

# Fast ion generation by combined RF-NBI heating in W7-X

M. Machielsen<sup>1,†</sup>, J.P. Graves<sup>1,2</sup>, H.W. Patten<sup>3</sup>, C. Slaby<sup>4</sup> and S. Lazerson<sup>4</sup>

<sup>1</sup>Ecole Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center (SPC), CH-1015 Lausanne, Switzerland

<sup>2</sup>York Plasma Institute, Department of Physics, University of York, York, Heslington YO10 5DD, UK

<sup>3</sup>Department of Statistics, University of Oxford, OX1 3LB Oxford, UK

<sup>4</sup>Max Planck Institute for Plasma Physics, 17491 Greifswald, Germany

(Received 30 October 2022; revised 20 March 2023; accepted 21 March 2023)

Good fast ion confinement in quasi-isodynamic stellarators still has to be demonstrated experimentally. In the absence of fusion alphas, auxiliary heating systems can be used to generate fast ions. For this purpose, W7-X has neutral beam injection and its newly installed ion cyclotron heating system. In practice, both systems can run simultaneously, the synergetic running of both is the focus of the current modelling development. With the upgrades to the SCENIC code described here, we find that, compared with pure neutral beam injection (NBI), the fast ion tail is significantly enhanced when radio frequency (RF) heating is enabled. But the effectiveness of heating and fast ion tail generation strongly depends on the plasma parameters and antenna frequency, which calls for further study, both numerical and experimental.

**Key words:** fusion plasma, plasma heating, plasma simulation

---

## 1. Introduction

Good alpha particle confinement is a requirement for a future stellarator reactor. This favourable property has yet to be demonstrated experimentally in W7-X. Due to the lack of fusion alphas, it is foreseen to generate fast ions using neutral beams, ion cyclotron resonance heating (ICRH) or both. It is these fast ions that function as ‘mock-up’ alpha particles for confinement tests. At an energy of 60 keV their Larmor radius divided by the minor radius is comparable to that of fusion alphas in a HELIAS reactor (assuming a scaled-up version of W7-X with 5 T magnetic field) (Drevlak *et al.* 2014). However, a perfect comparison will not be possible because particles in this energy range are still influenced by the radial electric field, while fusion alphas are not (Kolesnichenko *et al.* 2006; Faustin *et al.* 2016). NBI provides protons in this energy range, but only about a

† Email address for correspondence: [mike.machielsen@epfl.ch](mailto:mike.machielsen@epfl.ch)

third are at full energy. Therefore, heating by RF waves may aid an exploration in this regime. As the name ICRH suggests, it relies on a wave–particle resonance

$$\omega = n\omega_c + k_{\parallel}v_{\parallel}, \quad (1.1)$$

where  $\omega$  is the antenna frequency,  $\omega_c$  the cyclotron frequency,  $n$  the harmonic number,  $k_{\parallel}$  the parallel wavenumber and  $v_{\parallel}$  the parallel velocity. The cyclotron frequency depends on the particle position, so for a given  $\omega$ ,  $n$ ,  $k_{\parallel}$ ,  $v_{\parallel}$ , the positions where (1.1) holds describe a surface (in three dimensions) at constant magnetic field strength. The wave–particle interaction is strongest on this surface. Therefore, a common approximation is to assume that diffusion in velocity space only occurs where (1.1) holds, also known as resonant diffusion. In a Monte Carlo orbit following code this is done by applying a kick in velocity, only when a test particle crosses the resonant surface. As some particles reside in a sweet spot, they continue to receive kicks from the wave for an extended time, thereby reaching energies considerably higher than the bulk temperature. Therefore, a side effect of the bulk heating by ICRH is the generation of a fast ion population, if the conditions are suitable. The diffusion in velocity space is constrained by

$$\frac{dv_{\parallel}}{dv_{\perp}} = \frac{v_{\perp}}{v_{\parallel}} \frac{\omega - n\omega_c}{n\omega_c}, \quad (1.2)$$

where  $v_{\perp}$  is the perpendicular velocity (Kennel & Engelmann 1966). In pure RF heating schemes (so no NBI), the ion population that interacts with the wave typically starts with the same temperature as the background ions. So the numerator in (1.2) is nearly zero, and the kicks mainly increase the perpendicular velocity. Consequently, the high energy population that forms will have a lot of trapped ions, which are rapidly lost in the magnetic field of W7-X. One way to improve this is by using a so called three-ion scheme (Kazakov *et al.* 2015). In that case the power per ion is significantly enhanced, so higher energies can be obtained before the particles are lost. Another approach is to use neutral beams to seed a fast ion population, which is then further accelerated using RF heating. This is attractive for several reasons: the beams provide an extra 6.8 MW of power, in addition to a projected 1.5 MW of coupled RF power. The starting population is already above thermal energy, so it experiences less friction due to collisions as the ions are accelerated. And lastly, because the beam ions start with a considerable parallel velocity (and thus Doppler shift), the kicks are more isotropic, see (1.2). This results in a smaller trapped particle fraction compared with pure RF schemes.

Successful fast ion generation with combined RF-NBI heating was already demonstrated in stellarator experiments on Heliotron-J (Okada *et al.* 2016). Where a considerable enhancement of the fast ion tail was observed when ICRH was enabled. In order to aid experiments and extrapolate to future scenarios, extensive modelling of fast ions is a crucially important tool. In the present paper the code package SCENIC (Jucker 2010) is used to study such synergetic heating scenarios. It aims to self-consistently compute the magnetic equilibrium, the RF wave field and the fast ion distribution function. For this purpose the following codes are used respectively: VMEC (Hirshman & Whitson 1983) (or its anisotropic pressure version ANIMEC Cooper *et al.* 2009), LEMAN (Popovich, Cooper & Villard 2006) and VENUS-LEVIS (Pfefferlé *et al.* 2014). These three are iterated until a converged self-consistent state is reached. However, modelling of combined RF-NBI poses additional challenges over pure RF and pure NBI scenarios, which will be discussed in detail in § 2. In short, the thermal ions cannot simply be ignored because they provide a source of fast ions due to RF heating. A pragmatic way to deal with this is to simulate both fast and thermal markers (test particles) of the minority

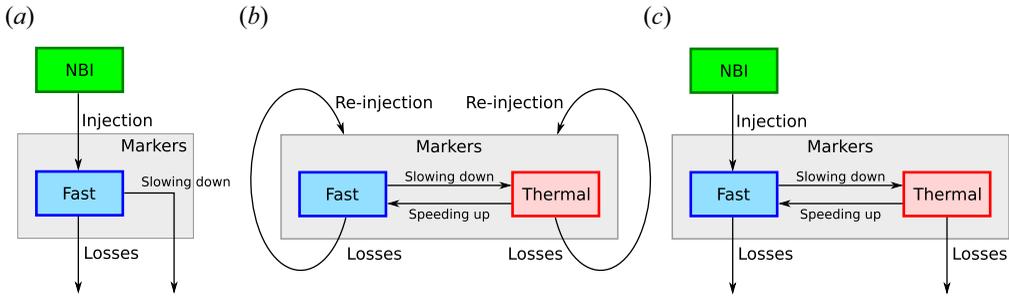


FIGURE 1. Schematic overview of modelling in VENUS-LEVIS, inflow at the top, outflow at the bottom. The population in the grey box is tracked by VENUS-LEVIS. (a) pure NBI, (b) pure RF, (c) RF-NBI.

species. Transport of thermal ions in many experiments far exceeds the predicted values by neo-classical transport. This ‘anomalous’ transport is thought to be caused by turbulence (Freidberg 2007, § 6.7). The true nature of this transport is very complex, but it is often well approximated as a diffusive process. The anomalous diffusion coefficient is typically of the order of  $1 \text{ m}^2 \text{ s}^{-1}$  (Wesson 2004, § 4.16). This type of diffusion is implemented as a Monte Carlo operator in VENUS-LEVIS, which has recently been made compatible with stellarator geometry (§ 2). Using this diffusion model, a steady state sets in after approximately 200 ms for pure NBI (§ 3). Enabling RF leads to a collisionless power deposition which strongly depends on the chosen parameters (§ 4). These parameters also have a major effect on the fast ion tail formation. If chosen correctly, the number of fast ions is enhanced considerably with RF-NBI (§ 5). The results are discussed in § 6 and summarised in § 7. Note that the goal of this study is not to perform an extensive parameter scan, but instead the focus is on identifying the key aspects that make a scheme suitable for fast ion generation. Hence, this paper describes the new feature (the radial diffusion model) in detail, and exploits it for an application where it is needed, in particular the RF-NBI scheme of W7-X.

## 2. Upgrade to VENUS-LEVIS

The fast ion population can be accurately modelled by sampling it with markers, and following their trajectories. In the case of pure NBI heating, the markers are born at high energy, and slow down on the background due to collisions. Since the main interest is in the fast ions, markers can be removed from the simulation as soon as they thermalise.<sup>1</sup> A balance is reached when the inflow of markers due to the beam equals the outflow due to losses and thermalisation.

Contrary to pure NBI, where the markers are born fast, and are mainly slowing down, RF provides a mechanism to speed up (thermal) markers. For this reason the thermalised markers are not removed from the simulation. In addition, in the absence of the real fuelling source (NBI), the particle balance has to be achieved in an ad hoc way. To ensure this balance in VENUS-LEVIS, the markers are ‘re-injected’ by sampling from the initial distribution function whenever they are lost, see figure 1.

<sup>1</sup>This happens when their energy drops below a predefined threshold, typically a few times the (local or central) electron temperature. A consequence of removing these thermal markers is that plasma dilution due to the beam species is not accounted for.

When modelling both heating methods simultaneously, the real particle source is known. But because RF is involved the thermal markers need to be included.<sup>2</sup> It has been found that this approach leads to an unphysical accumulation of thermalised beam ions, because the inflow far exceeds the losses in the simulation (Patten 2019).

This imbalance is expected as VENUS-LEVIS is designed for fast ion modelling. The guiding centre equations it solves are valid for both fast and thermal particles, however, the collision operator is incomplete for thermal ions. The transport experienced by these thermal ions, which includes turbulence, is complicated. In order to avoid modelling this complex physics, a radial diffusion model has been incorporated.

