

This is a repository copy of *The modular Weyl–Kac character formula*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/193106/>

Version: Published Version

---

**Article:**

Bowman, Christopher [orcid.org/0000-0001-6046-8930](https://orcid.org/0000-0001-6046-8930), Hazi, Amit and Norton, Emily (2022) The modular Weyl–Kac character formula. *Mathematische Zeitschrift*. ISSN 1432-1823

<https://doi.org/10.1007/s00209-022-03084-7>

---

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

PAPER • OPEN ACCESS

## LOCUST-GPU predictions of fast-ion transport and power loads due to ELM-control coils in ITER

To cite this article: S.H. Ward *et al* 2022 *Nucl. Fusion* **62** 126014

View the [article online](#) for updates and enhancements.

You may also like

- [Jump stabilization and landing control by wing-spreading of a locust-inspired jumper](#)  
Avishai Beck, Valentin Zaitsev, Uri Ben Hanan et al.
- [Design and demonstration of a bio-inspired flapping-wing-assisted jumping robot](#)  
Ngoc Thien Truong, Hoang Vu Phan and Hoon Cheol Park
- [Verification and validation of the high-performance Lorentz-orbit code for use in stellarators and tokamaks \(LOCUST\)](#)  
S.H. Ward, R. Akers, A.S. Jacobsen et al.

# LOCUST-GPU predictions of fast-ion transport and power loads due to ELM-control coils in ITER

S.H. Ward<sup>1,2,3,\*</sup>, R. Akers<sup>2</sup>, L. Li<sup>4</sup>, Y.Q. Liu<sup>5</sup>, A. Loarte<sup>3</sup>,  
S.D. Pinches<sup>3</sup>, A. R. Polevoi<sup>3</sup>, R.G.L. Vann<sup>1</sup> and M.A. Van Zeeland<sup>5</sup>

<sup>1</sup> York Plasma Institute, Department of Physics, University of York, York YO10 5DD, United Kingdom

<sup>2</sup> Culham Centre For Fusion Energy, Culham Science Centre, Abingdon, OX14 3DB, United Kingdom

<sup>3</sup> ITER Organization, Route de Vinon-sur-Verdon, CS 90 046, 13067 St. Paul Lez Durance Cedex, France

<sup>4</sup> College of Science, Donghua University, Shanghai 201620, China

<sup>5</sup> General Atomics, PO Box 85608, San Diego, CA 92186-5608, United States of America

E-mail: [samuel.ward@york.ac.uk](mailto:samuel.ward@york.ac.uk)

Received 14 January 2022, revised 30 July 2022

Accepted for publication 8 September 2022

Published 5 October 2022



CrossMark

## Abstract

Lorentz-orbit code for use in stellarators and tokamaks (LOCUST)-graphics processing unit has been applied to study the fast-ion transport and loss caused by resonant magnetic perturbations in the high-performance  $Q = 10$  ITER baseline scenario. The unique computational efficiency of the code is exploited to calculate the impact of the application of the ITER edge-localised modes (ELM)-control-coil system on neutral beam heating efficiency, as well as producing detailed predictions of the resulting plasma-facing component power loads, for a variety of operational parameters—the applied fundamental toroidal mode number  $n_0$ , mode spectrum and absolute toroidal phase of the imposed perturbation. The feasibility of continually rotating the perturbations is assessed and shown to be effective at reducing the time-averaged power loads. Through careful adjustment of the relative phase of the applied perturbation in the three rows of coils, peak power loads are found to correlate with reductions in neutral beam injection (NBI) heating efficiency for  $n_0 = 3$  fields. Adjusting the phase this way can increase total NBI system efficiency by approximately 2%–3% and reduce peak power loads by up to  $0.43 \text{ MW m}^{-2}$ . From the point of view of fast-ion confinement,  $n_0 = 3$  ELM control fields are preferred overall to  $n_0 = 4$  fields. In addition, the implementation of 3D magnetic fields in LOCUST is also verified by comparison with the SPIRAL code for a DIII-D discharge with ITER-similar shaping and  $n_0 = 3$  perturbation.

Keywords: tokamak, ITER, 3D fields, RMP, LOCUST, fusion, energetic particles

(Some figures may appear in colour only in the online journal)

\* Author to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

## 1. Introduction

The ITER tokamak aims to operate in high-confinement mode (H-mode) [1] to achieve high-gain ( $Q \sim 10$ ), stationary ( $\sim 500$  s) plasma discharges [2]. However, the additional power expelled during the type-I edge-localised modes (ELM) typically observed in H-mode plasmas poses a risk to plasma performance and can significantly reduce the lifetime of plasma-facing components (PFC), limiting the operational availability of ITER-scale devices [3]. To mitigate or suppress ELMs, ITER is equipped with a set of ELM-control coils (ECCs) that impose resonant magnetic perturbations (RMP) onto the plasma, breaking the underlying tokamak equilibrium axisymmetry.

Experimental and computational studies have suggested that RMPs lead to increased transport of fast ions [4]. Particles are predicted to be mostly affected where the RMP amplitude is greatest, often at the edge of the plasma, where the neutral beam injection (NBI) deposits part of its intense, anisotropic source of fast ions. Hence many studies corroborate the impact that the application of ECCs has on NBI ions in particular—specifically, the resulting reduction in beam ion confinement, NBI heating efficiency and associated increase in power flux to PFCs. All of these impacts must be mitigated if ITER is to achieve its aims.

The distribution of the PFC power flux is thought to be dictated by the influence on fast-ion orbits by the various magnetic structures that exist within the RMP field [5–9]. These structures are influenced in turn by the design and operation of the ECCs [10]—that is, the geometry of each coil and current it carries—in combination with the plasma’s response to the externally imposed RMP field [11, 12], which may reduce [13] or amplify the RMP field [4] and resulting transport [12]. Therefore, the potential impact that ITER’s high-power NBI system (a 33 MW flux of 1 MeV deuterons in the initial baseline configuration [14]) poses to the power loads on PFCs [15] (designed to tolerate up to  $10 \text{ MW m}^{-2}$  maximum at the divertor, less elsewhere [16]) could be aggravated or mitigated depending on how the ECC system is operated. This potential optimisation is why the ongoing study of ECC operation, plasma response models, and their combined influence on the PFC power flux is necessary to support the refinement of the ITER research plan [2].

