

Transient versus steady-state solutions: a qualitative study

D. Van Eester ^{1,†}, E.A. Lerche ^{1,2}, E. Pawelec ² and E. Solano ³

¹Laboratory for Plasma Physics, ERM-KMS, B-1000 Brussels, Belgium

²JET, Culham Science Centre, Abingdon OX14 3DB, UK

³Laboratorio Nacional de Fusión, CIEMAT, 28040 Madrid, Spain

(Received 23 October 2023; revised 26 January 2024; accepted 26 January 2024)

In view of the ultimate goal of producing long-lasting quasi-stationary discharges required for future fusion power stations, the numerical study of steady-state solutions of equations describing the particle and energy balance rightfully gets ample attention. Transient states may, however, differ significantly from the steady state ultimately reached and will – in practice – impact on the actual fate of the discharge. Using brutally simple models, the present paper highlights a number of aspects to illustrate this dynamics. It e.g. shows the different signature of wave and beam heating, potentially giving room to transiently trigger desirable effects that may allow us to better steer a discharge.

Keywords: plasma heating, fusion plasma, plasma simulation

1. Introduction

Future power stations relying on fusion will have superconducting coils allowing us to have long discharges; see e.g. the literature on the DEMO reactor (Federici, Maviglia & Holden 2022). Hence, a lot of modelling attention is devoted to studying steady-state solutions. However, most present-day tokamaks still have copper coils and hence have a limited flat top duration; see e.g. Mailloux *et al.* (2022). Moreover, even if it can be demonstrated through sufficiently sophisticated modelling that a steady state can potentially be maintained, it is first needed to reach that state, passing through intermediate stages where various phenomena occur that may prevent reaching of the desired equilibrium. Abundant experimental evidence shows that high-performance scenarios require fine tuning of the timing of the particle (gas, pellets) as well as energy injection (auxiliary heating systems) to properly enter H-mode and to avoid deleterious effects caused by too high impurity concentrations or MHD (magnetohydrodynamics) activity; see e.g. Garzotti *et al.* (2019). Also ‘landing’ high-performance H-mode discharges requires special care: switching off the heating causes the temperature to dwindle while the density stays high, potentially causing a runaway process of increased radiation resulting in a disruption.

If the high-performance phase lasts longer than a few times the slowest of the characteristic time scales, then a steady state can potentially be reached. If that is not

† Email address for correspondence: d.van.eester@fz-juelich.de

the case, the discharge dynamics is at least partly governed by transient effects. Various types of models are adopted to study the dynamics of discharges. The most realistic ones properly account for the time-dependent aspects of all the equations. Properly doing so requires intensive computational effort, however. A procedure often exploited is to solve the steady-state version of the reigning equations while partly accounting for the time dependence by introducing the actual source strength of the various sources at each snapshot in time. More crudely even, the sources are often simply assumed to be constant when the focus of the study is not on them. The goal of the present paper is to draw attention to the potential role and importance of intrinsic or externally forced/triggered transient effects and basic features such as the location of the power deposition (e.g. through the fact that diffusion is bi-directional so that it spreads heat towards the core as well as towards the edge). Although much more sophisticated transport as well as wave-particle interaction dynamics models are available and could in principle be used instead – some examples are Jaeger *et al.* (2001); Hellsten *et al.* (2004); Brambilla & Bilato (2009); Coster *et al.* (2010); Jucker *et al.* (2011); Bourdelle (2015); Bertelli *et al.* (2016); Breslau *et al.* (2018); Romanelli *et al.* (2020); Angioni (2021) and Hoelzl *et al.* (2021)) – the 2 simple models exploited in the present paper have purposely been brutally stripped of everything that is not or less of interest for the illustration of aspects of the transient dynamics: a single population Fokker–Planck equation illustrating kinetic aspects of the dynamics of energetic subpopulations and a single fluid transport equation highlighting macroscopic aspects. These two equations suffice to crudely grasp how the plasma kinetic profiles change under the influence of particle and heat sources, the first giving insight into how distributions are differently affected when heated through different means and how – under the crude approximation that particles are glued to magnetic surfaces – power is Coulomb collisionally redistributed to the non-directly heated populations on each magnetic surface, and the latter showing how the powers ending up in the thermal subpopulations are redistributed across magnetic surfaces to set up the macroscopically measured temperature profiles. Adopting more accurate models would undoubtedly enhance the realism of the description, but would deviate the focus of the paper. Although the individual steady-state solutions of the equations solved here do not depend on the path followed to reach that steady state because specific inputs are assumed to be prescribed so that the solutions are unique, the evolution of an actual experimental discharge often critically depends on the details of the path chosen, yielding a different nonlinear response (exploiting turbulence comes to mind) and yielding to a higher-performance regime, e.g. on the different timings of the switch-on or -off of the heating systems or the stimulated response when modulating the power or particle source. A better understanding of the transient physics can potentially help to better design discharges and steer them to a desired corner in configuration space in the future.

2. Simple Fokker–Planck model and transient effects in the setting up of a particle population

A very simplified version of the Fokker–Planck equation to obtain the distribution function F_o of a population of interest is solved as a function of the parallel ($v_{//}$) and perpendicular (v_{\perp}) velocity. We use $v = |\vec{v}| = |v_{\perp}\vec{e}_{\perp} + v_{//}\vec{e}_{//}| = |v_{\perp}^2 + v_{//}^2|^{1/2}$ in which \vec{e}_{\perp} and $\vec{e}_{//}$ are unit vectors with directions defined w.r.t. the direction of the confining magnetic field \vec{B}_o . It is assumed that heating close to the magnetic axis is studied, allowing to neglect the variation of the magnetic field so that trapping is of little importance (strictly, the variation of the magnetic field strength is very modest but non-zero close to the centre so a tiny fraction of the particles is actually trapped) and the homogeneous

plasma Fokker–Planck equation as provided in Stix (1992) can be adopted; for explicit expressions see e.g. Van Eester (1994). One gets

$$\frac{\partial F_o}{\partial t} = \frac{\partial F_o}{\partial t} \Big|_{\text{coll}} + \frac{\partial F_o}{\partial t} \Big|_{\text{RF}} + S - L, \tag{2.1}$$

where L is a loss term and S is a particle source term (a localised two-dimensional (2-D) Gaussian with a finite width at a prescribed beam velocity $(v_{\perp,o}, v_{\parallel,o})$ when describing a beam or a 1-D Gaussian at a prescribed birth velocity v_{birth} when describing fusion-born particles). If there is a finite source then a non-zero loss needs to be included in the description to ensure a steady state can be reached; $L = F_o/\tau$ is taken, where τ is a characteristic loss time. In absence of a particle source, the particle sink is omitted. The Coulomb collision operator can be written as

$$\begin{aligned} \frac{\partial F_o}{\partial t} \Big|_{\text{coll}} = & \left[3\Psi + v \frac{\partial \Psi}{\partial v} \right] F_o + v_{\perp} \left[-2\Phi + \frac{1}{v} \frac{\partial \Theta}{\partial v} + \Psi + \frac{1}{v_{\perp}^2} [\Theta + v^2 \Phi] \right] \frac{\partial F_o}{\partial v_{\perp}} \\ & + v_{\parallel} \left[-2\Phi + \frac{1}{v} \frac{\partial \Theta}{\partial v} + \Psi \right] \frac{\partial F_o}{\partial v_{\parallel}} \\ & + [\Theta + v_{\parallel}^2 \Phi] \frac{\partial^2 F_o}{\partial v_{\perp}^2} + [\Theta + v_{\perp}^2 \Phi] \frac{\partial^2 F_o}{\partial v_{\parallel}^2} - 2v_{\perp} v_{\parallel} \Phi \frac{\partial^2 F_o}{\partial v_{\perp} \partial v_{\parallel}}, \end{aligned} \tag{2.2}$$

where

$$\left. \begin{aligned} \Theta &= \sum_b \frac{\Gamma^{a/b}}{v} G(v/[2^{1/2} v_{\text{th},b}]) \\ \Phi &= \sum_b \frac{\Gamma^{a/b}}{2v^3} [\epsilon(v/[2^{1/2} v_{\text{th},b}]) - 3G(v/[2^{1/2} v_{\text{th},b}])] \\ \Psi &= \sum_b \frac{\Gamma^{a/b}}{v v_{\text{th},b}^2} \frac{m_a}{m_b} G(v/[2^{1/2} v_{\text{th},b}]) \\ \Gamma^{a/b} &= \sum_b N_b \frac{Z_a^2 Z_b^2 e^4 \ln \Lambda^{a/b}}{A_a^2 4\pi \epsilon_o^2 m_p^2} \\ G(x) &= [\epsilon(x) - x\epsilon'(x)]/[2x^2]. \end{aligned} \right\} \tag{2.3}$$

In the above ‘a’ refers to the examined ion population and the sum ‘b’ is on the different Maxwellian background electron and ion species, $\epsilon(x)$ is the error function and the quote refers to the derivative with respect to the error function’s argument; $v_{\text{th},b}$ is the thermal velocity; $\ln \Lambda^{a/b}$ is the Coulomb logarithm and N_b is the particle density of population ‘b’. The symbols m_{α} , Z_{α} and A_{α} refer to the mass, atomic number and ion mass of species α while m_p is the proton mass; ϵ_o is the vacuum permittivity. Finally, the RF diffusion operator appearing in the Fokker–Planck equation is

$$\frac{\partial F_o}{\partial t} \Big|_{\text{RF}} = \nabla \cdot \mathbf{D} \cdot \nabla F_o, \tag{2.4}$$

where

$$\mathbf{D} = \sum_N \frac{\pi}{8|k_{\parallel}|} \frac{Z_a^2 e^2}{A_a^2 m_p^2} \delta(v_{\parallel} - \frac{\omega - N\Omega}{k_{\parallel}}) \mathbf{a}_N \mathbf{a}_N^*, \tag{2.5}$$

with

$$\left. \begin{aligned} \mathbf{a}_N &= \Theta_N \frac{k_{//}}{\omega} \left[\left(\frac{\omega}{k_{//}} - v_{//} \right) \mathbf{e}_{\perp} + v_{\perp} \mathbf{e}_{//} \right] \\ \Theta_N &= E_+ J_{N-1}(k_{\perp} v_{\perp} / \Omega) + E_- J_{N+1}(k_{\perp} v_{\perp} / \Omega). \end{aligned} \right\} \quad (2.6)$$

In the above, and similar to what was done for the particle source, the delta function is replaced by a smooth, localised Gaussian or – following Stix – the delta function is removed altogether by formally integrating over the poloidal angle. The perpendicular wavenumber k_{\perp} used in the RF diffusion operator is associated with the fast wave; in practice, a proper guess is provided by running a hot plasma wave equation solver such as TOMCAT (Van Eester & Koch 1998) and picking the value of $k_{\perp, \text{fast}}$ at the position where the absorption is maximal. Likewise, an estimate for the polarisation $E_+/E_- = [E_{\perp,1} + iE_{\perp,2}]/[E_{\perp,1} - iE_{\perp,2}]$ is obtained from the same code. In an iterative process, the magnitude of the electric field is adjusted to ensure the power density is as externally prescribed. In the above, e is the (magnitude of the) electron charge, $\omega = 2\pi f$ where f is the antenna frequency, Ω is the cyclotron frequency and $k_{//}$ is the parallel wave vector component of the dominant mode in the antenna spectrum.