Note also that VENUS-LEVIS does not include self-collisions. Meaning, markers only experience a friction due to collisions with plasma species other than their own. This is justified as long as the marker species is a minority, where ‘species’ in this context refers to distinct populations which can be of the same particle type. For example, a hydrogen plasma with hydrogen beams will have electrons, thermal hydrogen and a high energy tail population which can be modelled as its own hydrogen ‘species’. VENUS-LEVIS would then be employed to only resolve this fast H population, which is indeed a small fraction of the total number of plasma particles, so neglecting self-collisions is justified. If thermal markers are not removed from the simulation this leads to an underestimation of the collisional drag. However, as will be shown later on, the particle fraction of markers is still only a few per cent in the examined cases.

### 2.1. *The RF-NBI model, step by step*

Typically, when simulating NBI using a Monte Carlo orbit tracer such as VENUS-LEVIS, all markers are spawned in the plasma at the start of the simulation, contrary to a continuous injection. This is allowed since every marker moves independently, and the procedure provides sufficient information to reconstruct the steady-state beam profiles.<sup>3</sup> They are then traced until they are either lost, or thermalised. Along the way the state of the markers is saved a number of times (e.g. every millisecond). From these data the beam profiles can be reconstructed.

To make the plasma quasi-neutral the beam species density profile is then subtracted from that of the background ion species’. These profiles are then passed on to LEMan. One limitation of the model is that the steady-state concentration of the beam is unknown at the start of the simulation, so there is no way to ensure the final beam concentration matches that of an experiment. Measurements provide the plasma profiles of all the species, and the injection rate of the beam is fixed too. Which leaves just one free parameter, the radial diffusion rate. This coefficient can be varied to find a best fit of the beam species concentration to that of the experiment (such a scan is not performed here).

The problem with the single injection (all markers spawned at once) is that it ends when all markers have been lost, thermalised or a given time is exceeded. The last state of the markers is therefore empty or nearly empty. For this reason a different approach is used when RF is also involved: markers are continuously injected and lost, and the simulation is run until particle balance sets in. Once this state is reached the profiles can be computed, LEMan is then run and VENUS-LEVIS can resume its run using the last marker state. This contains both freshly injected and slowed down beam ions.

<sup>2</sup>We do not rule out that a more efficient method exists, in which the thermal markers are kicked out and replaced by a known source term of fast ions that come from acceleration by the RF wave. This would result in better statistics. However, in this work a more pragmatic approach is chosen: the full minority distribution function is included (as in pure RF), but an ad hoc sink is introduced in the form of a diffusion model.

<sup>3</sup>This does not reproduce the true time dependence of the beam density profile, but that is irrelevant as these profiles are combined in a later stage anyway.

However, continuous injection has some challenges as well: the final number of confined markers (in steady state) is not yet known at the start of the simulation. Also the necessary simulated time needed to reach equilibrium is unknown. This is incompatible with VENUS-LEVIS because it is designed to model a fixed number of markers for a given time. A workaround is to introduce the concept of ‘cycles’. A cycle represents one run of VENUS-LEVIS, for a given fixed time (e.g. 100 ms). Instead of generating a marker list for just one cycle, it is made to include all future cycles as well. If after completing cycle  $j$  the simulation is not yet deemed converged, cycle  $j + 1$  can be launched by simply shifting the starting index in the marker list, which determines the freshly injected markers during cycle  $j + 1$ . After some time the simulation contains markers from cycles  $j + 1, j, j - 1, \dots$ , but most of the old ones will be lost (inactive). This leads to a larger, but still easily manageable memory footprint. In addition, the computational overhead associated with keeping track of the extra inactive markers is negligible.<sup>4</sup>

For simplicity, in this work an isotropic pressure profile is used to compute the magnetic equilibrium, so using VMEC, not ANIMEC. In addition, it is run only once, instead of including it in the iterative loop. In summary, the modelling procedure is as follows:

- (i) compute magnetic equilibrium using VMEC;
- (ii) compute NBI deposition (i.e. where the beam particles ionise) with VENUS-LEVIS;
- (iii) run VENUS-LEVIS (for multiple cycles) until steady state is reached;
- (iv) generate plasma profiles of the beam species and adjust the background density to make it quasi-neutral;
- (v) run LEMan to generate the wave field;
- (vi) repeat step 3, but now with ICRH enabled; and
- (vii) repeat steps 4-6 until converged.

## 2.2. Diffusion operator

In order to include the effect of collisions on markers, VENUS-LEVIS uses Monte Carlo operators. These are assumed to be independent. So the effects of pitch-angle scattering, energy kicks, interaction with the wave field, etc. can be treated separately. For related work see (Boozer & Kuo-Petravic 1981; Eriksson & Helander 1994), however, note that VENUS-LEVIS is not orbit averaged. VENUS-LEVIS also has the option to enable the effect of radial diffusion on markers, which too is implemented as a Monte Carlo operator. Now adopting the same diffusion operator as used by ASCOT, which is valid for any flux surface shape, including stellarators (Hirvijoki & Kurki-Suonio 2012). Its derivation can be summarised as follows:

the term responsible for radial diffusion can be written as

$$\left(\frac{\partial f}{\partial t}\right)_d = \nabla \cdot (D\nabla f). \quad (2.1)$$

With  $f$  the distribution function of the marker species and  $D$  the diffusion coefficient. The subscript  $d$  indicates that this is just the contribution of diffusion. Defining a generic moment as follows:

$$\langle h \rangle = \int_{\Omega} fh \, dV, \quad (2.2)$$

<sup>4</sup>Cleaner solutions exist, but those would require a considerable rewrite of an otherwise well working code.

where the integration is over the entire plasma volume  $\Omega$ . The time derivative, due to diffusion alone, is

$$\left(\frac{d\langle h \rangle}{dt}\right)_d = \int_{\Omega} \left(\frac{\partial f}{\partial t}\right)_d h dV = \int_{\Omega} h \nabla \cdot (D \nabla f) dV. \tag{2.3}$$

Integrating by parts (twice) gives

$$\left(\frac{d\langle h \rangle}{dt}\right)_d = \int_{\Omega} f \nabla \cdot (D \nabla h) dV + \int_{\partial \Omega} D (h \nabla f - f \nabla h) \cdot \mathbf{n} dS = \langle \nabla \cdot (D \nabla h) \rangle. \tag{2.4}$$

The surface integral appears due to the divergence theorem, but  $f$  is assumed to vanish at the plasma boundary, so this surface term disappears. Note that this assumption will be easily satisfied later on as a delta function will be used for  $f$ . Now the following coordinate system  $(u^1, u^2, u^3)$  with  $u^1$  the radial flux label is used. Substitution of  $h = u^1$  or  $h = (u^1)^2$  in (2.4) yields, respectively, the following in curvilinear coordinates:

$$\left(\frac{d\langle u^1 \rangle}{dt}\right)_d = \left\langle \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^i} (\sqrt{g} D g^{1i}) \right\rangle, \tag{2.5}$$

$$\left(\frac{d\langle (u^1)^2 \rangle}{dt}\right)_d = 2 \left\langle \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^i} (\sqrt{g} D u^1 g^{1i}) \right\rangle, \tag{2.6}$$

where the Einstein summation convention is applied,  $g^{ij}$  is the metric tensor and  $\sqrt{g}$  the Jacobian determinant. The time evolution of the variance  $\sigma^2$  can be written as

$$\left(\frac{d\sigma^2}{dt}\right)_d = \left(\frac{d}{dt} \left[ \langle (u^1)^2 \rangle - \langle u^1 \rangle^2 \right]\right)_d = \left(\frac{d\langle (u^1)^2 \rangle}{dt}\right)_d - 2 \langle u^1 \rangle \left(\frac{d\langle u^1 \rangle}{dt}\right)_d. \tag{2.7}$$

Initially, a particle is located in just one point  $u^1(0)$ , so its distribution function is  $\delta(u^1 - u^1(0))$ . Note that this Dirac delta is weighted such that  $\langle 1 \rangle = 1$ . Thus, for  $t = 0$  the volume in (2.6) can be eliminated

$$\left. \begin{aligned} \left(\frac{d\langle u^1 \rangle}{dt}\right)_d \Big|_{t=0} &= \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^i} (\sqrt{g} D g^{1i}), \\ \left(\frac{d\langle (u^1)^2 \rangle}{dt}\right)_d \Big|_{t=0} &= \frac{2}{\sqrt{g}} \left( u^1(0) \frac{\partial}{\partial u^i} (\sqrt{g} D g^{1i}) + \sqrt{g} D g^{11} \right), \\ \left(\frac{d\sigma^2}{dt}\right)_d \Big|_{t=0} &= 2 D g^{11}, \end{aligned} \right\} \tag{2.8}$$

where the right-hand side is also evaluated at  $t = 0$ . From this the stochastic radial kick is constructed

$$\begin{aligned} \Delta u^1 &= \Delta t \left(\frac{d\langle u^1 \rangle}{dt}\right)_d \Big|_{t=0} + R \sqrt{\Delta t \left(\frac{d\sigma^2}{dt}\right)_d \Big|_{t=0}} \\ &= \frac{1}{\sqrt{g}} \frac{\partial}{\partial u^i} (\sqrt{g} D g^{1i}) \Delta t + R \sqrt{2 D g^{11} \Delta t}, \end{aligned} \tag{2.9}$$

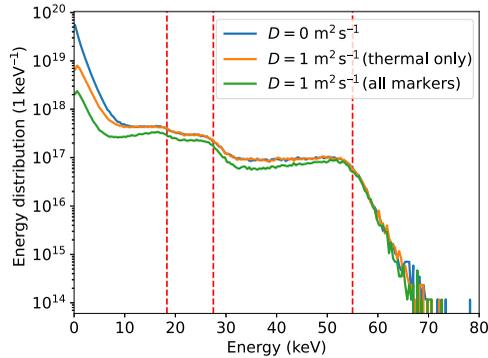


FIGURE 2. Comparison of the energy distribution for three cases: no diffusion, diffusion on thermal markers only or diffusion on all markers. The data represent the NBI slowing down distribution using base case parameters after 200 ms, normalised such that the integral over energy gives the total number of H ions in the plasma.

with  $R$  a randomly selected plus/minus sign. For this derivation no assumption has been made on the magnetic equilibrium, nor on the spatial dependence of  $D$ . But for simplicity, a homogeneous diffusion coefficient has been chosen in this work. Along the same lines as laid out above, a Monte Carlo operator for diffusion in the angular directions  $u^2$ ,  $u^3$  can also be derived. However, this type of diffusive transport in the angular variables is unimportant compared with the normal transport channels in those directions, hence it is omitted in VENUS-LEVIS.

