power loads are also higher, though they are still within the design tolerances of the first-wall panels. Furthermore, rotation along the contour can be seen to potentially reduce the time-averaged power load by a significant fraction - as relatively high power loads are observed at a select few phases, especially in the $n_0 = 3$ field. Furthermore, these phases also correlate well with those that exhibit maximum global losses, making optimisation simple in the case where the ECC system is fixed.

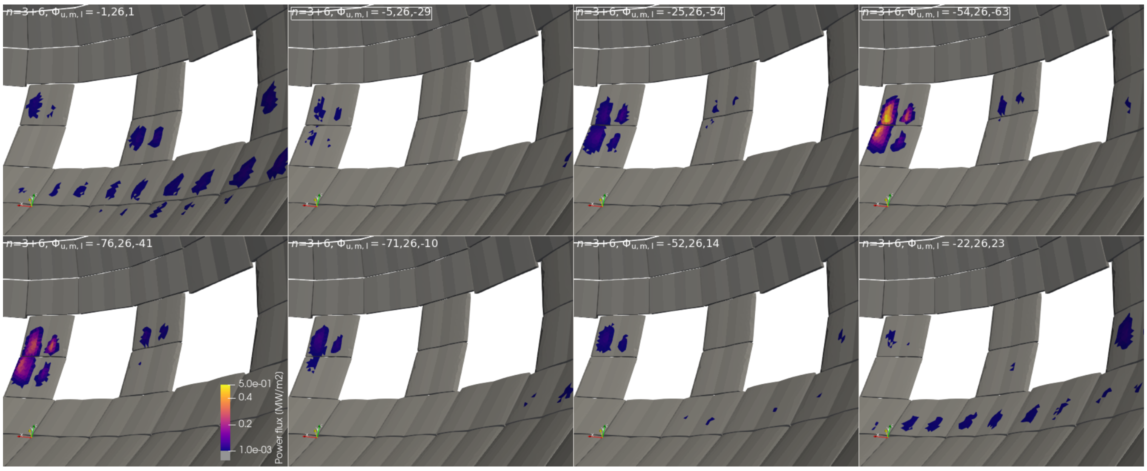

Figure 5.27: Same view and XPD as used in figure 5.21 but for RMP rotating along XPD isocontour. The relative current phase of the upper and lower coil rows is varied between eight points at which XPD is constant. Figure reproduced from [153].

The divertor region is again shown in figures 5.29 and 5.30. On the whole, the time-averaged divertor power loads are worse during XPD contour rotation. Interestingly, for both mode numbers, these power loads can be lower than those experienced at some phases where absolute phase is varied. Despite this, the correlation between divertor power loads and global losses is less clear. Nevertheless, the power loads do not exceed design limits.

5.3 Verification with the diagnostic neutral beam

The DNB is a single, low-power neutral beam system designed for active diagnosis of plasma dynamics through charge exchange recombination spectroscopy (CXRS) at the point where

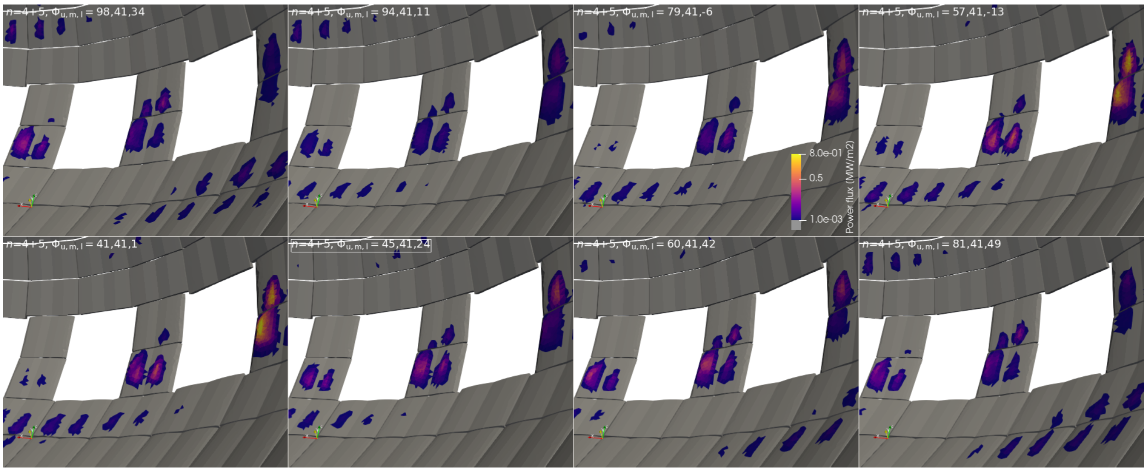

Figure 5.28: Same view and XPD as used in figure 5.22 but for RMP rotating along XPD isocontour. Figure reproduced from [153].

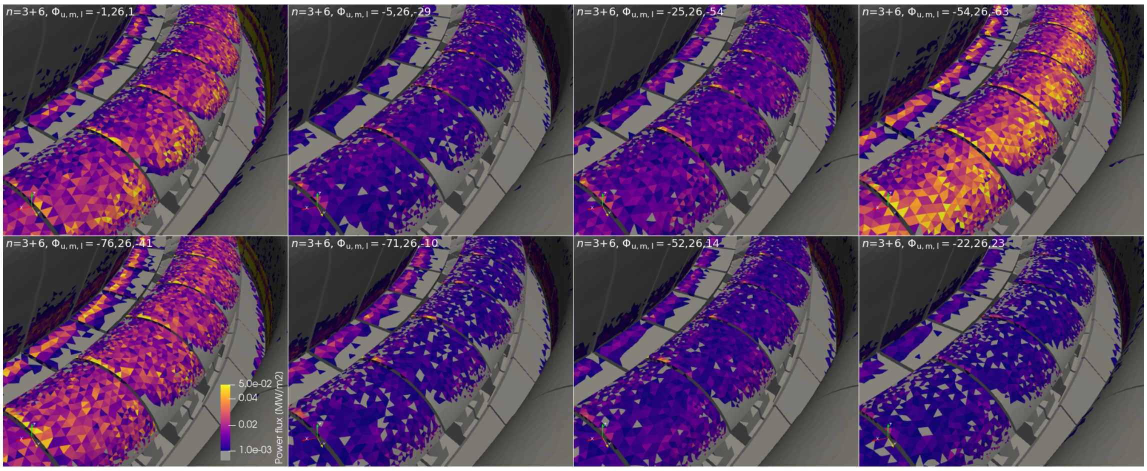

Figure 5.29: Same view and XPD as used in figure 5.23 but for RMP rotating along XPD isocontour and reduced ECC current. Figure reproduced from [153].

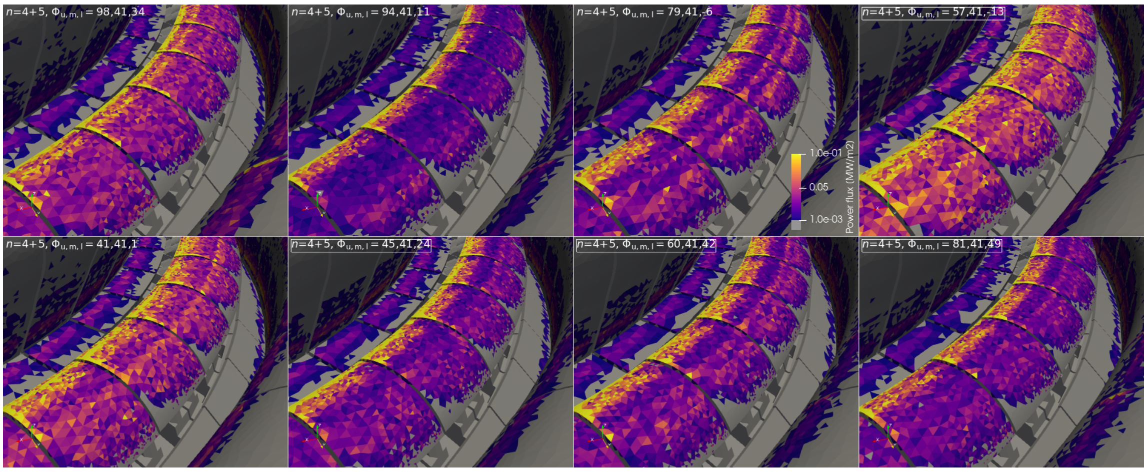

Figure 5.30: Same view and XPD as used in figure 5.24 but for RMP rotating along XPD isocontour and reduced ECC current. Figure reproduced from [153].

Table 5.3: Geometric and operational parameters describing the DNB and HNB systems in ITER. The ITER HNB system comprises two beams whilst the DNB is solitary.

	DNB	HNB
Initial particle pitch	-0.18	0.60
Energy [MeV]	0.1	1 or 0.87
Power [MW]	2 for 0.1s every 1.4s	33.3
RMS power [MW]	0.13	33.3
Species	H	D or H

the beam is ionised upon entry into the plasma. Table 5.3 compares key parameters of the DNB with HNB equivalents. The beam geometry and operational characteristics make the system likely to be resolved by detectors sensitive to fast-ion phase-space, such as the ITER FILD system. The low injection energy of the beam makes it likely that deposited ions are restricted to the plasma edge, potentially making its confinement highly sensitive to edge perturbations, such as those from the ECC system.

In this section the impact of the ECC system on the DNB is evaluated. As well as acting as an initial indication of the enhancement or degradation of signals originating from the beam for future diagnostic studies, the calculation of the beam confinement is specifically used here to judge the suitability of the DNB system as a sensor for experimentally verifying the methods of RMP field optimisation studied previously. This is repeated across the ITER cases in table 5.1 to determine which, if any, would provide a suitable low-risk testbed for pre-emptively determining optimal ECC configuration prior to the high-power ITER baseline scenario. As transport studies of the ITER DNB system are not yet present in the literature, the risk of power loss to the PFCs is also determined.

5.3.1 Implementation of the ITER DNB system

To supply `BBNBI` with the necessary injector geometry required to calculate the fast-ion deposition created by the DNB, the DNB geometry was determined before being amalgamated into the ITER machine description database in `IMAS`. The geometry of the DNB is similar to that of the HNB system. The injector face is comprised of nested grids; individual beamlets, which each inject a single collimated Gaussian beam of neutrals, including halo [158], arrange to form beam groups, which are arranged in grids on the injector face. Focussing of the beamlets and beamlet groups is performed independently, with different focal lengths in the vertical and horizontal planes. The IDSs previously used to describe the HNB system geometry could not be re-purposed because some elements of the geometry differ: the aforementioned focal lengths are different between beam systems, the DNB cannot adjust to inject on or off-axis and the DNB beam line originates and aims differently to the HNBs.

The geometry was instead derived from CAD drawings of the injection system. For simplicity, the injector face was assumed to be planar. The injection power was set to the RMS value in table 5.3. However, it is noted that, after the simulations herein, it was determined that the beam slowing down time is much shorter ($\sim$ms) than the beam pulse length ($\sim$100ms). Consequently, any measured absolute PFC power can be considered up to 15x higher than displayed here - as well as being pulsed. These simulation results remain valid due to the lack of self-interaction amongst fast ions, which means the marker weights may be scaled freely to account for these effects whilst leaving trajectories unaffected. Nevertheless, the loss power flux after scaling is still very small and poses little to the PFCs.

