Clearing the road for high-fidelity fast ion simulations in full three dimensions

T. Kurki-Suonio[1,†]**, K. Särkimäki**[1]**, J. Varje**[1]**, S. Äkäslompolo**[2]**, J. Kontula**[1]**,
P. Ollus[1]**, M. Bécoulet**[3]**, L. Chôné**[1]**, Y. Liu**[4] **and P. Vincenzi**[5]

[1]Department of Applied Physics, Aalto University, FI-00076 AALTO, Finland
[2]Max-Planck-Institut für Plasmaphysik, Teilinstitut Greifswald, D-17491, Germany
[3]CEA, IRFM, F-13108 Saint-Paul-lez-Durance, France
[4]General Atomics, San Diego, CA 92121, USA
[5]Consorzio RFX, I-35127 Padova, Italy

High-energy ions, such as fusion alphas and ions from external heating, can be very sensitive to any non-axisymmetric features in the confining magnetic field due to their collisionless nature. Since understanding the confinement properties of these ions is crucial for ITER (the first fusion reactor currently under construction in Cadarache, France) and beyond, it is of ultimate importance that the predictive simulations are accurate and free of numerical distortions. Adding the third dimension comes at substantial computational cost, calling for new kinds of approaches and computational platforms. In this contribution we discuss what new features, even new physics, the non-axisymmetry brings with it and how one could cope with the ever-increasing demands on both memory and CPU resources. In the end, a few simulation examples with a varying level of non-axisymmetry are given.

Key words: fusion plasma

---

1. Introduction

The fusion community has plenty to thank mathematicians for, but the very foundations of magnetic confinement fusion lie on a theorem in algebraic topology known as the hairy ball theorem: the only topology allowing a vector field (such as magnetic field) to lie tangent to the surface at all points is a torus. Therefore, in both tokamaks and stellarators the plasma is confined on concentric, (topologically) toroidal flux surfaces, each in principle defined by a single field line. Consequently, (in equilibrium) no pressure or potential difference can exist on a given flux surface, but gradients can exist between flux surfaces. This makes magnetic fusion devices essentially one-dimensional – at least in the first approximation. Basic magnetohydrodynamic (MHD) theory that tokamak equilibria are based on indeed consists of ordinary differential equations with the flux surface coordinate as its variable.

Introduction of the divertor geometry brought with it a new region where, in principle, this assumption no longer holds: the field lines   outside the magnetic

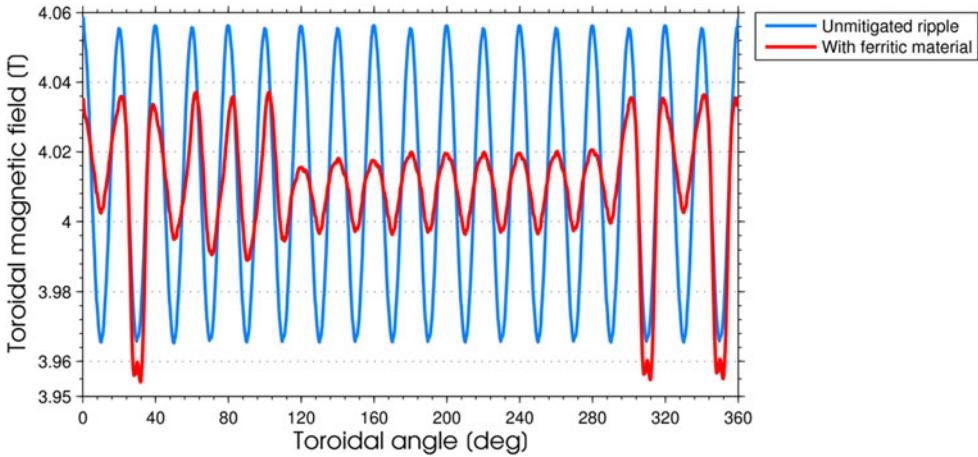† Email address for correspondence: taina.kurki-suonio@aalto.fi

FIGURE 1. The toroidal magnetic field strength at the outboard midplane separatrix as a function of the toroidal angle. The blue curve correspond to the unmitigated TF-ripple while the red curve includes the effects due to ferritic inserts and TBMs. The ripple mitigation is found very effective except at the beam ports, where the shape of the inserts had to be compromised, and the TBMs are evident as finger-like, double-ended structures at three locations.

separatrix intersect the divertor plates, allowing gradients also along the field lines. Therefore, modern fluid codes operating at the plasma edge have to be (at least) two-dimensional.

   All of the above holds for the plasma bulk particles. However, high performance tokamaks and stellarators feature a particle species the analysis of which requires higher dimensionality across the plasma: energetic ions, introduced either by external heating, such as neutral beam injection (NBI) and ion cyclotron resonance heating (ICRH), or fusion reactions, exhibit significant radial excursions from their nominal flux surface due to magnetic drifts. Therefore the fast ion physics has to be (at least) two-dimensional, covering the poloidal cross-section of a tokamak.

   Until recently, most tokamaks could be considered axisymmetric to a high degree. The dominant factor breaking the axisymmetry has been the toroidal field (TF) ripple, but with reasonably small machines and high number of TF coils its magnitude has been small. The situation with ITER is already different. A device twice the size of joint European torus (JET) in linear dimension will have only half of the TF coils (18) compared to JET (32). The resulting TF ripple would be too high and it will be partially mitigated by ferritic inserts (FI), to be installed in the double wall structure right in front of each coil, thus reducing the magnetic field strength there. However, the poloidal extent of the inserts is limited, which makes the field structure more complicated. Another factor breaking the axisymmetry in ITER is introduced by the so-called TBMs (test blanket modules) that will test tritium breeding. These massive modules are also made of ferritic material and, thus, will perturb the magnetic configuration. ITER will have six TBMs, installed in pairs in three different ports. The most dramatic symmetry-breaking effect in ITER is expected from the edge localized modes (ELM) control coils (ECC), to be installed in three toroidal rows around the torus. All these external contributors to the total field are discussed and their effect simulated in Kurki-Suonio *et al.* (2016*a*). Figure 1 illustrates how each affects the toroidal field strength at the outer midplane separatrix.

However, with the exception of the very periphery of the plasma, ITER will still be practically an axisymmetric device. This is not the case for the new (or, depending on the viewpoint: old) technology catching up with the tokamaks: the stellarators are inherently non-axisymmetric. This comes at high computational cost. In fact, already the optimized three-dimensional (3-D) design of the Wendelstein 7-X (W7-X) stellarator became possible only with the advent of the supercomputers.

The traditional fusion plasma physics has mainly been concerned with what happens in the confined plasma. However, with energy-producing fusion reactors in sight, more attention is now given also to what comes out of the plasma. It is the plasma-facing material components that will, to a high extent, decide if fusion power plants will be economical. When calculating the particle and power distributions to the reactor walls, one has to take into account that not only is the magnetic field non-axisymmetric but so is the first wall, particularly in present-day tokamaks and ITER. In fact, simulations of trace-impurity injection experiments in ASDEX Upgrade revealed that, for the deposition profile, the 3-D features of the first wall were more important than the magnetic non-axisymmetry (Miettunen *et al.* 2012).

Non-axisymmetry will naturally have some effect on all plasma constituents, but energetic ions are the most critical ones: the collision frequency drops quickly with increasing energy and, therefore, the fast ions faithfully follow any field perturbation. In this contribution we address the challenges faced in attempts to reliably follow energetic (E = tens of keV to MeV) ions in a truly three-dimensional, non-axisymmetric environment. The examples are mostly ones obtained with the ASCOT code, but reference to other 3-D fast ion codes, such as OFMC (Tani *et al.* 1981) (very similar to ASCOT) and LOCUST (Akers *et al.* 2012) and VENUS-LEVIS (Pfefferlé *et al.* 2014) (both with somewhat different conventions) is made when appropriate. In particular, it should be noted that a modern version of OFMC (Shinohara *et al.* 2003) was the pioneer in simulating non-trivial 3-D features due to ferromagnetic components.

2. Three-dimensional physics and implications for particle following

Fusion plasmas exhibit four classes of energetic particles:

(i) The 3.5 MeV fusion alphas are born with a practically isotropic velocity distribution (beam–target fusion compromises this slightly), thus featuring ions on both trapped and passing orbits.