The impact of diffusion is studied in a pure NBI scenario, see figure 2. Here, thermal markers are defined as having an energy less than 1.5 times the half-radius electron temperature, i.e.  $E < 1.5T_e$  ( $s = \rho^2 = 0.5$ ), and ‘fast’ is anything above that. The diffusion mostly effects low energy markers, but even in the range 20–50 keV a small but noticeable effect is seen. Above 50 keV the effects are negligible, as seen by the convergence of the lines around this energy in figure 2. To explain this effect, note the following: in the absence of RF there is no mechanism for markers to speed up (apart from the random collisions that lead to some markers increasing their energy). Thus the lower energy markers have been in the plasma for longer, during which diffusion has had a chance to kick out of some of them. For the freshly injected markers on the other hand, diffusion has not had the time. Hence the losses accumulate at the low energy range.

Besides the direct effect of diffusion on fast ions, there is also an indirect contribution, namely, the thermal concentration affects the dielectric properties of the plasma. And so in scenarios involving RF heating, the wave propagation and polarisation are affected. This in turn modifies the RF power per ion. As was already touched on in § 2, VENUS-LEVIS is designed for modelling fast ions, not thermal ions. This deficiency is remedied by adding a diffusion operator. However, to be consistent with the assumption that collisions are accurately modelled for fast ions, the diffusion is only applied to the thermal ions in the rest of this work. To be explicit, the following expression is adopted for the diffusion coefficient:

$$D(E) = \begin{cases} D_0, & E < 1.5T_e (s = 0.5), \\ 0, & E \geq 1.5T_e (s = 0.5). \end{cases} \quad (2.10)$$

Case	$N_e$ ( $10^{19} \text{ m}^{-3}$ )	$T_e$ (keV)	$T_i$ (keV)	$D_0$ ( $\text{m}^2 \text{ s}^{-1}$ )
Base case	5.0	3.0	1.5	1.0
$N \times 2$	10.0	3.0	1.5	1.0
$N \times 3$	15.0	3.0	1.5	1.0
$T \times 0.5$	5.0	1.5	0.75	1.0
$T \times 2$	5.0	6.0	3.0	1.0
$D_0 \times 0.5$	5.0	3.0	1.5	0.5
$D_0 \times 2$	5.0	3.0	1.5	2.0

TABLE 1. Description of the various test cases used. The density  $N$  and temperature  $T$  are on-axis values. For the profiles used see [Appendix A](#). The magnetic equilibrium is generated using base case parameters.

### 3. Neutral beam injection

Using the newly upgraded version of the code, the capabilities of fast ion generation in W7-X are assessed for several NBI and RF-NBI test cases, and their parameters are listed in [table 1](#). The plasma consists of a deuterium background, with H beams. Balanced injection is used (injectors 3, 4, 7, 8), which results in a total inflow of approximately  $1.2 \times 10^{21}$  protons per second. This amounts to a total injected power of approximately 6.84 MW. See Rust *et al.* (2011) for more information on the W7-X NBI system. Note that D can be swapped with He-4 with little change to the wave physics: in cold plasma two species with the same charge to mass ratio, and the same plasma frequency, make equal contributions to the dielectric tensor. Because of quasi-neutrality, the density of He-4 would be half that of D, therefore their plasma frequency is the same. However, it should be mentioned that the contributions are not exactly the same in hot plasma due to the difference in thermal velocity.

#### 3.1. Steady-state base case

After two cycles of 100 ms a steady state is reached, meaning loss rate  $\approx$  birth rate, see [figure 3](#). Note that the birth rate is less than the injection rate because of shine through. After one cycle of 100 ms, steady state is almost reached, but the amount of new markers still exceeds the number of lost ones. In the second cycle, 1 990 656 are born, and 1 918 915 are lost, demonstrating that a balance is reached. Note that perfect balance is not required, as multiple iterations of VENUS-LEVIS will be run in more complete simulations. After reaching this balance, moments of the slowing-down distribution function can be computed, such as the density as shown in [figure 4\(a\)](#), and the anisotropic pressure. As described in (Jucker 2010), the parallel and perpendicular temperatures can then be derived from this pressure, see [figure 4\(b\)](#). The central hydrogen concentration is approximately 7.3 %. To make the plasma quasi-neutral, this profile is subtracted from that of deuterium.

In order to use the marker population in the computation of the wave field, it is first split into a thermal background (isotropic Maxwellian) and a fast part (modified bi-Maxwellian Dendy *et al.* 1995). This splitting is unrelated to that used for the diffusion; it is chosen to find a reasonable fit to the marker population using a combination of the two aforementioned distribution functions. In this particular case ‘fast’ was defined as having  $E > 6T_e(s)$ , so using the local electron temperature. This criterion has also been used to generate [figure 4](#). The distribution of particle pitch  $\lambda = v_{\parallel}/v$  is seen in [figure 5](#).

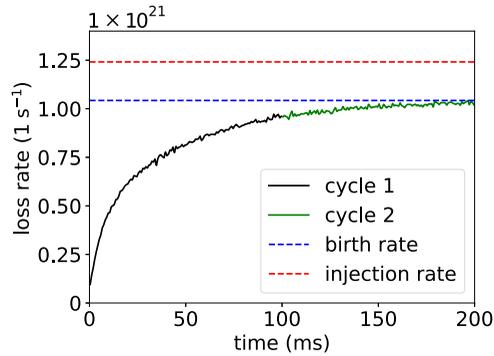


FIGURE 3. Loss rate of beam particles vs time. Using base case parameters, [table 1](#).

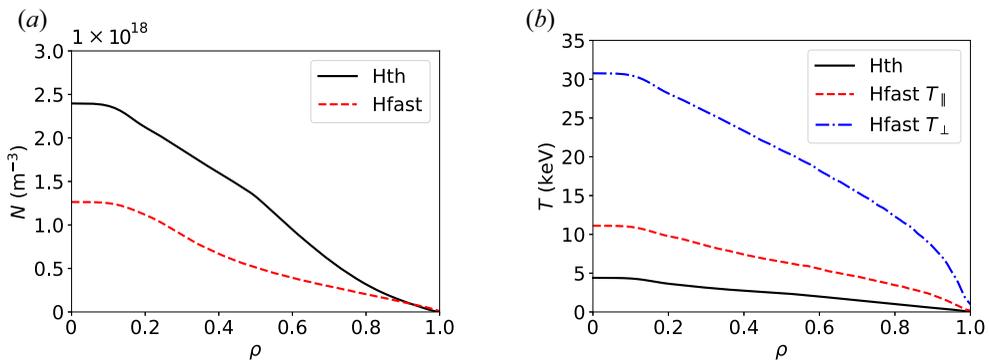


FIGURE 4. Hydrogen plasma profiles vs radial flux label  $\rho$ . The coordinate  $\rho$  is defined as the square root of the normalised toroidal flux. Base case parameters used. (a) Density, (b) temperature.

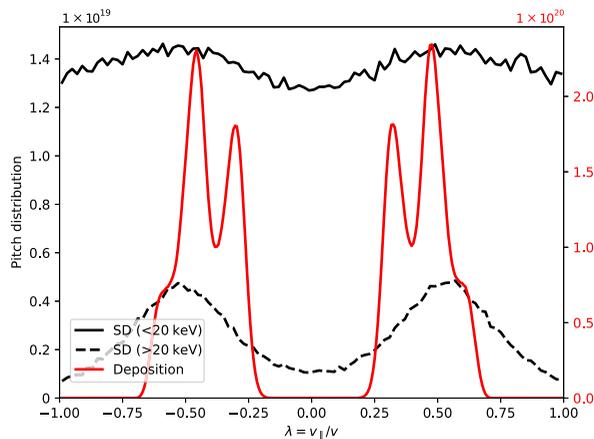


FIGURE 5. Particle pitch distribution for the slowing down (SD) distribution after 2 cycles (200 ms), normalised such that the integral over pitch gives the total number of H ions in the plasma. In red, with corresponding y-axis on the right, the deposition of particles for one whole cycle. The symmetry is a result of using balanced beams.

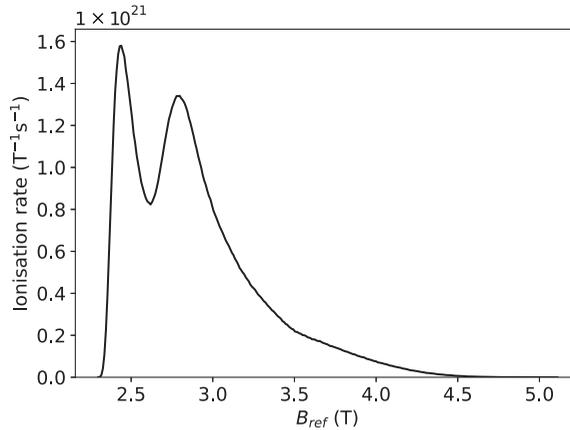


FIGURE 6. Distribution of  $B_{ref} = E/\mu$  values for the freshly injected beam ions, normalised so that the integral over  $B_{ref}$  gives the total ionisation rate. For comparison, the value of the magnetic field strength on axis is approximately 2.43 T in front of the antenna.