Figure 5.31 shows the coded DNB geometry and resulting deposition, compared to equivalents produced by HNB1 and HNB2 injecting on and off-axis respectively. Though no prior examples are available to compare to, the beam lines shown here, drawn from the injector grid to the point of tangency, align well with CAD drawings of the beam line in plane and elevation views. As expected, the deposition is much more edge-localised compared with the HNB system. The difference in pitch space between the HNB and DNB is much larger than the precision of the ITER FILD system studied elsewhere [159]. Combined with the difference in injection energy, this shows that it is likely that the DNB will be distinguishable from the HNB from the point of view of a FILD system - assuming that the detector field of view is sufficiently wide, the loss signal is strong enough and data can be gathered over at least one DNB pulse period.

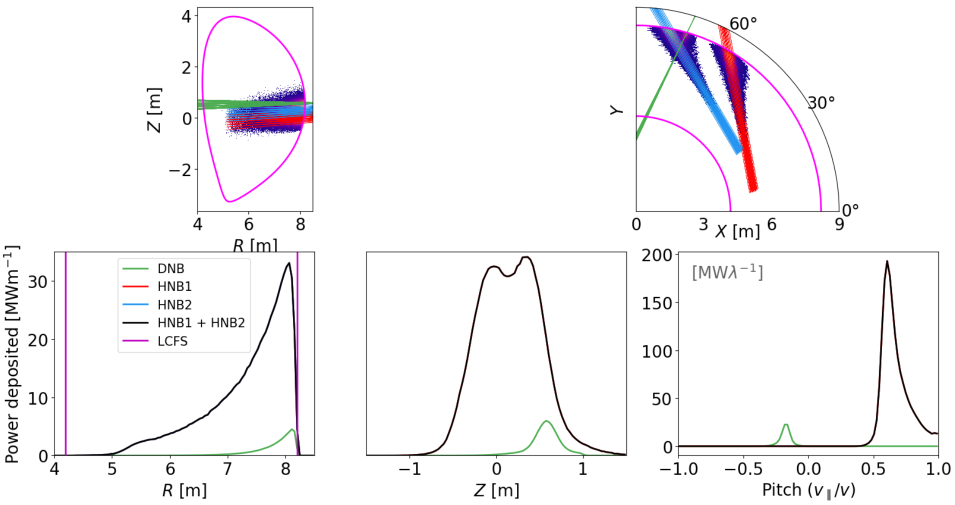

Figure 5.31: Geometry (top row) and deposition (top and bottom row) calculated by the DNB system in the $Q = 10$ ITER baseline scenario. Top left shows the elevation view, including the beam lines from the DNB (green, top), HNB1 (red, bottom) and HNB2 (blue, middle). Top right shows the plan view of a single quadrant of the device, showing the different tangency radii between beam types. Bottom row shows the deposited power from the DNB and combined HNB system along different projections. The concentration of the DNB at the edge of the plasma can be seen more clearly on the left, and the effect of the DNB geometry on the ion starting pitch can also be seen on the right to be much different to that of the HNB system. This may make it possible for FILDs to distinguish and identify fast ions specifically from the DNB system. Figure reproduced from [153].

5.3.2 DNB transport

To understand whether the DNB is a suitable tool for optimising the RMP field, the sensitivity of DNB transport to changes in the RMP is first established. As variation of the absolute RMP phase has been shown herein to be the optimal method for maximising global fast-ion confinement and minimising time-averaged power loads, the DNB confinement was measured over multiple values of absolute phase. To estimate the maximum transport, and likely maximum diagnostic signal, fast-ion losses were calculated during maximum ELM suppression - again through adjustment of the relative phase of the coil rows and maximisation of the ECC current. In addition, a more realistic scenario was studied in which the current was set to the lower value in table 5.1. To calculate the variability in transport over the full range of pitch and energy space, simulations were performed for longer time periods (500ms cut-off time) and with increased fidelity (2^{18} markers).

Normalised variation in the distribution function, measured over four phases within a rotation cycle, is shown for $n_0 = 3$ and $n_0 = 4$ fields in figures 5.32 and 5.33 respectively. The variation in transport is very regular in pitch-energy space across all scenarios and cases. Every case shows significant relative modulation of the fast-ion density near the injection location, where ions are born trapped. Because particles are injected from the DNB at such low pitch, near $\lambda \approx 0$, they can spread over a larger range of pitches than would be possible with a beam of the same energy injected at a high absolute pitch. And so in the time taken for the markers to thermalise, the beam spreads to cover the entire pitch domain, increasing its chances of detection by diagnostics - depending on the diagnostic weight function. In some cases, markers are significantly transported in pitch space - leading to a change in the trapped/passing fraction. $n_0 = 4$ fields are mostly likely to lead to this type of transport, and it is typically more likely at higher ECC currents. In these cases, modulation of the passing ion density is typically strongest at lower energies, as markers have had more time to scatter and populate this region.

To examine this transport more quantitatively, the DNB distribution function is plotted in figures 5.34 and 5.35, as a function of pitch and energy respectively, after being integrated over all other dimensions. The modulation in pitch in figures 5.32 and 5.33 can now be seen more clearly in figure 5.34 to correspond to a change in the pitch angle relaxation rate, likely caused by perturbed orbits accessing different regions of the plasma compared with the axisymmetric orbit. In some cases, the strength of this relaxation effect can be controlled by the perturbation phase. Interestingly however, fast-ion redistribution and fast-ion confinement do not always correlate; the fast-ion density in the $n_0 = 3$ variation of case 3 fluctuates very little whereas it produces the largest losses. A possible explanation is that fast-ion loss rate has become saturated by the RMP in a region of space, up to where the perturbation penetrates no further. However, this is unlikely, as the same behaviour occurs at lower RMP currents, where the fast-ion confinement is considerably improved (see figure 5.37). Instead, these losses, shown to remain high even at 10kAt in figure 5.16, may instead be too prompt to be affected by changes in the perturbation. In some cases, such as case 5 and case 7, and $n_0 = 3$ fields with low currents, there is a measured increase in the fast-ion density near the point of injection without a corresponding decrease elsewhere.

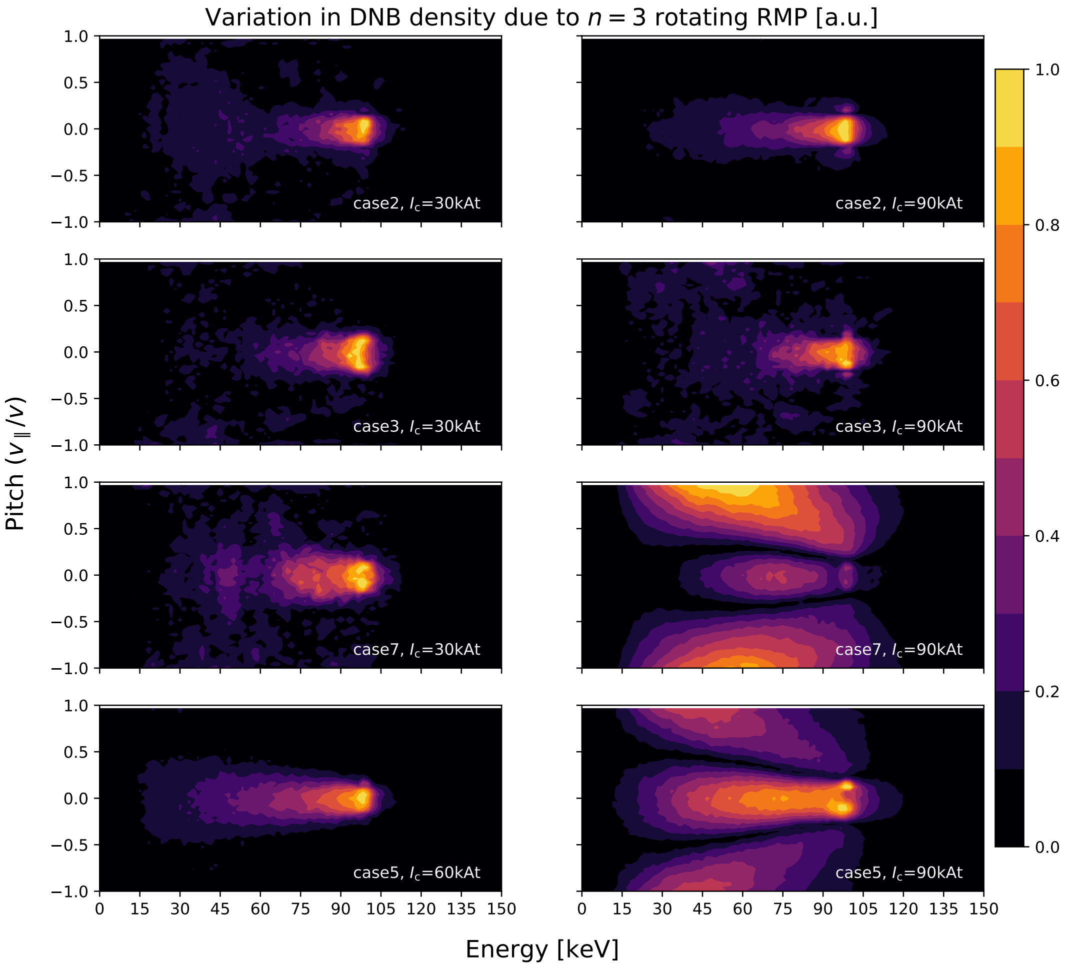

Figure 5.32: Normalised, absolute variation of the fast-ion distribution function measured over four phases during rotation of an applied $n_0 = 3$ RMP. The distribution function is integrated over all real space. Hot areas indicate regions of DNB ion phase-space which are relatively sensitive to rotations of the applied RMP field. Figure reproduced from [153].

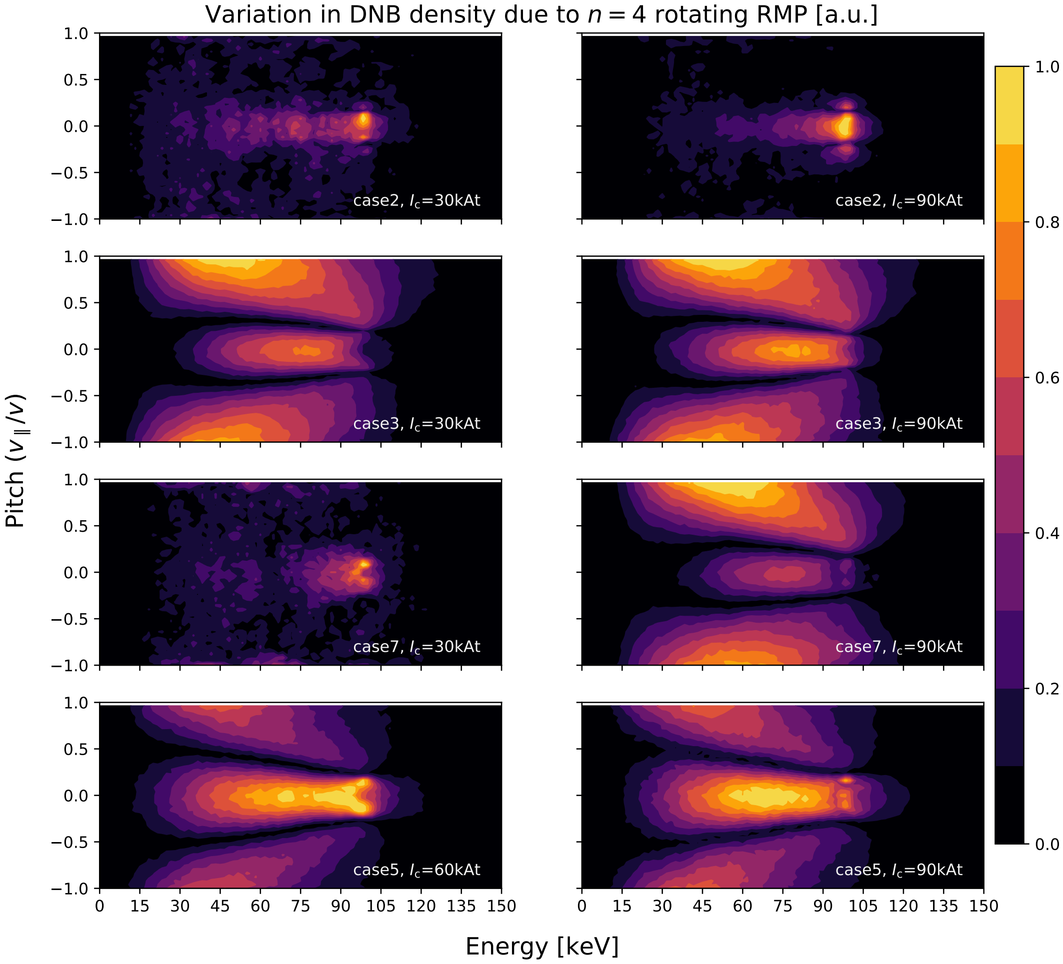

Figure 5.33: The $n_0 = 4$ equivalent to figure 5.32. The overall transport pattern remains the same, but the modulation is enhanced here. Figure reproduced from [153].