(ii) Beam ions are injected at a fixed angle, producing trapped ions at the edge and passing ones closer to the centre – depending on the injection angle.

(iii) ICRH produces ions up to MeV range by pumping perpendicular energy to them. Thus all these fast ions are on trapped orbits.

(iv) Runaway electrons on strongly passing orbits can be generated in the context of a disruption.

To be exact, fusion devices feature also a fifth, very important class of energetic particles: the 14.1 MeV fusion neutrons. However, in this work our interest is only in the particles confined by the magnetic field, and mainly ions.

According to Emmy Noether's famous theorem, for any continuous symmetry there exists a conserved quantity. In the case of particles moving in an axisymmetric tokamak magnetic field, this quantity is the canonical toroidal momentum, $P_\phi = mv_\parallel + q\Psi$, where Ψ is the toroidal flux. Conservation of P_ϕ implies that the orbit of a (collisionless) charged particle closes upon itself in the poloidal plane, i.e. even
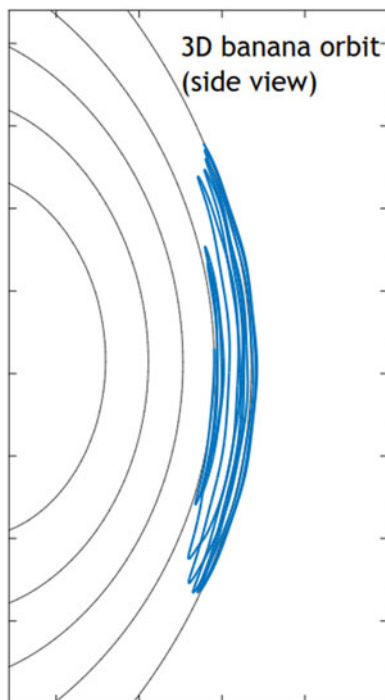
FIGURE 2. An illustration of the orbit of a deeply trapped ion in non-axisymmetric magnetic field. Since the particle is precessing toroidally, the bounce points can take place at different values of major radius, leading to a 'wandering' banana.

though due to magnetic drifts the particle moves radially between flux surfaces, at given poloidal angle it always ends up at the same surface. Therefore, ideally, even the energetic particles with their wide drift orbits are perfectly confined in an axisymmetric magnetic field. Consequently, they could be simulated with reduced models. The fastest models simply assign each test particle a given orbit topology and allow collisions to move the particle from one topology to another. The simulations are then carried out in, for instance, (Ψ, E, μ) space, where Ψ gives the radial position and E and μ define the orbit topology. Here, μ is the magnetic moment.

When the toroidal symmetry is broken by any of the features described in the Introduction, we lose the benefits of Noether's theorem: the drift orbits are no longer guaranteed to close upon themselves in the poloidal plane, resulting in net radial motion as illustrated in figure 2. Physically, this is readily seen for trapped particles: while for axisymmetric field the reflection point of a trapped particle stands at a given value of major radius, $R_{\mathrm{mirror}} = $ const., in a non-axisymmetric configuration the magnetic field value needed for reflecting the particle lies at values of major radius that vary with the toroidal angle, $R_{\mathrm{mirror}} = R_{\mathrm{mirror}}(\phi)$, leading to the orbit behaviour of figure 2. Thus, the absolute confinement of charged particles is lost and, in truly non-axisymmetric devices, it is necessary to follow the orbits of test particles.

The severity at which the non-axisymmetry affects the particle confinement is not universal but strongly depends on the orbit topology: the passing particles, with their small radial excursions, relatively speaking, are less influenced than the trapped orbits. Furthermore, the trapped orbits can be affected differently depending on their

pitch. For instance, for a pure TF ripple, particles with very small pitch can even be trapped between two adjacent TF coils and get rapidly lost by the (almost) vertical grad-B drift. Particles capable of moving full toroidal revolution but having their reflection points on the low-field side (LFS) experience a stronger ripple than those getting reflected on the high-field side (HFS). Thus their random walk process is faster but, on the other hand, their orbit widths are smaller than those reflected on HFS. Introducing ripple mitigation that would suppress the TF ripple on the LFS only might thus alter the ripple-diffusion pattern in the phase space. With increasing complexity in the magnetic field configuration it is thus very important to diligently follow the test particles.

Following the full gyro-orbits of energetic particles can be computationally very expensive. For instance, when following 3.5 MeV fusion alphas in ITER, the gyro-frequency is of the order of 10^8 Hz. This implies a time step of approximately 0.1–1 ns in order to fully resolve the gyro-motion, while the slowing-down time is approximately half a second. However, resolving the full Larmor motion is not necessary when the gradients in the magnetic background are shallow enough not to bring significant changes in the field strength during one gyro-orbit. In such a situation, it is sufficient to follow only the guiding centres (GC) of the gyro-orbits. The condition for the validity of the GC approach is commonly expressed as $\rho_L/L_B \ll 1$, where ρ_L is the Larmor radius and L_B the gradient length of the magnetic field. When following the guiding centres, simulating even the fusion alphas is computationally not an issue with modern super computers. In a device of ITER size, following guiding centres of fusion alphas for the full slowing-down time takes approximately 100 s per particle. As for the number of markers needed to get statistically reliable results, if one is interested only in zero-dimensional numbers (e.g. total power lost), already of the order of $N = 10^5$ markers is enough. However, were one interested in the peak power load on the first wall, given in MW m^{-2}, convergence of the results requires at least $N = 10^6$ (Kurki-Suonio *et al.* 2016*a*).

In cases with sufficient axisymmetry so that the orbits in the poloidal plane remain almost unchanged, it is possible to speed up the simulation of fast particles by accelerating the collisional time scale. In this approach, the guiding centres of particles are followed, but a single orbit is taken to stand for a multitude (10–100) of orbits when evaluating the effect of Coulomb collisions. This is a valid approximation when the collision time is 10–100 times longer than the bounce time.

When simulating strongly perturbed magnetic fields the validity of the guiding-centre approach has to be carefully considered: with non-axisymmetric magnetic fields it is not sufficient to check if the Larmor radius is sufficiently smaller than the standard gradient length of the magnetic field, but now one has to also make sure that the change in the magnetic field strength in the parallel direction is not too large during one Larmor orbit, $\Delta B/B \ll 1$, where ΔB is the change in the magnetic field strength along the field line during one gyration period. If the parallel gradients are sufficiently strong, one has to revert to following the full gyro orbits (GO), which increases the computation time by a factor of 10–100. However, this problem can be partially circumvented in situations where the 3-D effect is radially limited, like the TF ripple. In such cases a hybrid method could be used: the guiding centres can be followed in the essentially axisymmetric field of the core plasma, but gyro-orbit would be adopted upon entering the non-axisymmetric part of the plasma. This method is naturally of limited use for energetic particles with very wide banana widths that would cover both the axisymmetric and the non-axisymmetric regions, but it has been successfully applied to calculate more precisely the power load distribution on the

wall when the 3-D effects have not prohibited using GC approach: the GC markers are followed until approximately one Larmor radius from the wall. At that point also the gyro-orbit of the marker is followed, allowing detection of the exact deposition location on the wall. If the gyro-orbit makes no encounter with a plasma-facing component, the GO following is dropped when the marker has receded far enough from the wall and only GC following is continued.