Despite the importance of this topic, it remains challenging to study computationally due to the extreme scale of the ITER device. Kinetic fast-ion codes which aim at accurately resolving PFC power loads must cope with a huge spatiotemporal domain (1 s slowing down time,  $\sim 800 \text{ m}^3$  plasma volume), model expansive yet intricate wall geometries, and resolve fine structures in RMP fields. This bottlenecks attempts to systematically simulate the ITER system without significant computational resources, making the optimisation problem difficult. However, the Lorentz-orbit code for use in stellarators and tokamaks (LOCUST) code [17] is a novel kinetic fast-ion algorithm that is designed to make reactor-scale systems tractable on desktop hardware. While this could potentially provide a way of discovering new physics, in this

context LOCUST also enables the routine study of ITER at high fidelity, allowing for precise optimisation—for example through the minimisation of localised, component-specific power loads.

This paper aims to evaluate methods of mitigating reductions in NBI heating efficiency and localised PFC power loads due to ECCs by predicting the related fast-ion transport in ITER. The distribution of lost NBI power is calculated for applied  $n_0 = 3$  and  $n_0 = 4$  ECC waveforms, which can be oscillated to rotate the RMP toroidally while maintaining ELM suppression. Here,  $n_0$  is used to indicate the fundamental toroidal mode number of the applied perturbation. To resolve the PFC power at the component level, we use the LOCUST code to model fast-ion dynamics in the presence of detailed and realistic models of the first wall and RMP field. Because of this, we have also verified the implementation and convergence of the 3D magnetic field model in LOCUST.


After describing and testing the 3D magnetic field model in section 2, the ITER ECC system, and the physical model used to represent it, is described in section 3. Section 4 then presents the results of studying this model with LOCUST, including the measured fast-ion transport and PFC power loads at different RMP phases. Finally, section 5 presents a summary and outlook of the work.

## 2. 3D magnetic fields in LOCUST


LOCUST is a high-performance kinetic fast-ion code. It utilises hardware acceleration via programmable graphics processing units (GPU) and tuned software algorithms to massively parallelise the calculation of fast-ion slowing-down trajectories at speeds unobtainable with traditional central processing unit (CPU) codes. LOCUST achieves this by assigning fast-ion markers to individual GPU threads, which each solve the particle Lorentz equation of motion and periodically apply a Fokker–Planck collision operator. This allows for the efficient, full-orbit tracking of large populations (millions) of markers over long periods of time (seconds). For this reason, the code is uniquely suited for calculating fast-ion transport and detailed PFC power loads in ITER.

To ensure accuracy when extrapolating to ITER, we tested the implementation of 3D magnetic fields in LOCUST by applying it to an already operating device where the impact of RMPs on NBI fast-ion confinement has been assessed experimentally and by other modelling tools. The code has previously been shown [17] to compare well with other fast-ion codes and experiment in a variety of scenarios with axisymmetric plasmas, including those with ITER-similar shapes such as DIII-D shot #157418 [7]. Here we study DIII-D shot #157418 again, including the  $n_0 = 3$  RMP applied during the discharge.

Using the same plasma input data from [7], the corresponding results from LOCUST were compared to the simulations presented in [7] by SPIRAL [18]. SPIRAL is a widely accepted code that is regularly used in comparisons with experiment. Though SPIRAL’s speed and reliance on conventional CPU hardware makes it unable to study ITER at high



**Figure 1.** The normalised, element-wise difference— $\delta f(\epsilon, \lambda)$ —between  $f(\epsilon, \lambda)$  calculated with and without a pure  $n = 3$  RMP field, where  $f(\epsilon, \lambda)$  is the distribution function integrated over all dimensions except energy and pitch. Filled contours are calculated by SPIRAL while the white intermediate contours are calculated by LOCUST. A perfect match corresponds to white contours perfectly defining the boundaries between the filled contours. In both simulations the RMP leads to transport in the same region of phase space, while quantitative discrepancies are limited to high energies and are caused by physics unrelated to the 3D magnetic field model.



**Figure 2.** The element-wise difference,  $\delta f(\lambda)$ , expressed as a fraction of  $f_{2D}(\lambda)$  calculated without the RMP field. Included are results from LOCUST simulations using a simple 2D outline for the axisymmetric wall and a defeatured CAD model for the 3D wall. The small difference at high pitch can be explained by differences in the electron drag—not differences in the 3D magnetic fields.

fidelity without significant computational resources, its wide use makes the code very well-suited for verification exercises. LOCUST used a higher precision version of the perturbed field and plasma response calculated by M3D-C1 [19], generated on a grid with 3.7 mm spacing to ensure convergence. LOCUST used the same neutral beam deposition as SPIRAL, which was generated by beams injected both co and counter to the plasma current, including all three energy components as well as the effects of the plasma displacement on the temperature and density profiles due to the applied RMP.

Figures 1 and 2 show the changes in the steady-state fast-ion distribution function,  $f$ , caused by the RMP. The distribution function is cropped to include only the plasma edge, in regions satisfying  $\rho > 0.77$ , where  $\rho$  represents the


square root of the normalised toroidal flux. Though this cut-off is different in [7], where  $\rho > 0.7$ , the small volume within  $0.7 < \rho < 0.77$ , and the resulting influence of finite bin-widths, means manual tuning was required to capture this sensitive region. Figure 1 shows the difference between  $f$  calculated with and without the 3D magnetic field,  $f_{3D}(\epsilon, \lambda) - f_{2D}(\epsilon, \lambda)$ , as a function of energy  $\epsilon$  and particle pitch angle  $\lambda = v_{||}/v$  measured against the plasma current. There is good qualitative agreement between the codes; both show transport in the same region of phase space, corresponding to edge-localised, co-passing particles. Of these particles, those with energies above 40 keV—half injection energy—are affected somewhat differently between the simulations. Figure 2 illustrates this more clearly, showing the RMP-induced transport as a fraction of the axisymmetric distribution function,  $[f_{3D}(\lambda) - f_{2D}(\lambda)]/f_{2D}(\lambda)$ . The effect of these high-energy markers is a 10% point difference in the transport within  $0.8 < \lambda < 1.0$ ,  $\epsilon > 40$  keV. This could be because SPIRAL, unlike LOCUST, extends the electron density profile past the last-closed flux surface, generating increased drag at higher energies for edge-localised markers. Likewise, LOCUST did not include plasma rotation, impurities or electric fields. Both 2D and 3D representations of the first wall were used in LOCUST to gauge the effects of wall model variations. Compared with SPIRAL, which uses a 3D representation, the variations are small.

Despite the small difference found in the results, these results show that there are no errors in LOCUST's implementation of 3D magnetic fields.