This simplified ‘locally homogeneous’ model is clearly inadequate to properly describe the interaction between waves and particles well away from the core, especially in compact machines that do not have a large aspect ratio. However, for the purpose of illustrating the importance of transient effects, Stix’s model is sufficiently rich. A Crank–Nicolson time stepping system (see the [Appendix](#)) is adopted to solve the time-dependent Fokker–Planck equation. At a set of velocity grid points the finite difference method is exploited to write the Fokker–Planck equation as a linear system of the form $\partial \mathbf{F} / \partial t = \mathbf{A} \cdot \mathbf{F} + \mathbf{B}$, where the vector \mathbf{F} contains the values of the distribution function F_o at all the grid points and \mathbf{B} is the ‘source’ term, independent of F_o .

The plasma is assumed to consist of a balanced D and a T bulk population, supplemented with a D beam and a H minority. The density adopted is $8 \times 10^{19} \text{ m}^{-3}$ and the background temperatures are 7 keV for the electrons and 8 keV for the ions. The corresponding slowing down time is 750 ms. The wave–particle interaction for 3 example populations will be treated: a 6% D beam population with launch velocity $v_{\perp,o} = 2.41 \times 10^6 \text{ m s}^{-1}$ and $v_{//,o} = 2.76 \times 10^6 \text{ m s}^{-1}$ (i.e. $E_o = 110 \text{ keV}$ and angle $\arctan(v_{//,o}/v_{\perp,o})$ of 48°), a D bulk population and a 3% H minority population. The radio frequency (RF) frequency is 51 MHz and a single core $k_{//}$ of 9 m^{-1} was chosen. Core RF heating requires a magnetic field strength of $B_o = 3.35T$, the D population being heated at $N = 2$ while H is heated at $N = 1$, where N is the cyclotron harmonic. Corresponding to the adopted parameters, $k_{\perp} = 75 \text{ m}^{-1}$. The electric field strength is adjusted to ensure a prescribed steady-state power density of 0.6 MW m^{-3} is reached for the H minority and D majority while the beam has a somewhat more modest absorption of 0.2 MW m^{-3} ; from an assessment of the wave fields for the adopted parameters, the relative polarisation of E_+ (mostly responsible for the heating thermal ions) and E_- (responsible for heating fast ions) was chosen to be 1/4. [Figure 1](#) depicts the normalised RF diffusion coefficient for the chosen parameters and for fundamental cyclotron H heating as well as for second harmonic heating of a D population. The normalisation factor is determined by the assumed RF power density; the electric field amplitude needs to be adjusted to match an imposed RF power density. The top plot depicts the dependence when expressed in term of the perpendicular velocity; the bottom plot has the same data but for the corresponding energy.

The first example treated is the D majority heated at its second harmonic. The local energy density $v_{\perp} E F_o$ in $(v_{\perp}, v_{//})$ space is depicted in \log_{10} scale and normalised to the

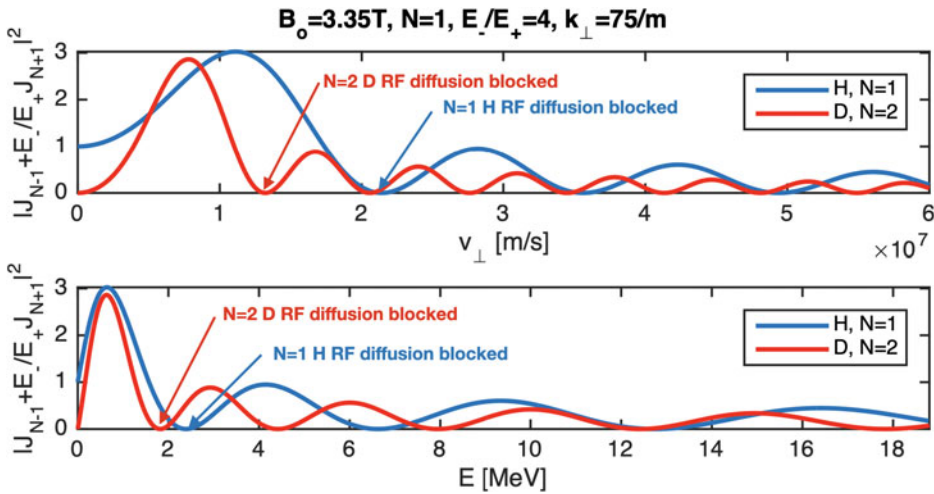


FIGURE 1. The RF diffusion coefficient for $N = 1 H$ and $N = 2 D$. Note the arrows indicating the first zeros of the RF diffusion coefficients, corresponding to the velocities beyond which the particles cannot diffuse.

maximal value in figure 2; the data were clipped at a fraction of 10^{-5} to exclude entirely negligible contributions so the log plot contour levels vary from -5 (dark colour in the contour plot) to 0 (bright colour in the plot). The first observation that can be made is that it takes some time to develop a strong tail. After 60 ms and when resonant interaction occurs for all $v_{//}$, where the distribution is significant, the tail stretches only up to approximately $v_{\perp} = 8 \times 10^6 \text{ m s}^{-1}$. After 100 ms the tail has still not reached its maximal extent. For the adopted power density, the tail formed by second harmonic heating stretches up to the first zero of the RF diffusion coefficient ($v_{\perp} \approx 1.3 \times 10^7 \text{ m s}^{-1}$ corresponding to approximately 1.8 MeV) where further diffusion is blocked; $D_{\text{RF}} \propto |E_+ J_{N-1} + E_- J_{N+1}|^2$, where N is the cyclotron harmonic and the argument of the Bessel functions is $k_{\perp} v_{\perp} / \Omega$. The upper limit of the integration domain was purposely chosen well beyond the zero of the diffusion coefficient to underline that the v_{\perp} up to which the population extends is not artificially imposed by the integration limits but results from vanishing of the diffusion. Increasing the power density does not increase the maximum velocity reached by the tail particles; it merely causes the tail to be more populated so that a gradually more pronounced secondary maximum forms at high energy. When the maximal velocity has been reached the evolution is not finalised yet: Coulomb collisional interaction causes the tail to broaden by pitch angle scattering while the tail still gains extra energy on a slower time scale and a shallow maximum forms in the high density domain. Aside from the fact that second harmonic heating is most efficiently heating particles that already have a high energy, the high energy of the tail is also due to the fact that the tail is constituted of very few particles. Binning the particles in bins of various velocity ranges (not shown here but illustrated later) reveals that – although the distribution seems to have reached a steady state after 2 seconds – there is still a slow evolution. It takes a few more seconds before the distribution is truly stationary. Reaching an actual equilibrium takes most time for the interaction of the weak tail with the electron population; on account of their small mass, electrons and ions occupy very different regions of velocity space. Since the tail is slowing down, providing energy and particles to the lower energy region, also the rest of the distribution is mildly affected as long as the high energy tail has not fully stabilised.




FIGURE 2. Temporal evolution of the D majority energy density locally in velocity space. Six different times were picked with respect to the switching of the RF heating: 60, 100 ms, 0.5, 1, 2 and 3 s. In each of the panels, the perpendicular velocity (horizontal axis) varies from 0 to $16 \times 10^6 \text{ m s}^{-1}$ while the parallel velocity varies from -8 to $+8 \times 10^6 \text{ m s}^{-1}$.

The hydrogen minority is heated at the fundamental cyclotron frequency. The evolving energy density is provided in [figure 3](#). By comparison, the D bulk thermal population is much more weakly affected; it is modified indirectly by some of the fast tail population slowing down and increasing the bulk temperature but hardly by direct RF heating. Because of the difference in mass and hence in Larmor radius, the first zero of the RF diffusion coefficient is now at $v_{\perp} \approx 2.2 \times 10^7 \text{ m s}^{-1}$. Compared with the D second harmonic heating, the equilibrium is more quickly reached for the H minority. While the D tail is still not fully in equilibrium after several seconds (the density and energy in the thermal subpopulation still getting input from the slowing down tail), the H distribution hardly budes after 1 s. [Figure 4](#) provides some details on how the steady state is reached. Panel (a) shows the power balance, illustrating that a very high wave absorption (approximately 2/3 of the launched power) is reached as soon as the RF power is switched on (second time point) while the Coulomb collisional relaxation takes somewhat more time to set in. In panel (b) it can be seen that the equilibrium of the H minority is reached much quicker ($\approx 200 \text{ ms}$) with the ion background than with the electrons ($\approx 1 \text{ s}$). Moreover, the interaction with the electrons builds up very gradually, a direct consequence of the fact that it takes time to build up a high energy tail. In panel (b) it is also clear that the minority tail ultimately passes much more energy on to the electrons than to the ions, a result well known from steady-state analysis.

The 110 keV D beam distribution function heated at its second cyclotron harmonic and depicted in [figure 5](#) reveals a behaviour somewhat in between that of the first 2 shown examples. To enable reaching a steady state when beam particles are constantly injected, a loss term is required. A (constant) characteristic loss time of $\tau = 300 \text{ ms}$ – typical for the adopted densities and bulk temperatures in the actual JET Baseline experiments on which the adopted parameters are loosely inspired (Garzotti *et al.* 2019; Maggi *et al.* 2023) – was chosen both for the energy and the particle loss. Because a large fraction of the particles are already preheated and hence have velocities well beyond that of the thermal subpopulation, most of the particles in the beam distribution equally take less time to reach their equilibrium state: After 1 s it essentially does not change anymore. A key difference which has its repercussions on MHD activity but which is a well-known characteristic of a beam population is that the local source gives rise to a ‘bump on tail’ region, allowing us to excite rather than damp waves locally in a subregion of velocity space. Note that the distribution is again blocked at the first zero of the combination of Bessel functions underlying the RF diffusion coefficient.

[Figure 6\(a\)](#) provides the D beam power balance as a function of time. At the outset the incoming power exceeds the Coulomb collisional redistribution to the background. While the tail gradually forms, the Coulomb collisional sink as well as the energy loss term become more significant. Unlike what was observed for fundamental heating, it takes approximately a second before the RF absorbed power is significant and able to match the incoming power. Panel (b) shows how the collisional power is subdivided: most the power flows to the electrons but much less than for the case of the H minority. Because the tritons are heavier than the deuterons, their thermal velocity is smaller so that the slowing down fast tail particles start interacting with the D bulk before they start slowing down on the T bulk. Hence more of the tail energy flows to the D thermal population than to the T thermal population (recall that the same density has been adopted for the D as for the T bulk). This is true for any high energy tail in a mixture of balanced D and T majorities. Similar to the case of the H minority, the fraction collisionally passed on to the electrons gradually increases while the tail is gradually building up. Because the slowing down from the source at $v_{\perp,o} = 2.41 \times 10^6 \text{ m s}^{-1}$ and $v_{\parallel,o} = 2.76 \times 10^6 \text{ m s}^{-1}$ (110 keV) to the thermal energy of the bulk ions takes time, it now requires time for the tail




FIGURE 3. Temporal evolution of the H minority energy density locally in velocity space. Six different times were picked with respect to the switching of the RF heating: 60, 100 ms, 0.5, 1, 2 and 3 s. In each of the panels, the perpendicular velocity (horizontal axis) varies from 0 to $27 \times 10^6 \text{ m s}^{-1}$ while the parallel velocity varies from -10 to $+10 \times 10^6 \text{ m s}^{-1}$. The data were clipped at a fraction of 10^{-5} so the log plot contour levels vary from -5 (dark colour in the contour plot) to 0 (bright colour in the plot).

to reach its equilibrium with the background ions as well. Just as for the H minority case, reaching the equilibrium with the electrons takes longer. Various processes are mixed here: the natural slowing down and pitch angle scattering of the beam particles from their injection location to form a distribution that is the sum of a thermal subpopulation