In these cases, DNB losses in 2D and 3D fields are comparable, and increasing the RMP field amplitude can even reduce losses. Hence it may be that the 3D field is improving the DNB confinement in some regions of phase space. Looking at the energy distribution of fast ions in figure 5.35, it can also be seen that the perturbation alters the prompt loss rate - as evident in the difference in fast-ion density at the injection energy. Again, this variation can be controlled by the perturbation phase in some cases.

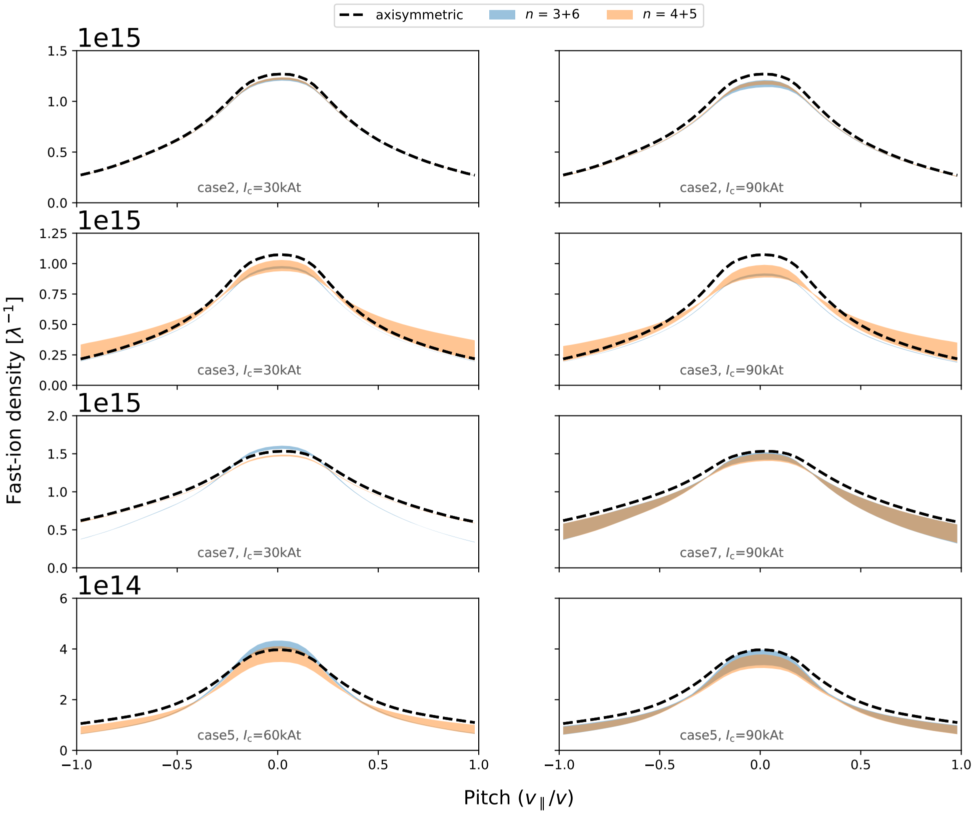

Figure 5.34: Steady-state distribution function of the ITER DNB in a rotating RMP field plotted as a function of pitch. The distribution function in an axisymmetric field without the RMP is also plotted for comparison in black. Figure reproduced from [153].

Whilst DNB transport may be sensitive to the RMP, a correlation between this transport and that measured for the HNB system would enable the DNB to act as a proxy in experiments. Figure 5.36 shows the DNB distribution function evaluated at $\lambda \sim 0$ as the perturbation is rotated. Though there is some stability in the relationship between density and phase, for the $n_0 = 3$ field for example, it is much more sensitive to RMP phase than the HNB losses shown in figure 5.19. Overall, the correlation between the DNB phase-space

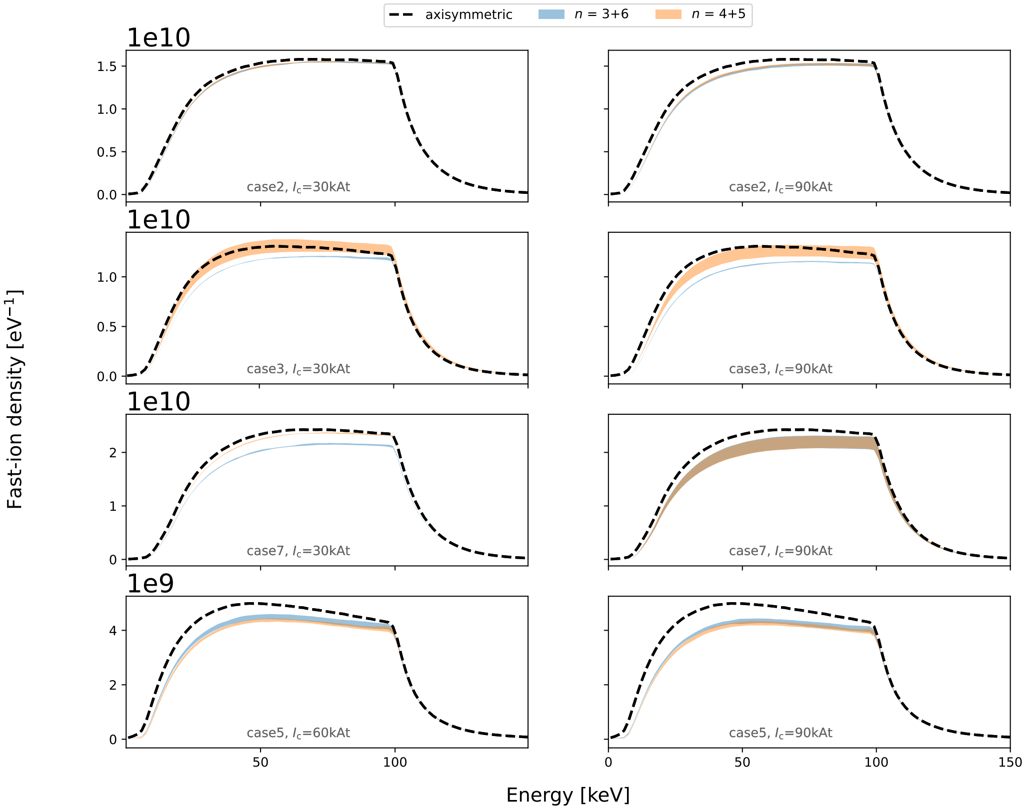

Figure 5.35: Steady-state distribution function of the DNB in a rotating RMP field plotted as a function of energy. The distribution function in an axisymmetric field without the RMP is also plotted for comparison in black. Figure reproduced from [153].

transport and HNB losses plotted in figures 5.2 and 5.19 is unclear.

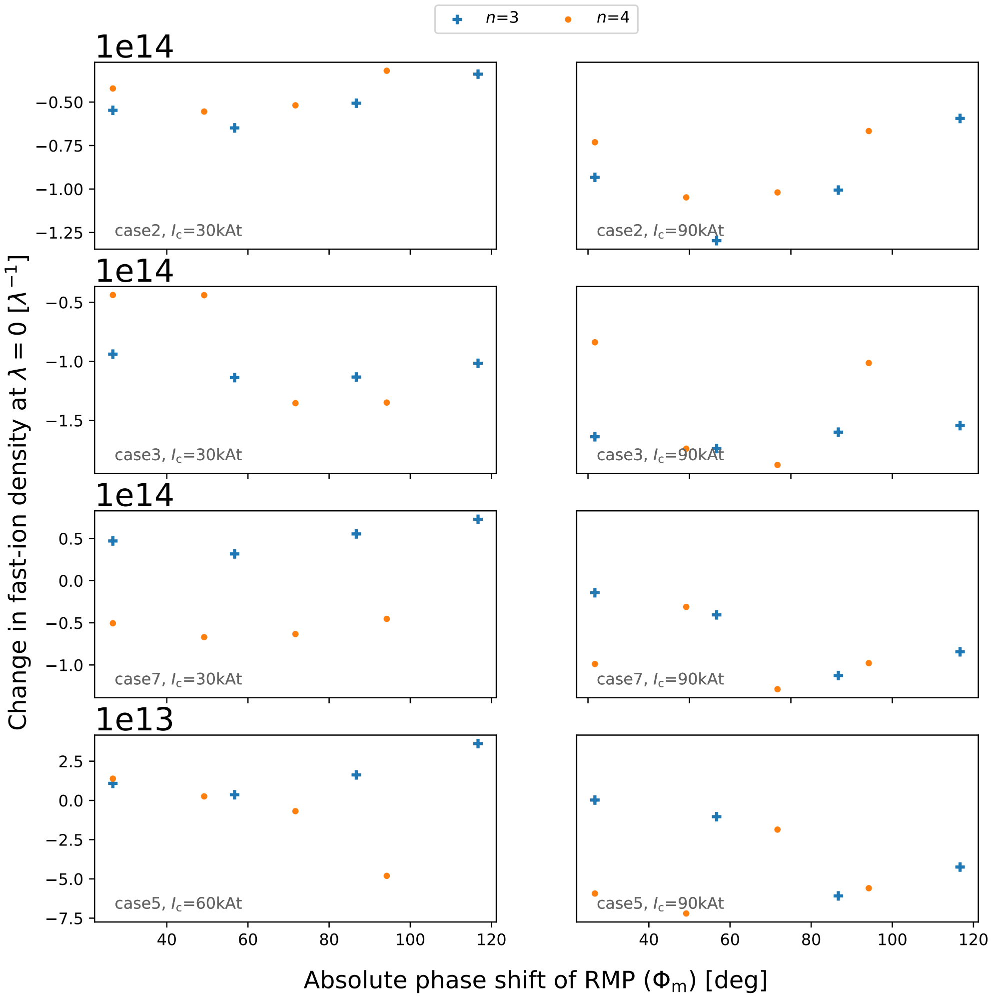

Figure 5.36: DNB fast-ion distribution function evaluated at $\lambda \sim 0$ and plotted against RMP absolute phase as represented by middle coil row phase. Figure reproduced from [153].