### 3. ITER ELM-control coil fields

#### 3.1. The ITER ECC system

The ITER ECC system is illustrated in figure 3. It consists of three rows of nine, regularly spaced window frame coils, with centres starting at  $\phi = 30^\circ, 26.7^\circ$  and  $30^\circ$  for the upper, middle and lower rows respectively. Each coil is independently



**Figure 3.** Top left: view of the ITER ECC system (not to scale). The three rows of rectangular coils are responsible for imposing 3D perturbations onto the plasma by each passing a current defined by equation (1). 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 RMP field. Bottom: distribution of ECC coils in poloidal ( $\theta$ ) and toroidal ( $\phi$ ) 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 RMP spectrum.




**Figure 4.** Current distribution (magenta) defined by equation (1) for  $n_0 = 3$  and  $n_0 = 4$  RMPs in ITER ECCs (black quadrangles where height represents current passed), 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 magenta current profile. In both cases, the finite number of coils means the higher harmonic significantly contributes to the total waveform.

powered and consists of six windings passing a maximum of 15 kA each (90 kAT total magnetomotive force) [20]. To impose RMPs of a given fundamental toroidal mode number,  $n_0$ , each ECC will operate with a current defined by


$$I(\phi_{\text{coil}}) = I_0 \cos(n_0[\phi_{\text{coil}} - \Phi] - \omega t), \quad (1)$$

where  $\phi_{\text{coil}}$  is the toroidal location of the coil centre,  $\omega$  (5 Hz [3]) is a rotational frequency and  $\Phi$  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,  $I_0$ ,  $n_0$ ,  $\Phi_{u,m,1}$ ,  $\omega$  are the controllable degrees of freedom, subject to the operational limitations and a desired level of ELM suppression. It is worth noting that the toroidal spectrum of an RMP imposed by coils of finite width and

number is often largely composed of harmonics higher than the fundamental toroidal mode number,  $n_0$ , where the spectrum  $n$  is made up of  $n_0, n_1$  etc. For example, the ratio of amplitudes of the first and second harmonics are approximately 70%–85% for  $n = 3, 6$  and 88%–95% for  $n = 4, 5$  in ITER, depending on the coil geometry (which differs between rows). Importantly, these additional harmonics, which sometimes rotate counter to the coil current profile and fundamental harmonic, lead to variations in the poloidal spectrum throughout rotation—potentially altering the fast-ion dynamics and resulting PFC footprint. Typical spectra for ITER are illustrated in figure 4. Included are only the first two harmonics, as these by far have the largest amplitude in  $n_0 = 3$  and  $n_0 = 4$  mode spectra.



**Figure 5.** 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. 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 10 cm, with both decreasing approximately linearly. Losses saturate at roughly 3 cm; to give some margin for error, a grid spacing of 1 cm subsequently then chosen for all simulations.




**Figure 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 under-dome pipes, while 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 in figure 10.

### 3.2. RMP and plasma model


The 3D magnetic fields created by the ITER ECC system used herein have been calculated previously [21, 22]. The plasma response was calculated by MARS-F [23, 24] using plasma parameters determined by ASTRA [25] under various assumptions of plasma Prandtl number (ratio of toroidal momentum to thermal diffusivity in the core) and ratio of toroidal momentum to thermal confinement times— $(\tau_\phi/\tau_E)$ . Together, these affect the plasma rotation at the core and edge respectively, in turn influencing the plasma's ability to screen the external perturbation. 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. If anything, this amounts to linearly scaling the total fast-ion losses [26]. The fields used here are based on equilibria calculated assuming  $\tau_\phi/\tau_E = 2$  and a Prandtl number of 0.3 as expected from turbulent transport

simulations for ITER [27] using TGLF [28], which correspond to a smaller X-point displacement (XPD) amplitude, due to the low plasma resistivity and relatively high absolute value of the toroidal rotation [29] of 10–30 kRad  $s^{-1}$ , which is low when normalised to the Mach number in ITER. Like in [21, 22], the XPD, the plasma displacement normal to the flux surfaces at the X-point, is taken as a metric for ELM suppression [30].

The fields calculated by MARS-F are decomposed into toroidal harmonics for each coil row and are accessible via `mhd_linear` interface data structures (IDS) in the ITER integrated modelling and analysis suite (IMAS) [31]. `mhd_linear` IDSs are storage containers native to ITER's IMAS software ecosystem that enforce a data schema, allowing for the storage and exchange of simulation data between physics codes. This allows for rapidly simulating an arbitrary level of ELM suppression and ECC configuration by linearly rescaling, phase-shifting and re-combining the fields from individual coil rows prior to reading by LOCUST—without



**Figure 7.** Plasma data representing the  $Q = 10$  ITER DT scenario in MARS-F and subsequent LOCUST calculations. Included here are the safety factor  $q$ , electron and ion temperatures  $T_e$  and  $T_i$ , electron density  $n_e$  and rotation profile  $\omega_\phi$  plotted against normalised poloidal flux  $\psi_n$ . The safety factor and rotation play a crucial role in determining the plasma response to RMPs, while the plasma density and temperature dictate the location of NBI-deposited fast-ions and their collision frequencies. The high plasma density leads to many fast-ions being deposited at the plasma edge, where the RMP amplitude is strongest.