The validity of the GC approach is not absolutely clear in all situations. Therefore we have done a few tests where the same ASCOT simulations were carried out first with GC following and then by GO following. In the original tests we found that GO following gave 25 %–50 % lower wall power loads than GC. This observation sparked an investigation of the actual guiding-centre transformation (A. Brizard, private communication 2017). A consistent GC description requires both the equations of motion and the collision operator to be transformed from the particle frame to the GC frame. In fast ion modelling, it is customary to use the standard particle collision operator even when applying GC following, but ASCOT now has a genuine GC collision operator (Brizard 2004; Hirvijoki *et al.* 2013). This transformation is carried out to the zeroth order in the formal expansion parameter, while the transformations for the equations of motion, assumed more critical for energetic ions, are first order. Also the initial transformation from the particle position to the GC location was transformed to the first order. However, the magnetic moment and parallel momentum have still been taken at their zeroth-order (particle) value. We now have fixed this, including the first-order terms even in the magnetic moment and parallel momentum, so that the initial transformation now reads (see Cary & Brizard (2009) and, more recently, Lanthaler *et al.* (2017))

$$ \boldsymbol{X} = \boldsymbol{x} - \boldsymbol{\rho}, \quad p_\parallel^{GC} = p_\parallel + p_\parallel^1, \quad \text{and,} \quad \mu^{GC} = \mu + \mu^1, \qquad (2.1a\text{--}c) $$

where

$$ p_\parallel^1 = -p_\parallel \boldsymbol{\rho} \cdot \boldsymbol{\kappa} + \frac{m\mu}{q}(\tau_B + \boldsymbol{a_1} : \nabla \boldsymbol{b}) \qquad (2.2) $$

$$ \mu^1 = \boldsymbol{\rho} \cdot \left(\mu \nabla \ln B + \frac{p_\parallel^2}{mB}\boldsymbol{\kappa} \right) - \frac{\mu p_\parallel}{qB}(\tau_B + \boldsymbol{a_1} : \nabla \boldsymbol{b}). \qquad (2.3) $$

Here, $\boldsymbol{X}$ and $\boldsymbol{x}$ are the guiding-centre and particle positions, respectively, $\boldsymbol{\rho}$ stands for the gyro-vector, $p_\parallel$ is the parallel momentum and μ the magnetic moment. The magnetic curvature is $\boldsymbol{\kappa} = \boldsymbol{b} \cdot \nabla \boldsymbol{b}$, the field line twist is $\tau = \boldsymbol{b} \cdot \nabla \times \boldsymbol{b}$, both given by the magnetic unit vector $\boldsymbol{b}$. The dyadic tensor $\boldsymbol{a_1}$ facilitates the compact form of the above equations and is nicely outlined in, e.g. Lanthaler *et al.* (2017).

At first this seemed to correct the inconsistency observed in the locations of the banana turning points as shown in figure 3(*a*): with the consistent application of the first-order transformation appeared to bring the GC turning point to coincide with the GO one. Unfortunately, repeating the test for a multitude of orbits it was discovered that this was not a generic fix: figure 3(*b*) shows several GC orbits corresponding to the same particle but assuming a different original gyro-angle. In addition, in red the figure shows an orbit where the gyro-orbit is followed but the GC transformation is carried out at each time step. The result is not a simple line but, rather, a helical orbit like the gyro-orbit but with significantly smaller radius. This study revealed the original gyro-angle does matter, as shown in figure 4 that reveals how, for the same
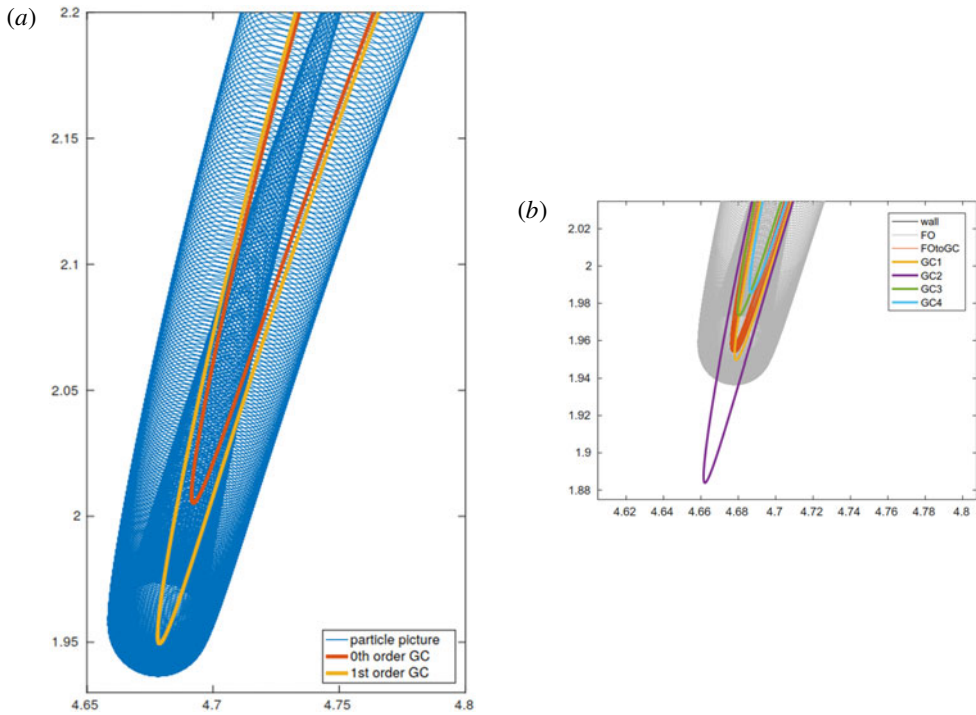
FIGURE 3. A test ion trajectory followed in three different ways: full gyro-orbit (blue), with zeroth-order guiding-centre transformation (red) and with first-order guiding-centre transformation (red). In (*b*), the guiding centers of the same particle but with different initial gyro phases are also shown.

gyro-orbit, one gets different GC paths depending on the gyro-angle. Therefore there always can be a slight difference between the GC and GO turning points.

This inconsistency is not very likely to jeopardize the validity of the GC results. For instance, for fusion alpha particles the original gyro-phase is random to begin with. The fast ion species most sensitive to this inaccuracy is the counter-injected beam ions that could be lost during their first orbit – same naturally holds for fusion alphas born on counter-propagating orbits but, unlike beam ions, their birth probability in the cold edge is very low. Nonetheless, efficient alternatives to GC approach are worth investigating. In fact, the fairly recent GPU-based code LOCUST (Akers *et al.* 2012) relies entirely on following full gyro-orbits. The argument is based on the fact that the length of a time step is not all that matters. For fusion alphas the time step in GC following is of the order of $\Delta t \sim 0.1$ μs, but the conventional fourth-order Runge–Kutta scheme with 5th-order error correction requires 6 look-ups for the magnetic field. On the contrary, using the Boris method in integrating the gyro-orbit the time step is limited to $\Delta t < 1$ ns, but only 1 look-up for the magnetic field is needed. Careful comparison between the different integrators in different situations is thus called for.

3. High-fidelity representation of 3-D magnetic field

The accurate representation of the magnetic field is essential for fast ion orbit-following studies, and it is an twofold issue: how the field is evaluated within
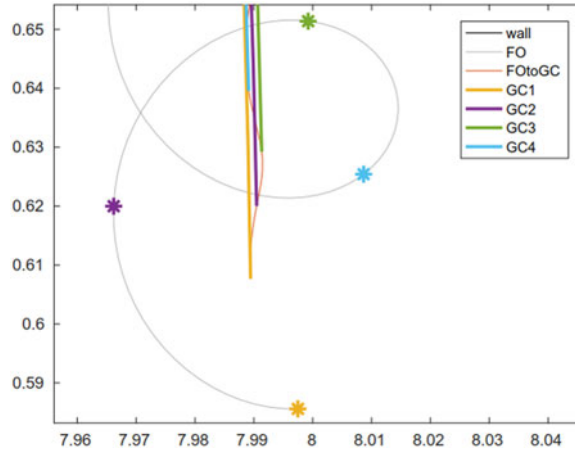
FIGURE 4. A gyro-orbit with corresponding guiding centres started at four different gyro-phases leading to slightly different values of the magnetic moment. Also the trajectory obtained by carrying out the GC-transformation at each step of the gyro-orbit following is shown.

the orbit-following code and how the magnetic field data are assembled. We shall first discuss the evaluation issue, followed by how a complicated 3-D field can be assembled and what potential pitfalls are involved.