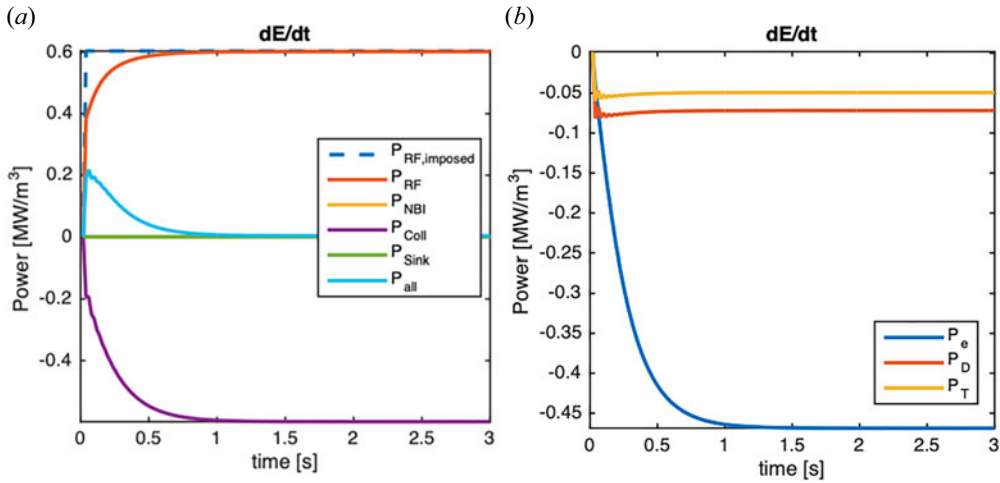


FIGURE 4. Temporal evolution of the H minority power balance (a) and Coulomb collisional redistributed power loss to the background species (b).

and a ‘slowing down’ distribution (Gaffey (1976) analytically solved the equation for the latter, adopting a collisional operator that is lossy, artificially but purposely preventing the thermal subpopulation forming – thereby removing it from the solution – and enabling him to omit an explicit loss term), the RF acceleration pulling particles to higher energies and the competition between the RF and collisional dynamics.

Figure 7 – depicting the partial density $2\pi \int_{v_{bot}}^{v_{top}} dv_{//} dv_{\perp} v_{\perp} F_o$ and the partial energy density $2\pi \int_{v_{bot}}^{v_{top}} dv_{//} dv_{\perp} v_{\perp} mv^2/2F_o$ in a set of velocity bins $v_{bot} \leq v \leq v_{top}$ – illustrates that the distribution of a D beam heated at its second harmonic cyclotron layer is by far most populated in the lowest energy regions (thermal + beam source) while the energy density has significant contributions from the RF-induced tail. The density and energy density in the region where the source lies reaches its steady state most quickly while regions only populated after collisions have been allowed to populate them reach their steady state later. Among the more significantly populated density and energy density regions, the thermal region reaches its steady-state value the slowest: it takes time before the competition between the RF net acceleration and the Coulomb collisional slowing down are in equilibrium, the bulk temperature having been increased above the temperature in the absence of RF power because of the extra energy fed into this region. The thermal population is also affected by second harmonic heating, be it indirectly: by Coulomb collisions. For that reason, it takes time (in this case approximately 2 seconds) before this global equilibrium has been reached.

The present computation is done imposing an electric field strength corresponding to the steady-state power balance. It was shown that it takes time to set up a high energy tail that is in equilibrium with the bulk ions and electrons. Because second harmonic heating primarily affects high energy particles, most of the power is collisionally passed on to the electrons and the reason for the slower convergence is because the tiny amount of particles in the tail is sensitive to small changes in local power density. In reality – the machine vessel being metallic so acting as a Faraday cage – the incoming RF power has to be balanced at all times, so also immediately after powering the RF antennas. For a given distribution at a given time, this requires the electric field strength to be amplified to ensure this is the case. The electric field strength responds to the total power, not to the

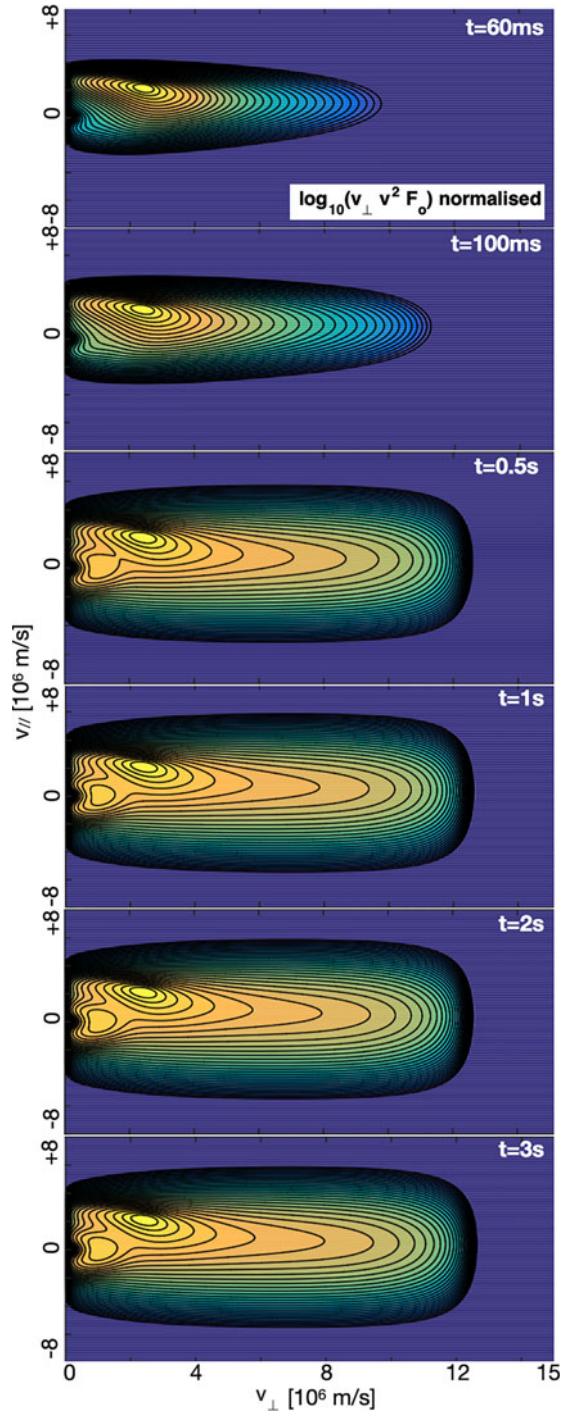


FIGURE 5. Temporal evolution of the D beam energy density locally in velocity space. Six different times were picked with respect to the switching of the RF heating: 60, 100 ms, 0.5, 1, 2 and 3 s. In each of the panels, the perpendicular velocity (horizontal axis) varies from 0 to $15 \times 10^6 \text{ m s}^{-1}$ while the parallel velocity varies from -8 to $+8 \times 10^6 \text{ m s}^{-1}$. The same colour coding as before is used. The source location is easily recognised as the location where the power density is maximal: the bright yellow spot at $(v_{\perp}, v_{\parallel}) \approx (2.41, 2.76) 10^6 \text{ m s}^{-1}$.

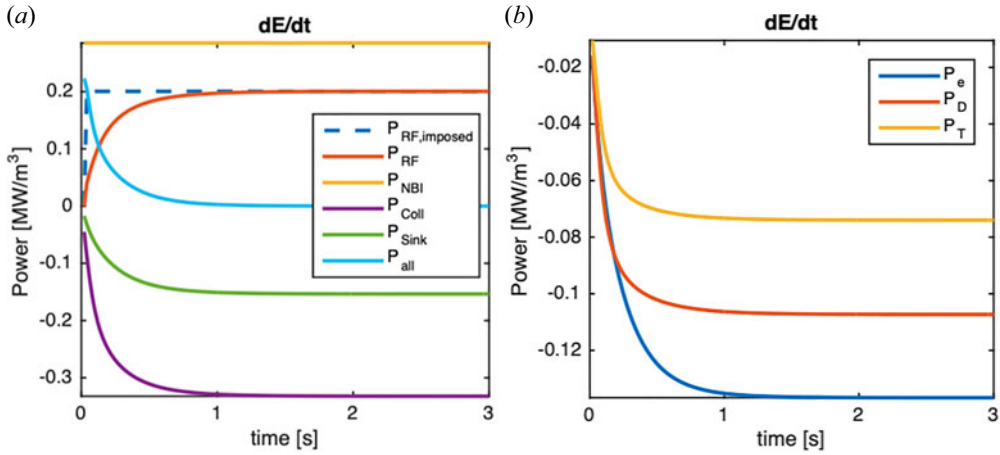


FIGURE 6. Temporal evolution of the $N = 2$ RF heated D beam power balance (a) and Coulomb collisional redistributed power loss to the background species (b).

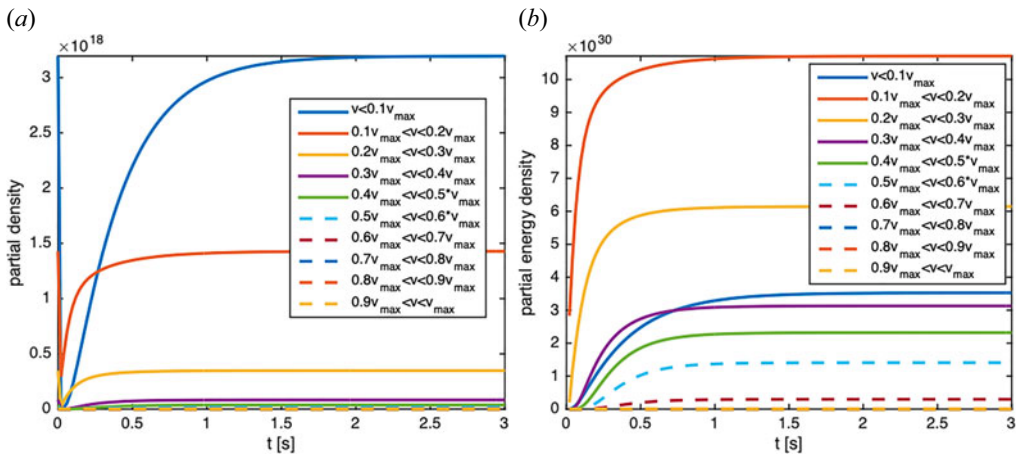


FIGURE 7. Temporal evolution of the partial density $2\pi \int_{v_{bot}}^{v_{top}} dv_{//} dv_{\perp} v_{\perp} F_o$ (a) and partial energy $2\pi \int_{v_{bot}}^{v_{top}} dv_{//} dv_{\perp} v_{\perp} mv^2/2F_o$ (b) of the RF heated D beam up to a steady state being reached. The lower ('bot') and upper ('top') velocity limits of the various velocity bins over which integrations are performed are depicted in the legend of the panels; v_{max} is the largest $|v|$ of the adopted grid.

partial power density passed on to a species that is damped. As amplification of the field potentially renders parasitic effects such as collisional damping or sheath rectification in the edge more efficient, such amplification should be avoided. Since harmonic heating is inefficient at the outset, a fundamental heating scheme of a minority is most suited to avoiding such amplification. As was shown in figure 4(a), the RF power density of a minority heated at its fundamental cyclotron harmonic matches the RF power density corresponding to the steady state already soon after the RF power is switched on. Adding a small minority at a proper concentration to avoid electric field amplification at the start is indeed experimentally known to be beneficial in the transient state preceding the actual steady state. Needless to say that, when solving the steady-state equation to grasp the wave-particle dynamics, this effect of reduced absorption efficiency necessitating a

growth of the electric field amplitude and the risk it poses for allowing parasitic edge damping goes unnoticed. The reason for the very efficient damping is that $N = 1$ heating works well even at very low temperature since it mainly interacts with the thermal population.