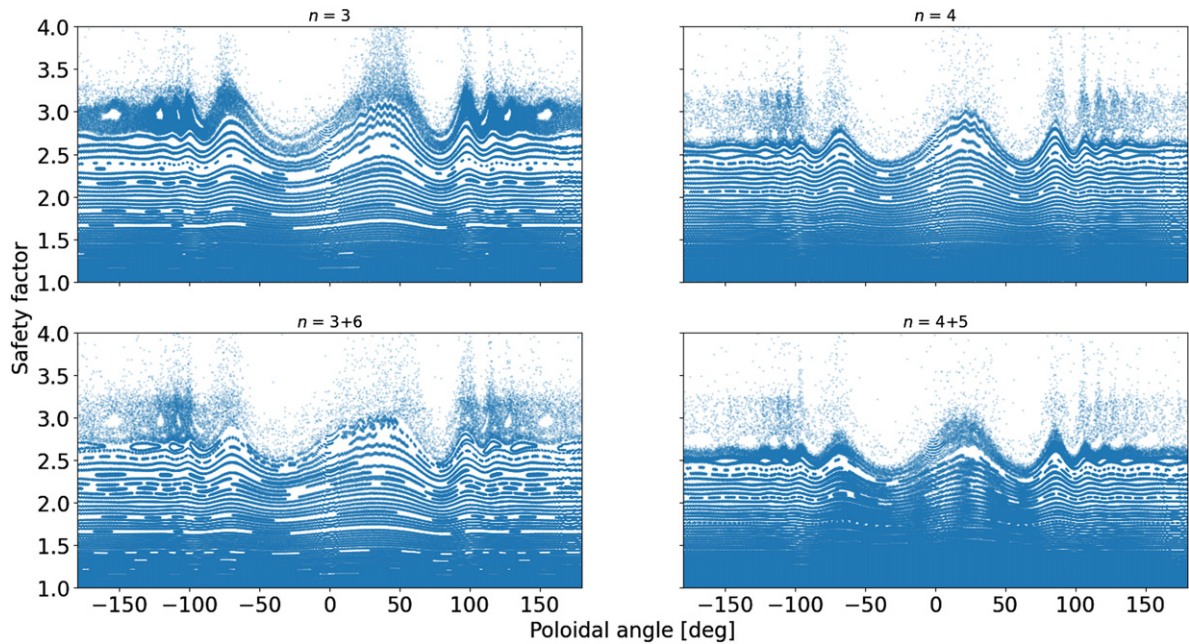
**Figure 8.** Global fast-ion losses to PFCs as a percentage of deposited beam power ( $\approx 33$  MW) for different middle coil row phases,  $\Phi_m$ , while maintaining relative phases between all coil rows. This reflects the situation in which the RMP is rotated to spread the fast-ion power loads. 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 passing maximum current ( $\Phi_m = 26.7^\circ$ ). 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,  $\Phi$  remains as defined in equation (1)—that is, in the ITER machine coordinate system. Hence, absolute phase is varied through  $360^\circ/n_0$ .

re-running MARS-F. The total field is then generated by interpolating to a Cartesian  $R - Z$  grid, where high-order splines are used near the separatrix region to resolve the fine structures created by thin current sheets. As this may be different to methods used in other fast-ion codes, figure 5 shows a simple convergence test to gauge the required grid precision when including the first two harmonics. The NBI power loss was calculated over multiple simulations in which only the perturbation grid spacing changed.  $\nabla \cdot \mathbf{B}$  was evaluated and averaged over all points within the plasma, before being scaled by one of two lengths:  $B$  on-axis and either the Boris integrator 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  $\mathcal{O}(\Delta x_{\mathbf{B}})$ . Convergence of fast-ion losses begins below 10 cm, however it takes until  $\approx 3$  cm for fast-ion losses to stabilise. Therefore, a grid size of 1 cm was chosen.

A similar convergence test was also performed for particle time step, with the time distribution of losses converging for a Boris time step of  $\approx 1$  ns (approximately 25 steps of 1 cm per gyroperiod for a 1 MeV deuteron). Finally, losses were found to near saturation after  $\approx 30$  ms, which was then used as a time cut-off in all simulations. It was therefore assumed that the time-dependent fast-ion losses due to an ECC system oscillating at 0–5 Hz could be well-approximated by separate simulations using static RMP fields at discrete phase intervals; i.e. we approximate  $\omega = 0$ .





**Figure 9.** Poincaré maps for fields at the same  $\Phi_m$  but different toroidal mode numbers and spectra, plotted against safety factor  $q$  and poloidal angle  $\theta$ . The maps are evaluated at  $\phi = 60^\circ$ . The patterns near  $\theta = 0^\circ$  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 the presence of the second harmonic changes the field structure at the edge.

For consistency, the same axisymmetric equilibrium and plasma parameters, such as temperature and density profiles, used in MARS-F from ASTRA, were also used here by the collision operator in LOCUST. No specific model was used to extend the plasma into the scrape-off layer. For simplicity, and due to observing a negligible effect on high-energy fast-ion losses, impurities were ignored. The BBNBI [32] IMAS actor was used to calculate a realistic fast-ion deposition from the heating neutral beam system into the axisymmetric equilibrium plasma. Finally, a volumetric mesh (57.6 M 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 under-dome 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 6, which shows the PFC geometry and its individual components highlighted.

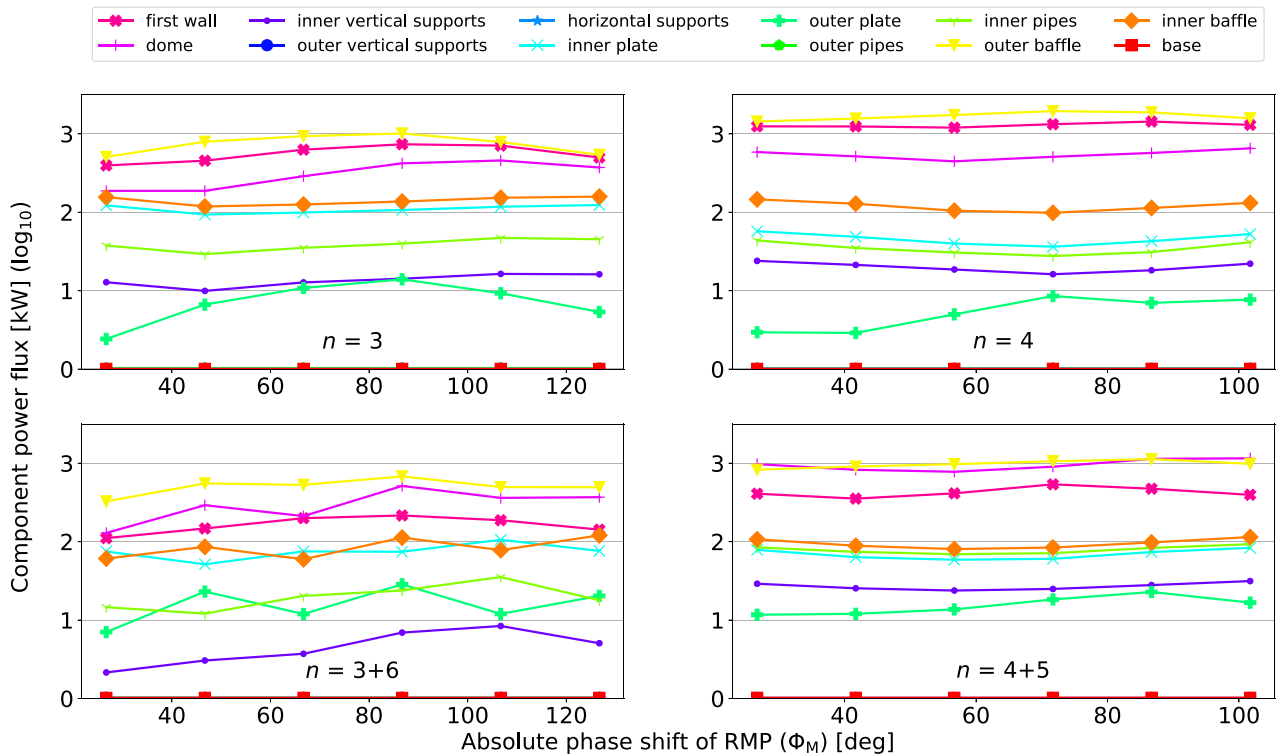
#### 4. Fast-ion power loads in FPO discharges