Finally, to receive a signal in the first place, a detector at the first wall must intercept lost DNB ions. The DNB losses to the PFCs were therefore calculated and are shown in figure 5.37 for two values of current - maximum and the reduced values in table 5.1. The overall fast-ion loss rate is steady, and it is much flatter than that exhibited by the HNB system - even at maximum RMP current. Where there is appreciable variation, in the $n_0 = 3$ variation of case 3, it does not clearly correlate to the equivalent HNB losses. And ultimately, the loss rate, especially at low RMP current, is extremely small. This makes it unlikely that any diagnostics adhering to the geometry used in this CAD model will detect lost DNB ions.

Hence, from the data shown here, it seems unlikely that the DNB is suitable for ECC system optimisation.

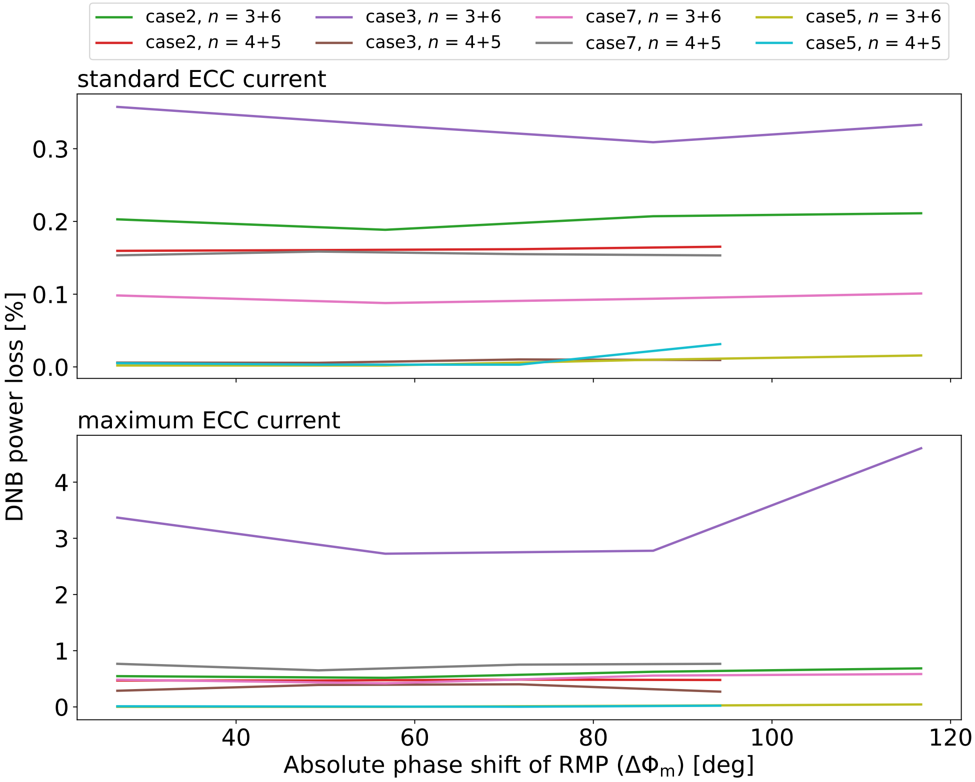

Figure 5.37: Global losses to PFCs in rotating RMP fields operated at maximum current and the reduced current shown in table 5.1. Figure reproduced from [153].

5.4 Summary

In this chapter the fast-ion workflow was used to quantitatively estimate the impact of externally imposed RMPs by the ITER ECC system on fast-ion transport and confinement.

In chapter 5.1, `LOCUST` was applied to study global fast-ion losses and local power fluxes for the $Q = 10$ ITER scenario with maximum RMP applied and optimised to achieved best ELM control. Depending on the symmetry of the RMP applied ($n_0 = 3$ or $n_0 = 4$) and absolute phase applied, the RMP fields were shown to reduce overall NBI heating efficiency by $\approx 2 - 12\%$ (0.7-4.0MW), though each of the two ITER heating neutral beams is affected in turn if the RMP is rotated. The toroidal mode number was found to be the most sensitive parameter for controlling fast-ion confinement, with $n_0 = 4$ modes found to be consistently worse for both global fast-ion confinement and PFC power loads. By adjusting the absolute phase of the RMP, NBI heating efficiency can be increased by 1.7-3.2%, whilst the fast-ion power loads can be redistributed. As such, rotating the RMP is useful to reduce the RMS power flux to most components. This is particularly true for the first-wall, where mid-plane tiles diametrically opposite to the HNB injection ports are subject to relatively intense power loads for some RMP phases. The optimisation of local power fluxes versus global NBI losses is deeply impacted by the n_0 of the applied RMP. For $n_0 = 3$ RMP fields, both local power fluxes and global NBI losses are directly correlated; in which case the optimal solution may be simply to fix the absolute RMP phase at the point where the first middle coil carries maximum current - as both global fast-ion losses and PFC power flux are minimised here. For $n_0 = 4$, use of RMP rotation may provide a better approach to optimisation, depending on the local power fluxes to optimise, since there is no straightforward correlation between local power fluxes and global NBI losses. However, it should be noted that, even in the worst-case scenario, the resulting power loads are typically a factor of 5-50 lower than the design value of the affected panel of the first-wall and divertor components in ITER.

In chapter 5.2, the unique speed and scalability of the workflow was used to massively extend the methodology used in chapter 5.1 to study additional types of discharge, levels of ELM suppression and methods of ECC operation. The response to changes in XPD created by scaling ECC current and modifying the relative phase of the ECC rows was measured. It was found that changes in the ECC current could lead to significant redistributions of fast-ion losses, the nature of which changes across ITER cases. This may have ramifications for future planned diagnostics designed around the monitoring or sensing of fast-ion power loads and losses - such as FILDs or infra-red cameras. Therefore, simulations covering a range of ECC currents and discharge scenarios will be necessary when optimising their placement. By measuring the fast-ion confinement at different levels of XPD, it was concluded that, as well as ELM suppression, XPD also correlates well with global fast-ion losses - despite the 3D nature of the system geometry. Whilst the gradient of this correlation changed between cases, its remained strong throughout, signifying that XPD forms a robust combined metric for both HNB heating efficiency and ELM suppression. Based on this, two methods of operating the ECCs, equivalent from the point of view of ELM suppression, were then compared: one in which the absolute phase of the perturbation was varied and one in which the perturbation was oscillated by changing the relative coil phase. Despite the same

XPD, there was a strong preference for the former method with respect to global fast-ion confinement. In this situation, the low-power discharges exhibited much less variation in total losses compared with the baseline scenario, making them relatively inappropriate from the point of view of pre-emptive optimisation if the RMP is chosen to be fixed at a single absolute phase. However, if the ITER baseline demands maximum ELM suppression, then these measurements showed that the best way of optimising the absolute phase may simply to operate first at lower ECC current, as the correlation between global fast-ion losses and first-wall power fluxes found at maximal ELM suppression remains at reduced ECC currents.

In chapter 5.3, a model for the DNB was implemented in `IMAS` and used in the fast-ion workflow. To test the feasibility of diagnosing the HNB fast-ion confinement through sensing of the DNB population, the confinement and transport of the beam was studied over the full slowing-down time. In doing so, the regions of fast-ion phase space most sensitive to variations in the RMP absolute phase were determined for potential future use with diagnostics designed to resolve the fast-ion distribution function - such as reciprocating FILDs. This variation was then quantified before being found to correlate poorly with the changes in global HNB confinement. Furthermore, the absolute fractions of DNB ions lost to the PFCs were found to be extremely small across the entire rotation cycle even at maximal ELM suppression.

Chapter 6

Summary

6.1 Conclusions

In this work, a high-fidelity physics workflow was assembled, tested and used to propose and verify methods of operating the ELM-control coil system in ITER, so as to maximise fast-ion confinement and thus optimise overall NBI heating efficiency and PFC power loads.

Software to enable the assembly, testing and automatic execution of this workflow was created. This achievement has demonstrated the feasibility of including the novel `LOCUST` fast-ion code as a component in integrated models. Whilst the workflow assembled here was constructed purely from Python, `LOCUST` was also adapted to read data from `IMAS` IDSs as part of this work, allowing additional `IMAS` workflows in the future to access the performance offered by `LOCUST`.

`LOCUST` was verified against experimental data and validated against prominent and reliable fast-ion codes, showing agreement to within expected Monte Carlo noise and theoretical limitations. This activity led to the discovery and removal of bugs in some of the codes used in this comparison, including ones at the centre of widely used integrated modelling platforms. It was found to be important to employ tests which populate as wide a region of fast-ion phase-space as possible, in order to maximise test coverage of implemented models, as well as the boundaries between them; for example, the plasma near and outside the LCFS and the implementation of the PFC wall were found to exacerbate finite Larmor radius effects which may be disregarded in certain codes. Achieving this coverage when comparing codes and experiments together was found to require studying tokamak discharges in which fast-ions are injected counter to the plasma current. Convergence tests were performed for the final workflow, as well as calibration of the 3D perturbation components. The effects of including additional toroidal harmonics in the 3D field and modelling plasma collisions were shown to be important when evaluating the fast-ion transport.

The assembled and tested workflow was then deployed to studying fast-ion transport driven by 3D fields in ITER. Overall, this transport, and the resulting power loads, was found to be greatly affected by the arrangement of the perturbation. The final positions of lost markers were highly correlated with the perturbation amplitude at the plasma edge as well as fast-ion pitch, confirming the mutual influence of orbit topology and perturbation structure on fast-ion transport.

The effects of varying the ELM-control coil operating parameters on fast-ion confinement were also studied. Ramping the current in the ECCs was found to steadily increase the fast-ion losses in most cases, however the exact behaviour varied across toroidal mode number and the type of plasma. This meant that neutral beam heating efficiency could fall by up to 12% at maximum ECC current, depending on the scenario. Hence it was recommended that, throughout the ITER experimental campaigns, ECC current should be minimised. It is likely that ITER operates below the maximum possible ECC current - as low as 45kAt [72]. At this point, fast-ion losses decrease, and the heating efficiency in the worst case scenario rises by ∼3% points. This is in line with the losses estimated in the ITER baseline scenario at maximum ECC current, which were found to be acceptable.

The impact of the relative phase between ECC rows on the fast-ion confinement was also evaluated. In combination with the effects of ECC current, it was found that the X-point displacement amplitude metric for ELM suppression also acts as an approximate measure of fast-ion confinement. For a particular plasma scenario and applied RMP spectrum, the metric was sufficiently robust that the Pearson correlation between XPD and loss rate of fast ions to the first wall was 0.84 (0.69 for the outer divertor target and baffle). This may signal that the conditions for ELM suppression mostly affect passing particles, which tend to be lost due to the chaotic motion of their guiding-centres which diffuse out of the stochastic layer at the plasma edge. This makes sense from the point of view of density pump-out, where ELM suppression correlates with the onset of stochasticity at the pedestal edge. Though it remains to be seen how much the perturbed X-point geometry directly dictates fast-ion confinement; given the lack of fast ions reaching the X-point, even those on loss orbits, it is likely small. The correlation between XPD and loss rate was not perfect, as losses varied whilst ECC system was guided through a configuration space which should maintain XPD. This could be due to the influence of higher toroidal harmonics, which perturb the XPD but are not accounted for in the calculation of XPD by `MARS-F`. Nevertheless, this discovery, which remains relatively robust despite the realism of the geometry studied, offers potential for the future development of reduced models. In this capacity, the reduction in HNB heating efficiency could be approximately calculated throughout the entire ECC configuration space by simulating the fast-ion loss rate at the two points of maximum and minimum XPD for a given plasma configuration and applied mode spectrum.