Injectors 4,8 are more tangential, creating peaks at  $\lambda = \pm 0.46$ , while 3,7 are more radial, resulting in peaks at  $\lambda = \pm 0.30$  (figure 5). The exact pitch of the peaks in figure 5 depends on the magnetic field. After some time has passed, these peaks have become smoothed due to pitch-angle scattering, which is most effective at low energy. This explains why the ‘low’ energy ( $< 20$  keV) part of the distribution is nearly isotropic in pitch. Eventually the distribution reaches a steady state which is well described by a Maxwellian plus a high energy population. This low density, high temperature population will never become isotropic no matter the waiting time, because the particle source itself is anisotropic.

Note that the initial NBI population consists overwhelmingly of passing particles, see figure 6, which shows the distribution of fresh beam ions in terms of their reflection field  $B_{ref} = E/\mu$ . Where  $E$  is the kinetic energy, and  $\mu = mv_{\perp}^2/(2B)$  the magnetic moment. A particle is classified as passing when  $B_{ref} > \max(B)$ , and trapped otherwise, where the maximum is to be taken along its full trajectory. In other words, reflection would occur at  $B = B_{ref}$ , but if  $B$  is always less than  $B_{ref}$  this bounce point is never encountered, hence it is passing. Figure 4(b) shows that the perpendicular temperature of the fast population exceeds its parallel temperature. This does not contradict the finding that most particles are passing; indeed  $v_{\perp} > |v_{\parallel}|$  on average, which results in the anisotropy, yet the ions are mostly passing because of the injection sites. Which are at high magnetic field strength, meaning small  $\mu$ , and thus large  $B_{ref}$ .

Note that in all of the runs in this work the same high mirror (KJM) configuration has been used. The reason for this choice is because the high mirror scenario results in better core NBI confinement than in low/standard mirror (Patten 2019). To clarify, if the empirical goal is simply RF bulk heating, one may be better off using the low mirror because the magnetic field strength varies less in toroidal direction. Suppose on-axis heating is used in front of the antenna, then ions will still resonate near the axis at all toroidal angles. So the share of thermal ions in the core which can interact with the wave is maximised. In the high mirror, on the other hand, the resonant surface even leaves the plasma in the triangular cross-section. However, remember that bulk heating is not the aim of this work. The goal is to maximise fast ion generation for advanced studies, such as testing confinement, exciting Alfvén waves and transport suppression, which can only be

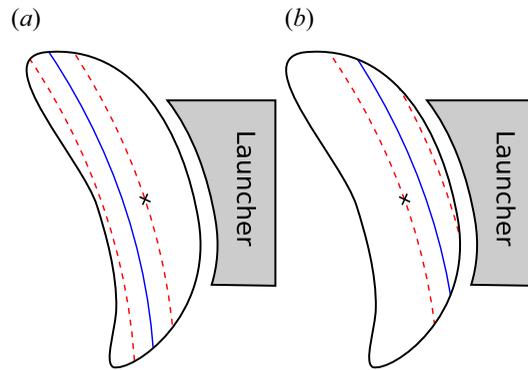


FIGURE 7. Simple cartoon to illustrate where the resonance is. The thermal ions resonate near the cold resonance ( $\omega = \omega_c$ ), indicated in blue. The red dashed lines indicate where the beam ions resonate, and the  $\times$  highlights the magnetic axis. In panel (a) the antenna frequency is shifted up, in (b) it is shifted down.

achieved if the NBI ions are well confined, and if these ions can be further accelerated by RF.

#### 4. The ICRF wave

##### 4.1. Doppler shift

The wave–particle resonance is given by (1.1). In this work the antenna frequency is chosen such that the hydrogen resonates at the fundamental cyclotron frequency, i.e.  $n = 1$ . In addition electrons will resonate at  $n = 0$  and deuterium at  $n = 2$ . The latter being a hot plasma effect, which LEMan can now simulate due to its recent upgrades, detailed in (Patten 2019) and (Machielsen, Graves & Cooper 2021). Including hot plasma effects is important in order to get an accurate estimate of the power fractions of the various species.

The thermal ions will resonate near the cold resonance,  $\omega = \omega_c$ , but freshly injected beam ions experience a considerable Doppler shift, as illustrated in figure 7. As the wave enters the plasma from the low-field side it will first encounter a resonance with the beam, then the cold resonance and then another resonance with the beam. This Doppler shift itself cannot be controlled in any feasible way. On the other hand, an operator can adjust the antenna frequency up or down (i.e. ‘shifting’ it) so that one of the beam resonances goes through the core.<sup>5</sup> In the simulation the antenna frequency is set indirectly by choosing a resonant magnetic field  $B_{\text{res}} = m\omega/q$ .

Choosing the optimal antenna frequency for fast ion generation is not obvious. Ideally, the beam ions should resonate in the core, where the confinement is best. However, the magnetic field strength on axis depends on toroidal angle. Secondly, there is a spread of parallel velocities figure 5, and  $k_{\parallel}$  is also not uniform in space. The optimal antenna frequency should follow from an extensive simulation (or experimental) scan, but a rough estimate can be obtained by assuming the following values:  $\lambda = \pm 0.4$  and full beam energy (55 keV), hence  $v = 3.2 \times 10^6 \text{ m s}^{-1}$ . The parallel wavenumber is computed by LEMan, the on-axis value near the antenna is found to be  $k_{\parallel} = 14 \text{ m}^{-1}$ . This results in a frequency shift of approximately 2.9 MHz (up or down), corresponding to a shift in resonant magnetic field of 0.19 T. For the chosen magnetic configuration the field strength

<sup>5</sup>The antenna shift is defined with respect to the default case, in which the antenna frequency is chosen such that the cold resonance is on axis in front of the antenna. In practice it may be more convenient to leave the antenna frequency fixed, and instead modify the magnetic field strength.

on axis in front of the antenna is approximately 2.43 T, meaning  $B_{\text{res}} = (2.43 \pm 0.19)$  T, and the antenna frequency should be  $\omega/(2\pi) = (37.1 \pm 2.9)$  MHz. As a starting point, a value of  $B_{\text{res}} = 2.5$  T is used unless specified otherwise. In order to evaluate the dielectric tensor, LEMan assumes that locally in each point a single parallel wavenumber dominates (or set of wavenumbers). To replicate dipole phasing, both  $k_{\parallel}$ ,  $-k_{\parallel}$  are used. See Faustini (2017, § 4.5.1) for a description of the antenna model. The LEMan simulations in this work used a radial resolution of 200 elements and  $30 \times 300$  (poloidal  $\times$  toroidal) Fourier modes.

#### 4.2. Wave field

Several different plasma scenarios are compared, and their parameters summarised in [table 1](#). With respect to the base case, an increased density leads to a lower power fraction of H. This is because the injection rate is kept the same, while the background density has increased, i.e. there is a lower H density fraction. It should be mentioned that the lower minority concentration leads to an improved wave polarisation at the cyclotron resonance. However, the H concentration in the base case was not yet high enough to significantly deteriorate the wave polarisation, therefore lowering the concentration results in less minority heating. An increased diffusion rate leads to a lower thermal concentration, which results in a somewhat increased fast H power fraction. If instead the temperature is raised, the damping on electrons and D is more effective. This is not surprising, because D only absorbs via second harmonic heating, which is a finite Larmor radius (FLR) effect. In addition, the increased temperature results in a broader resonance layer (Doppler broadening due to larger  $T_{\parallel}$ ), which could also explain part of the increased power fraction of D. Electrons also absorb more power, in part because transit time magnetic pumping plays a bigger role as temperature is increased.

The power fractions in [figure 8](#) only indicate the direct wave heating. Subsequently, the power is redistributed over the plasma species via collisions. This collisional power is not computed by LEMan, but is obtained by the follow-up VENUS-LEVIS run, see e.g. [figure 15](#).

It can be seen in [figure 9](#) that the RF power deposition is not peaked on axis. There are multiple reasons for this: the ion cyclotron resonance happens near a surface where  $\omega = \omega_c$ . This surface gets close to the magnetic axis in front of the antenna, but departs from the magnetic axis at other toroidal angles, meaning that, at those other angles, the power is mainly deposited off axis. Secondly, the aforementioned surface cuts through the plasma, and resonance therefore does not just occur on axis but also at other radial locations (all the way to the plasma edge). And lastly, as can be seen from [figure 10](#), [figure 11](#), the wave amplitude is not at all homogeneous. It would be a coincidence if a peak in  $|E^+|$  was exactly aligned with the magnetic axis.

As can be seen from the electric field in [figure 12](#), [figure 13](#) the wave propagates less far in the high density case (stronger absorption). The heating of the plasma is therefore also more localised near the antenna.

### 5. Fast ion distribution

In order to compute the effect of the ICRF wave on the markers, VENUS-LEVIS needs the correct wave amplitude. To this end VENUS-LEVIS is run first for a number of time steps to measure the absorbed wave power. The field is then rescaled such that this power matches the expected coupled power (assumed to be 1.5 MW here), multiplied by the power fraction of the H species. The simulation is restarted after this step. Later in the

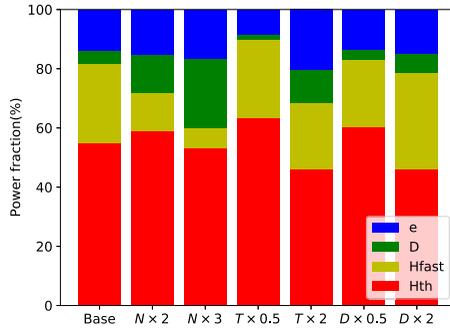


FIGURE 8. Fraction of the RF power directly deposited on the various plasma species. A distinction is made between the thermal (Maxwellian) hydrogen and the fast hydrogen.