The simplest method of mitigating localised fast-ion power loads in ITER is to adjust the absolute phase of the applied RMP while maintaining the requirements for ELM control (i.e. the relative phase). Experiments suggest that rotation of RMP fields modulates the intensity of fast-ion power flux on a given PFC—for instance, as measured by a fast-ion loss detector (FILD) [7, 33]. While this suggests that the absolute RMP phase may be used to control where the power flux

lands, it cannot be excluded that the absolute phase may also modulate the global loss rate, since NBIs are located at fixed toroidal positions. Therefore, one may choose to optimise a given absolute RMP phase according to some combination of maximising total NBI heating efficiency and minimising peak PFC power fluxes to avoid localised hot spots. In both cases, an optimal RMP phase could be determined and fixed following the optimisation criteria. However, if large localised power fluxes on PFCs are unavoidable, then it may be beneficial to vary the absolute RMP phase in time to reduce the time-averaged power loads as quantified by their root mean square (RMS). The downside of this choice is the potential thermal cycling of components and reduction in average NBI efficiency over the RMP cycle. Ultimately, the best case scenario is if both criteria can be optimised simultaneously—that is, if peak power flux correlates with global losses. Determining this is a key aim of this study.

Given their higher fusion power and stored energy, and their importance to the ITER mission, discharges from the fusion power operation (FPO) stage are prioritised for investigation. As the plasma response to a given ECC perturbation is predicted to be similar in both the  $Q = 5$  and  $Q = 10$  stages of the  $I_p = 15$  MA deuterium–tritium baseline FPO scenario, we focus here on studying the high-performance  $Q = 10$  stage. Figure 7 shows the relevant plasma data for this discharge, including the  $q$  and rotation profiles which impact plasma response and the associated edge stochasticity.

After adjusting the relative phase between the rows of ECCs, the RMP is rotated by oscillating the current in each ECC, while the relative current phase between each ECC row is fixed to maintain ELM suppression. To quantify the fast-ion



**Figure 10.** Power loads to various tokamak first-wall components as a function of absolute toroidal perturbation phase for different RMP mode spectra at 90 kA coil current amplitude for two-harmonic simulations. Components which receive 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 6. Where traces do not follow the same pattern exhibited by global losses in figure 8, this is because the field redistributes fast-ions among components.


transport in the likely extreme case, when ELM suppression is maximal, the optimal upper and lower row phases (those which maximise XPD) relative to the middle coil row were taken from [21, 22]. All coils also passed the maximum current amplitude of 90 kA. To generate a general yet realistic distribution of deposited neutral beam ions, both heating neutral beams included in the reference baseline (HNB1 and HNB2 [34]) were used, with HNB1 injecting off-axis and HNB2 injecting on-axis. The total injected NBI power is split equally between the beams— $\approx 16.7$  MW each. In the future, ITER foresees to increase the NBI power for heating and current drive by installing a third heating beam, HNB3. However, this was not considered in our studies. In the  $Q = 10$  scenario, the calculated shine-through losses were extremely small and thus assumed to be negligible throughout.

#### 4.1. Fast-ion transport


Figure 8 shows the total measured power lost to PFCs from both beams as the RMP is rotated. This is calculated by summing the power flux to all PFCs including the effects of fast-ion collisions over  $\approx 30$  ms, which was found to be the saturation time for PFC power flux. Approximately 33 000 ( $2^{15}$ ) markers (values greater than  $\approx 8000$  ( $2^{13}$ ) were sufficient for estimating global losses) were tracked over different values of absolute perturbation phase, represented in figure 8 by middle row phase,  $\Phi_m$ , while relative toroidal phase between coil rows ( $\Delta\Phi_u - \Delta\Phi_m$  and  $\Delta\Phi_l - \Delta\Phi_m$ ) was maintained. In the limit

where the toroidal spectrum consists purely of the fundamental mode ( $I(n) \sim \delta(n - n_0)$ ), the XPD remains constant as the perturbation rotates. However, in reality the inclusion of additional toroidal harmonics introduces a dependency on  $\Phi_m$ . To see whether this has a noticeable effect on the fast-ions, we toggle the inclusion of the second harmonic,  $n_1$ . In the case where the sideband 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. 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 ECCs.

Over a rotation cycle, the total measured losses vary by approximately  $\pm 2.8\%$ – $3.2\%$  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. While other work observes roughly double this variation for an individual beam in the  $n = 3$  case [10], the toroidal separation of the beams means each is impacted by the perturbation in turn as it moves past, with losses from one beam lagging the other. Therefore the discrepancy between our study and [10] is likely smaller. Furthermore, the magnitude of global losses reflects those calculated in [15], and the absolute phases corresponding to minimal and maximal losses align well with those predicted in [10] ( $\Phi \approx 22^\circ$  and  $82^\circ$  respectively [35]). Interestingly, these phases also align across toroidal mode numbers and spectra,



**Figure 11.** Two views of an  $n = 3 + 6$  perturbation evaluated at  $\psi_n = 0.99$  with lost markers plotted at their final poloidal angle  $\theta$  and toroidal angle  $\phi$  locations, coloured according to final pitch  $\lambda$ . The particle loss pattern adheres to the perturbation as it is rotated toroidally.




**Figure 12.** 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  $\phi = 90^\circ$ , while 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 in grey. Shown in the lower left of each frame are the coordinate axes which can be ignored.

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  $\phi = 58^\circ$  and  $74^\circ$  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, illustrated by the Poincaré plots of figure 9.

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  $\approx 1$  MW). This is the result of the magnitude of the first harmonic being lower for a given level of current in the ECC coils, compared to that produced by the coils assuming that they would produce a single harmonic. The impact of the second harmonic on edge magnetic field structure is shown in



**Figure 13.** Same view used in figure 12 but for an  $n_0 = 4$  RMP field. A similar behaviour to the  $n_0 = 3$  field can be seen.