ASCOT represents the magnetic field with cylindrical coordinates, (R, z, ϕ), and the total field is a superposition of an axisymmetric component and arbitrary 3-D magnetic field. The axisymmetric component, giving the 2-D plasma equilibrium and evaluated from the poloidal flux, is represented by bi-cubic splines, whereas the 3-D component, consisting of all magnetic field components B_R, B_z and B_ϕ, is represented by tricubic splines. However, while the axisymmetric component is divergence free, this is not guaranteed to be the case for the 3-D component. Significant divergence corresponds to a magnetic monopole and, thus, would cause particle orbits to drift unphysically. Therefore minimizing the divergence is important. The requirement of vanishing divergence could be trivially fulfilled if the input data would consist of the vector potential instead of the field itself, and this approach should be adopted in the future, if possible.

An alternative, and divergence free, way of evaluating the magnetic field would be to use Fourier representation, i.e. decompose the magnetic field into toroidal modes. The Fourier representation is not used in ASCOT, but it is used in codes such as LOCUST and VENUS-LEVIS. The drawback of this method is that the magnetic field evaluations would get computationally increasingly expensive as the number of modes included increases. This is not an issue for the toroidal ripple or ELM control coils which both produce perturbations with only a few significant modes. However, a much wider range of the spectrum is required to include non-periodic features introduced by, for instance TBMs or even FIs which, to avoid interference with the NBI ducts, are not identical in each of the 18 ITER sectors. If the complete magnetic field can be represented with only few modes, the Fourier approach is probably faster than the spline interpolation: memory access is usually the bottleneck in modern computers and storing spline coefficients requires much more memory than the Fourier approach.

Setting up the magnetic field is non-trivial in three dimensions since the data have to have a sufficient resolution to capture the magnetic perturbations also in the

toroidal direction. Resolution is rarely an issue with the axisymmetric component, which is usually provided by the equilibrium codes such as VMEC. Neither is it an issue when the magnetic field perturbation is caused by the external coils as these perturbations can be accurately and efficiently calculated by, e.g. using tools based on the Biot–Savart law, such as BioSaw (Äkäslompolo, Koskela & Kurki-Suonio 2015*b*). The difficulties with resolution arise when including perturbations caused by components made of ferritic material that get magnetized in the reactor, like FIs and TBMs in ITER. To calculate their effect, a finite element method is applicable. Unfortunately, this is not a trivial process for a device of ITER size where the smallest details on magnetized components are on a millimetre scale. One way to alleviate the problem is to simplify the internal structure of the ferromagnetic component, but if the whole point of the study is evaluate the effect of the component as realistically as possible, this should not be taken too far. A simplification process aiming at retaining relevant features while eliminating some of the smallest structures is detailed in Äkäslompolo *et al.* (2015*a*). Even with such preprocessing of the input data, when applying today's finite element method (FEM) solvers such as the commercial COMSOL package used for ASCOT backgrounds, the total field cannot be reproduced at sufficient accuracy in a single run but one has to rely on a multi-step process outlined in Äkäslompolo *et al.* (2015*a*).

However, simply calculating the perturbation field due to coils or magnetized material is not necessarily enough since the plasma adjusts to external perturbations. The effect of this plasma response can only be calculated with dedicated MHD codes such as JOREK and MARS-F which both have their strengths and weaknesses. The plasma response is known to 'heal' magnetic islands inside the plasma, but at the very periphery it has been found to sustain a narrow, stochastic region that can have the opposite effect on fast ion confinement. For ASCOT studies, two different MHD codes have been used. JOREK (Orain *et al.* 2013) is a nonlinear code, making the computations CPU-intensive, that has full tokamak geometry, including the X-point, but uses reduced MHD model and as such omits the possible toroidal component of the plasma response. MARS-F (Liu *et al.* 2000), on the other hand, does not use the reduced MHD, but it lacks the X-point geometry of JOREK. Furthermore, it is a linear code, which makes the computations fast but raises the concern on the validity of results when the perturbations are large. We have done a preliminary comparison between the response given by JOREK and MARS-F, showing that while there are differences in the details, the features leading to changes in fast ion deposition appeared similarly in both cases (Särkimäki *et al.* 2018).

Figure 5 illustrates the role of the plasma response for the ITER 15MA baseline scenario when the ELM control coils are operated at somewhat exaggerated current value. While the island in the core plasma almost disappear, at the edge there remains a substantial region where no field lines are able to carry out even a single poloidal turn. Thus energetic ions entering/born in this region run a high probability of being promptly lost.

4. Fast particle interactions

While the emphasis in this contribution is on the accurate following of the trajectories of almost-collisionless particles, the interest typically is in phenomena that require addressing time scales that involve interactions with other particles, most notably the background plasma. The first problem to arise here comes from the 2-D nature of the plasma equilibrium when evaluating the effect of Coulomb collisions:
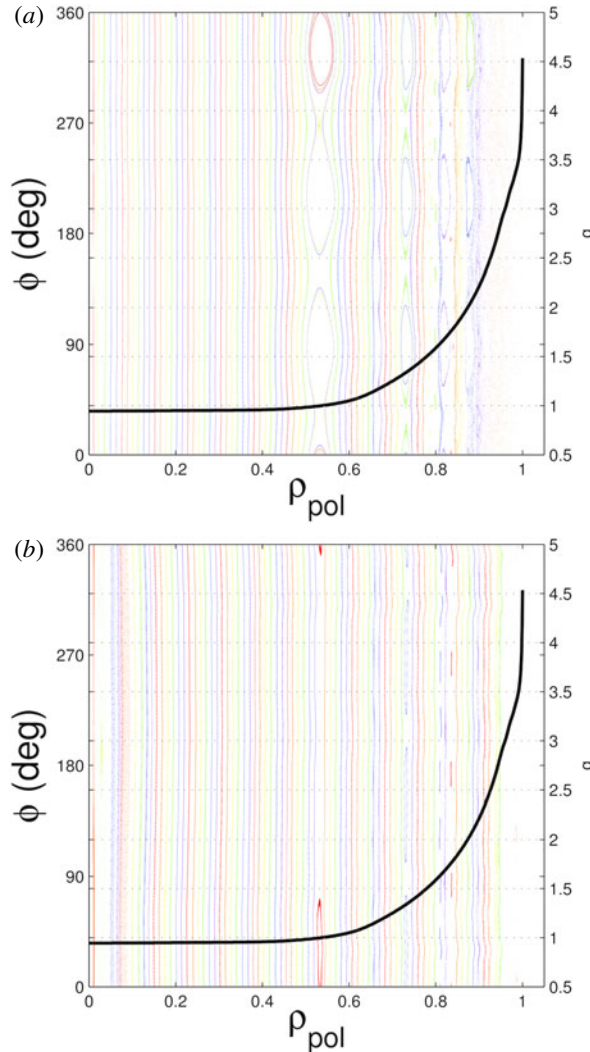
FIGURE 5. Toroidal Poincaré puncture plots of the magnetic field lines in the ITER 15 MA baseline scenario when the effect of FIs, TBMs and ECCs are included. (*a*) In vacuum approximation, (*b*) including the plasma response.

in the case of strong toroidal deformations the kinetic profiles do not follow the actual radial deviations of the flux surfaces. However, considering the experimental uncertainties in the kinetic profiles this is not likely to be a serious issue.

Fusion reactions in a tokamak reactor can take place via three different channels, if you will: in thermonuclear fusion the fusion rate is calculated directly from the plasma kinetic profiles and, thus, only involves the issue discussed above for the Coulomb collisions. However, in today's tokamaks the fusion reactions are dominated by beam–plasma collisions in the case of beam-heated plasmas, or by reactions between plasma and ICRH-generated ions. The source profiles of the fast ions involved in the latter reaction channels are not toroidally symmetric and, thus, require full 3-D spatial profiles to accomplish. This is usually accomplished as a 2-step process. In

the case of ASCOT, first a slowing-down simulation is carried out for ions generated by BBNBI (Asunta *et al.* 2015) (beams) or ASCOT-RFOF (Johnson *et al.* 2011), and the resulting fast ion distributions are fed to the AFSI (Sirén *et al.* 2018) fusion source code to calculate the fusion production. Also two 3-D fast ion distributions can be made interact with each other, but this channel is typically small.

The third type of interactions, relevant for fast ions at least at the plasma periphery, are the charge-exchange (CX) reactions with the neutrals. These are the most difficult to accurately model. Not only is the distribution of the neutrals, those being immune to the magnetic structures, inherently three-dimensional, but this distribution is very poorly known. Even the most sophisticated fluid codes, such as SOLPS (Schneider *et al.* 2006), that model neutrals in three dimensions provide only 2-D distributions.