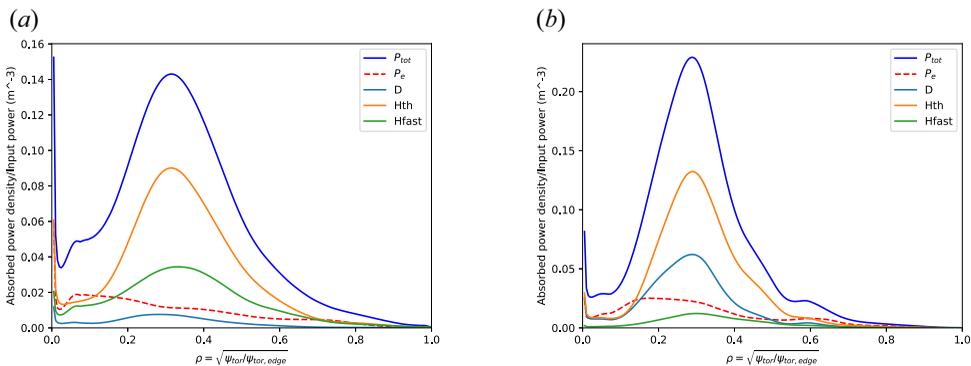


FIGURE 9. Flux surface averaged power density vs radial position. (a) Base case, (b) density  $\times 3$ .

simulation the field is rescaled at regular time intervals,<sup>6</sup> but since these are typically minor adjustments the simulation is not reset to the state of the last time interval.

The addition of RF provides a clear improvement in the fast ion tail over pure NBI, see figure 14 (and appendix A for all cases). However, the effectiveness strongly relies on the plasma parameters. In the high density case the collisionality is too high, preventing a significant tail from forming. The opposite is true for the case with increased temperature (reduced collisionality).

The binning in figure 14 has been truncated at 150 keV, where all distributions disappear into noise. This can be improved in the future by using more markers, which reduces the noise and higher energies will be resolved. However, this would require excessive marker counts. Alternatively, the effect of noise can be mitigated by using markers with adaptive weights (Schuster *et al.* 2021). The basic idea is to continuously split and delete markers to increase the resolution in the high energy range. Another way to improve the energy resolution is by replacing the thermal population by an analytical particle source term, so only fast ions need to be modelled. In this method the thermal beam ion population is assumed to be Maxwellian, at some given concentration and temperature known from experiments. Using this distribution and the wave field, an expression for the fast ion particle source can be computed in principle. In other words, the left most scheme in

<sup>6</sup>Typically 1 ms intervals, so much larger than the time steps used for updating the markers.

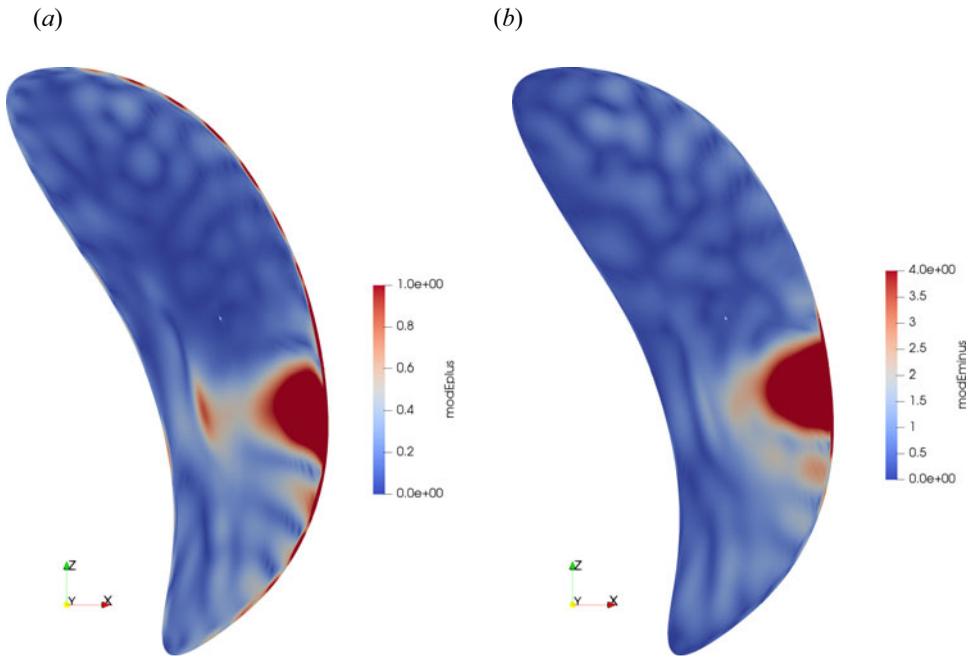


FIGURE 10. A plot of the amplitudes of the left handed (*a*) and right handed (*b*) polarised component of the electric field of the wave in front of the antenna (base case). Note,  $E^+$  is mainly responsible for heating H. (arbitrary units used)

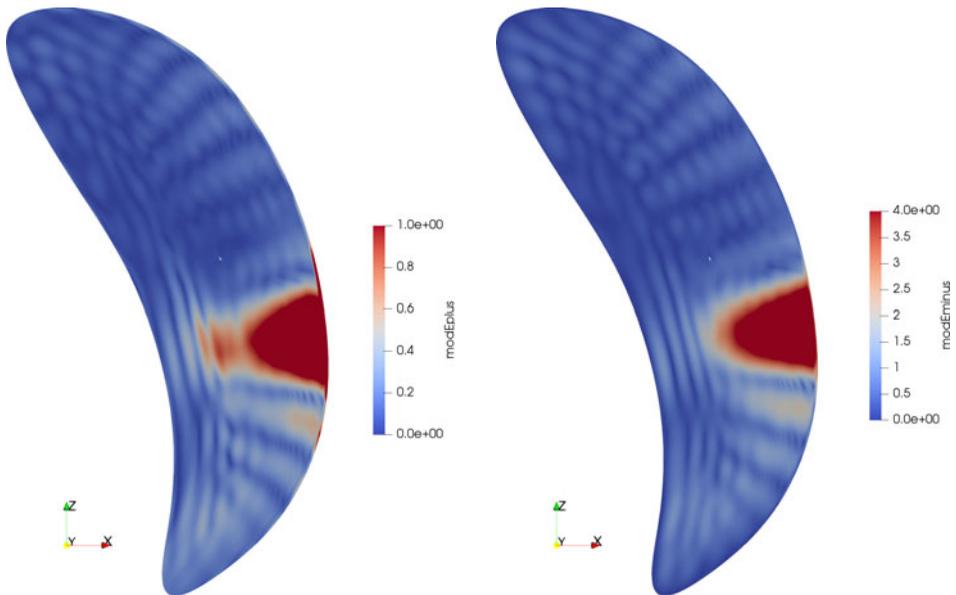


FIGURE 11. Plot of the left and right handed polarised components of the wave field (density  $\times 3$ ).

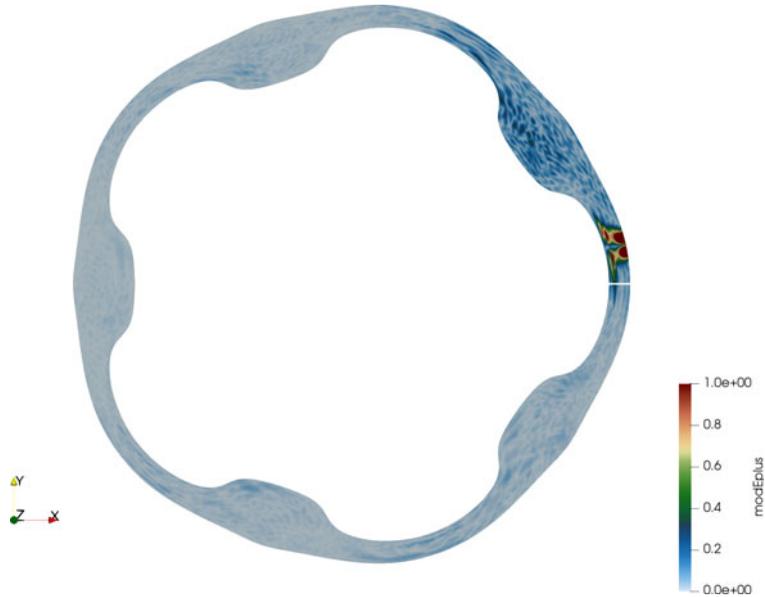


FIGURE 12. Cross-section of the plasma, top-down view, colour indicates  $|E^+|$  in arbitrary units. The hot spot is located just in front of the antenna.

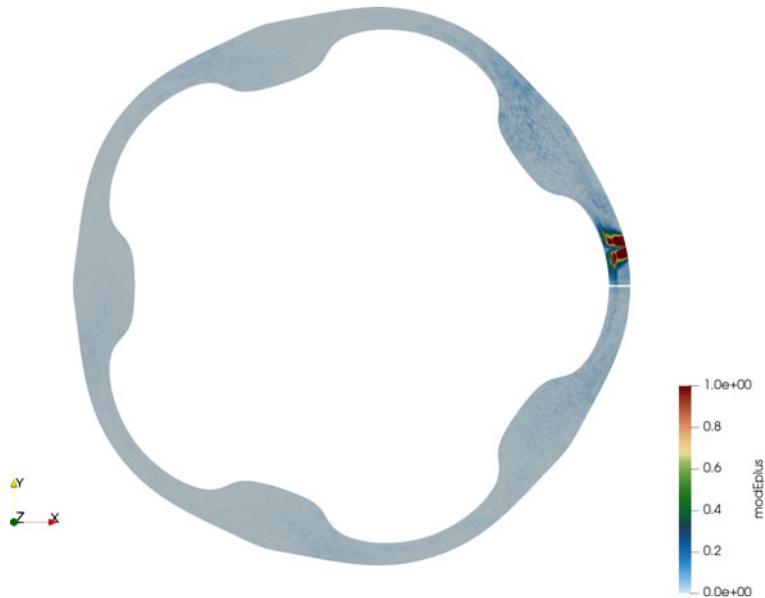


FIGURE 13. Same, but for density x3 case.

figure 1 would be used, instead of that on the right. Except that it will be supplemented by an influx of markers computed based on the RF heating of the bulk population. As the number of thermal markers vastly outnumbers the fast ones, this approach would potentially lead to the biggest improvement. However, deriving such a source term is far from trivial. To be clear, these three approaches are not mutually exclusive, a combination of them can be used.