**Figure 14.** 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 9. Because of this, and the fact that pure  $n = n_0$  fields are unrealistic, only fields including the second harmonic will be the focus of most of our studies.


#### 4.2. Power loads

It is now important to determine whether it is possible to optimise overall NBI heating efficiency while minimising localised PFC power fluxes. Figure 10 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 non-realistic  $n = n_0$  fields, where rotation does not affect the field structure. Hence the rotation


and spectrum of an RMP each acts not only to scale the fast-ion losses globally but also to redistribute them among PFCs in ITER.

It is therefore important to examine the spatial distribution of lost particles, which is plotted in figure 11 for a single phase over the perturbation amplitude evaluated at the plasma edge. While we do not aim to completely explore the loss or redistribution mechanisms, we start by noting the influence of particle orbit topology by highlighting the strong correlation between loss location and particle pitch angle  $\lambda$  ( $v_{\parallel}/v$ , as measured against plasma current direction, which is typically clockwise in ITER). This correlation between pitch and loss location persists as the perturbation phase changes.

Markers are ultimately lost due to the enhancement of radial diffusion within the stochastic layer at the plasma edge. This predominantly affects passing particles, which are typically



**Figure 15.** Same view as figure 14 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 part of the first wall.




**Figure 16.** Fast-ion power loads shown from the divertor looking upward towards the wall panels 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 power loads typically located on the side facing the neutral beam ports.


lost along the inboard divertor leg and strike the inboard divertor targets. Most fast ions in ITER are deposited near to 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. While 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 flux region to strike outboard divertor targets.

The loss pattern that is created, 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 FIELDS detect oscillations in fast-ion losses as RMPs are rotated in experiments.

To identify hotspots within the footprint, and determine whether they persist, redistribute, or fluctuate, the number of markers was increased to approximately two million ( $2^{21}$ ), 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 12 over the same phases used for figure 8. 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 11. The same can be said for the  $n = 4 + 5$  field; figure 13 shows similar



**Figure 17.** Same view and wall model used above in figure 16 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 16 above for  $n_0 = 3$  fields. The additional power loads in the upper region of the first wall can also be seen.




**Figure 18.** 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 number. 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.

patterns, except a new footprint is introduced on the outboard ceiling first-wall panels.


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 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 14 and 15 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, the power loads here do not persist throughout the rotation, meaning that, for both mode numbers, first-wall RMS power loads could be reduced

by RMP rotation. The only persistent loads are limited to specific divertor regions shown below.

Further close-ups of these first-wall loads are displayed in figures 16 and 17, 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 whose surface normal is anti-parallel to the beam direction. While not negligible, the maximum power loads of approximately  $0.5\text{--}0.7\text{ MW m}^{-2}$  are tolerable for the first-wall panels in these locations. The wall panels in question correspond to blanket modules 14 and 15 [36] which are designated ‘enhanced heat flux’ panels [37], designed to accommodate up to  $\approx 5\text{ MW m}^{-2}$  (compared to ‘normal heat flux’ panels rated for  $1\text{--}2\text{ MW m}^{-2}$ ). As can be seen, the peak power loads can



**Figure 19.** 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 8. The cassette gaps on the outer baffle and divertor dome are subject to the highest power loads.



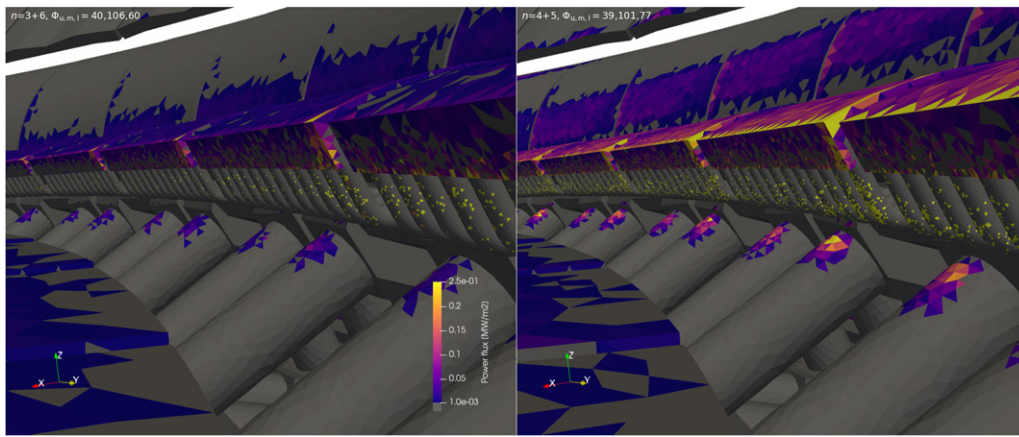
**Figure 20.** Same view as figure 19 but for an  $n_0 = 4$  RMP field. Unlike  $n_0 = 3$  fields, peak power loads are located in multiple regions, such as cassette gaps, the outer baffle and inboard side of the dome. In some of regions, peak power loads persist throughout the RMP cycle.

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 16 and 17 is plotted at various RMP phases in figure 18. In both cases, rotation of the RMP can reduce the RMS peak first-wall power load, by  $0.37 \text{ MW}^{-2}$  and  $0.43 \text{ MW}^{-2}$  for  $n_0 = 3$  and  $n_0 = 4$  fields respectively. It should be noted however that, for  $n_0 = 3$  fields, the peak power flux correlates well with the global fast-ion losses shown in figure 8, so that the optimum phase to minimise global NBI losses automatically ensures the lowest peak power load thus making RMP rotation unattractive as a scheme to decrease peak power loads. However, the same cannot be confidently said for  $n_0 = 4$  fields, where the correlation between global

NBI losses and peak power loads is less clear. This is corroborated when considering the divertor region, shown in figures 19 and 20. Like the first wall, peak power loads to the divertor dome and outer baffle oscillate with global losses in the  $n = 3 + 6$  case. Whereas in the  $n = 4 + 5$  case, for example on the dome, peak power loads only redistribute and can be high even when global losses are low. For example, at the phase ( $\Phi_m = 41^\circ$ ) which leads to minimum global NBI losses, the RMP causes significant power fluxes to reach the inboard side of the dome structure and outer baffle.