5. High-fidelity representation of the first wall 3-D features

Fast ion losses are not important only because they would make plasma heating less effective but also because they could cause local damage at the plasma-facing components. Therefore care has to be taken also when reconstructing the first wall and divertor regions so that they contain all relevant components at sufficient resolution. Naturally, this comes at the cost of memory, but it is not so obvious that a wall with high resolution would significantly degrade the performance of the simulation code. However, this becomes an issue when simulating species with variable weight factors, such as alpha particles, where a single marker generated in the highly reactive core can represent orders of magnitude larger population of real-life alpha particles than a marker launched at the plasma periphery. Upon reaching a plasma-facing component we refer to these ions as 'monster ions', since while they contribute a very high peak power load on the wall, their statistical significance is small due to the low escape probability from the centre. The situation is worsened as the resolution of the wall is increased, since the power carried by the marker is divided by ever-smaller surface area. Therefore, when increasing the wall resolution, also the number of markers should be correspondingly increased, making the simulations very CPU-intensive.

In an axisymmetric simulation, the boundary is typically represented as a segmented 2-D contour of either the last closed flux surface itself, leading to anomalously high loss rate, or the approximate shape of the limiters and the first wall. The requirements for memory and calculations to evaluate collisions with these segments are trivial, as the representation typically consists of at most some hundreds of elements, requiring a few kilobytes of data. However, in the 3-D case, the boundary typically consists of some structure based on detailed 3-D computer assisted design (CAD) drawings. Depending on the level of detail, a triangle mesh representation can consist of up to tens of millions of triangle elements, requiring gigabytes of memory. There are efficient algorithms, often originating from the 3-D graphics field, for searching for intersections between trajectories and triangles, including tree representations such as octrees or bounding volume hierarchy (BVH) trees, as used by ASCOT. However, the CPU requirements for evaluating the collisions are still greater compared to the 2-D case.

The 3-D nature of the first wall has proven to play a key role in the distribution of the wall power load, as will become clear in the next section that gives examples on fast ion simulations in non-axisymmetric environments.

6. Examples of 3-D simulations with ASCOT

To illustrate how different 3-D features of a fusion device affect the results we show results from three different ASCOT simulations. We start with simulations for DEMO,

which is anticipated to have a lot simpler construction than ITER, which we will present right after. The section is finished with the ultimate 3-D device, the stellarator, with simulations of the W7-X beams. In each case, the markers corresponding to beam ions are generated using the BBNBI code (Asunta *et al.* 2015) while the fusion alpha markers are generated with the AFSI code (Sirén *et al.* 2018). The hybrid method, described above, is used in all cases.

6.1. *Fast ion power loads in DEMO*

Fast ions play a key role in any fusion reactor. Hence, it is no surprise that tools like ASCOT are utilized in the design phase. A good example is the design of the European DEMO, where ASCOT has been used to study various fast ion phenomena and the results are utilized to optimize the design. The current European design for DEMO will have up to 50 MW NBI heating in the flat-top phase, while during ramp-up it may need much more auxiliary power (Vincenzi *et al.* 2017). At the same time, it is expected to use different construction material (EUROFER-97) than ITER (CuCrZr). While more radiation resilient, EUROFER-97 has lower thermal conductivity and, consequently, is more vulnerable with excessive power loads. Therefore, while the first wall power limit in ITER is 4.7 MW m^{-2}, in DEMO it is only 1 MW m^{-2} (Wenninger *et al.* 2017). Against such low tolerance for power losses, the small number of TF coils (18) appears worrisome. It must be considered that this 1 MW m^{-2} includes all contributions, not just that of fast ions. In a reactor, the fast ion contribution can be comparable or even smaller than losses caused by radiation as shown in Wenninger *et al.* (2017). Fast ion power loads do not, however, pose a serious issue since this DEMO version is designed to have a sizable outer gap, 23 cm, see figure 6, so that even with the limited number of TF coils the (unmitigated) ripple strength at the separatrix reaches only 0.8 %, as compared to the (unmitigated) ITER ripple of 1.1 %. Possible utilization of ferritic inserts, such as in ITER, are being considered. At this stage of the study, the first wall of DEMO is assumed not quite axisymmetric but symmetric with 18 sectors. Even in the final design, it is assumed not to have nearly as many features as ITER.

Figure 7 shows the mesh used in calculating the DEMO magnetic field with the COMSOL code package.

ASCOT was used to simulate both the fusion alphas and the beam ions. For the beams the latest DEMO NBI reference design (Sonato *et al.* 2017) was used, where the power is 16.8 MW per injector, and the acceleration voltage is 800 keV. The injectors consist of 20 modules with 60 beamlets each. In simulations with unmitigated ripple, the power loss due to alphas was approximately 450 kW and due to beams 50 kW. Introducing the ferritic inserts at their design mass reduced the beam losses to practically zero and even the alpha power down to 30 kW. It thus might not be necessary to implement as massive inserts as foreseen now, or it might be possible to even further reduce the number of TF coils. In Varje *et al.* (2017) the effect of ferritic inserts with reduced mass has been investigated. Figure 8 shows the wall load distribution for the fusion alphas. The peak powers are observed at the first wall, not at the divertor, but they do not reach 1 MW m^{-2}.

It is interesting to note that the calculated DEMO fast ion power loads are actually smaller than in ITER baseline scenario. In addition to the large plasma–wall gap, other things contributing to the smallness of the load are the high plasma current, 19.6 MA compared to the 15 MA in ITER, and less steep pedestal. However, it should be kept in mind that these simulations were carried out for MHD-quiescent plasmas.
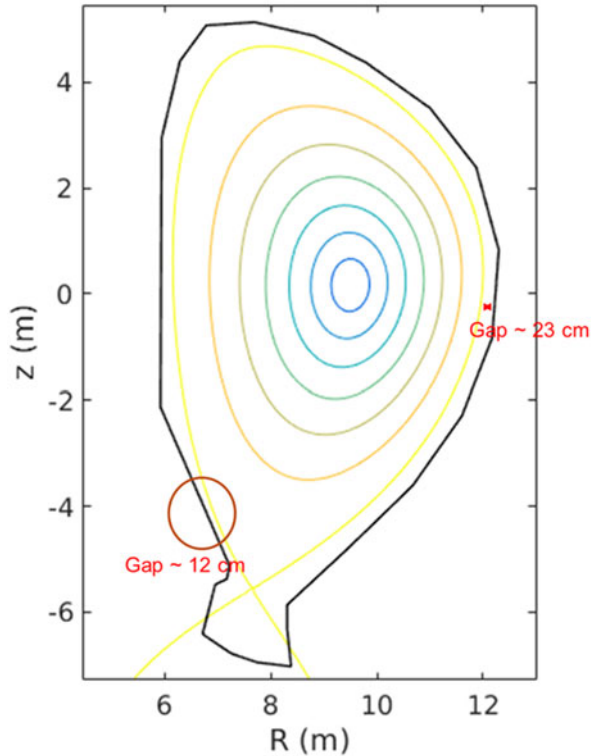
FIGURE 6. Poloidal cross-section of a DEMO design indicating the closest approach of the plasma to the wall at the high-field side and the sizable outer gap at the low-field side.

Furthermore, a study to have ECC to mitigate ELMs in DEMO is being carried out and, as we shall see, introduction of the perturbation generated by ECCs can alter the confinement, in particular of the beam ions.