The workflow was then deployed to estimate the largest possible impact that ECCs would have on fast-ion confinement in the ITER baseline scenario. The reduction in heating efficiency was estimated at 2-12%, depending on the toroidal mode number, spectrum and phase of the perturbation. Of these parameters, toroidal mode number was found to be the most sensitive, with $n_0 = 4$ fields consistently leading to higher losses. However, for a given RMP spectrum, losses could be varied by 1.7-3.2% by adjusting the RMP phase without significantly affecting ELM suppression.

Fast-ion losses were also calculated in the situation where ELM suppression could be reduced. In this case, two parameter spaces were explored which both correspond to the same XPD: the absolute perturbation phase and position along XPD contours in coil configuration space. This study was expanded to include other ITER scenarios as well as the baseline. It was found that optimising the ECCs by varying the absolute phase typically leads to lower minimum and average losses. However, more optimal points along the XPD contour may be found if additional harmonics are taken into account in its calculation. Furthermore, it was found that moving along the XPD contour, unlike absolute phase, produces the same fast-ion response in most ITER operating scenarios - making experimental verification using low-power scenarios more straightforward. On the other hand, a relatively low-risk method for finding the optimal absolute RMP phase in the ITER baseline was found to be possible if the ECC current was lowered.

The power lost from the plasma to the PFCs, and ways to mitigate it, were also determined. PFC power loads in the ITER baseline were shown to be high yet tolerable. The power densities are largest at the first wall and focussed on panels diametrically opposing the HNB injection ports. The oscillation of the RMP to reduce the time-averaged power loads in any given spot was also studied. In the baseline scenario at lower XPD, power loads were lower when varying the absolute perturbation phase, as opposed to varying position along the XPD contour at increased current. The RMS power loads during this absolute RMP rotation were found to be $\approx$0.44MWm^{-2} lower than at peak for both $n_0 = 3$ and $n_0 = 4$ fields. For an $n_0 = 3$ field, this strategy may not be necessary, as it was found that the peak and average power loads correlate well with peak reductions in HNB heating efficiency. Therefore, the optimal solution may be to choose $n_0 = 3$ RMP fields whereby a single absolute phase is selected, at which both global and local fast-ion losses are optimised together.

Lastly, the diagnostic neutral beam was implemented, and the first simulations of DNB fast-ion transport were performed. The feasibility of using the DNB to experimentally verify the findings herein during low-power discharges was assessed. The regions of fast-ion phase space that are most sensitive to changes in RMP phase were determined. Transport in phase-space was found to be uncorrelated with HNB confinement. Furthermore, DNB losses to the first wall were also found to be extremely low, implying that detection of these losses with FILD systems will be difficult in ITER. Hence alternative methods may be required to verify the optimisations recommended here - for example, the use of infra-red cameras to observe the power loads predicted by `LOCUST`. Conversely, the direct detection of HNB ions themselves may prove adequate.

6.2 Future study

As a result of the work performed here, the question of fast-ion confinement in ITER has been elaborated, whilst new questions and avenues for ongoing research have been presented.

The workflow assembled herein now provides a platform upon which new models of fast-ion physics in ITER may be continuously integrated (CI). This is specifically possible due to the data translation and automation offered by `LOCUST_IO`; abstracting `LOCUST` in Python to leverage of Python's native test-driven developer utilities, such as Pytest; existing CI services, such as GitHub actions; and workflow tools, such as Luigi or those in `IMAS`. This ultimately makes the creation and automation of physics-informed CI easier. The adaptation of `LOCUST` to read from `IMAS` IDSs also means that CI can be extended to the workflow level, facilitating the quick and safe assembly of complex workflows through regression tests which specifically cover the integration of physics components.

The software infrastructure assembled here also enables users to leverage `LOCUST`'s speed in new ways to reach even higher levels of fidelity. For example, `LOCUST` can now be used as a component in the virtual development cycle of new tokamak components or systems, such as diagnostics. This could be trivially expanded to account for variations in upstream physics models, for example the 3D field, by running many variations of a simulation. Indeed this general methodology could be applied in any context - for example, an ensemble of `LOCUST` simulations could be performed, each using variations of a given input from different upstream models, the results from which may then be combined. For instance, PFC power loads calculated using a 3D field model that is accurate at the plasma edge may be combined with the core fast-ion density calculated using a 3D field more accurate in the core. Another context in which the automation provided here offers improved precision, outside of individual simulations, is the performance of many unrelated simulations. For example, independent parallelised simulations could be used to scan parameters, as was demonstrated here, or the workflow could feasibly be deployed iteratively as a hill-climbing tool, performing a local search in an arbitrary configuration space. Once this combination of physics component integration and GPGPU hardware is paired with novel numerical algorithms, for example those which ease the adjoint problem [160], a new regime of fast-ion physics fidelity will be accessible, potentially bringing the discovery of new physics.

As well as the addition of new software infrastructure on top of that assembled herein, there is also the possibility for swapping or improving sub-components. For example, the plasma data calculated by `ASTRA`, used in `MARS-F` and subsequently `LOCUST`, feature no model for representing the plasma outside the LCFS. The finite density at the plasma edge is however important to the edge beam deposition, the confinement of which is extremely sensitive to the 3D field. Therefore, physics models which extrapolate the plasma generated by fixed-boundary codes should be included in future incarnations of the workflow. Another factor which may potentially affect this edge beam deposition is the 3D displacement of the plasma by the perturbations - the magnitude of which can reach centimetres [161]. This is especially true for ITER, given that most ions are born close to the trapped-passing boundary, where orbit topology is highly sensitive to collisional transport and thus the plasma composition. Therefore, it should be a priority that this 3D plasma perturbation is taken into account in future beam models.

The thorough testing of `LOCUST` has prepared it for the study of many future reactor-scale devices. The inclusion of various tokamak geometries, such as spherical devices, allows `LOCUST` to be confidently deployed amongst novel designs for future reactor-scale devices - for example STEP and SPARC/ARC. Whilst it is difficult to verify the accuracy of the complete workflow when simulating ITER plasmas, it will soon become easier after the addition of a middle ECC row on the DIII-D device [162], which has already significantly contributed to the study of ELM-suppressed ITER-like plasmas.

The final workflow was extremely capable of performing extensive parameter scans, but there still exist regions which could not be explored due to computational constraints. For example, the variation in beam injection geometry was not fully explored. However, it should be noted that the simulations carried out in this work were performed on relatively modest GPU clusters - a maximum of 5 nodes at any one time and each node comprising only eight Nvidia P100 cards (or equivalent) - whereas it is not uncommon to see GPU systems featuring hundreds of equivalent cards. In this case, exploring the entire operating space of the ECC system is possible.

Whilst XPD was shown to be a relatively robust metric for fast-ion losses, as well as ELM suppression, the correlation does not always persist across plasmas and spectra. It would be useful to study exactly why this is the case - exactly how much the topology at the X-point influences fast-ion confinement. Using this reasoning to formulate reduced models, and the limits within which they remain accurate for a given tokamak system, would prove extremely useful for including in larger, slower, and more complex legacy integrated models. The decrease in the required computation would enable these workflows to very quickly determine the average fast-ion confinement of a given tokamak/plasma design across its operational space.

The predicted divertor power loads, which show an increased flux between cassette gaps, suggests that power loads from fast ions could be greatly impacted by PFC misalignment or deformation. Therefore, future studies may wish to include physics models for the evolution of the PFC surface [163], which can be re-meshed using the `IMAS` GGD [135] and iteratively re-entered into the simulation with `LOCUST` - resembling the methodology employed in [164]. This type of workflow could guide the iterative configuration and placement of components close or invasive to the plasma, such as FILDs or radio frequency antenna - which should aim to be included in future versions of the CAD model studied here.

The placement of FILDs in particular may also be informed by these simulations, which show that the ECC-driven fast-ion power flux to some PFCs may only appear after a threshold RMP current has been passed. Above this threshold, FILDs may receive a near-constant signal despite re-scaling of the ECC current - depending on the location of the FILD.

After showing the feasibility of varying the absolute and relative phases of the ECC rows in this work, future methods of operating the ECC system whilst maintaining ELM suppression may also be explored. For example, it may prove advantageous to combine the two approaches - rotating along the contour of relative phase between the rows whilst the absolute coil phase is also varied. If the ratio of the frequencies of these two independent cycles is tuned correctly, then the ECC system may be operated in a way where, at any given time over the course of an ITER discharge, the ECC system configuration, and resulting PFC power load footprint, will be unique. The automation enabled by the workflow assembled here would also allow the iterative determination of any arbitrary ergodic trajectory in the ECC operation space, along which the local fast-ion losses are minimised.

Finally, it is becoming clearer that innovation is not only required to demonstrate fusion but also deploy it at scale. It is hoped that the work herein demonstrates the potential for high-performance physics codes to inform and accelerate the development of fusion reactors which avoid technological lock-in, work synergistically at the grid level, and provide power in a manner that is accessible to everyone everywhere.

Chapter 7

Appendices

.1 Guiding-centre equations of motion

The guiding-centre equations of motion as derived in [16] that are solved by `LOCUST` are included here for reference:

$$\dot{\chi} = \frac{qB}{m}, \tag{1}$$

where χ represents gyro-phase,

$$\dot{\mu} = 0 \tag{2}$$

$$\dot{v_{\parallel}} = \frac{q}{mB^{\star}_{\parallel}} \mathbf{B}^{\star} \cdot \mathbf{E}^{\star} \tag{3}$$

$$\dot{\mathbf{R}} = v_{\parallel} \frac{\mathbf{B}^{\star}}{B^{\star}_{\parallel}} + \mathbf{E}^{\star} \times \frac{\hat{\mathbf{b}}}{B^{\star}_{\parallel}} \tag{4}$$

and where parallel terms are vector components in the direction of the magnetic field $\hat{\mathbf{b}} = \mathbf{B}/B$. The modified potentials $\mathbf{E}^{\star}$ and $\mathbf{B}^{\star}$ are given by

$$\mathbf{E}^{\star} \equiv -\frac{\partial \mathbf{A}}{\partial t} - \nabla \Phi^{\star} \tag{5}$$

$$\mathbf{B}^{\star} \equiv \nabla \times \mathbf{A}^{\star} \tag{6}$$

where

$$\Phi^{\star} \equiv \Phi + \frac{\mu B}{q} \tag{7}$$

$$\mathbf{A}^{\star} \equiv \mathbf{A} + \frac{m v_{\parallel}}{q}\hat{\mathbf{b}}. \tag{8}$$

.2 `TRANSP` settings

A noteworthy conclusion of this study is the environment required to accurately compare codes on scales that did not disguise discrepancies. For reference, included in table 1 below are the namelist settings required by `TRANSP` to achieve this fidelity.

Table 1: `TRANSP` namelist settings.

setting	value	comment
`GOOCON`	20	low orbit acceleration
`AVGTIM`	0.0	no statistical smoothing
`DXBSMOO`	0.0	no statistical smoothing
`NZONE_FB`	40	—
`NZONE_FP`	40	—
`NZONE_NB`	200	typically 20-60 *with* smoothing
`NZONES`	200	—
`XBMBND`	1.5	previously locked to 1.3
`NPTCLS`	10^7	1.7×10^5 in `ASCOT`, 1.3×10^5 in `LOCUST`
`WGHTA`	20.0	—

.3 Coulomb logarithm

The two definitions of the Coulomb logarithm, $\ln(\Lambda)$, used by `LOCUST` in this study are included here for reference. The complete expressions below describe the Coulomb logarithm in `NUBEAM`, whilst removing the terms labelled by tildes yields the definition used in `ASCOT`.