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 changes in shadowing



**Figure 21.** 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. While 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.

and regions with discontinuities, such as gaps between divertor cassettes, cause higher power loads due to the low angle of incidence. Rotation is unlikely to reduce these power loads between the cassette domes in the  $n_0 = 4$  case, again making optimisation less straightforward. However, it should be noted that the dome and outer baffle are designed to handle power loads of  $5 \text{ MW m}^{-2}$ , while the highest power loads from NBI losses on these components predicted above are only  $0.1 \text{ MW m}^{-2}$  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 interaction with the plasma, such as the under-dome cooling pipes; figure 21 shows these for the two RMP phases which exhibit the highest power loads for  $n_0 = 3$  and  $n_0 = 4$ . The loads rarely exceed  $0.25 \text{ MW m}^{-2}$ , making them of little concern for the pipes themselves. Again, these loads are sensitive to the local geometry—in this case, the curved edges of the inner support legs and gaps between divertor cassette gaps, which leads to the highest loads on the endmost cooling pipes. The precise value of the power loads on the individual surface triangles of the under-dome pipes are subject to uncertainties due to a lack of marker statistics, however the presence of more marker hits in the  $n = 4 + 5$  case signifies an increased power load compared to  $n_0 = 3$ .

## 5. Summary

The LOCUST code has been applied to studying fast-ion transport in ITER. In particular, transport due to the ITER ECC system was calculated as the system was cycled to toroidally rotate the imposed RMP for ELM control. Initially, the 3D magnetic field model in LOCUST was verified against the SPIRAL code by studying fast-ion transport in an ITER-similar RMP experiment at DIII-D. Results from both codes are sufficiently close to conclude that the 3D magnetic field model in LOCUST is implemented correctly. For its convergence, centimetre perturbation grid sizes and multiple toroidal

harmonics were found to be highly important. Further efforts to validate the 3D magnetic field models used in various fast-ion codes against existing RMP experiments, such as those carried out by the ITPA-EP group, remains an important open activity. The benefits of this type of activity are clear, as the results herein exhibit differences and similarities to equivalent studies performed with other fast-ion codes and 3D magnetic field models.

After testing, 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  $\approx 2\% - 12\%$  ( $0.7-4.0 \text{ MW}$ ), 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 while increased by  $1.7\%-3.2\%$ , while 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 this case the optimal solution may be simply to fix the absolute RMP phase at the point where the first middle coil passes 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 that we have modelled, the resulting power loads are typically a factor of 5–50 lower than the design value of the affect panel of the first wall and divertor components in ITER. In practice, the results presented here may be confirmed in ITER by the use of the relatively low-power diagnostic neutral beam in conjunction with available diagnostics [26] before high NBI power experiments are carried out with RMP fields at significant current levels in the plasma and the ECCs.

The sensitivity of power loads to orbit topology makes it important to study the effects of 3D plasma edge displacement ( $\sim$ cm s) on the initial neutral beam distribution. This is especially true given that most fast-ions are born close to the trapped-passing boundary. Likewise, the increased flux between cassette gaps suggests that, in reality, power loads from fast ions could be greatly affected by PFC misalignment and a realistic geometry for these components, integrated with the evaluation of the power loads [38], and the resulting effects on the PFCs [39], such as the system employed in [40], could be a natural continuation of the work reported here.

In summary, clear gains in NBI heating efficiency when RMPs are applied for ELM control in ITER can be achieved with relatively minor adjustments to the ECC system operation—namely, adjustments to the absolute phase of the applied RMP. The results here show that, if these adjustments are made, then, even at the maximum level of current in the ECCs for ELM suppression in the ITER  $Q = 10$  baseline scenario, fast-ion power loads are very unlikely to exceed the engineering limits of the PFCs for  $n_0 = 3$  and  $n_0 = 4$  fields in ITER. However, if ELM suppression can be achieved with a smaller applied RMP, new avenues emerge for further optimisation of ECC operation for minimising NBI global losses and local power fluxes—namely, the lowering of ECC current amplitude and the adjustment of relative phase between upper and lower coil rows [26]. This problem is highly suited for study with LOCUST, as well as other novel algorithms, such as the recent backward Monte Carlo approach to solving the adjoint problem [41].

## Acknowledgments

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 Viking Cluster, which is a high performance compute facility provided by the University of York. We are grateful for computational support from the University of York High Performance Computing service, Viking and the Research Computing team. Additional thanks are also owed to Lucía Sanchís and Todd Evans for their help and valuable discussions. 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 LEGAL6411249v1. 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.

## ORCID iDs

S.H. Ward  <https://orcid.org/0000-0002-0641-2589>  
 Y.Q. Liu  <https://orcid.org/0000-0002-8192-8411>  
 A. Loarte  <https://orcid.org/0000-0001-9592-1117>  
 S.D. Pinches  <https://orcid.org/0000-0003-0132-945X>  
 R.G.L. Vann  <https://orcid.org/0000-0002-3105-2546>  
 M.A. Van Zeeland  <https://orcid.org/0000-0002-7911-2739>

## References

- [1] Wagner F. *et al* 1982 Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak *Phys. Rev. Lett.* **49** 1408
- [2] ITER Organization 2018 ITER research plan within the staged approach ITER Technical Report ITR-18-003
- [3] Loarte A. *et al* 2014 Progress on the application of ELM control schemes to ITER scenarios from the non-active phase to dt operation *Nucl. Fusion* **54** 033007
- [4] Evans T.E. 2015 Resonant magnetic perturbations of edge-plasmas in toroidal confinement devices *Plasma Phys. Control. Fusion* **57** 123001
- [5] de Rover M., Schilham A.M.R., Montvai A. and Cardozo N.J. 1999 Test particle transport in perturbed magnetic fields in tokamaks *Phys. Plasmas* **6** 2443–51
- [6] Evans T.E., Joseph I., Moyer R.A., Fenstermacher M.E., Lasnier C.J. and Yan L.W. 2007 Experimental and numerical studies of separatrix splitting and magnetic footprints in DIII-D *J. Nucl. Mater.* **363–365** 570–4
- [7] Van Zeeland M.A. *et al* 2015 Fast ion transport during applied 3D magnetic perturbations on DIII-D *Nucl. Fusion* **55** 073028
- [8] Sanchis L. *et al* 2018 Characterisation of the fast-ion edge resonant transport layer induced by 3D perturbative fields in the ASDEX upgrade tokamak through full orbit simulations *Plasma Phys. Control. Fusion* **61** 014038
- [9] Särkimäki S. 2020 Efficient and rigorous evaluation of fast particle losses in non-axisymmetric tokamak plasmas *Nucl. Fusion* **60** 036002
- [10] Sanchis L. *et al* 2021 Optimizing beam-ion confinement in ITER by adjusting the toroidal phase of the 3D magnetic fields applied for ELM control *Nucl. Fusion* **61** 046006
- [11] Varje J. *et al* 2016 Effect of plasma response on the fast ion losses due to ELM control coils in ITER *Nucl. Fusion* **56** 046014
- [12] Särkimäki S., Varje J., Bécoulet M., Liu Y. and Kurki-Suonio T. 2018 Mechanics of ELM control coil induced fast particle transport in ITER *Nucl. Fusion* **58** 076021
- [13] Nardon E. *et al* 2011 Strike-point splitting induced by external magnetic perturbations: observations on jet and mast and associated modelling *J. Nucl. Mater.* **415** S914–7
- [14] Hemsworth R.S. *et al* 2017 Overview of the design of the ITER heating neutral beam injectors *New J. Phys.* **19** 025005
- [15] Akers R. *et al* 2016 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
- [16] Pitts R.A., Kukushkin A., Loarte A., Martin A., Merola M., Kessel C.E., Komarov V. and Shimada M. 2009 Status and physics basis of the ITER divertor *Phys. Scr.* **T138** 014001