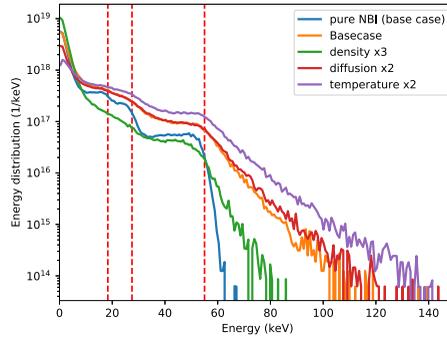


FIGURE 14. Comparison of the H particle energy distributions, normalised such that the integral (i.e. summing over bin values times bin widths) gives the total number of H particles in the whole plasma. The vertical lines indicate the injection energy (full, half, third).

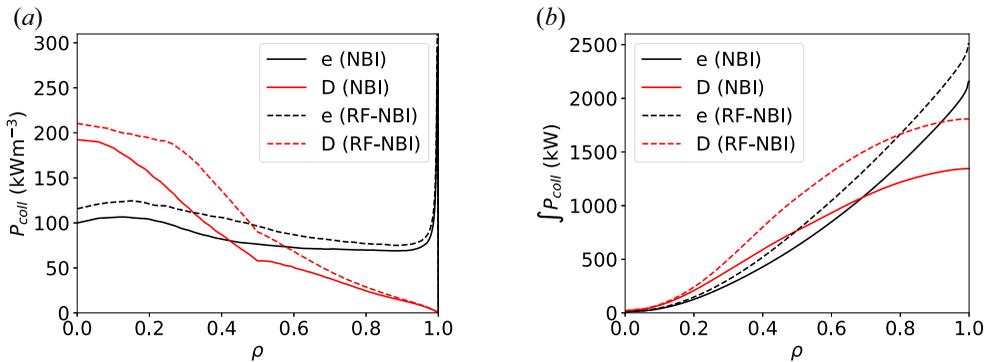


FIGURE 15. (a) Collisional power density vs radial position (base case). For comparison, the NBI-only case (second cycle) has been plotted too. (b) Total (cumulative) collisional power within given flux surface.

VENUS-LEVIS also computes the collisional power on the background species (electrons and deuterons), [figure 15](#). This collisional power is dominated by the beam, meaning that adding RF does not significantly affect this profile. Also observe the spike in electron power near the plasma edge. This is because the beam temperature is much larger than that of the background plasma. A more detailed explanation is found in [Appendix B](#). This spike is an artefact emerging from the choice of plasma profiles, however, it constitutes only a minor fraction of the total power as seen from [figure 15\(b\)](#). Hence its effect on the core physics is negligible.

So far the antenna frequency has been set to 38.1 MHz for simplicity, however, it can be varied in the range of 25–38 MHz ([Ongena \*et al.\* 2020](#)). The antenna frequency dictates  $B_{\text{res}}$ , which strongly affects which ions resonate and where, see (1.1). The more  $B_{\text{res}}$  deviates from the on-axis value (2.43 T), the more the resonant surfaces are displaced. Which means that only particles with larger  $|v_{\parallel}|$  can still resonate in the core where the density is high. As a result, the larger the shift, the more power goes to the fast H vs the thermal H, [figure 16](#). This is a trade-off, the antenna frequency should be shifted to optimally accelerate fast ions present in the core even further. But shifting it too far means that the slower ions do not get a chance to reach this optimal level. This was also investigated in [Patten \(2019, § 5.2.3\)](#). Note that  $B_{\text{res}} = 2.7$  T requires an antenna frequency

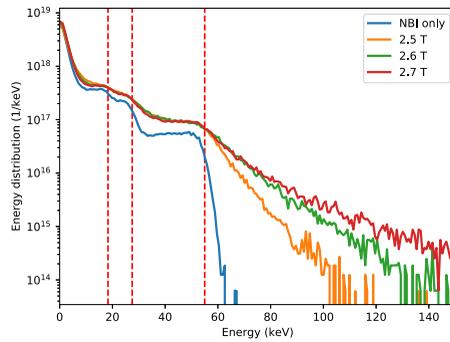


FIGURE 16. H particle energy distributions for the base case, but repeated with different resonant magnetic field strengths.

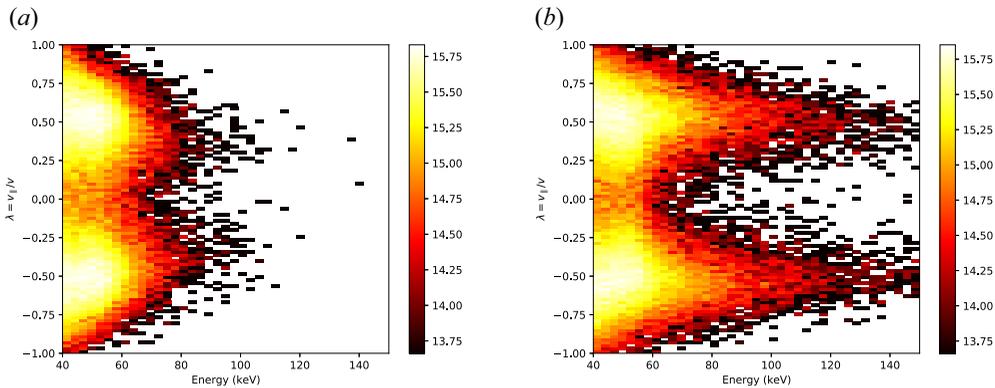


FIGURE 17. Fast ions binned in pitch and energy. The colour indicates  $\log_{10}$  of the total number of particles per bin. Base case used; (a)  $B_{\text{res}} = 2.5$  T, (b)  $B_{\text{res}} = 2.7$  T.

of 41.2 MHz, which is outside of the accessible range, so this is a more academic test case. In a real experiment the coil currents can be reduced, while maintaining the same frequency shift, e.g.  $B_{\text{res}} = 2.6$  T, while having  $B = 2.33$  T on axis in front of the antenna. Another option would be to shift the antenna frequency down instead of up. This is yet another variation that should be investigated. Note that the resonant surface at  $\omega = \omega_c$  is surrounded by two Doppler shifted surfaces, as illustrated in figure 7. Since the injection is from the outboard side, the wave first resonates with the beam ions, then with the thermal H and then with the beam again. By shifting the frequency down, the beam ions in the core encounter the wave last, which may lead to less acceleration of beam ions. Also note that, by shifting  $\omega$  too much (up or down), only two of the three resonances remain in the plasma (near the antenna).

The particle pitch is also investigated, and found to be anisotropic, see figure 17. It is mostly similar to the pitch in figure 5, except with enhanced energy. Particles with low pitch do exist, but they are lost rapidly, see figure 18. The losses can also be binned in terms of toroidal Boozer angle, see figure 19. It is nearly, but not fully, fivefold periodic.

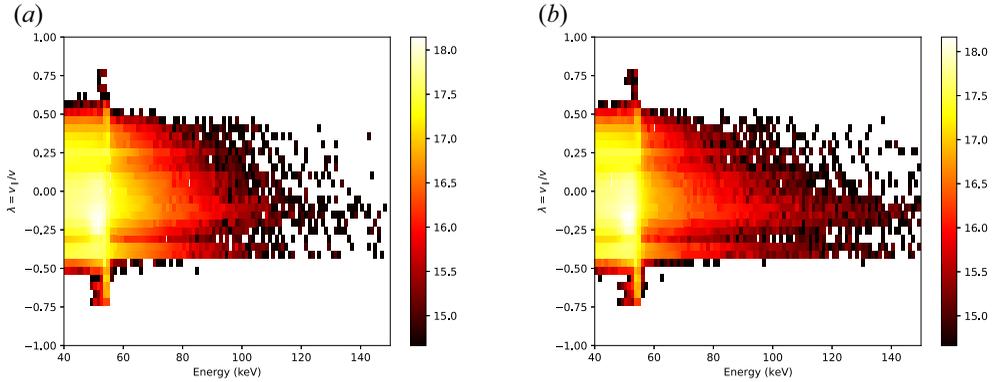


FIGURE 18. Fast ion losses binned in pitch and energy, for the base case. The colour of a bin indicates  $\log_{10}$  of the lost particles per second; (a)  $B_{\text{res}} = 2.5$  T, (b)  $B_{\text{res}} = 2.7$  T.

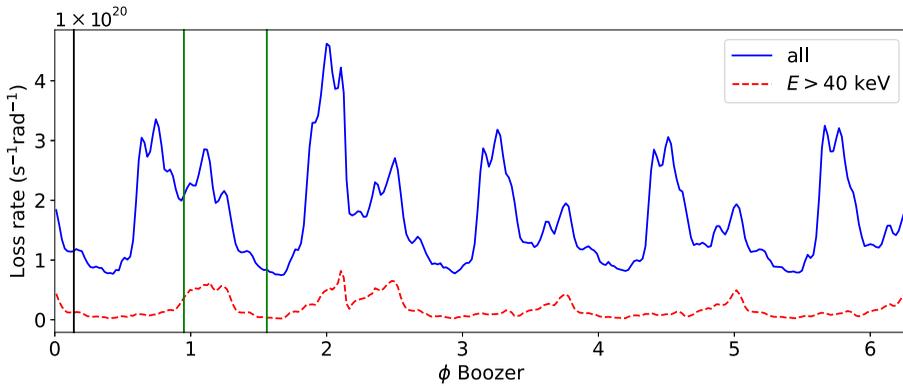


FIGURE 19. Particle loss rate vs toroidal Boozer angle  $\phi$ , normalised such that the integral over  $\phi$  gives the total loss rate in particles/second. The black vertical line indicates the position of the antenna, and the green lines the positions where the NBI deposition is strongest.

## 6. Discussion

The current modelling approach presents a number of notable strengths and weaknesses, summarised below.

Strong points:

- (i) Sampling particles and tracing their orbits is an effective method to deal with the high-dimensional problem associated with the evolution of the distribution function.
- (ii) The model finds a self-consistent solution for the fast ion distribution function and RF wave field in the inherently three-dimensional plasma of a stellarator.
- (iii) VENUS-LEVIS includes finite orbit width effects.
- (iv) The wave code uses a localised antenna model in three dimensions.
- (v) RF heating on all species is computed, including FLR effects. Based on the H power fraction, the electric field amplitude is rescaled in VENUS-LEVIS.