The Coulomb logarithm, for a test particle species i colliding with background plasma species p, is defined as

$$\ln(\Lambda_{\mathrm{ip}}) \equiv \ln\left(\frac{\mathrm{b}_{\mathrm{max,i}}}{\mathrm{b}_{\mathrm{min,ip}}}\right) \tag{9}$$

where the b terms denote the maximum and minimum impact parameters. $b_{\mathrm{max,i}}$ is defined as

$$b_{\mathrm{max,i}} = \left[\sum_{\mathrm{p}} \frac{\omega_{\mathrm{p}}^2 + \widetilde{\Omega_{\mathrm{p}}^2}}{\frac{T_{\mathrm{p}}}{m_{\mathrm{p}}} + \frac{\widetilde{2E_{\mathrm{i}}}}{m_{\mathrm{i}}}} \right]^{-\frac{1}{2}}, \tag{10}$$

with ω_{p} the plasma frequency,

$$\omega_{\mathrm{p}}^2 = \frac{n_{\mathrm{p}} q_{\mathrm{p}}^2}{m_{\mathrm{p}} \epsilon_0}, \tag{11}$$

and Ω_{p} the cyclotron frequency,

$$\Omega_{\mathrm{p}} = \frac{q_{\mathrm{p}}}{m_{\mathrm{p}}} B. \tag{12}$$

$b_{\mathrm{min,ip}}$ is defined as

$$b_{\mathrm{min,ip}} = \max\left(b_{\mathrm{min,ip}}^{\mathrm{quantum}}, b_{\mathrm{min,ip}}^{\mathrm{classical}} \right), \tag{13}$$

the largest minimum impact parameter out of the classical expression and the expression with quantum corrections:

$$b_{\mathrm{min,ip}}^{\mathrm{classical}} = \frac{q_{\mathrm{i}} q_{\mathrm{p}}}{4\pi\epsilon_0 \mu_{\mathrm{ip}} u_{\mathrm{ip}}^2} \tag{14}$$

$$b_{\mathrm{min,ip}}^{\mathrm{quantum}} = \frac{\hbar}{2\mu_{\mathrm{ip}} u_{\mathrm{ip}}} \widetilde{e^{-\frac{1}{2}}}, \tag{15}$$

where both define the effective relative velocity u_{ip} as

$$u_{\mathrm{ip}}^2 = \frac{3T_{\mathrm{p}}}{m_{\mathrm{p}}} + \frac{2E_{\mathrm{i}}}{m_{\mathrm{i}}} \tag{16}$$

and μ_{ip} the reduced mass

$$\mu_{\mathrm{ip}} = \frac{m_{\mathrm{i}} m_{\mathrm{p}}}{m_{\mathrm{i}} + m_{\mathrm{p}}}. \tag{17}$$

List of references

[1] D. D. Clayton. *Principles of stellar evolution and nucleosynthesis.* University of Chicago press, 1983.

[2] S. Atzeni and J. Meyer-ter Vehn. *The physics of inertial fusion: beam plasma interaction, hydrodynamics, hot dense matter*, volume 125. OUP Oxford, 2004.

[3] W. Commons. Binding energy curve - common isotopes, 2022.

[4] N. Peacock. Procs. 3rd conf. on plasma phys. and controlled nucl. *Fusion, Novosibirsk*, 2, 1968.

[5] T. Nicholas, T. Davis, F. Federici, J. Leland, B. Patel, C. Vincent, and S. Ward. Re-examining the role of nuclear fusion in a renewables-based energy mix. *Energy Policy*, 149:112043, 2021.

[6] V. Masson-Delmotte, et al. *Global Warming of 1.5 OC: An IPCC Special Report on the Impacts of Global Warming of 1.5 C Above Pre-industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty.* World Meteorological Organization Geneva, Switzerland, 2018.

[7] Lazard's Levelized Cost of Energy Analysis. Technical Report October, Lazard, 2020.

[8] World Energy Outlook 2021. Technical report, IEA, 2021.

[9] T. H. Rider. A general critique of inertial-electrostatic confinement fusion systems. *Physics of Plasmas*, 2(6):1853, 1995.

[10] T. H. Rider. Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium. *Physics of plasmas*, 4(4):1039, 1997.

[11] S. E. Wurzel and S. C. Hsu. Progress toward fusion energy breakeven and gain as measured against the lawson criterion. *Physics of Plasmas*, 29(6):062103, 2022.

[12] E. ITER et al. Plasma confinement and transport. *Nuclear Fusion*, 39(12 ITER physics basis):2175, 1999.

[13] I. Organization. Iter research plan within the staged approach. *ITER technical report ITR-18-003*, 2018.

[14] J. R. Cary and A. J. Brizard. Hamiltonian theory of guiding-center motion. *Reviews of modern physics*, 81(2):693, 2009.

[15] H. Goldstein, C. Poole, and J. Safko. Classical mechanics, 2002.

[16] R. G. Littlejohn. Variational principles of guiding centre motion. *Journal of Plasma Physics*, 29(1):111, 1983.

[17] A. Lal and S. M. Badiger. An analogy of a magnetic mirror in mechanics. *European journal of physics*, 26(5):883, 2005.

[18] J. Wesson and D. J. Campbell. *Tokamaks*, volume 149. Oxford university press, 2011.

[19] P. Helander, R. Akers, and L.-G. Eriksson. On neutral-beam injection counter to the plasma current. *Physics of plasmas*, 12(11):112503, 2005.

[20] J. Hedin, T. Hellsten, L.-G. Eriksson, and T. Johnson. The influence of finite drift orbit width on icrf heating in toroidal plasmas. *Nuclear Fusion*, 42(5):527, 2002.

[21] R. Hemsworth, et al. Overview of the design of the iter heating neutral beam injectors. *New Journal of Physics*, 19(2):025005, 2017.

[22] F. Hinton and R. D. Hazeltine. Theory of plasma transport in toroidal confinement systems. *Reviews of Modern Physics*, 48(2):239, 1976.

[23] K. Ghantous, N. Gorelenkov, H. Berk, W. Heidbrink, and M. Van Zeeland. 1.5 d quasilinear model and its application on beams interacting with alfvén eigenmodes in diii-d. *Physics of Plasmas*, 19(9):092511, 2012.

[24] M. Podesta et al. A reduced fast ion transport model for the tokamak transport code TRANSP. *Plasma Physics and Controlled Fusion*, 56(5):055003, 2014.

[25] M. Podesta, M. Gorelenkova, N. Gorelenkov, R. White, P. Bonofiglo, F. Poli, A. Teplukhina, J. Yang, M. Cecconello, and M. Vallar. Development of a reduced model for energetic particle transport by sawteeth in tokamaks. *Plasma Physics and Controlled Fusion*, 64(2):025002, 2021.

[26] T. Kurki-Suonio, O. Asunta, E. Hirvijoki, T. Koskela, A. Snicker, T. Hauff, F. Jenko, E. Poli, and S. Sipilä. Fast ion power loads on iter first wall structures in the presence of ntms and microturbulence. *Nuclear Fusion*, 51(8):083041, 2011.

[27] C. Cheng. Kinetic extensions of magnetohydrodynamics for axisymmetric toroidal plasmas. *Physics reports*, 211(1):1, 1992.

[28] M. Fitzgerald, J. Buchanan, R. J. Akers, B. N. Breizman, and S. Sharapov. Halo: A full-orbit model of nonlinear interaction of fast particles with eigenmodes. *Computer Physics Communications*, 252:106773, 2020.

[29] S. Chandrasekhar. Stochastic problems in physics and astronomy. *Reviews of modern physics*, 15(1):1, 1943.

[30] M. N. Rosenbluth, W. M. MacDonald, and D. L. Judd. Fokker-planck equation for an inverse-square force. *Physical Review*, 107(1):1, 1957.

[31] F. L. Hinton. Collisional transport in plasma. *Handbook of Plasma Physics*, 1:147, 1983.

[32] A. H. Boozer. Monte carlo collision operators for use with exact trajectory integrators. *Physics of Plasmas*, 9(10):4389, 2002.

[33] A. Boozer et al. Monte Carlo evaluation of transport coefficients. *The Physics of Fluids*, 24(5):851, 1981.

[34] A. Galeev and R. Sudan. *Basic plasma physics.* Number v. 1 in Handbook of plasma physics. North-Holland Pub., 1983.

[35] W. Heidbrink and R. White. Mechanisms of energetic-particle transport in magnetically confined plasmas. *Physics of Plasmas*, 27(3):030901, 2020.

[36] W. Heidbrink, J. M. Park, M. Murakami, C. Petty, C. Holcomb, and M. Van Zeeland. Evidence for fast-ion transport by microturbulence. *Physical review letters*, 103(17):175001, 2009.

[37] S. Pinches, et al. The role of energetic particles in fusion plasmas. *Plasma physics and controlled fusion*, 46(12B):B187, 2004.

[38] N. N. Gorelenkov, S. Pinches, and K. Toi. Energetic particle physics in fusion research in preparation for burning plasma experiments. *Nuclear Fusion*, 54(12):125001, 2014.

[39] S. Pinches, I. Chapman, P. W. Lauber, H. Oliver, S. Sharapov, K. Shinohara, and K. Tani. Energetic ions in ITER plasmas. *Physics of Plasmas*, 22(2):021807, 2015.

[40] G. Fu and J. Van Dam. Excitation of the toroidicity-induced shear alfvén eigenmode by fusion alpha particles in an ignited tokamak. *Physics of Fluids B: Plasma Physics*, 1(10):1949, 1989.

[41] S. Pinches, et al. Observation and modelling of fast ion loss in jet and asdex upgrade. *Nuclear fusion*, 46(10):S904, 2006.

[42] M. Fitzgerald, S. Sharapov, P. Rodrigues, and D. Borba. Predictive nonlinear studies of tae-induced alpha-particle transport in the q= 10 ITER baseline scenario. *Nuclear Fusion*, 56(11):112010, 2016.

[43] M. Van Zeeland, et al. Alfvén eigenmode stability and fast ion loss in DIII-D and ITER reversed magnetic shear plasmas. *Nuclear Fusion*, 52(9):094023, 2012.

[44] I. Chapman, et al. Empirical scaling of sawtooth period for onset of neoclassical tearing modes. *Nuclear Fusion*, 50(10):102001, 2010.

[45] B. Geiger et al. Fast-ion transport and neutral beam current drive in ASDEX upgrade. *Nuclear Fusion*, 55(8):083001, 2015.

[46] B. Geiger, et al. Fast-ion transport in the presence of magnetic reconnection induced by sawtooth oscillations in asdex upgrade. *Nuclear Fusion*, 54(2):022005, 2014.

[47] A. Snicker et al. Power loads to ITER first wall structures due to fusion alphas in a non-axisymmetric magnetic field including the presence of MHD modes. *Nuclear Fusion*, 53(9):093028, 2013.

[48] I. Chapman, et al. The physics of sawtooth stabilization. *Plasma Physics and Controlled Fusion*, 49(12B):B385, 2007.

[49] M. Garcia-Munoz, et al. Fast-ion losses induced by elms and externally applied magnetic perturbations in the asdex upgrade tokamak. *Plasma Physics and Controlled Fusion*, 55(12):124014, 2013.

[50] R. White. Chaos in trapped particle orbits. *Physical Review E*, 58(2):1774, 1998.

[51] G. Saibene, et al. Toroidal field ripple effects on h-modes in jet and implications for ITER. In *Proc. 34th EPS Conf. on Plasma Physics Controlled Fusion (Warsaw, Poland, 2007).* 2007.

[52] R. Boivin, S. Zweben, and R. White. Study of stochastic toroidal field ripple losses of charged fusion products at the midplane of tftr. *Nuclear fusion*, 33(3):449, 1993.