The effects of the synergy between wave and beam heating differ depending on the wave heating scheme chosen. Figure 8 depicts the evolution of a H beam heated at its fundamental cyclotron frequency. Whereas the energy density peak around the birth velocity remains localised at that location when the RF heated beam tail forms for harmonic heating (as can be seen in figure 5), the spread observed when adopting a fundamental heating scheme is more pronounced. A similar spread occurs in the thermal region, the waves stretching the whole population over a larger perpendicular velocity region. While high energy (MeV) tails are formed, a large number of the particles remain in the energy range of 10 s of keV, filling the gap between the beam source and the actual thermal subpopulation. Figure 9 illustrates the temporal evolution towards the steady state. Unlike what was the case for minority heating, it now takes time before the RF absorbed power transferred to the beam matches the imposed launched power. This is because the beam and wave heatings were started simultaneously so that there is a very modest thermal particle density to transfer the power to early after the switch-on, where $N = 1$ is most efficient. When the RF heating is switched on later than the beam – allowing the beam to first populate the thermal region through slowing down of beam particles – then the RF heating becomes more efficient much more quickly.

All previous computations were done assuming that the electric field can potentially interact with all particles of the distribution. The underlying idea – that is very commonly adopted – is that although the wave–particle interaction is a resonant interaction selecting a specific $v_{//}$ from the condition that $N\Omega + k_{//}v_{//} = \omega$, the fact that Ω is location dependent allows us to satisfy the resonance condition somewhere along the orbit of any particle in the population intended to be heated ($\omega \approx N\Omega$). This is not necessarily true, however. It can readily be shown that the biggest possible distance away from the magnetic axis for the cyclotron resonance to cut the orbit for a given $v_{//}$ is $R_o|(1 - v_{//}k_{//}/\omega)^{-1} - 1|$. It is assumed poloidal field and drift effects are absent so that B_o simply varies as $1/R$ while particles are glued to magnetic surfaces. The frequency was chosen to place the cold plasma cyclotron resonance at the magnetic axis. With these assumptions, any magnetic surface intersecting the midplane at more than 16 cm from the axis will always ensure the resonance is met somewhere on the orbit for all (passing) particles of an RF heated D Maxwellian when recalling that the parallel diffusion caused by the RF waves is modest and simply assuming that most of the distribution falls inside $|v_{//}| < 3v_{th}$. Recall that the adopted simplified model ignores trapped particle effects; more strictly, a fraction of sufficiently deeply trapped particles will be out of reach for RF heating. Specifically when aiming at heating very close to the magnetic axis – as was done here to justify omitting trapping effects – the Ω variation is modest and the condition is *not* met except for a very small subset of the population. Doing so the transient behaviour of the distribution being set up and the differing roles of the RF heating and the Coulomb collisional relaxation become even more evident. This is illustrated in figure 10 for the D beam heated at its second harmonic layer. To avoid spurious numerical problems, the delta function $\delta(v_{//} - [\omega - N\Omega]/k_{//})$ – selecting the parallel velocity at which the resonant interaction takes place – has been smoothed, replacing it by a Gaussian with half-width $4 \times 10^4 \text{ m s}^{-1}$ but identical integral across the resonance turned into a quasi-resonance. A value of $\xi = N\Omega/\omega = 1.02$ was chosen for the present example but the electric field strength was adjusted to have the same RF power density as in previous examples ($P_{\text{RF,beam}} = 0.2 \text{ MW m}^{-3}$). As the number of particles able to profit from RF heating is




FIGURE 8. Temporal evolution of the H beam energy density locally in velocity space. Six different times were picked with respect to the switching of the RF heating: 60, 100 ms, 0.5, 1, 2 and 3 s. In each of the panels, the perpendicular velocity (horizontal axis) varies from 0 to $27 \times 10^6 \text{ m s}^{-1}$ while the parallel velocity varies from -10 to $+10 \times 10^6 \text{ m s}^{-1}$. The same colour coding as before is used.

small, the RF tail is formed very quickly when heating at or very close to the magnetic axis. In energy it has already reached its full extent after just 60 ms, be it that the tail is not yet very populated at that time. Further populating the higher energy regions by RF waves happens more slowly, and the collisional relaxation to the final state happens even slower. The overall result very much depends on the type of heating and the type of population: since fundamental cyclotron heating affects the whole distribution but dominantly the thermal subpopulation, the RF tail of a minority heated at $N = 1$ forms slower than a tail at $N = 2$, the competition between RF heating and collisional relaxation being more pronounced. The relative timing of the switch-on of the RF and the beam reveals different

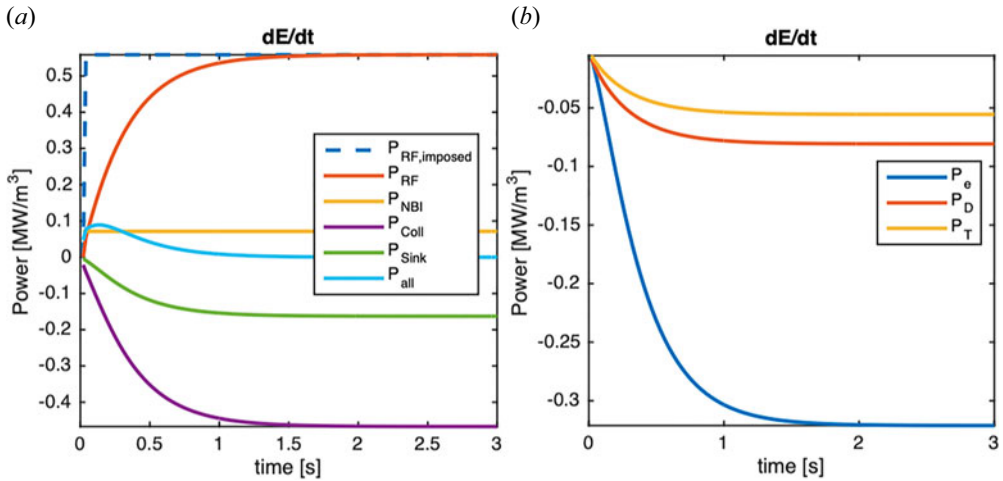


FIGURE 9. Temporal evolution of the $N = 1$ RF heated H beam power balance (a) and Coulomb collisional redistributed power loss to the background species (b).

transient effects. For the here adopted simple model, the same steady state is ultimately always reached but in the actual experiment clever timing allows us to minimise unwanted effects that can negatively influence the fate of the discharge and is exploited to nonlinearly reach higher performance.

Especially for orbits close to the plasma centre for which the magnetic field strength varies little, it is not always possible to satisfy the resonant condition underlying wave–particle interaction. Figure 11 shows the power density corresponding to the steady-state solution for an imposed electric field strength of $E_+ = 0.6$ and $E_- = 2.4 \text{ kV m}^{-1}$ but varying the resonant parallel velocity via the value of the local cyclotron frequency while adopting the smoothed delta function treatment. Because the source velocity of the beam is in the co-direction, the RF power density peaks at a value below $N\Omega/\omega = 1$ i.e. on the low field side. The heating is non-zero in a narrow interval only: in the equatorial plane, the dominant beam heating occurs at $\xi = 0.94$ so $x = (1/\xi - 1)R_o = 9 \text{ cm}$ on the low field side and the RF heating of the D beam is finite in the interval $-17 < x < +33 \text{ cm}$ provided the considered magnetic surface stretches up to these values. At the magnetic axis and for the given magnetic field and electric field strength, the RF power density is only 0.39 MW m^{-3} . While the mean parallel energy of the beam $\langle mv_{\parallel}^2/2 \rangle / \langle 1 \rangle$ where $\langle \dots \rangle = \int d\mathbf{v} \dots F_o$ is hardly affected (it merely increases from 11 keV at $\xi = 0.9$ and 1.1 to 17 keV when $\xi = 0.94$; recall that the adopted electron temperature is 7 keV and the ion temperature 8 keV so that the mean energy of the bulk population is $E_{th} = 3/2 \langle kT \rangle$ which is slightly smaller than 12 keV), its mean perpendicular energy $\langle mv_{\perp}^2/2 \rangle / \langle 1 \rangle$ reaches values up to 105 keV. This latter value is neither a good measure for the bulk nor for the tail temperature, the consequence of the fact that the tail is weakly populated but carries a non-negligible fraction of the energy of a population heated at its second cyclotron harmonic. Note the secondary maximum in both the power density and the perpendicular energy, highlighting the (modest) role the thermal population plays when the cyclotron resonance passes through the region in v_{\parallel} that is most densely populated.

Modulation of the RF wave power level is a routinely exploited tool to determine the experimental power deposition profile; see e.g. Mantica *et al.* (2011, 2021): since net wave heating requires satisfying of the condition $\omega = N\Omega + k_{\parallel}v_{\parallel}$, cyclotron heating gives rise




FIGURE 10. Temporal evolution of the D beam energy density locally in velocity space adopting the fact that the resonance is localised around $v_{\parallel, \text{res}} = (\omega - N\Omega)/k_{\parallel}$, here adopting $N = 2$ and $N\Omega_D/\omega = 1.02$. The same 6 different times as before were picked.




FIGURE 11. Steady-state RF power density (a) and corresponding mean perpendicular energy (b) for a series of $N\Omega/\omega$ values scanning the resonant velocity over the beam distribution function.

to well localised power deposition and hence is an excellent means to provide a suitable heat source for transport studies linking predictions to experimental evidence. As the temperature responds to the change of RF power, Fourier analysing the electron and ion temperature response to the modulated power allows us to determine the experimental power deposition profile under the influence of direct heating and collisional relaxation and to confront predictions with experimental data. Modulation of wave RF power equally allows ‘sawtooth pacing’ – see e.g. Lennholm *et al.* (2021) and Lerche *et al.* (2017) – as it destroys the stabilising effect of high energy populations and results in forced sawtooth crashes, which mitigates the impact of giant sawteeth and flushes high-Z particles – responsible for deflated or even hollow temperature profiles through radiation – out of the core. Figure 12 depicts the power collisionally transferred to the various species in response to a 1 Hz modulation of the RF power source with an 80/20 duty cycle. Because the typical time scale on which electrons and ions react is different, the response to differing heating schemes has a different signature. Note, for example, that the power flowing to the ions for the case of a H minority heated at its fundamental cyclotron frequency quickly increases but then gradually decreases while the power flowing to the electrons simply increases steadily as long as the RF power source is on. Because the adopted modulation period is too short compared with the total needed to reach a steady state, the collisional power levels never reach their asymptotic level. For $N = 2D$ majority heating, the ion power levels approach saturation but the electron response shows no sign of it either. To get a quantitative assessment of the differing response of the temperature to a specific modulated heat source, a Fourier analysis of the temperature is often made. In the legends of figure 12 the phase of the $n = 1$ Fourier component of the electron and ion temperature response (oscillating at the modulation frequency i.e. proportional to $\exp[i\omega_{\text{mod}}t]$ where $\omega_{\text{mod}} = 2\pi f_{\text{mod}} = 2\pi$) to the imposed RF power modulation has been included.

3. Simple transport model

Strictly speaking, the whole dynamics of hot tokamak plasmas should be studied globally in terms of distribution functions interacting with one another in the whole vessel. Numerically, this is a daunting task. In practice, the influence of particle and heat sources




FIGURE 12. Impact of RF power modulation on the various plasma constituents for fundamental cyclotron minority heating (a) and second harmonic majority heating (b).

is typically modelled assuming the description of the whole process can be captured by focussing on a set of sub-processes which can be studied individually but should be solved as a coupled set. Because the time scale on which RF waves vary is much shorter than the other characteristic times on which the plasma changes, one typically solves the wave equation as a driven problem, the RF electric field oscillating at the rhythm imposed by the generator. Hence it makes sense to distinguish between a ‘fast’ time scale on which the waves vary, and a slower one on which the distribution is set up and in which the effect of the RF field shows up as the net effect averaged over an RF period. Naively assuming the particles do not deviate too much from a given magnetic surface, it is customary to describe the setting up of the particle distributions on isolated magnetic surfaces. On each

surface, the distribution functions evolve under the influence of Coulomb collisions. After Coulomb collisional redistribution, a final power density profile is set up. The fraction of the power that flows to the various thermal subpopulations constitutes the heat source term needed in the energy transport equation – subject of the present section – allowing us to determine how the (thermal) energy diffuses across magnetic surfaces, setting up the temperature profile.