Weak points:

- (i) Geometric limitations: the antenna model is tied to the magnetic equilibrium. This equilibrium is computed by VMEC, so it assumes nested magnetic flux surfaces, and

runs up to the LCFS only. Moreover, VMEC is not included in the self-consistent loop, meaning  $B_0$  is assumed to be fixed. Lastly, the NBI injection geometry is simplified compared with that in ASCOT/BEAMS3D.

- (ii) Antenna model is an approximation of dipole phasing.
- (iii) A steady-state radial electric field is not included.
- (iv) The diffusion model requires an additional input, namely, the diffusion coefficient. However, this *ad hoc* parameter can be used to match the minority concentration to that measured in an experiment.
- (v) Mode conversion is not included in the wave code.

#### General comments:

VENUS-LEVIS does not use any type of bounce averaging, it solves the full four-dimensional guiding centre equations. This means that the model includes finite orbit width effects, which are important for getting a realistic estimate of the particle losses and collisional power deposition profile. Coulomb collisions are included as well in the form of a Monte Carlo operator.

Regarding the full-wave code, the antenna is implemented with a given radial, poloidal and toroidal extension ('localised in three dimensions'). Additionally, the dielectric tensor in the hot model makes no approximation on the size of the Larmor radius compared with the wavelength of the electric field.

Compared with the previous RF-NBI work with SCENIC (Patten 2019; Patten *et al.* 2020), several features have been enhanced/added: the use of cycles to resume NBI runs, and the radial diffusion model in order to reach particle balance. In addition, the plasma profiles have been updated and the aforementioned hot plasma model is enabled. Additionally, the NBI deposition in the new version of SCENIC is more complete, taking into account the beam focussing.

Some further improvements could be made in the future. Most notably, on the subject of the geometric effects. It should be mentioned that this is not a fundamental limitation, in the sense that the code could be modified to take in a more general magnetic geometry that extends to the metal wall. Fast ion hot spots on the wall depend rather sensitively on the magnetic configuration (Äkäslompolo *et al.* 2018). In order to get good agreement with experiment on this front, these limitations should be addressed within the SCENIC package. Recently, however, strides have already been made to compute the fast ion wall loads for RF minority heating. This was done by taking the markers from VENUS-LEVIS that crossed the LCFS, and importing them into ASCOT. These were then traced to the actual metal wall (Slaby *et al.* 2022).

A steady-state radial electric field arises from the ambipolarity condition. However, for simplicity it has not been included in these runs. The radial electric field (ion root) typically helps to confine fast ions because of the resulting  $E \times B$  drift, see e.g. Kolesnichenko *et al.* (2006); Faustin *et al.* (2016). Therefore, by omitting this effect the fast ion population is likely underestimated. Especially for trapped ions in the NBI range of energies, which might otherwise be heated by the RF wave. It should be noted that the radial electric field depends on the plasma profiles, so it will be different for every case in table 1.

It should furthermore be noted that in these simulations the placement of the antenna differs somewhat with respect to the real antenna position. In reality it is offset in toroidal direction with respect to the bean section, by approximately  $8^\circ$ , away from the injection points of the beams. In the model it has been offset  $8^\circ$  in the other direction, i.e. closer to the beam injection sites.

## 7. Summary

The radial diffusion operator in VENUS-LEVIS has been generalised to a stellarator geometry. This, together with some other code improvements, enables the marker count to saturate. That is, the particle inflow due to NBI matches the outflow due to losses. Several different scenarios involving synergetic RF-NBI heating have been investigated. The best performance (in terms of fast ion generation) was obtained with a low collisionality plasma, i.e. low background density and high temperature. In addition, for the RF-NBI scheme it is crucial to have the right amount of antenna frequency shift. However, given the vastness of the parameter space, further study is needed to properly optimise this scenario for W7-X.

## Acknowledgements

The Piz Daint (CSCS, Switzerland) supercomputer facilities were used for the simulations presented in this research. Numerical calculations were supported by a grant from CSCS under project ID s1087.

*Editor P. Helander thanks the referees for their advice in evaluating this article.*

## Funding

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 – EUROfusion). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. The project was also supported in part by the Swiss National Science Foundation.

## Declaration of interests

The authors report no conflict of interest.

## Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

## Appendix A. Plasma parameters

The plasma profiles for the base case are found in [figure 20](#). The other scenarios from [table 1](#) use the same profile shapes, just rescaled. These are assumed to be typical profiles for W7-X. The comparison of the energy distributions of [figure 14](#) is expanded to include more scenarios, see [figure 21](#).

## Appendix B. Collisional power

As explained before, VENUS-LEVIS uses Monte Carlo operators to implement collisions. The kick in energy due to Coulomb collisions can be written as follows:

$$\Delta E = - \sum_s 2v_E^s \Delta t \left[ E - \left( \frac{3}{2} + \frac{E}{v_E^s} \frac{dv_E^s}{dE} \right) T_s \right] + 2R \sqrt{\sum_s v_E^s T_s E \Delta t}, \quad (\text{B1})$$

where  $s$  is the background species (so e, electrons or D, deuterons in this case) and  $R$  is a

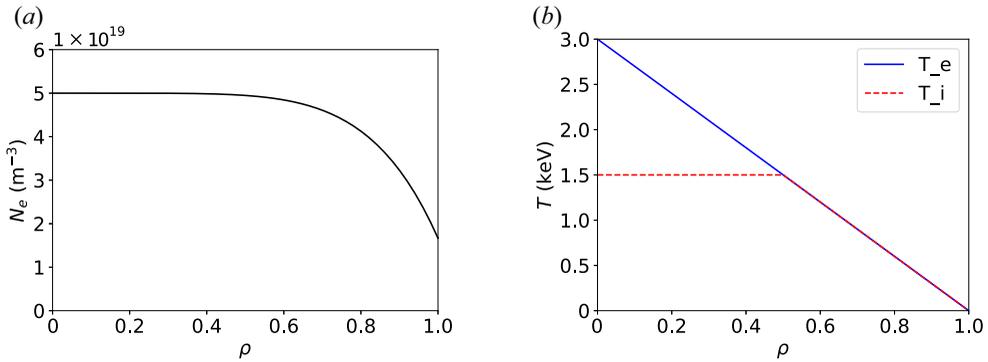


FIGURE 20. Plasma profiles vs flux label  $\rho$ , the latter being defined as the square root of the normalised toroidal flux. (a) Electron density, (b) temperature.

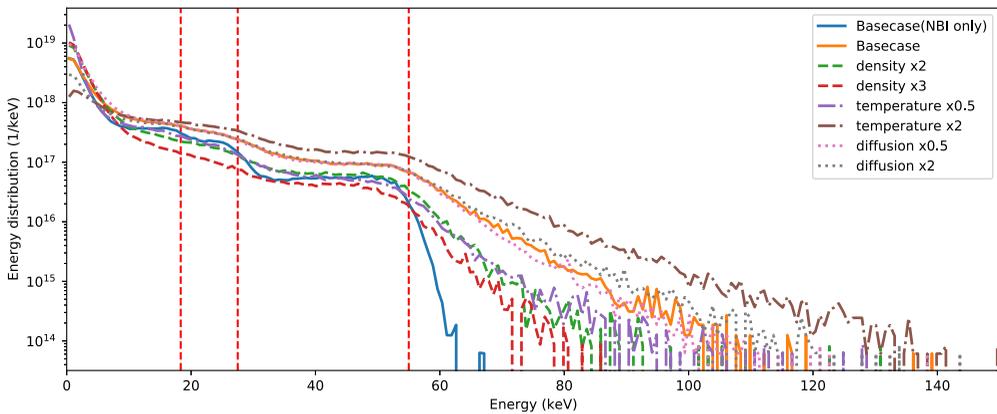


FIGURE 21. Comparison of the H particle energy distributions.

random plus/minus sign. The frequency is given by

$$v_E^s = \frac{c_s \sqrt{m_s}}{\sqrt{2}} \frac{N_s}{T_s^{3/2}} \frac{\Psi(v/v_{T_s})}{v/v_{T_s}}, \tag{B2}$$

with  $v_{T_s}$  the thermal velocity of species  $s$ , and

$$c_s = \frac{q^2 q_s^2 \log \Lambda_s}{4\pi \epsilon_0^2}, \tag{B3}$$

where  $q$  is the particle charge, and  $\log \Lambda_s$  is the Coulomb logarithm. The function  $\Psi$  is defined as

$$\Psi(x) = \frac{\text{erf}(x) - \frac{2}{\sqrt{\pi}} x e^{-x^2}}{2x^2}. \tag{B4}$$

The power transfer to the background by one marker, during one time step, is computed

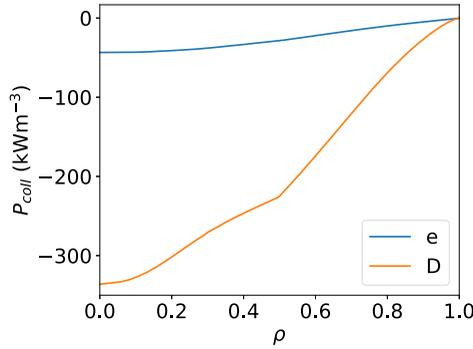


FIGURE 22. Collisional power to background species, plotting only the contribution of  $-\sum_s 3v_E^s G(v/v_{T_s}) T_s$ . Base case parameters, RF-NBI scenario.