[53] G. J. Kramer, R. White, R. Nazikian, and H. Berk. Fusion-born alpha particle ripple loss studies in ITER. *Princeton Plasma Physics Laboratory Preprint.–2009.–March.–PPPL-4363*, 2009.

[54] R. J. Goldston, R. White, and A. H. Boozer. Confinement of high-energy trapped particles in tokamaks. *Physical review letters*, 47(9):647, 1981.

[55] T. Kurki-Suonio, et al. ASCOT simulations of fast ion power loads to the plasma-facing components in ITER. *Nuclear Fusion*, 49(9):095001, 2009.

[56] T. Kurki-Suonio et al. Protecting ITER walls: fast ion power loads in 3D magnetic field. *Plasma Physics and Controlled Fusion*, 59(1):014013, 2016.

[57] F. Wagner, et al. Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the asdex tokamak. *Physical Review Letters*, 49(19):1408, 1982.

[58] A. Loarte, et al. Progress on the application of elm control schemes to iter scenarios from the non-active phase to dt operation. *Nuclear Fusion*, 54(3):033007, 2014.

[59] E. Daly, et al. Update on design of the iter in-vessel coils. *Fusion Science and Technology*, 64(2):168, 2013.

[60] S. Ward, R. Akers, A. Loarte, L. Li, Y. Liu, S. Pinches, A. Polevoi, R. Vann, and M. V. Zeeland. Locust-gpu predictions of fast-ion transport and power loads due to elm-control coils in iter. *Nuclear Fusion*, 62(12):126014, 2022.

[61] T. Evans, D. Orlov, A. Wingen, W. Wu, A. Loarte, T. Casper, O. Schmitz, G. Saibene, M. Schaffer, and E. Daly. 3D vacuum magnetic field modelling of the ITER ELM control coil during standard operating scenarios. *Nuclear Fusion*, 53(9):093029, 2013.

[62] T. Evans. Resonant magnetic perturbations of edge-plasmas in toroidal confinement devices. *Plasma Physics and Controlled Fusion*, 57(12):123001, 2015.

[63] T. Evans, I. Joseph, R. Moyer, M. Fenstermacher, C. Lasnier, and L. Yan. Experimental and numerical studies of separatrix splitting and magnetic footprints in diii-d. *Journal of Nuclear Materials*, 363:570, 2007.

[64] L. Li, Y. Liu, A. Loarte, S. Pinches, A. Polevoi, Y. Liang, and F. Zhong. Modeling 3d plasma boundary corrugation and tailoring toroidal torque profiles with resonant magnetic perturbation fields in iter. *Nuclear Fusion*, 59(9):096038, 2019.

[65] L. Li, Y. Liu, A. Loarte, S. Pinches, A. Polevoi, and F. Zhong. Elm control optimization for various iter scenarios based on linear and quasi-linear figures of merit. *Physics of Plasmas*, 27(4):042510, 2020.

[66] Y. Liu, A. Kirk, Y. Gribov, M. Gryaznevich, T. Hender, and E. Nardon. Modelling of plasma response to resonant magnetic perturbation fields in mast and iter. *Nuclear Fusion*, 51(8):083002, 2011.

[67] Y. Liu, A. Kirk, and E. Nardon. Full toroidal plasma response to externally applied nonaxisymmetric magnetic fields. *Physics of Plasmas*, 17(12):122502, 2010.

[68] S. Kim, et al. Nonlinear modeling of the effect of n= 2 resonant magnetic field perturbation on peeling-ballooning modes in kstar. *Nuclear Fusion*, 60(2):026009, 2020.

[69] M. Wade, et al. Advances in the physics understanding of elm suppression using resonant magnetic perturbations in diii-d. *Nuclear Fusion*, 55(2):023002, 2015.

[70] Q. Hu, R. Nazikian, B. Grierson, N. Logan, D. Orlov, C. Paz-Soldan, and Q. Yu. Wide operational windows of edge-localized mode suppression by resonant magnetic perturbations in the diii-d tokamak. *Physical Review Letters*, 125(4):045001, 2020.

[71] Q. Hu, et al. Predicting operational windows of elms suppression by resonant magnetic perturbations in the diii-d and kstar tokamaks. *Physics of Plasmas*, 28(5):052505, 2021.

[72] Q. Hu, J.-K. Park, N. Logan, S. Yang, B. Grierson, R. Nazikian, and Q. Yu. Nonlinear two-fluid modeling of plasma response to rmps for the elm control in the iter baseline. *Nuclear Fusion*, 61(10):106006, 2021.

[73] F. Waelbroeck. Theory and observations of magnetic islands. *Nuclear Fusion*, 49(10):104025, 2009.

[74] R. Fitzpatrick. Bifurcated states of a rotating tokamak plasma in the presence of a static error-field. *Physics of Plasmas*, 5(9):3325, 1998.

[75] F. Waelbroeck. Shielding of resonant magnetic perturbations in the long mean-free path regime. *Physics of plasmas*, 10(10):4040, 2003.

[76] M. Becoulet, et al. Screening of resonant magnetic perturbations by flows in tokamaks. *Nuclear Fusion*, 52(5):054003, 2012.

[77] D. Pfefferlé, C. Misev, W. A. Cooper, and J. P. Graves. Impact of rmp magnetic field simulation models on fast ion losses. *Nuclear Fusion*, 55(1):012001, 2014.

[78] E. Nardon, et al. Edge localized mode control experiments on MAST using resonant magnetic perturbations from in-vessel coils. *Plasma Physics and Controlled Fusion*, 51(12):124010, 2009.

[79] S. R. Haskey, M. J. Lanctot, Y. Liu, J. M. Hanson, B. 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, 2014.

[80] M. De Rover, A. Schilham, A. Montvai, and N. L. Cardozo. Test particle transport in perturbed magnetic fields in tokamaks. *Physics of Plasmas*, 6(6):2443, 1999.

[81] M. A. Van Zeeland et al. Fast ion transport during applied 3D magnetic perturbations on DIII-D. *Nuclear Fusion*, 55(7):073028, 2015.

[82] K. McClements, et al. The effects of resonant magnetic perturbations on fast ion confinement in the mega amp spherical tokamak. *Plasma Physics and Controlled Fusion*, 57(7):075003, 2015.

[83] K. McClements, K. Tani, R. Akers, Y. Liu, K. Shinohara, H. Tsutsui, and S. Tsuji-Iio. The effects of resonant magnetic perturbations and charge-exchange reactions on fast ion confinement and neutron emission in the mega amp spherical tokamak. *Plasma Physics and Controlled Fusion*, 60(9):095005, 2018.

[84] J. Varje et al. Effect of plasma response on the fast ion losses due to elm control coils in ITER. *Nuclear Fusion*, 56(4):046014, 2016.

[85] K. Särkimäki et al. Mechanics of elm control coil induced fast particle transport in ITER. *Nuclear Fusion*, 58(7):076021, 2018.

[86] L. Sanchis, et al. Optimizing beam-ion confinement in iter by adjusting the toroidal phase of the 3d magnetic fields applied for elm control. *Nuclear Fusion*, 61(4):046006, 2021.

[87] L. Sanchis, et al. Characterisation of the fast-ion edge resonant transport layer induced by 3d perturbative fields in the asdex upgrade tokamak through full orbit simulations. *Plasma Physics and Controlled Fusion*, 61(1):014038, 2018.

[88] K. Särkimäki. Efficient and rigorous evaluation of fast particle losses in non-axisymmetric tokamak plasmas. *Nuclear Fusion*, 60(3):036002, 2020.

[89] Y. Liu, L. Li, A. Loarte, S. Pinches, and A. Polevoi. Drift orbit islands of energetic particles due to 3d fields in iter. *Nuclear Fusion*, 61(10):106029, 2021.

[90] H. E. Mynick. Transport of energetic ions by low-n magnetic perturbations. *Physics of Fluids B: Plasma Physics*, 5(5):1471, 1993.

[91] L. Krlín and P. Cahyna. Particle diffusion in a system of magnetic islands in tokamaks in fully hamiltonian approach. *Czechoslovak Journal of Physics*, 56(2):B111, 2006.

[92] P. Cahyna and L. Krlín. Full hamiltonian description of the interaction of particles with magnetic islands in tokamak. *Czechoslovak Journal of Physics*, 56(4):367, 2006.

[93] K. Shinohara, A. Bierwage, Y. Suzuki, J. Kim, G. Matsunaga, M. Honda, and T. Rhee. Estimation of orbit island width from static magnetic island width, using safety factor and orbit pitch. *Nuclear Fusion*, 58(8):082026, 2018.

[94] K. Shinohara, A. Bierwage, A. Matsuyama, Y. Suzuki, G. Matsunaga, M. Honda, S. Sumida, and J. Kim. Efficient estimation of drift orbit island width for passing ions in a shaped tokamak plasma with a static magnetic perturbation. *Nuclear Fusion*, 60(9):096032, 2020.

[95] L. Yu, E. Xue, D. Zhang, S. Zheng, X. Zhang, J. Huang, E. Li, Z.-M. Sheng, and J. Fu. Simulation of the loss of passing fast ions induced by magnetic islands in east tokamak plasmas. *AIP Advances*, 11(2):025020, 2021.

[96] K. He, Y. Sun, B. Wan, S. Gu, M. Jia, and Y. Hu. Full-orbit simulation of fast ion loss under resonant magnetic perturbations in the east tokamak. *Nuclear Fusion*, 61(1):016009, 2020.

[97] Y. Liu, et al. Toroidal modeling of plasma response and resonant magnetic perturbation field penetration. *Plasma Physics and Controlled Fusion*, 54(12):124013, 2012.

[98] D. Ryan, et al. 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, 2015.

[99] K. He, Y. Sun, B. Wan, S. Gu, and M. Jia. Roles of primary and sideband resonances in the confinement of energetic passing ions under resonant magnetic perturbations. *Nuclear Fusion*, 60(12):126027, 2020.

[100] M. Heyn, I. B. Ivanov, S. Kasilov, W. Kernbichler, A. Loarte, V. Nemov, and A. M. Runov. On the confinement of passing alpha particles in a tokamak-reactor with resonant magnetic field perturbations shielded by plasma currents. *Nuclear fusion*, 52(5):054010, 2012.

[101] E. Nardon, et al. Strike-point splitting induced by external magnetic perturbations: observations on jet and mast and associated modelling. *Journal of nuclear materials*, 415(1):S914, 2011.

[102] C. Su, S. Chen, M. Mou, W. Guo, X. Chen, and C. Tang. Effect of resonant magnetic perturbations with plasma response on particle orbit characters. *Physica Scripta*, 96(3):035604, 2021.

[103] O. Asunta, S. Äkäslompolo, T. Kurki-Suonio, T. Koskela, S. Sipilä, A. Snicker, M. García-Muñoz, et al. Simulations of fast ion wall loads in asdex upgrade in the presence of magnetic perturbations due to elm-mitigation coils. *Nuclear Fusion*, 52(9):094014, 2012.

[104] R. Akers, B. Colling, J. Hess, Y. Liu, A. Turner, S. Äkäslompolo, J. Varje, K. Särkimäki, S. Pinches, and M. Singh. High fidelity simulations of fast ion power flux driven by 3d field perturbations on iter (iaea-cn–234), international atomic energy agency (iaea). Technical report, 2016.

[105] Y. Xu, L. Li, Y. Hu, Y. Liu, W. Guo, L. Ye, and X. Xiao. Numerical simulations of nbi fast ion loss with rmps on the east tokamak. *Nuclear Fusion*, 60(8):086013, 2020.