The previous section dealt with describing the evolution of the distribution function of a heated ion population under the influence of wave or combined wave and beam heating in velocity space. The present section discusses the evolution of the radial profile of the temperature, naively assuming the temperature is constant on magnetic surfaces and that the plasma can be described as a single fluid. Such an elementary model was preferred to highlight that a lot of the dynamics of the plasma can qualitatively be understood from rather simple models. Of course, more sophisticated models are required to get a quantitatively correct description. The simplest possible yet fairly realistic density and energy evolution equation is a set of 2 diffusion-type equations with a constant diffusion D and convection V coefficient for both density N and energy $E = 3NkT/2$. Assuming the plasma cross-section is circular without a Shafranov shift, such a set reads

$$\left. \begin{aligned} \frac{\partial N(\rho, t)}{\partial t} &= \frac{1}{\rho} [\rho [D_N N' + V_N N]]' + S_N \\ \frac{\partial E(\rho, t)}{\partial t} &= \frac{1}{\rho} [\rho [D_E E' + V_E E]]' + S_P, \end{aligned} \right\} \quad (3.1)$$

i.e.

$$\left. \begin{aligned} \rho \frac{\partial N(\rho, t)}{\partial t} &= \rho D_N N'' + [D_N + \rho V_N] N' + [V_N + \rho] N + \rho S_N \\ \rho \frac{\partial E(\rho, t)}{\partial t} &= \rho D_E E'' + [D_E + \rho V_E] E' + [V_E + \rho] E + \rho P, \end{aligned} \right\} \quad (3.2)$$

in which the primes represent the radial derivatives ($' = \partial/\partial\rho$; ρ is the minor radius), S_N and $S_E = P$ are the particle and energy sources and losses; P is the power density. The energy equation is transformed into an equation for the temperature evolution, accounting for the actual density

$$\begin{aligned} \rho N \frac{\partial T}{\partial t} &= -\rho T \frac{\partial N}{\partial t} + [\rho D_E N] T'' + [D_E N + \rho(2N'D_E + V_E N)] T' \\ &+ [D_E N' + V_E N + \rho(D_E N'' + V_E N')] T + \frac{\rho P}{3k/2}. \end{aligned} \quad (3.3)$$

The local fluxes $D_N N' + V_N N$ and $D_T T' + V_T T$ are imposed to be zero at the magnetic axis $\rho = 0$ and at the edge the density and temperature are assumed to be reduced to negligible values (10^{17} m^{-3} and 20 eV will typically be assumed but the actual values are of little importance since the gradient terms dominate over them). Numerically accounting for the time evolution can be done using the Crank–Nicolson scheme (Crank & Nicolson 1947); see the [Appendix](#). Since the adopted model is too simple to describe the very early phase of the discharge (breakdown and initial ohmic heating phase) the initial profiles are imposed. They qualitatively represent the density and temperature prior to the main heating and gas injection being switched on.

Two types of particle sources are considered: the direct gas or pellet injection sources, and the particle source corresponding to the injected beam. The source for the energy

equation contains 3 sub-sources: the power injected by the neutral beam injection (NBI), that by the ion cyclotron resonance heating (ICRH) waves and a radiation loss term. For a steady state to be reached, these externally imposed sources representing the incoming particles and energy have to be balanced at the edge $\rho = a_p$ (where a_p is the minor radius of the last closed magnetic surface) by the edge particle and energy fluxes

$$\left. \begin{aligned} [\rho(D_N N' + V_N N)]|^{a_p} &= - \int_0^{a_p} d\rho \rho S_N \\ [\rho(D_E E' + V_E E)]|^{a_p} &= - \int_0^{a_p} d\rho \rho P. \end{aligned} \right\} \quad (3.4)$$

The flux is zero at $\rho = 0$. If a different total source is imposed and if the convection is assumed to be zero while the diffusion coefficient is known, the gradient has to increase or decrease accordingly. Solving a single transport equation for the whole plasma is a drastic simplification but can more easily be justified for the density than for the temperature. When one injects gas or beams, atoms rather than ions or electrons are injected. Hence, for a given S_N - and for hydrogen-like species ($Z = 1$) - we have the same N_e as N_i . No need to treat them apart since ions and electrons are ‘born’ by the same ionisation, so for the same D and V one will get the same N . There is a difference between particle and energy sources, however. The total energy is the sum of the energies of all species. If the power is evenly spread over ions and electrons (and knowing their density is the same) while assuming the corresponding D and V are the same, the total energy is the same so the individual electron and ion temperatures need to be half. If we send in a neutral beam with a given injection velocity then - after ionisation but when the particles still share their birth velocity - the electron energy is small since $m_e \ll m_i$ in the sum $m_e v_e^2/2 + m_i v_i^2/2$ so $E_{\text{beam}} \approx m_i v_i^2/2$ to first approximation. That suggests the factor 2 does not need to be there as only the ions are heated so T_i is prescribed by P while $T_e = 0$. But equipartition cranks up T_e . If the transport coefficients are the same and the power is evenly divided among electrons and ions, this process will go on until $T_e = T_i$, again yielding half the temperature. By definition, the single-species transport equation assumes all species can be represented by the same temperature so that an equipartition term is not needed. Apart from the evident conclusion that simpler models necessarily lack potentially important physics, this states that D and V differ depending on the context in which they are used. The above conservation law linking the total power and the flux at the edge states that, for a given total power, the transport coefficients of the single-species and two-species model need to differ by a factor 2 if the same temperature is to be achieved in either model when the same total power is considered.

To study the importance of gradients of the transport coefficients the above equations have mildly been generalised allowing D values and V values to be position dependent. The equations become

$$\rho \frac{\partial N(\rho, t)}{\partial t} = \rho D_N N'' + [D_N + \rho(D'_N + V_N)]N' + [V_N + \rho V'_N]N + \rho S_N, \quad (3.5)$$

$$\begin{aligned} \rho N \frac{\partial T}{\partial t} &= -\rho T \frac{\partial N}{\partial t} + [\rho D_T N]T'' + [D_T N + \rho(D'_T N + 2D_T N' + V_T N)]T' \\ &+ [D_T N' + V_T N + \rho(D'_T N' + D_T N'' + V'_T N + V_T N')]T + \frac{\rho P}{3k/2}. \end{aligned} \quad (3.6)$$

3.1. Preliminary notes on power density profiles

While the RF power density profile is imposed in the present work, the beam profiles are computed accounting for the beam trajectory and the reigning density. Provided the shine through is low enough for the density at a given time, the beam is fired and the total beam power and source energy are imposed. The beam intensity decreases as

$$I = I_o \exp \left[- \int ds N(s(\rho)) \sigma \right], \quad (3.7)$$

where s is the distance travelled along a straight beam trajectory and σ is the ionisation cross-section. The decrease of the beam intensity thus critically depends on the density of the plasma the neutral beam travels through. When solving the transport equations, the beam source will be computed, accounting for the actual density profile at each step in time. Exploiting an approximate fit, Kazakov proposed a simplified expression for the beam intensity adopting $\sigma \approx 1.8 \times 10^{18} A_{\text{NBI}}/E_o$, where E_o is the beam birth energy in keV and A_{NBI} is the atomic mass of the species the beam consists of (Kazakov, Van Eester & Ongena 2015). Note the ionisation efficiency depends on the atomic mass A_{NBI} of the beam species type: lighter atoms penetrate better than heavier ones so a T beam deposition profile will be more concentrated near the edge than a D beam will. The angle of incidence α of the beam with respect to the major radius direction is imposed and the beam is assumed to be in the equatorial plane and launched from the low field side. Because of the tokamak's toroidal curvature, the beam either ends on the high field side (small α for dominantly perpendicular beam) or on the low field side (large α for dominantly parallel beam). For simplicity we assume the machine walls are simply at major radius $R_M = R_o + \rho_{\text{wall}}$ (low field side) and at $R_m = R_o - \rho_{\text{wall}}$ (high field side) i.e. a cylindrical machine is assumed when describing the beam source; R_o is the major radius and ρ_{wall} the minor radius of the vessel wall. The points where the beam trajectory intersects the machine wall on the low field side are given by

$$X = R_M \frac{[t^2 \pm 1]}{t^2 + 1}, \quad (3.8)$$

where $t = \tan \alpha$, and the points where the beam trajectory intersects the machine wall on the high field side are given by

$$X = \frac{R_M}{[t^2 + 1]} \left[t^2 \pm \frac{R_m}{R_M} \left[1 + t^2 \left(1 - \left(\frac{R_M}{R_m} \right)^2 \right) \right]^{1/2} \right]. \quad (3.9)$$

The former always has 2 real roots (one of them being the beam injection reference point), the latter roots are only physically relevant (real) when the angle is low enough. The corresponding Y values in both cases are

$$Y = -t(X - R_M). \quad (3.10)$$

Provided the density is known, the reduction of the beam intensity can now be computed. To account for the fact that the beam is not necessarily fired horizontally, a factor $C_{\text{nbi}} \geq 1$ is introduced to mimic the corresponding increase in length of the beam trajectory. For purely geometrical reasons 'normal/perpendicular' beams (small α) require higher densities than 'tangential/parallel' beams (α closer to 45°), the intersection point with the wall for the former being on the high and for the latter on the low field side




FIGURE 13. Beam trajectory and associated beam intensity for a constant density of $N_e = 10^{20} \text{ m}^{-3}$.

so that the former path is significantly shorter; perpendicular and parallel refer to the direction of the beam path with respect to the toroidal direction, the direction of the dominant component of the static magnetic field. The left panel of figure 13 depicts a few typical normal and tangential beams for $\alpha = 0, 20^\circ, 40^\circ$ and 60° while the right panel depicts the reduction in beam intensity as a function of the minor radius. Unlike what will be done later in the paper, a constant density of 10^{20} m^{-3} – as opposed to a density obtained by solving the particle transport equation – was considered here, merely for illustration purposes. Note that the tangential beams do not pass all ρ values on the innermost leg of the trajectory, potentially giving rise to an increased deposition when the beam is nearly parallel to the local magnetic surface and beyond which a discontinuous jump in the local intensity and hence in the power density profile is observed. Depending on the value of α this will give rise to a local increase of the temperature gradient (more or less steep depending on the local transport strength but geometrical in nature and not due to the transport itself; in the latter case one would refer to such an increase as a ‘transport barrier’) near the specific ρ where the beam trajectory reaches its innermost R . To clearly highlight the effect, figure 14 illustrates this for a somewhat reduced (constant) density of $2 \times 10^{19} \text{ m}^{-3}$ and $\alpha = 40^\circ$.

Radiation losses are often a concern in high-performance plasmas, especially in the presence of a metallic plasma facing wall. The radiation fraction is assumed to be a fixed fraction of the incoming power but either (i) has its own specific profile or (ii) has the same profile as that of the auxiliary heating. Gaussian profiles with an imposed width and location are adopted for the particle source, RF power density and – in case (i) – power loss density. The RF power density typically peaks near the plasma centre thanks to a proper choice of the frequency for a given magnetic field strength, while the particle source peaks near the edge as ionisation is most concentrated there. If localised, a loss source peaking near the edge is adopted. The latter is not universal but is e.g. typical for high-performance JET Baseline plasmas with dominant radiation from a low field edge radiating blob due to heavy, rapidly rotating tungsten impurity ions – most likely sputtered from the divertor – merging to the midplane under the influence of the centrifugal force (Garzotti *et al.* 2019).