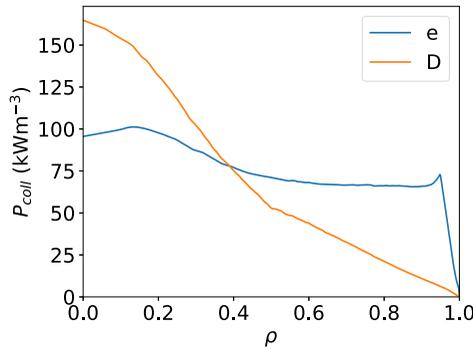


FIGURE 23. Collisional power to background species,  $\min(T_e) = 150$  eV, base case parameters, NBI-only, one cycle.

by taking the advective part of (B1)

$$P_{\text{coll}} = \sum_s 2v_E^s \left[ E - \left( \frac{3}{2} + \frac{E}{v_E^s} \frac{dv_E^s}{dE} \right) T_s \right]. \tag{B5}$$

Using simple algebra, the power can be written as

$$P_{\text{coll}} = \sum_s 2v_E^s \left[ E - \frac{3}{2} G \left( \frac{v}{v_{T_s}} \right) T_s \right], \tag{B6}$$

with

$$G(x) = 1 + \frac{1}{3} \frac{x\Psi'(x) - \Psi(x)}{\Psi(x)} = \frac{4x^3}{3\sqrt{\pi}e^{x^2}\text{erf}(x) - 6x}. \tag{B7}$$

In the studied scenarios the background temperature is just a few eV in the edge.<sup>7</sup> From figure 4(b) the beam temperature can be crudely estimated as 1 keV at  $\rho = 1$ . So for a

<sup>7</sup>Not perfectly zero at  $\rho = 1$  actually, for numerical reasons.

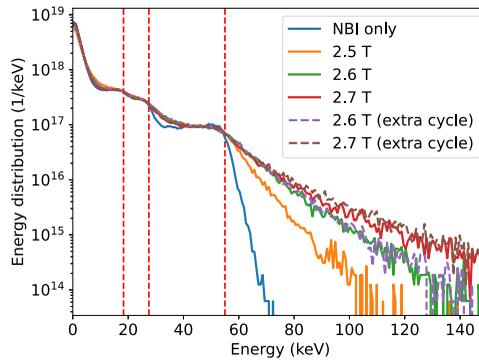


FIGURE 24. The H particle energy distributions for the base case, for different  $B_{res}$ . All runs with RF enabled used three cycles, except for the 2.6 T and 2.7 T cases which have been continued for one additional cycle as indicated by the legend.

typical beam marker with this energy, the  $E$  term in (B6) dominates in the edge

$$P_{\text{coll}} \approx \sum_s 2v_E^s T_H. \quad (\text{B8})$$

The excessive temperature ratio  $T_H/T_e$  causes the spike in the collisional power. No such peak is observed for deuterium because the factor  $\Psi(v/v_{T_D})/(v/v_{T_D}) \approx 0$  in the edge, so  $v_E^D \approx 0$ . To demonstrate that the contribution of  $E$  indeed causes the spike, see figure 22 which shows the collisional power without this term. To be clear, the  $E$  term has not been omitted from the simulation, only from the plot. The power in figure 22 is negative, meaning the background loses power to the H minority species. But this is expected. Since  $E$ ,  $v_E^s$ ,  $T_s$ ,  $G(v/v_{T_s})$  are all positive, the contribution of the  $E$  term to the power is positive, while the rest is negative. It can be argued how realistic this spike in the collisional power is. In reality, this effect is likely mitigated because of negative feedback: if the power really is so high, the background plasma will heat up, lowering the temperature ratio. This last case is also tested explicitly, by rerunning the base case with NBI-only for one cycle, but using a modified electron temperature profile where the minimal temperature is set to 150 eV. That is, the same profiles as in figure 20(b), but a slightly raised  $T_e$  for  $\rho > 0.95$ ; the ion temperature is left unchanged. Under these circumstances the spike is not encountered, see figure 23.

### Appendix C. Note on compute resources

The vast majority of compute time was consumed by the VENUS-LEVIS runs, which ran on 72 nodes on the CPU partition of Piz Daint, XC40, dual socket Intel Xeon E5-2695 v4. One cycle of VENUS-LEVIS took approximately 18 h. Six cycles (100 ms each) worth of marker data were pre-generated for each case in table 1, which is more than enough to reach a converged state. All the simulations start by finding the steady-state marker population for a pure NBI scenario. This usually takes two cycles, see figure 3. Using the plasma profiles, LEMan is run. After this VENUS-LEVIS is ran again for one more cycle. The additional cycles were found to be unnecessary for the base case, see figure 24, hence the number of cycles has been limited to three for all cases. This is because the H population is dominated by NBI. As was shown above, the addition of RF produces an enhanced fast ion tail, but this is at a relatively low concentration compared with the rest of the H. Its effect on the computation of the moments has therefore been found to

be minor. Thus, running LEMan does not produce a significantly different wave field, nor power fractions. The fourth cycle of VENUS-LEVIS thus uses essentially the same inputs, and no appreciable difference was observed. In conclusion, running one case consumed approximately 4000 node-hours.

## REFERENCES

- ÄKÄSLOMPOLO, S., DREVLAK, M., TURKIN, Y., BOZHENKOV, S., JESCHE, T., KONTULA, J., KURKI-SUONIO, T. & WOLF, R.C. 2018 Modelling of NBI ion wall loads in the W7-X stellarator. *Nucl. Fusion* **58** (8), 082010.
- BOOZER, A.H. & KUO-PETRAVIC, G. 1981 Monte carlo evaluation of transport coefficients. *Phys. Fluids* **24** (5), 851.
- COOPER, W.A., HIRSHMAN, S.P., MERKEL, P., GRAVES, J.P., KISSLINGER, J., WOBIG, H.F.G., NARUSHIMA, Y., OKAMURA, S. & WATANABE, K.Y. 2009 Three-dimensional anisotropic pressure free boundary equilibria. *Comput. Phys. Commun.* **180** (9), 1524–1533.
- DENDY, R.O., HASTIE, R.J., MCCLEMENTS, K.G. & MARTIN, T.J. 1995 A model for ideal  $m = 1$  internal kink stabilization by minority ion cyclotron resonant heating. *Phys. Plasmas* **2** (5), 1623–1636.
- 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.
- ERIKSSON, L.-G. & HELANDER, P. 1994 Monte carlo operators for orbit-averaged Fokker–Planck equations. *Phys. Plasmas* **1** (2), 308–314.
- FAUSTIN, J.M., COOPER, W.A., GRAVES, J.P., PFEFFERLÉ, D. & GEIGER, J. 2016 Fast particle loss channels in Wendelstein 7-X. *Nucl. Fusion* **56** (9), 092006.
- FAUSTIN, J.M.P. 2017 Self-consistent interaction of fast particles and icrf waves in 3d applications of fusion plasma devices. PhD thesis, Lausanne, EPFL.
- FREIDBERG, J.P. 2007 *Plasma Physics and Fusion Energy*. Cambridge University Press.
- HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26** (12), 3553–3568.
- HIRVIJOKI, E. & KURKI-SUONIO, T. 2012 Monte Carlo diffusion operator for anomalous radial transport in tokamaks. *Europhys. Lett.* **97** (5), 55002.
- JUCKER, M. 2010 Self-consistent icrh distribution functions and equilibria in magnetically confined plasmas. PhD thesis, EPFL.
- KAZAKOV, Y.O., EESTER, D.V., DUMONT, R. & ONGENA, J. 2015 On resonant ICRF absorption in three-ion component plasmas: a new promising tool for fast ion generation. *Nucl. Fusion* **55** (3), 032001.
- KENNEL, C.F. & ENGELMANN, F. 1966 Velocity space diffusion from weak plasma turbulence in a magnetic field. *Phys. Fluids* **9** (12), 2377.
- KOLESNICHENKO, Y.I., LUTSENKO, V.V., TYKHYY, A.V., WELLER, A., WERNER, A., WOBIG, H. & GEIGER, J. 2006 Effects of the radial electric field on the confinement of trapped fast ions in the Wendelstein 7-X and Helias reactor. *Phys. Plasmas* **13** (7), 072504.
- MACHIENSEN, M., GRAVES, J.P. & COOPER, W.A. 2021 ICRF modelling in 2d and 3d magnetic configurations using a hot plasma model. *Plasma Phys. Control. Fusion* **63** (9), 094002.
- OKADA, H., MURAKAMI, K., JINNO, Y., KOBAYASHI, S., KADO, S., NAGASAKI, K., MINAMI, T., YAMAMOTO, S., OHSHIMA, S., MUTOH, T., *et al.* 2016 Fast ion generation by combination heating of ICRF and NBI in Heliotron J. 26th IAEA Fusion Energy Conference Proceedings.
- ONGENA, J., MESSIAEN, A., KAZAKOV, Y.O., SCHWEER, B., STEPANOV, I., VERVIER, M., CROMBÉ, K., SCHOOR, M.V., BORSUK, V., CASTAÑO-BARDAWIL, D., *et al.* 2020 The ICRH system for the stellarator Wendelstein 7-X. *AIP Conf. Proc.* **2254** (1), 070003.
- PATTEN, H.W. 2019 Development and optimisation of advanced auxiliary ion heating schemes for 3d fusion plasma devices. PhD thesis, EPFL.
- PATTEN, H.W., GRAVES, J.P., COOPER, W.A., ERIKSSON, J. & PFEFFERLÉ, D. 2020 Identification of an optimized heating and fast ion generation scheme for the Wendelstein 7-X stellarator. *Phys. Rev. Lett.* **124** (15), 155001.

- PFEFFERLÉ, D., COOPER, W.A., GRAVES, J.P. & 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.
- POPOVICH, P., COOPER, W.A. & VILLARD, L. 2006 A full-wave solver of the Maxwell's equations in 3D cold plasmas. *Comput. Phys. Commun.* **175** (4), 250–263.
- RUST, N., HEINEMANN, B., MENDELEVITCH, B., PEACOCK, A. & SMIRNOW, M. 2011 W7-X neutral-beam-injection: Selection of the NBI source positions for experiment start-up. *Fusion Engng Des.* **86** (6–8), 728–731.
- SCHUSTER, C.U., JOHNSON, T., PAPP, G., BILATO, R., SIPILÄ, S., VARJE, J. & HASENÖHRL, M. 2021 Moment-preserving and mesh-adaptive reweighting method for rare-event sampling in monte-carlo algorithms. *Comput. Phys. Commun.* **267**, 108041.
- SLABY, C., MACHIELSEN, M., LAZERSON, S. & GRAVES, J.P. 2022 Simulation of radio-frequency heating and fast-ion generation in Wendelstein 7-X. *J. Phys.: Conf. Ser.* **2397** (1), 012006.
- WESSON, J. 2004 *Tokamaks*, 3rd edn. Clarendon Press.