[106] O. Meneghini, et al. Integrated modeling applications for tokamak experiments with OMFIT. *Nuclear Fusion*, 55(8):083008, 2015.

[107] M. Romanelli, et al. Jintrac: a system of codes for integrated simulation of tokamak scenarios. *Plasma and Fusion research*, 9:3403023, 2014.

[108] F. Imbeaux, et al. Design and first applications of the iter integrated modelling & analysis suite. *Nuclear Fusion*, 55(12):123006, 2015.

[109] S. Ward, R. Akers, A. Jacobsen, P. Ollus, S. Pinches, E. Tholerus, R. Vann, and M. V. Zeeland. Verification and validation of the high-performance lorentz-orbit code for use in stellarators and tokamaks (LOCUST). *Nuclear Fusion*, 61(8):086029, 2021.

[110] G. L. Delzanno et al. On particle movers in cylindrical geometry for particle-in-cell simulations. *Journal of Computational Physics*, 253:259, 2013.

[111] K. Tretiak and D. Ruprecht. An arbitrary order time-stepping algorithm for tracking particles in inhomogeneous magnetic fields. *Journal of Computational Physics: X*, 4:100036, 2019.

[112] E. Fehlberg. *Low-order classical Runge-Kutta formulas with stepsize control and their application to some heat transfer problems.* National aeronautics and space administration, Washington, D.C., 1969.

[113] J. R. Cash and A. H. Karp. A variable order Runge-Kutta method for initial value problems with rapidly varying right-hand sides. *ACM Transactions on Mathematical Software (TOMS)*, 16(3):201, 1990.

[114] J. R. Dormand and P. J. Prince. A family of embedded Runge-Kutta formulae. *Journal of computational and applied mathematics*, 6(1):19, 1980.

[115] D. Goeken and O. Johnson. Fifth-order Runge-Kutta with higher order derivative approximations. *Electronic Journal of Differential Equations*, 2(November):1, 1999.

[116] J. P. Boris and R. A. Shanny. *Proceedings: Fourth Conference on Numerical Simulation of Plasmas, November 2, 3, 1970.* Naval Research Laboratory, Washington, D.C., 1972.

[117] E. Hirvijoki et al. ASCOT: Solving the kinetic equation of minority particle species in tokamak plasmas. *Computer Physics Communications*, 185(4):1310, 2014.

[118] A. Pankin, D. McCune, R. Andre, G. Bateman, and A. Kritz. The tokamak Monte Carlo fast ion module NUBEAM in the national transport code collaboration library. *Computer Physics Communications*, 159(3):157, 2004.

[119] M. D. Kruskal. *Advanced Plasma Theory, Proceedings of the International School of Physics Course XXV.* Academic Press Inc., 1964.

[120] T. A. Wareing, J. M. McGhee, and J. E. Morel. ATTILA: A three-dimensional, unstructured tetrahedral mesh discrete ordinates transport code. *Transactions of the American Nuclear Society*, 75, 1996.

[121] H. Lütjens, A. Bondeson, and O. Sauter. The chease code for toroidal mhd equilibria. *Computer Physics Communications*, 97(3):219, 1996.

[122] L. Lao, H. S. John, R. Stambaugh, A. Kellman, and W. Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611, 1985.

[123] G. Huysmans and O. Czarny. MHD stability in x-point geometry: simulation of ELMs. *Nuclear Fusion*, 47(7):659, 2007.

[124] Y. 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.

[125] C.-M. Fransson, B. Lennartson, C. Breitholtz, A. Bondeson, and Y. Liu. Feedback stabilization of nonaxisymmetric resistive wall modes in tokamaks. ii. control analysis. *Physics of Plasmas*, 7(10):4143, 2000.

[126] S. Jardin et al. Multiple timescale calculations of sawteeth and other global macroscopic dynamics of tokamak plasmas. *Computational Science & Discovery*, 5(1):014002, 2012.

[127] O. Asunta, J. Govenius, R. Budny, M. Gorelenkova, G. Tardini, T. Kurki-Suonio, A. Salmi, and S. Sipilä. Modelling neutral beams in fusion devices: Beamlet-based model for fast particle simulations. *Computer Physics Communications*, 188:33, 2015.

[128] M. Schneider, L.-G. Eriksson, I. Jenkins, J. Artaud, V. Basiuk, F. Imbeaux, T. Oikawa, J.-E. contributors, et al. Simulation of the neutral beam deposition within integrated tokamak modelling frameworks. *Nuclear Fusion*, 51(6):063019, 2011.

[129] B. Geiger, et al. Progress in modelling fast-ion d-alpha spectra and neutral particle analyzer fluxes using FIDASIM. *Plasma Physics and Controlled Fusion*, 62(10):105008, 2020.

[130] F. Assous, P. Degond, and J. Segre. A particle-tracking method for 3d electromagnetic pic codes on unstructured meshes. *Computer Physics Communications*, 72(2):105, 1992.

[131] M. Matsumoto and T. Nishimura. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 8(1):3, 1998.

[132] G. V. Pereverzev and P. Yushmanov. Astra. automated system for transport analysis in a tokamak. 2002.

[133] C. Chrystal, B. Grierson, G. Staebler, C. Petty, W. Solomon, J. DeGrassie, K. Burrell, T. Tala, and A. Salmi. Predicting rotation for iter via studies of intrinsic torque and momentum transport in diii-d. *Physics of Plasmas*, 24(5):056113, 2017.

[134] L. Li, Y. Liu, A. Loarte, S. Pinches, A. Polevoi, and F. Zhong. Toroidal modeling of resonant magnetic perturbations in preparation for the initial phase of ITER operation. *Nuclear Fusion*, 60(1):016013, 2019.

[135] L. Kos, R. Pitts, G. Simič, M. Brank, H. Anand, and W. Arter. Smiter: A field-line tracing environment for iter. *Fusion Engineering and Design*, 146:1796, 2019.

[136] R. Lake. *Consequences of fast ion driven modes in MAST*. Ph.D. thesis, University of Warwick, 2013.

[137] M. Van Zeeland, et al. Measurements and modeling of alfvén eigenmode induced fast ion transport and loss in DIII-D and asdex upgrade. *Physics of Plasmas*, 18(5):056114, 2011.

[138] M. Van Zeeland et al. Modulation of prompt fast-ion loss by applied n=2 fields in the DIII-D tokamak. *Plasma Physics and Controlled Fusion*, 56(1):015009, 2013.

[139] J. Varje, K. Särkimäki, J. Kontula, P. Ollus, T. Kurki-Suonio, A. Snicker, E. Hirvijoki, and S. Äkäslompolo. High-performance orbit-following code ASCOT5 for Monte Carlo simulations in fusion plasmas. *arXiv preprint arXiv:1908.02482*, 2019.

[140] J. Breslau, M. Gorelenkova, F. Poli, J. Sachdev, and X. Yuan. TRANSP. Computer Software `https://doi.org/10.11578/dc.20180627.4`, 2018.

[141] C. Ordonez and M. Molina. Evaluation of the coulomb logarithm using cutoff and screened coulomb interaction potentials. *Physics of plasmas*, 1(8):2515, 1994.

[142] J. Krommes. An introduction to the physics of the coulomb logarithm, with emphasis on quantum-mechanical effects. *Journal of Plasma Physics*, 85(1), 2019.

[143] J. Harrison, et al. Overview of new MAST physics in anticipation of first results from MAST Upgrade. *Nuclear Fusion*, 59(11):112011, 2019.

[144] C. Michael, et al. Dual view FIDA measurements on MAST. *Plasma Physics and Controlled Fusion*, 55(9):095007, 2013.

[145] A. Sperduti, M. Cecconello, S. Conroy, and A. Snicker. Neutron rate estimates in MAST based on gyro-orbit modelling of fast ions. *Nuclear Fusion*, 61(1):016028, 2020.

[146] O. Jones. *Experimental fast-ion transport studies on the Mega-Amp Spherical Tokamak.* Ph.D. thesis, Durham University, 2015.

[147] A. Kolmogorov. Sulla determinazione empirica di una lgge di distribuzione. *Inst. Ital. Attuari, Giorn.*, 4:83, 1933.

[148] G. Kramer et al. A description of the full-particle-orbit-following spiral code for simulating fast-ion experiments in tokamaks. *Plasma Physics and Controlled Fusion*, 55(2):025013, 2013.

[149] M. Cecconello, W. Boeglin, D. Keeling, S. Conroy, I. Klimek, R. Perez, et al. Discrepancy between estimated and measured fusion product rates on MAST using TRANSP/NUBEAM. *Nuclear Fusion*, 59(1):016006, 2018.

[150] T. Evans, et al. Rmp elm suppression in diii-d plasmas with iter similar shapes and collisionalities. *Nuclear fusion*, 48(2):024002, 2008.

[151] T. Evans, et al. Elm suppression in helium plasmas with 3d magnetic fields. *Nuclear Fusion*, 57(8):086016, 2017.

[152] M. Singh, D. Boilson, A. Polevoi, T. Oikawa, and R. Mitteau. Heating neutral beams for iter: negative ion sources to tune fusion plasmas. *New Journal of Physics*, 19(5):055004, 2017.

[153] S. Ward, R. Akers, A. Loarte, L. Li, Y. Liu, S. Pinches, A. Polevoi, and R. Vann. Optimising fast-ion transport due to ELM-control coils in ITER discharges with LOCUST-GPU (in preparation). 2022.

[154] L. Sanchís. personal communication.

[155] R. Pitts, et al. Physics basis and design of the iter plasma-facing components. *Journal of Nuclear Materials*, 415(1, Supplement):S957, 2011. Proceedings of the 19th International Conference on Plasma-Surface Interactions in Controlled Fusion.

[156] R. Mitteau, et al. The design of the iter first wall panels. *Fusion Engineering and Design*, 88(6):568, 2013. Proceedings of the 27th Symposium On Fusion Technology (SOFT-27); Liège, Belgium, September 24-28, 2012.

[157] R. A. Pitts, A. Kukushkin, A. Loarte, A. Martin, M. Merola, C. E. Kessel, V. Komarov, and M. Shimada. Status and physics basis of the ITER divertor. *Physica Scripta*, T138:014001, 2009.

[158] W. Heidbrink. Fast-ion d α measurements of the fast-ion distribution. *Review of Scientific Instruments*, 81(10):10D727, 2010.

[159] M. Garcia-Munoz et al. Conceptual design of the ITER fast-ion loss detector. *Review of Scientific Instruments*, 87(11):11D829, 2016.

[160] E. Hirvijoki. Eliminating poor statistics in monte-carlo simulations of fast-ion losses to plasma-facing components and detectors. *arXiv preprint arXiv:1905.04952*, 2019.

[161] I. Chapman, et al. Assessing the merits of resonant magnetic perturbations with different toroidal mode numbers for controlling edge localised modes. *Nuclear Fusion*, 54(12):123003, 2014.

[162] D. Weisberg, C. Paz-Soldan, Y. Liu, and N. Logan. Optimizing multi-modal, non-axisymmetric plasma response metrics with additional coil rows on diii-d. *Nuclear Fusion*, 59(8):086060, 2019.

[163] E. Thorén, S. Ratynskaia, P. Tolias, R. A. Pitts, et al. The memos-u code description of macroscopic melt dynamics in fusion devices. *Plasma Physics and Controlled Fusion*, 63(3):035021, 2021.

[164] J. Coburn, et al. Energy deposition and melt deformation on the iter first wall due to disruptions and vertical displacement events. *Nuclear Fusion*, 62(1):016001, 2021.