Neutral beam injection cannot be used when the density is too low because of excessive shine through, yielding deleterious edge sputtering and impurity influx. Unless the beam energy is low or the beam power is huge, the contribution of the beam to the fuelling (particle source) is modest. Since the RF power density is proportional to the square of

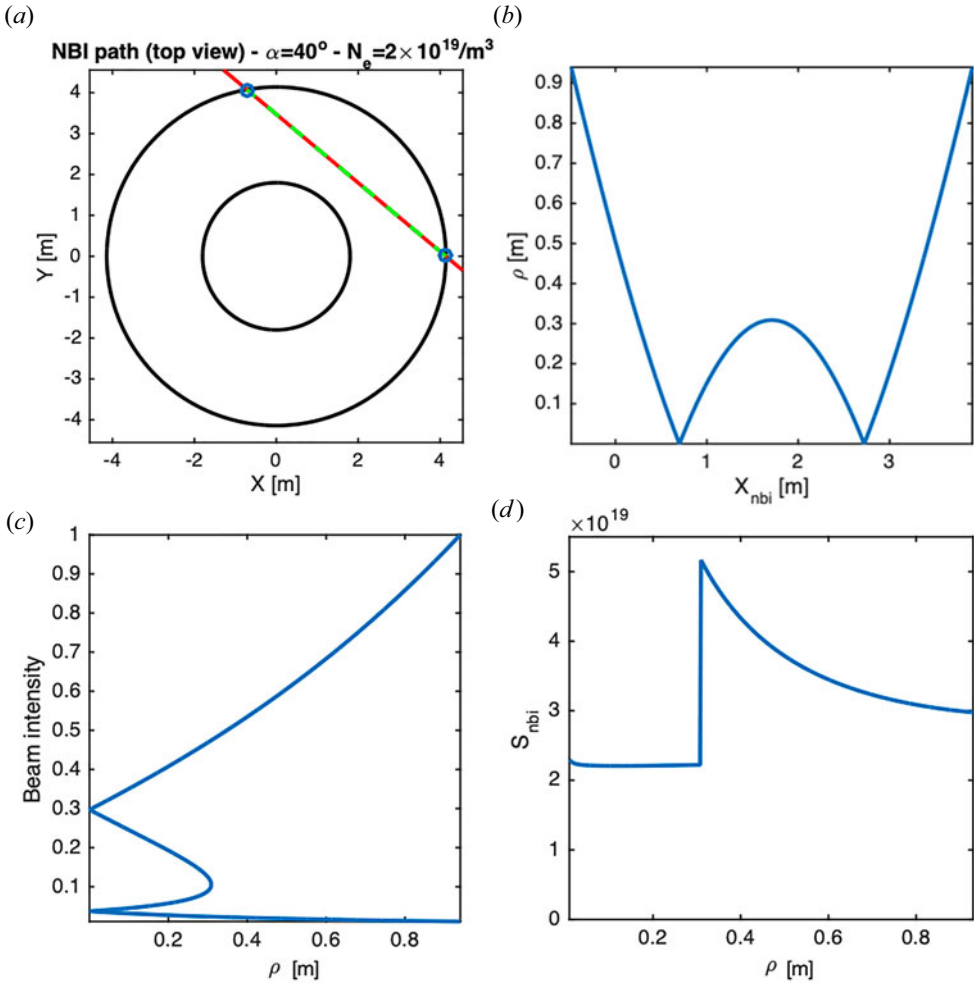


FIGURE 14. Beam trajectory. Top view of the tokamak in terms of the Cartesian coordinates (a) and minor radius in terms of the Cartesian X (b), associated beam intensity for a constant density of $N_e = 2 \times 10^{19} m^{-3}$ (c) and beam source (d).

the plasma frequency and hence to the density for a given electric field, core RF heating becomes inefficient at too low density as well. In presence of metallic vessel walls, the RF antenna power coupled to the plasma needs to be absorbed inside the plasma, if need be by the waves sloshing over the plasma multiple times before being damped. To ensure the power can be damped, the electric field amplitude increases until the absorbed and coupled powers match. As a consequence, the RF absorption is a weaker function of the density than what a local model suggests but it still makes sense to consider the power deposition as prescribed. Only when the density is very low such that the single pass absorption is lower than $\approx 30\%$ of the incoming power do parasitic damping mechanisms compete strongly with cyclotron and Cerenkov heating, making the absorption very dependent on the standing wave pattern of the global wave modes set up inside the vessel. For the present study, it is assumed RF-heating-related edge damping does not vary significantly and can be left out of the description.

3.2. Discussion of transient effects in transport

We now proceed to solving the simplified time-dependent transport equations accounting for combined beam and wave heating. As in the section on the setting up of the distribution and consistent with the settings adopted there, the parameters needed for the computations are very loosely inspired on a tokamak of the size of and available auxiliary heating like JET: the major radius is 2.97 m, the minor radius is 0.94 m. A D beam with source energy of 110 keV and total injected beam power of 30 MW is assumed and 3 MW of RF power is launched into the machine unless specified otherwise; $\alpha = 26.5^\circ$ (tangential beam with turning point at the high field side edge). Unless when radiation is purposely changed, the radiated power is 40 % of the incoming beam power. The Gaussian half-widths of the particle source, the RF source and the losses are 0.06, 0.1 and 0.1 m, respectively. These half-widths can indeed be assumed to be small since RF power allows us to deposit power in a narrow region with a width related to the Doppler shift of the population for the dominant antenna spectrum mode, the losses due to radiation tend to be localised in a mantle or – for heavy impurities such as tungsten so that centrifugal force pushes them to larger major radii on account of their toroidal rotation and mass – a low field side ‘blob’ close to the top of the pedestal and most of the ionisation of the gas happens close to the last closed flux surface. The RF power density peaks at $\rho = 0.1$ m, the losses at 0.9 m and the particle source at 0.93 m. The (constant) particle diffusion coefficient is $D_N = 0.3 \text{ m}^2 \text{ s}^{-1}$ and that of the energy is $D_E = 0.8 \text{ m}^2 \text{ s}^{-1}$; these brutally simple 1-fluid transport equation values are chosen by ‘reverse engineering’ to yield ballpark densities and temperatures characteristic of JET’s Baseline shots – see e.g. Mailloux *et al.* (2022) and Garzotti *et al.* (2019) – in steady state. Based on the shape of the collisionality (see the discussion at the end of this section) an extra diffusion contribution of $D_N = 1.0 \text{ m}^2 \text{ s}^{-1}$ localised at the pedestal was added with a gaussian halfwidth of 0.1 m. The gas/pellet injection is started immediately; the total particle source of combined beams and pellets is either $4 \times 10^{22} \text{ s}^{-1}$ (‘gas + pellets’) or $2 \times 10^{22} \text{ s}^{-1}$ (‘gas only’). The density needs to build up before the beams can be switched on. This is done self-consistently but a minimum time of 600 ms is allowed before checking if the shine through is low enough for the reached density. The RF power is switched on 400 ms after that. Figure 15 depicts the particle and energy sources when a steady state has been reached. Note that, while auxiliary power is mainly delivered by the neutral beam system, the contribution of the beam to the particle input is very modest: approximately 1/20. Since the actual breakdown cannot be modelled by the adopted simple model, an initial density and temperature profile have been assumed. They are L-mode type ohmically heated profiles with modest values at the edge and with a peak core value of $2 \times 10^{19} \text{ m}^{-3}$ and 2 keV for density and temperature, respectively.

Figure 16 shows the evolution towards the steady state; the thick red line is the final state reached after 3.6 s. The density is first strongly hollow. The void is gradually filled over time. When the beam intensity computation shows the shine through is low enough, the beams are powered and the temperature starts rising (no explicit ohmic power source has been considered so the temperature remains negligible as long as the beams are not switched on). Ultimately (fat red lines) a temperature of 8 keV and a core density of $7.2 \times 10^{19} \text{ m}^{-3}$ are reached. A density ‘pedestal’ is formed, not because the model contains a description of the H-mode but simply because the particle (ionisation) source is very close to the edge; a mild temperature pedestal is formed since the beams start depositing power close to the edge (the sharper the edge density gradient, the more the beam deposition is significant close to the edge). The rightmost panel depicts the evolution of the energy $3/2NkT$ which – at least for the adopted parameters – reaches a steady state much quicker than the individual densities and temperatures.




FIGURE 15. Source (a,b) and power (c,d) density (a,c) and corresponding integrated source (b,d) when the steady state is reached.




FIGURE 16. Simulation of the density, temperature and energy evolution to a steady state.

Figure 17 shows the temporal evolution of the density and temperature as a function of time. While the density near the edge rises very quickly but also reaches its steady-state value very soon, it takes time to allow for a gradual increase of the density in the core via diffusion. The temperature stays untouched as long as the auxiliary heating is not switched on. It first rises quickly when the beam is switched on when the density is high enough for the shine through to be small, and then shoots up further when the RF is switched on. There is a small temperature overshoot due to the beam heating but a more significant one due to

the wave heating. This overshoot occurs mainly for 2 reasons: (i) at the somewhat modest density at the start of the heating phase, the temperature response to a prescribed heat source is more pronounced. (ii) Specifically for heating schemes able to deposit power very locally, the associated temperature response is bigger, be it that the extent of the overshoot is mitigated if transport is faster (not shown). As was illustrated in the previous section, however, the overshoot is expected to be less pronounced when modelling a discharge more carefully: rather than directly ending up in the thermal subpopulation modelled by the transport equations, the energy flows to fast particle populations which subsequently slow down. This intermediate step causes a delay, smoothing the overshoot actually observed. After that, the temperatures near the core very gradually reduce to their steady-state level. An actual steady state is only reached after more than 3 seconds. The 2 overshoots suggest that the timing of the switch-on of the RF and NBI power can be optimised to transiently achieve as high as possible temperatures for a given maximally available auxiliary power. [Figure 18](#) illustrates this: a 200 ms delay between the NBI and RF switch-on times seems optimal to transiently reach a high temperature but the peak reached is not very different when choosing a time difference in the range from 200 to 600 ms. Separating the RF and beam switch-on times more decreases the peak at which combined wave and beam heating occurs, and illustrates that the peak is actually composed of 2 subpeaks, one associated with each heating mechanism. The steady-state value reached is – for the simple model adopted where e.g. the radiation fraction is prescribed – identical. In actual experiments, different choices yield different end results, however; see e.g. the *Nuclear Fusion* issue dedicated to the recent JET DT experiments (Maggi *et al.* 2023). One aspect in that dynamics is the RF-only peak that occurs before the beam is switched on and when the density is still low. This predicted peak occurs since a constant RF density has been imposed in the simple computation so at a lower density the model will predict a higher temperature to be reached (an example with reduced particle injection will be treated later; see [figure 23](#)). However, as was dealt with in some more detail in the first part of this paper which discussed transient kinetic effects (§ 2), core heating efficiency depends critically on which RF heating scheme is considered: some wave heating schemes are efficient already at low temperature while others are not. Since the RF power density scales with the density and – up to polarisation effects determining which species benefit most of the wave energy – with the square of the electric field, the core RF heating efficiency should then – at first sight – be lower at lower density. All RF power launched from the antenna has, however, to be absorbed inside the tokamak vessel. If the adopted wave absorption scheme is weak because the density is still low, the electric field magnitude has to grow in response to ensure the launched power will indeed be absorbed in the volume out of which the wave power cannot escape. This makes the density a less critical parameter and suggests the predicted RF-only temperature peak is at least partly artificial. The critical issue is the fate of launched power: when core damping is weak, the globally enhanced electric field favours parasitic edge damping; power deposited in the low density edge region potentially triggers enhanced sputtering. To avoid this, it is preferred to delay the switch-on of the ICRH until a significantly high density plasma target has formed, and the choice of a heating scheme that is potent even at low temperature is preferred.