6.2. *Fast ion power loads in ITER*

ASCOT has been used to carry out a comprehensive analysis of the effect of ferritic inserts and TBMs on the confinement fusion alphas and beam ions in ITER (Kurki-Suonio *et al.* 2016*a*). The study covered four scenarios, the inductive 15 MA baseline scenario, the 12.5 MA hybrid scenario, the partially non-inductive advanced 9 MA scenario and the half-field 7.5 MA scenario. A more detailed analysis of the results can be found in Kurki-Suonio *et al.* (2016*a*), but here we focus on how 3-D features affect the power load and distribution. The losses for each scenario are summarized in table 1, separating the power arriving at the divertor (desired location) from the power arriving at the first wall. Furthermore, results are given separately for simulations where the magnetic background was calculated in the vacuum approximation and where also the response of the plasma was included. While the losses to the first wall reveal the expected result, i.e. that further breaking the axisymmetry by the introduction of the TBMs increases the power loads, a more careful study of the results helps understanding how many different things can affect the power distribution in a tokamak reactor.
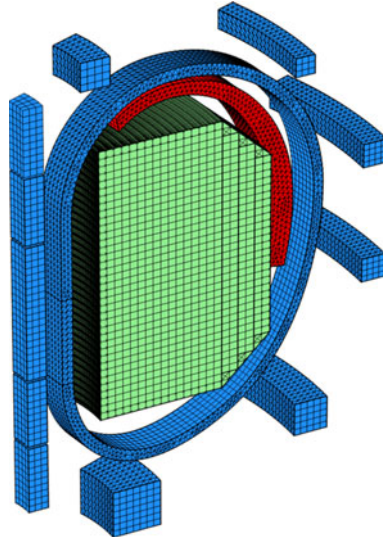
FIGURE 7. COMSOL mesh for magnetic field calculations with coils (blue), ferritic inserts (red) and plasma volume (green).
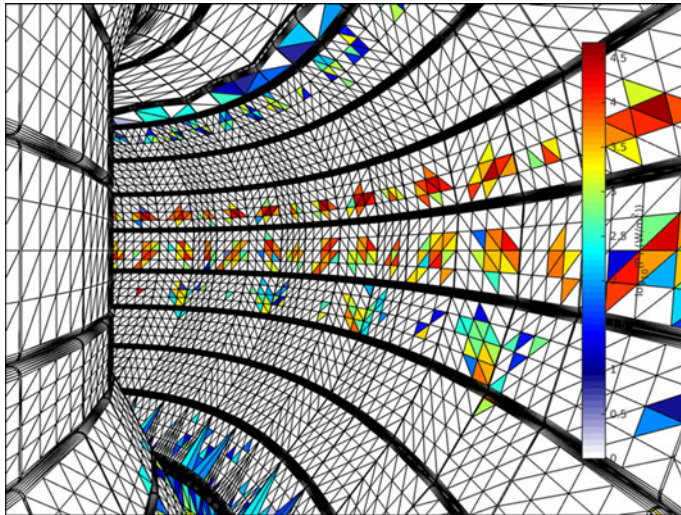

FIGURE 8. Illustration of the alpha power deposition profile on the first wall and divertor of the present DEMO design.

First, looking at the alpha loads, it is observed that the losses do not scale with the plasma current (which would give advanced scenario the highest alpha load) or fusion power (giving baseline the highest value), but the highest power loss is observed for the hybrid scenario. It turned out that triangularity of the hybrid plasma was somewhat stronger than for the other two fusion scenarios, reducing the plasma–wall gap at the outer midplane by approximately one Larmor radius of the fusion alphas. It is also important to realize that, in ITER, the plasma–wall gap is not a single number: the

Scenario	α wall load (kW)	α divertor load (kW)	NBI wall load (kW)	NBI divertor load (kW)
7.5 MA	$-$ / $-$	$-$ / $-$	12/13	0/3
$+$ TBM	$-$ / $-$	$-$ / $-$	19/19	0/3
9 MA	160/160	130/150	6/5	2/4
$+$ TBM	250/270	130/180	15/14	2/9
12.5 MA	510/530	190/190	3/3	1/1
$+$ TBM	580/640	190/210	7/8	1/3
15 MA	20/19	120/120	2/$-$	1/$-$
$+$ TBM	39/42	110/130	7/$-$	1/$-$
$+$ ECC	70/160	1900/1300	9/10	1150/1300

TABLE 1. Power loads to first wall and divertor in the different ITER scenarios: 15 MA baseline scenario with $P_{\text{fus}} = 85$ MW, 12.5 MA hybrid scenario with $P_{\text{fus}} = 50$ MW, 9 MA advanced scenario with $P_{\text{fus}} = 80$ MW and the 7.5 MA half-field scenario. The two numbers given for each case indicate 'no plasma response/with plasma response'.

separation between the plasma and the closest first wall element varies by up to 15 cm along the toroidal direction (see figure 2 in Kurki-Suonio *et al.* (2016*b*)).

The relative importance of the first wall structure and (basic) magnetic field structure is illustrated by figure 9 which shows the fast ions arriving at the LFS first wall both in the baseline and half-field scenarios (*a*), together with the toroidal field variation at the outboard midplane separatrix (*b*). The (mitigated) TF ripple is found to have its minima and maxima interchanged in the half-field scenario compared to the baseline case. This is because the FIs remain fully saturated even at the half of the ITER nominal TF field and, thus, overcompensate the ripple. However, the locations at which the ions intersect the wall are not altered, but the '2-humped' limiter structure remains intact even in the half-field case, indicating that the structure of the first wall is here more important.

According to table 1, neither (mitigated) ripple nor TBMs compromise the fast ion confinement in any of the ITER major operating scenarios. The situation is seen to change dramatically upon introducing the ELM control coils (ECC) in the baseline scenario: the power arriving at the divertor plates is increased by orders of magnitude even when the healing of the magnetic surfaces by plasma response is taken into account (Varje *et al.* 2016). The mechanism by which ECCs degrade the fast ion confinement is found to be different from TBMs: at the plasma periphery there appears a stochastic region that, unlike islands formed further in, is not healed by the plasma response. The stochastic field lines promptly lead energetic ions born at the edge to the divertor. The mechanisms involved are studied in more detail in Särkimäki *et al.* (2018).

It is thus important to realize that even in ITER which, in principle, represents an axisymmetric device, fast ion power loads cannot be estimated by simply considering their orbit widths, given by the nominal plasma current and particle energy, together with a nominal plasma–wall gap, but that it is necessary to include the 3-D features of both the magnetic field and the first wall.

6.3. *Beam ion power loads in Wendelstein 7-X*

In stellarators, axisymmetry is broken by default as illustrated in figure 10: even the poloidal cross-section of the plasma varies dramatically as a function of toroidal
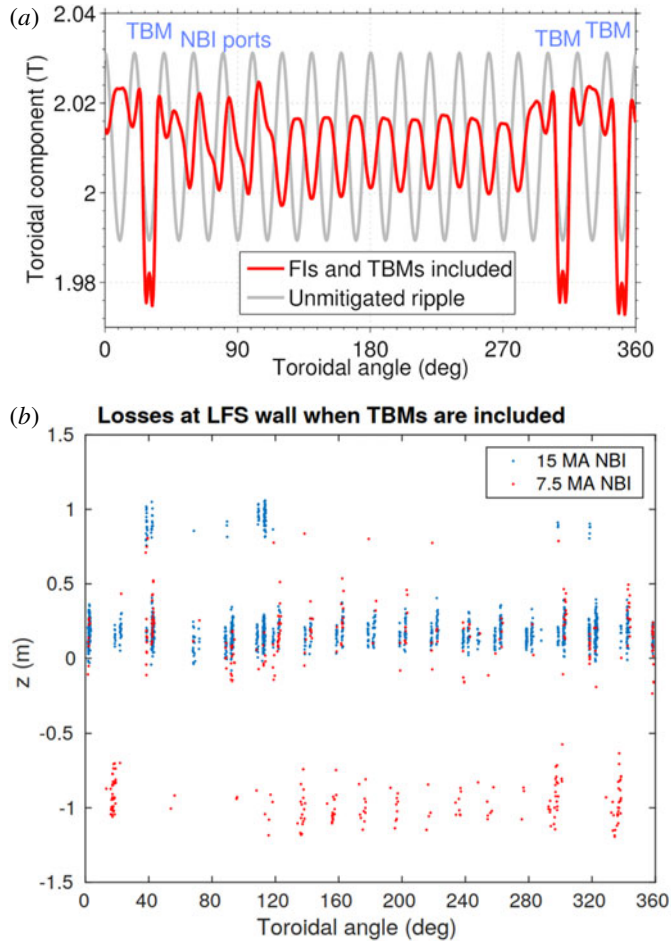
FIGURE 9. (*a*) The (reverse) ripple at outboard midplane for the half-field scenario, and (*b*) structure of losses for both 15 MA baseline and the half-field scenario with its reversed ripple. The wall structure is seen to dominate over the ripple structure of the magnetic field.