- [17] Ward S.H., Akers R., Jacobsen A.S., Ollus P., Pinches S.D., Tholerus E., Vann R.G.L. and Van Zeeland M.A. 2021 Verification and validation of the high-performance Lorentz-orbit code for use in stellarators and tokamaks (LOCUST) *Nucl. Fusion* **61** 086029
- [18] Kramer G.J., Budny R.V., Bortolon A., Fredrickson E.D., Fu G.Y., Heidbrink W.W., Nazikian R., Valeo E. and Van Zeeland M.A. 2013 A description of the full-particle-orbit-following spiral code for simulating fast-ion experiments in tokamaks *Plasma Phys. Control. Fusion* **55** 025013
- [19] Jardin S.C., Ferraro N., Breslau J. and Chen J. 2012 Multiple timescale calculations of sawteeth and other global macroscopic dynamics of tokamak plasmas *Comput. Sci. Disc.* **5** 014002
- [20] Daly K.I. *et al* 2013 Update on design of the ITER in-vessel coils *Fusion Sci. Technol.* **64** 168–75
- [21] Li L., Liu Y.Q., Loarte A., Pinches S.D., Polevoi A., Liang Y. and Zhong F.C. 2019 Modeling 3D plasma boundary corrugation and tailoring toroidal torque profiles with resonant magnetic perturbation fields in ITER *Nucl. Fusion* **59** 096038
- [22] Li L., Liu Y.Q., Loarte A., Pinches S.D., Polevoi A. and Zhong F.C. 2020 ELM control optimization for various ITER scenarios based on linear and quasi-linear figures of merit *Phys. Plasmas* **27** 042510
- [23] Liu Y.Q., Bondeson B., Fransson C.-M., Lennartson B. and Breitholtz C. 2000 Feedback stabilization of nonaxisymmetric resistive wall modes in tokamaks: I. Electromagnetic model *Phys. Plasmas* **7** 3681–90
- [24] Fransson C.-M., Lennartson B., Breitholtz C., Bondeson A. and Liu Y.Q. 2000 Feedback stabilization of nonaxisymmetric resistive wall modes in tokamaks: II. Control analysis *Phys. Plasmas* **7** 4143–51
- [25] Gregorij Pereverzev V. and Yushmanov P.N. 2002 ASTRA Automated System for Transport Analysis in a Tokamak
- [26] Ward S.H., Akers R., Loarte A., Li L., Liu Y.Q., Pinches S.D., Polevoi A. and Vann R.G.L. 2022 Optimising fast-ion transport due to ELM-control coils in ITER discharges with LOCUST-GPU *Nucl. Fusion* (in preparation)
- [27] Chrystal C., Grierson B.A., Staebler G.M., Petty C.C., Solomon W.M., deGrassie J.S., Burrell K.H., Tala T. and Salmi A. 2017 Predicting rotation for ITER via studies of intrinsic torque and momentum transport in DIII-D *Phys. Plasmas* **24** 056113
- [28] Staebler G.M., Kinsey J.E. and Waltz R.E. 2007 A theory-based transport model with comprehensive physics *Phys. Plasmas* **14** 055909
- [29] Li L., Liu Y.Q., Loarte A., Pinches S.D., Polevoi A. and Zhong F.C. 2019 Toroidal modeling of resonant magnetic perturbations in preparation for the initial phase of ITER operation *Nucl. Fusion* **60** 016013
- [30] Liu Y., Kirk A., Gribov Y., Gryaznevich M.P., Hender T.C. and Nardon E. 2011 Modelling of plasma response to resonant magnetic perturbation fields in mast and ITER *Nucl. Fusion* **51** 083002
- [31] Imbeaux F. *et al* 2015 Design and first applications of the ITER integrated modelling & analysis suite *Nucl. Fusion* **55** 123006
- [32] Asunta O., Govenius J., Budny R., Gorelenkova M., Tardini G., Kurki-Suonio T., Salmi A. and Sipilä S. 2015 Modelling neutral beams in fusion devices: beamlet-based model for fast particle simulations *Comput. Phys. Commun.* **188** 33–46
- [33] Van Zeeland M.A. *et al* 2013 Modulation of prompt fast-ion loss by applied  $n = 2$  fields in the DIII-D tokamak *Plasma Phys. Control. Fusion* **56** 015009
- [34] Singh M.J., Boilson D., Polevoi A.R., Oikawa T. and Mitteau R. 2017 Heating neutral beams for ITER: negative ion sources to tune fusion plasmas *New J. Phys.* **19** 055004
- [35] Sanchís L. 2021 personal communication
- [36] Pitts R.A. *et al* 2011 Physics basis and design of the ITER plasma-facing components *J. Nucl. Mater.* **415** S957–64
- [37] Mitteau R. *et al* 2013 The design of the ITER first wall panels *Fusion Eng. Des.* **88** 568–70
- [38] Kos L., Pitts R.A., Simič G., Brank M., Anand H. and Arter W.A. 2019 SMITER: a field-line tracing environment for ITER *Fusion Eng. Des.* **146** 1796–800
- [39] Thorén E. *et al* 2021 The MEMOS-U code description of macroscopic melt dynamics in fusion devices *Plasma Phys. Control. Fusion* **63** 035021
- [40] Coburn J. *et al* 2021 Energy deposition and melt deformation on the ITER first wall due to disruptions and vertical displacement events *Nucl. Fusion* **62** 016001
- [41] Hirvijoki E. 2019 Eliminating poor statistics in Monte-Carlo simulations of fast-ion losses to plasma-facing components and detectors (arXiv:1905.04952)