For comparison, and to illustrate the importance of having the proper diffusion coefficients, the temperature profile obtained in the absence of the enhanced pedestal temperature diffusion is shown in [figure 19](#). Qualitatively, the evolution looks similar but the temperature now reaches a higher steady-state value of 11.5 keV since part of the edge diffusion is absent. Recall that diffusion is a random walk process which causes the quantity studied at a given location at some reference time to be spread away from its initial location at later times, the square of the distance over which it spreads (both inwards and




FIGURE 17. Density and temperature temporal evolution at a few reference positions from the core to the edge.

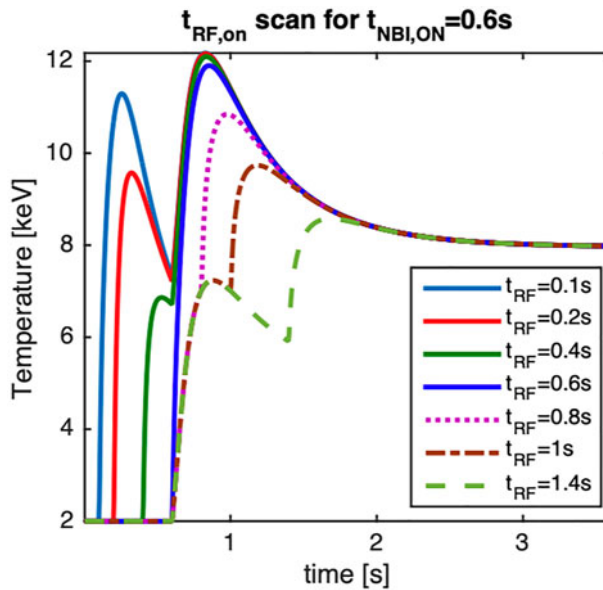


FIGURE 18. Scan of the RF switch-on time for fixed beam switch-on time.

outwards) being proportional to the diffusion coefficient D as well as the elapsed time: $\langle(\rho - \rho_{\text{ref}})^2\rangle = D(t - t_{\text{ref}})$. Enhanced energy diffusion tends to broaden the energy profile quicker, lowering the peak central temperature reached but – more importantly in this case – also transports energy quicker to the edge, across which it is lost. This spillage is masked when the steady state has been reached since local differences in fluxes then merely reflect differences in local source strength, an equilibrium being reached between in- and outgoing fluxes and the net equilibrium flux then having to be inward unless the integrated loss terms in the power source overtake the integrated heating term (a case which was artificially excluded here by fixing the imposed radiation fraction to be less than 1 at all times). Unless the transport coefficients are governed by other dynamics at the




FIGURE 19. Temporal evolution of the temperature profile when the fraction of the total radiated power with respect to the beam power stays constant at 0.4 while the radiation is localised near the pedestal as in figure 16 but the enhanced localised diffusion at that location is absent.

early stage (the here adopted model merely imposes them) power is preferably deposited well away from the edge to keep risk of spillage across the edge minimal.

With the adopted model and for transport equation coefficients that do not vary in time a steady state can be reached, as was shown in the previous figures. Often, the plasma conditions evolve, however. The rate at which that happens depends on the strength of the reigning processes. Figure 20 illustrates the temperature change when the fraction of the total radiated power with respect to the beam power evolves from 0.4 at the start of the discharge to 0.7 at its end and is localised near the pedestal (left panel) or is not localised near the edge but spread throughout the plasma, sharing the same power density profile as the beam up to the evolving factor. As expected, the temperature and hence the energy profiles now show a decreasing trend rather than approaching a steady state. It is the effective local power – the sum of the net power locally deposited and the energy carried into the region of interest per second by diffusion or convection – that modifies the temperature locally. The radiation being localised, the key difference between the middle panel in figure 16 and figure 19 is the behaviour close to the edge. The core heating is largely steered by well-localised ICRH, although it is also affected by the other mechanisms. However, at the edge, a more pronounced ‘pedestal’ is formed when the losses there are more modest. Also the temperature profile is fatter when the edge-localised losses are smaller; e.g. at $\rho = 0.6$ m, the temperature of the 2 cases looked at differs crudely by 1 keV. The edge temperature is less sensitive to power than that in the centre because of the volume effect: at the same density, the same amount of power causes a bigger temperature change closer to the core because the power is distributed over a smaller volume i.e. over fewer particles. Comparing the left and right panels, it is clear that the sensitivity to the radiation is very different depending on the shape of the radiation profile: when the radiation is at the same location as the heating, the impact is direct and

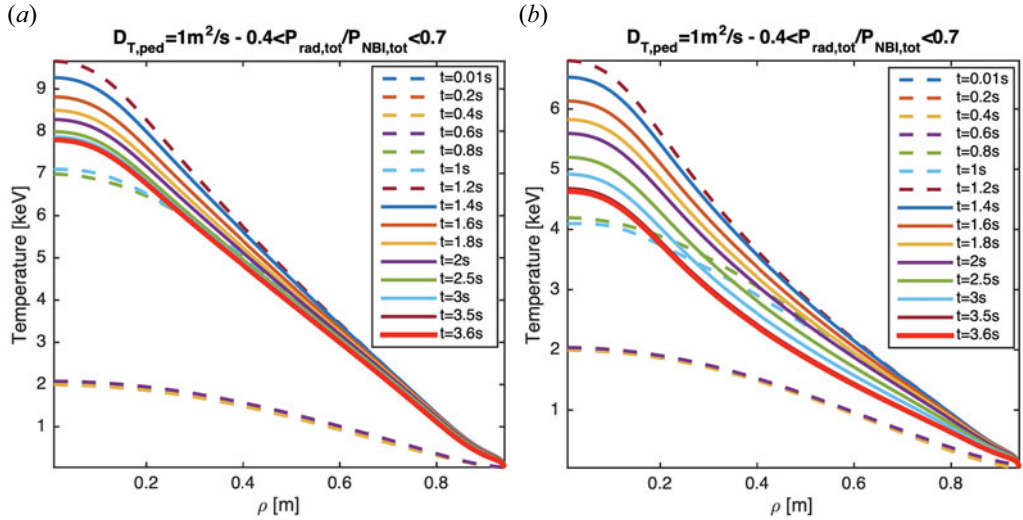


FIGURE 20. Temporal evolution of the temperature profile when the total radiated power fraction with respect to the power of the beam evolves from 0.4 at the start to 0.7 at the end of the discharge while the radiation is localised near the edge (a) as in the middle panel of the reference case [figure 16](#), or has the same fraction throughout the profile (b).

strong while, when the heating and loss occur at different locations, the heat first needs to diffuse to the loss region for the effect to be visible. One example one may think of is core-localised radiation when high-Z impurities can reach the core and hence directly reduce the efficiency of the local (dominantly RF) heating there. Another example is beam heating at high density and/or with a heavy beam ion type, causing the beam heating to be more dominantly located in the outer regions of the plasma (for example when considering T rather than D beams and early after the switch-on of the beam power). The beam deposition profile critically depends on the density and the transient behaviour of the density has a direct impact on the beam deposition and hence on the evolution of the discharge.

As the temperature drops, radiation will become gradually more pronounced. The radiation is imposed in the present, simple model while radiation actually strongly depends on temperature. This effect is only faked in the present modelling by increasing the radiation fraction. Radiation can further be increased for various reasons, for example when ELM (Edge Localised Mode) flushing dwindles, giving rise to high-Z impurities remaining in the discharge while also allowing the density to further increase, both effects producing enhanced collisionality and radiation.

To assess the relative importance of the wave versus the beam heating, the temperature profile evolution for reduced and enhanced wave and beam heating are presented next; the increased pedestal diffusion is omitted to highlight the bare effect of increased power. [Figure 21](#) shows the achieved temperature when the RF heating is switched off altogether (a), and when it is increased to 5 (b) instead of 3 MW. In the absence of RF power the temperature first reaches a value of 7 keV when the density is still fairly modest (the density is not replotted since it is identical to that in [figure 16](#)) but then drops to a mere 5 keV; recall there is 30 MW of beam power being launched into the plasma. Increasing the RF power to 5 MW yields a temperature overshoot to almost 12 keV, after which the core temperature settles at a value of 10 keV.

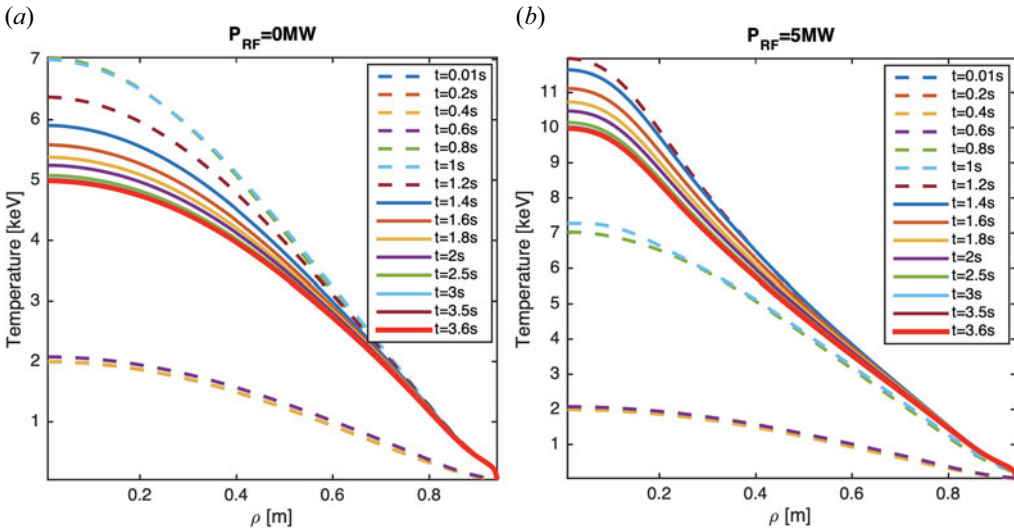


FIGURE 21. Temporal evolution of the temperature profile in absence of RF power (a) and for $P_{RF} = 5 \text{ MW}$ (b).