angle. Wendelstein 7-X is the world's largest and most advanced stellarator. The modular coil geometry of W7-X is a result of optimization process that was only made possible by the emergence of supercomputers. The optimization was carried out with respect to good neoclassical confinement, small Grad–Shafranov shift, small bootstrap current and, maybe most importantly, good confinement of fast ions at high beta. Although non-axisymmetric, it still features another kind of symmetry: it is stellarator symmetric (Dewar & Hudson 1998) with 5 toroidal periods. Wendelstein 7-X has altogether 70 coils (20 planar, 50 shaped), with 7 independent power sources, giving seven degrees of freedom. This gives flexibility in devising different magnetic configurations. One of the main goals of W7-X is to demonstrate good confinement of fast ions, which in W7-X are produced by ion heating, such as neutral beam injection. Confining beam ions would also show that fusion alphas could be well confined in a Helias-like reactor: the relative size of beam ion gyro-radius with
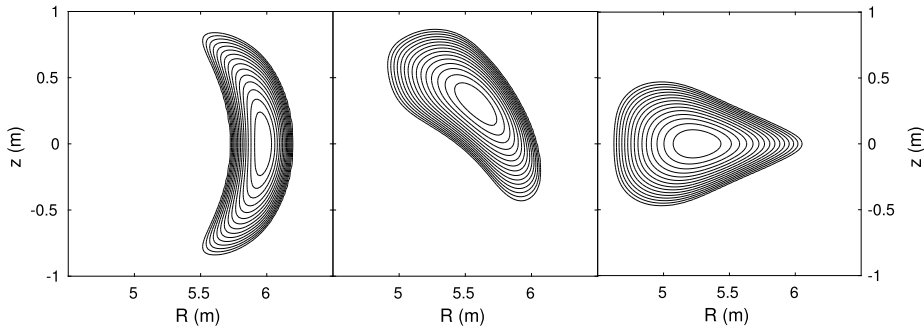
FIGURE 10. Flux surface contours of W7-X at three toroidal angles: $\phi = 0°$, $\phi = 18°$ and $\phi = 36°$ (Kontula 2017).

respect to the W7-X dimensions corresponds to the same ratio of alpha gyro-radius in a foreseen reactor.

Beam operation will start with two NBI injectors with two sources, each injecting hydrogen at 55 keV to hydrogen plasma. The power per source in this case is up to 1.7 MW. The beams are being commissioned, with first experiments foreseen for the summer of 2018. The beams in W7-X indeed pose a hard challenge for confinement: due to engineering constraints the beams are practically radial, and many ions are born on trapped orbits, vulnerable to losses due to the multitude of ripples in a stellarator. Beam ion losses have already been studied with the ANTS code and a simple wall (Drevlak *et al.* 2014). This work has been extended with ASCOT, using a detailed 3-D wall (4×10^6 triangles), for the W7-X reference magnetic configurations (Äkäslompolo *et al.* 2018).

The simulations indicate wall power losses even to some sensitive wall components (made of steel), including panels and poloidal closure, pumping slits, vacuum vessel and ports. Even in the optimized scenarios, beam losses are substantial, several per cent, compared to tokamaks where they are almost negligible. The configuration with high mirror ratio appears superior for fast ion confinement, with lowest lost power fraction. Furthermore, most power goes to the high heat capacity components, divertor parts and heat shields, but non-negligible amount still lands on some vulnerable parts (figure 11). Since none of the reference scenarios appear ideal for fast ion confinement, hunt for the perfect (beam) scenario is ongoing, facilitated by the flexible coil system.

7. Conclusions on computational requirements due to 3-D structures and future prospects

The increased dimensionality can affect the computer time requirements in two ways. Firstly, the physics dictate limitations to the simulation, increasing CPU time requirements. The 3-D fields must be sampled at sufficiently high resolution to capture the physical effects, requiring shorter time steps. Likewise, collisions with the detailed 3-D wall need to be evaluated at sufficiently high resolution, complicated by the drastically larger number of elements in the representation.

Secondly, the higher resolution and consequently larger data size can degrade performance on modern supercomputers. The limited amount of memory available on a computer node can limit the number of simultaneous processes, requiring undersubscription and limiting the utilization of the full performance of the CPU.
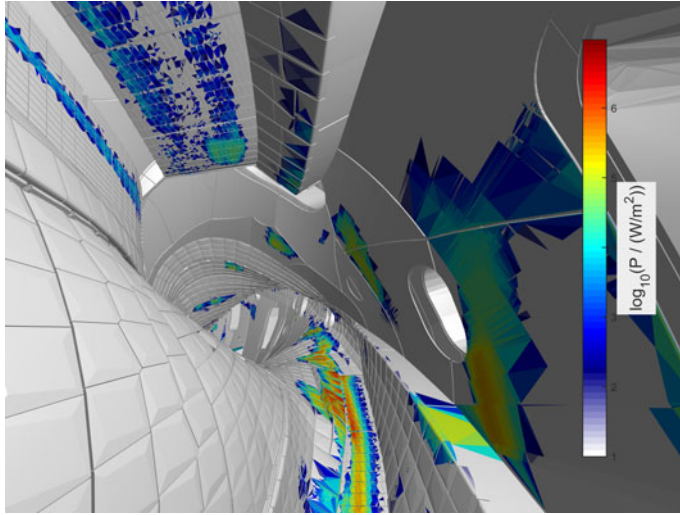
FIGURE 11. The NBI wall loads in W7-X in a specific magnetic configuration. The data are for the high statistics run described in Äkäslompolo *et al.* (2018).

This can be alleviated by shared memory in frameworks such as OpenMP or MPI 3. The increased memory footprint also causes issues with cache utilization, which is key for getting optimal performance from modern CPUs.

With the introduction of GPU and many core architectures to scientific computing, the above memory issues become even more pronounced due to higher number of cores but lower amounts of memory per core. When utilizing these architectures it becomes increasingly necessary to adapt software to make full use of the hardware. In particular, GPU and many cores allow for efficient use of array operations. However, the limitations of these architectures in terms of available memory together with the need for finer resolved computations, call for the use of parallelization techniques, such as domain decomposition, in order to spread the load among the nodes and cores. The problem is then one of optimizing the memory load versus the communication time between units. In addition, making efficient use of the available computing power necessitates to apply load balancing techniques in order to avoid underused memory and idle cores.

Finally, the development of numerical methods consistent with the mathematical framework of Lagrangian mechanics can allow for more efficient and faithful simulations. The so-called variational integrators, which are built by discretizing the variational principle and then deriving a set of discrete equation from it, self-consistently preserve the invariants of motion with a bounded error. Indeed, as per Noether's theorem of classical mechanics, this method yields discrete conservation equations tied to the equations of motions in a self-consistent manner. As a consequence, such integrators have very good conservation properties over long simulation times, and also allow the use of longer time steps compared to *ad hoc* integrators (Kraus 2013).

Acknowledgements

REFERENCES

ÄKÄSLOMPOLO, S., ASUNTA, O., BERGMANS, T., GAGLIARDI, M., GALABERT, J., HIRVIJOKI, E., KURKI-SUONIO, T., SIPILÄ, S., SNICKER, A. & SÄRKIMÄKI, K. 2015a Calculating the 3D magnetic field of ITER for European TBM studies. *Fusion Engng Des.* **98–99**, 1039–1043; Proceedings of the 28th Symposium on Fusion Technology (SOFT-28), arXiv:1506.00659.

ÄKÄSLOMPOLO, S., DREVLAK, M., TURKIN, Y., BOZHENKOV, S., JESCHE, T., KONTULA, J., KURKI-SUONIO, T. & WOLF, R. 2018 Modelling of NBI ion wall loads in the w7-x stellarator. *Nucl. Fusion* **58** (8), 082010.

ÄKÄSLOMPOLO, S., KOSKELA, T. & KURKI-SUONIO, T. 2015b Biot savart law integrator biosaw. Arxiv.org, arXiv:151101629.

AKERS, R., VERWICHTE, E., MARTIN, T., PINCHES, S. & LAKE, R. 2012 GPGPU Monte Carlo calculation of Gyro-Phase resolved fast ion and n-state resolved neutral deuterium distributions. In *Proceedings of the 39th EPS Conference on Plasma Physics*, Europhysics Conference Abstracts, vol. 36F, P5.058. European Physical Society.

ASUNTA, O., GOVENIUS, J., BUDNY, R., GORELENKOVA, M., TARDINI, G., KURKI-SUONIO, T., SALMI, A. & SIPILÄ, S. 2015 Modelling neutral beams in fusion devices: beamlet-based model for fast particle simulations. *Comput. Phys. Commun.* **188**, 33–46.

BRIZARD, A. J. 2004 A guiding-center fokker–planck collision operator for nonuniform magnetic fields. *Phys. Plasmas* **11** (9), 4429–4438.

CARY, J. R. & BRIZARD, A. J. 2009 Hamiltonian theory of guiding-center motion. *Rev. Mod. Phys.* **81**, 693–738.

DEWAR, R. L. & HUDSON, S. R. 1998 Stellarator symmetry. *Physica* D **112** (1–2), 275–280; Proceedings of the Workshop on Time-Reversal Symmetry in Dynamical Systems.

DREVLAK, M., GEIGER, J., HELANDER, P. & TURKIN, Y. 2014 Fast particle confinement with optimized coil currents in the W7-X stellarator. *Nucl. Fusion* **54** (7), 073002.

HIRVIJOKI, E., BRIZARD, A., SNICKER, A. & KURKI-SUONIO, T. 2013 Monte Carlo implementation of a guiding-center fokker-planck kinetic equation. *Phys. Plasmas* **20** (9), 092505.

JOHNSON, T., SALMI, A., STEINBRECHER, G., ERIKSSON, L.-G., HELLSTEN, T., HÖÖK, L., SCHNEIDER, M. & CONTRIBUTORS, I.-T. 2011 Library for rf interactions in orbit following codes. In *AIP Conference Proceedings*, vol. 1406, pp. 373–376. AIP.

KONTULA, J. 2017 Neutral beam injection simulations in the wendelstein 7-x stellarator. Master's thesis, Aalto University, School of Science.

KRAUS, M. 2013 Variational integrators in plasma physics. PhD thesis, Technical University of Munich.

KURKI-SUONIO, T., ÄKÄSLOMPOLO, S., SÄRKIMÄKI, K., VARJE, J., ASUNTA, O., CAVINATO, M., GAGLIARDI, M., HIRVIJOKI, E., PARAIL, V., SAIBENE, G. *et al.* 2016a Effect of the european design of tbms on iter wall loads due to fast ions in the baseline (15 ma), hybrid (12.5 ma), steady-state (9 ma) and half-field (7.5 ma) scenarios. *Nucl. Fusion* **56** (11), 112024.

KURKI-SUONIO, T., SÄRKIMÄKI, K., ÄKÄSLOMPOLO, S., VARJE, J., LIU, Y., SIPILÄ, S., ASUNTA, O., HIRVIJOKI, E., SNICKER, A., TERÄVÄ, J. *et al.* 2016b Protecting iter walls: fast ion power loads in 3d magnetic field. *Plasma Phys. Control. Fusion* **59** (1), 014013.

LANTHALER, S., PFEFFERLÉ, D., GRAVES, J. & COOPER, W. 2017 Higher order larmor radius corrections to guiding-centre equations and application to fast ion equilibrium distributions. *Plasma Phys. Control. Fusion* **59** (4), 044014.

LIU, Y., BONDESON, A., FRANSSON, C.-M., LENNARTSON, B. & BREITHOLTZ, C. 2000 Feedback stabilization of nonaxisymmetric resistive wall modes in tokamaks. I. Electromagnetic model. *Phys. Plasmas* **7** (9), 3681–3690.

MIETTUNEN, J., KURKI-SUONIO, T., MAKKONEN, T., GROTH, M., HAKOLA, A., HIRVIJOKI, E., KRIEGER, K., LIKONEN, J., ÄKÄSLOMPOLO, S. & THE ASDEX UPGRADE TEAM 2012 The effect of non-axisymmetric wall geometry on ^{13}C transport in ASDEX Upgrade. *Nucl. Fusion* **52** (3), 032001.

ORAIN, F., BÉCOULET, M., DIF-PRADALIER, G., HUIJSMANS, G., PAMELA, S., NARDON, E., PASSERON, C., LATU, G., GRANDGIRARD, V., FIL, A. *et al.* 2013 Non-linear magnetohydrodynamic modeling of plasma response to resonant magnetic perturbations. *Phys. Plasmas* **20** (10), 102510.

PFEFFERLÉ, D., COOPER, W., GRAVES, J. & MISEV, C. 2014 Venus-levis and its spline-fourier interpolation of 3d toroidal magnetic field representation for guiding-centre and full-orbit simulations of charged energetic particles. *Comput. Phys. Commun.* **185** (12), 3127–3140.

SÄRKIMÄKI, K. J. V., BÉCOULET, M., LIU, Y. & KURKI-SUONIO, T. 2018 Mechanics of elm control coil induced alpha particle transport. *Nucl. Fusion* **58**, 076021.

SCHNEIDER, R., BONNIN, X., BORRASS, K., COSTER, D., KASTELEWICZ, H., REITER, D., ROZHANSKY, V. & BRAAMS, B. 2006 Plasma edge physics with b2-eirene. *Contrib. Plasma Phys.* **46** (1–2), 3–191.

SHINOHARA, K., KAWASHIMA, H., TSUZUKI, K., URATA, K., SATO, M., OGAWA, H., KAMIYA, K., SASAO, H., KIMURA, H., KASAI, S. *et al.* 2003 Effects of complex magnetic ripple on fast ions in jft-2m ferritic insert experiments. *Nucl. Fusion.* **43** (7), 586.

SIRÉN, P., VARJE, J., ÄKÄSLOMPOLO, S., ASUNTA, O., GIROUD, C., KURKI-SUONIO, T., WEISEN, H. & JET CONTRIBUTORS 2018 Versatile fusion source integrator afsi for fast ion and neutron studies in fusion devices. *Nucl. Fusion* **58** (1), 016023.

SONATO, P., AGOSTINETTI, P., BOLZONELLA, T., CISMONDI, F., FANTZ, U., FASSINA, A., FRANKE, T., FURNO, I., HOPF, C., JENKINS, I. *et al.* 2017 Conceptual design of the demo neutral beam injectors: main developments and R&D achievements. *Nucl. Fusion* **57** (5), 056026.

TANI, K., AZUMI, M., KISHIMOTO, H. & TAMURA, S. 1981 Effect of toroidal field ripple on fast ion behavior in a tokamak. *J. Phys. Soc. Japan* **50** (5), 1726–1737.

VARJE, J., AGOSTINETTI, P., KURKI-SUONIO, T., SNICKER, A., SONATO, P., SÄRKIMÄKI, K. & VINCENZI, P. 2017 Effect of 3d magnetic perturbations on fast ion confinement in the European demo. In *44th EPS Conference on Plasma Physics*, Europhysics Conference Abstracts, vol. 41F, P2.147. European Physical Society.

VARJE, J., ASUNTA, O., CAVINATO, M., GAGLIARDI, M., HIRVIJOKI, E., KOSKELA, T., KURKI-SUONIO, T., LIU, Y., PARAIL, V., SAIBENE, G. *et al.* 2016 Effect of plasma response on the fast ion losses due to ELM control coils in ITER. *Nucl. Fusion* **56** (4), 046014.

VINCENZI, P. *et al.* 2017 EU DEMO transient phases: main constraints and heating mix studies for ramp-up and ramp-down. *Fus. Engng. Design* **123**, 473–476.

WENNINGER, R., ALBANESE, R., AMBROSINO, R., ARBEITER, F., AUBERT, J., BACHMANN, C., BARBATO, L., BARRETT, T., BECKERS, M., BIEL, W. *et al.* 2017 The demo wall load challenge. *Nucl. Fusion* **57** (4), 046002.