Decreasing the NBI power level from 30 to 25 MW only very slightly drops the peak and ultimate values reached, as can be seen by comparing figure 22(a) with figure 16(b). Similarly, increasing the power to 35 MW only mildly increases the temperatures. Somewhat more surprisingly, even when adding an extra 15 MW of power to have $P_{NBI} = 45 \text{ MW}$ still hardly affects the results. The maximal temperature reached is approximately 11.5 keV while the steady-state maximum is 9 keV. There are 2 reasons one can immediately point at to help explain this: (a) the beam deposition profile is very wide, giving rise to a much broader spread temperature profile than that achieved by RF heating. (b) In high density plasmas, raising the temperature requires much power, especially in the presence of point (a). The present model is, however, too simple to capture the full dynamics of the heating process since a single temperature evolution equation is solved, although the dynamics of electrons and ions, and the dynamics of the minority, beam and majority species can be very different. As was illustrated in § 2 when solving a(n equally simple) Fokker–Planck equation, fast subpopulations tend to preferentially dump their power Coulomb collisionally on the electrons. The critical energy

$$E_c(eV) = 14.8kT_e(eV)A_f \left[\sum_i \frac{N_i Z_i^2}{N_e A_i} \right]^{2/3}, \tag{3.11}$$

(the energy at which a test particle transfers as much energy to electrons as to ions when slowing down through Coulomb collisions) being proportional to the electron temperature, such wave-induced electron heating cranks up the critical energy, subsequently allowing more of the beam energy to be transferred to the ions, of benefit for producing neutrons. It also reduces the friction, making the heating – a competition between auxiliary power reaching a population and the Coulomb collisional passing on of power to other populations – more efficient i.e. causing a larger temperature increase for a given power and density. Whereas a more sophisticated model ought to be used to study the dynamics in more detail, (a) is an argument that holds true in any case, even if the here adopted model is overly simple. When seeking to reach high temperatures allowing ion populations to




FIGURE 22. Temporal evolution of the temperature profile for $P_{\text{NBI}} = 25$ MW (a), $P_{\text{NBI}} = 35$ MW (b) and $P_{\text{NBI}} = 45$ MW (c).

fuse, it seems beneficial to include RF wave heating in the heating mix to allow power to be deposited in a small volume where it allows us to reach high temperatures. A few MW of wave power deposited near the centre allows us to push the core temperature up. Broadening the temperature profile is, in contrast, achieved more efficiently by beam heating since this type of heating relies on ionisation and is characterised by a broad deposition profile. Harvesting a maximum of neutrons via thermal fusion reactions (the temperatures described here represent those of the bulk, not that of high energy populations) evidently benefits from both an increased and broadened profile. The fact that diffusion carries power away from its source both towards the core as to the edge suggests adopting heating schemes able to localise power where it is needed has clear benefits.

The energy evolution (rightmost) panel in figure 16 indicates that opting for more modest densities allows us to reach higher temperatures. This is illustrated in figure 23, which shows the density and temperature profile evolution when the particle source strength is cut to half its size, $2 \times 10^{22} \text{ s}^{-1}$ rather than $4 \times 10^{22} \text{ s}^{-1}$. The core density reached now only is $5.8 \times 10^{19} \text{ m}^{-3}$, but the steady-state temperature is over 12 keV while the maximal temperature reached in the temperature overshoot is almost 16 keV. When it comes to harvesting neutrons, reducing the density is not the first option taken since the fusion yield scales linearly with the densities of the 2 interacting ion populations and hence scales with the square of the electron density. It may, however, be beneficial to lower the density if the associated rise in temperature allows for reaching of much higher fusion cross-section values and/or ensures neutrons are also harvested from populations with somewhat higher energies (e.g. the $D-T$ fusion cross-section favours particles with energies around 100 keV).

Modulation is one of the tools often used to help obtain more detailed knowledge of specific phenomena. Heat transport studies e.g. rely on localised power modulation to obtain the experimental power deposition profile; see e.g. Mantica *et al.* (2011, 2021). Figures 24 and 25 illustrate the impact of modulating the particle source between 0.1 and 1.5 s, the NBI power between 1.8 and 2.3 s and the RF power between 2.4 and 2.9 s. For all modulations, a modulation frequency of 20 Hz was adopted. The particle source is doubled when modulating (this modulation either representing an increase in gas opening or the firing of extra pellets at a desired rate) while the powers are halved (in high-performance discharges, the heating systems are often running at their maximal performance so further increasing the power is impossible). As can be expected, the effect




FIGURE 23. Density and temperature evolution for a reduced gas injection scenario.

of the particle source modulation is dominantly seen in the edge; the amplitude of the modulation is reduced going deeper into the plasma and is – on the scale of the present drawing – even unnoticeable close to the plasma core. As long as the neutral beams are not fired, the temperature remains unchanged (recall that as its impact is small compared with that of the other heating sources, ohmic heating was omitted in the present computation) but the temperature change is immediately observed when the beams are switched on. Because the beam power deposition profile is wide, the effect of the beam modulation is observed throughout the plasma but is rather modest in amplitude, even for a beam power modulation with an amplitude of 15 MW. In contrast the modulation of the RF power – the RF power deposition being well localised – is only visible in the core. Note that a change of RF power by 1.5 MW causes a much bigger oscillation of the temperature than the 15 MW change of the beam power, suggesting wave heating is a much better tool for affecting the local temperature and for that reason – since it allows the heating to be very localised – a better tool to perform heat transport studies. Figure 25 illustrates why pellets are beneficial for ELM triggering: The increased source rate not only sharply raises the edge density but also its spatial gradient. Seeing ELMs as relaxation phenomena when a too steep gradient occurs, firing pellets helps to ensure this. The present model does, however, not shed any light on the detailed dynamics of ELMs.

Understanding ELM behaviour requires a study in itself and necessitates insights well beyond the scope of the present qualitative modelling; see e.g. Hoelzl *et al.* (2021). More generally, it speaks for itself that the here adopted model is adequate for highlighting specific effects but does not give details on how or why they happen. Moreover, it sidesteps quite a few potentially important ones. It is e.g. known that MHD activity (often manifesting itself by bursts on time scales vastly shorter than those on which convection and diffusion occur) is a key player in high-performance dynamics; hence the note on edge-localised modes. Furthermore, heating the plasma via beams or ICRH waves unavoidably creates high energy populations, the orbits of which can have widths that are non-negligible with respect to the minor radius, an effect likely to have an impact e.g. on the fate of the freshly ionised beam particles close to the plasma edge. Finally, and most obvious since it involves the parameters studied in this section, changes of the density and temperature have a big impact on the plasma’s resistivity and collisionality, the former being important for the current diffusion (and hence the safety factor) and both




FIGURE 24. Temporal evolution of the density and temperature at 5 positions in the plasma: $\rho = 0, 0.24, 0.49, 0.74$ and 0.92 m when modulating the particle source from $t = 0.1$ to $t = 0.5$, the beam source from $t = 1.8$ to 2.3 s and the RF source from 2.4 to 2.9 s.




FIGURE 25. Temporal evolution of the density and its spatial derivative at 5 positions in the edge of the plasma: $\rho = 0.85, 0.87, 0.89, 0.91$ and 0.93 m.




FIGURE 26. Typical evolution of the collisionality (in \log_{10} scale) as a function of time and minor radius.

being important for the diffusivity of the plasma, here taken as an external parameter. For example, neoclassical diffusion involves both the electron–ion collisionality and the safety factor, the latter indirectly depending on the current density itself controlled by the resistivity. Figure 26 depicts a typical evolution of the electron–ion collision frequency $\nu_{ei}[\text{Hz}] = 0.15 \times 10^5 N_e [10^{19} \text{ m}^{-3}] / (T[\text{keV}])^{3/2}$ profile and illustrates that the collisionality is higher in the early stages of the discharge than when it reaches its steady state; likewise, the current diffusion is very rapid when the temperature is very low but slows down significantly later on in the discharge. The time scales on which collisionality and resistivity vary at the start of the discharge make accurate modelling of that phase tough; it is the reason why it has been sidestepped here. The initial density and temperature considered are simply imposed; they are qualitatively relevant for the ohmic heating phase just prior to the main heating being switched on. The figure also shows that the pedestal collisionality is roughly an order of magnitude larger than the central value, underlining why neoclassical effects are thought to be important for understanding the pedestal. As has been illustrated in the earlier given examples, including different profiles yields significantly different results so, ideally, the dynamics should be modelled self-consistently. The procedure of externally imposing D , V and the loss profile was adopted to keep the adopted procedure simple but excludes obtaining more than handwaving results. The proper evaluation of the transport coefficients is a rich research topic in its own right.

4. Conclusions

In the present paper two simple models were exploited to illustrate the role of transient effects on tokamak discharges. Although it is clear that simple models at best provide a qualitative and not a quantitatively correct description of the physics, stripping almost all of the less essential aspects allows us to highlight that transient solutions often differ significantly from the ultimate steady state reached. Whereas the simple models used here yield the same steady state irrespective of how it was reached, experimental evidence

suggests that careful tuning of the transient phase allows reaching of differing steady-state regimes, likely under the influence of nonlinear phenomena.

The fusion yield involves the densities of the fusing populations but also depends sensitively on the relative velocities of fusing particles so that the temperature or energy – and more specifically the shape and amplitude of the distribution close to the maximum of the fusion cross-section – is a key parameter as well. Tail formation takes time. Different mechanisms occur on different time scales. For the adopted core parameters, wave heating is typically faster than collisional effects. Fundamental cyclotron heating deforms the whole population and hence is slower than second harmonic heating, but reaching an equilibrium between heating and collisions takes time. Inside a metallic shell, the electric field is forced to grow to ensure all incoming power is damped and hence weak wave damping in the core indirectly favours parasitic edge damping, which needs to be avoided as it gives rise to sputtering and hence impurity influx. Because harmonic heating is inefficient in the thermal region, it is better exploited when a preheated population (such as a beam) is available. Lacking such a population it is better to initiate the RF heating using fundamental cyclotron minority heating, which is efficient even in cold plasmas. Heating very close to the magnetic axis excludes particles with a too high parallel velocity from being heated by RF waves while – for a given power density – tail formation is faster. Well away from the core all particles can be accelerated but the power density is smaller and hence the tail energies reached are smaller. Because of their differing mass, electrons and ions occupy a very different region in velocity space. As a result, collisional cross-talk between ions is typically faster than cross-talk between electrons and ions. Seeking ion heating, exploiting heating methods that directly affect the ions is therefore beneficial.

Using a simple model to compute the beam deposition profile, it was shown that different densities yield different power depositions; also heavier beam ions tend to deposit more power near the edge. Diffusion being a random walk process, it both distributes particles or energy towards the core and to the edge, suggesting that transient temperature effects benefit most from power deposited well away from the edge of the plasma. The fact that diffusion is bi-directional is masked when merely looking at the steady-state solution. That state being reached, the integrated power (gain and loss) is linked to the flux at the edge and a net steady-state power gain simply requires a net negative (inward) flux across the edge. Leaving aside the actual breakdown, which requires more sophisticated modelling, the initial density profile is hollow and it takes time before the density gradient becomes positive throughout the plasma. For the adopted parameters, the core temperature typically somewhat overshoots its steady-state value before it relaxes to its ultimate value. When the sources and losses are not at the same location, it takes time (on the diffusion time scale) to globally feel the impact of the imbalance and move to a steady state. Localised heating near the core allows for the reaching of high temperatures, both transiently and in steady state. Seeking to minimise undesired effects, fusion power stations will likely primarily exploit thermonuclear reactions i.e. will dominantly generate neutrons from the reactions involving thermalised populations. Although a broader energy profile is clearly beneficial, the present study suggests that the impact of core-localised ICRH heating in the highest density region where also the bulk temperature is highest and hence where most neutrons will be harvested is often more important than that of widely spread NBI heating, in particular if neutrons are partly harvested through beam–bulk fusion reactions; NBI heating – on the other hand – ensures the volume in which neutrons can be harvested is broadened. As formation of the density profile starting from a neutral population requires ionisation, particle source modulation mainly impacts on the edge density while its impact is almost invisible deeper inside the plasma. Also, the density gradient is sensitive to source modulation (think of firing pellets, a technique often used

to trigger ELMs by forcing the edge gradients across the threshold value). Because of its localised nature, ICRH power modulation is much more efficient in locally affecting the temperature than NBI power modulation. Operating at reduced particle throughput allows the reaching of higher temperatures at the expense of lower densities reached.

5. End remark

The here adopted models by far do not do justice to the rich physics of wave–particle and particle–particle interactions accounted for in sophisticated modelling codes. Strictly, a multi-fluid transport model making the distinction between electrons and ions, and between bulk, minority and impurity ions should be used. The diffusion and convection coefficients themselves depend on the density and temperature, an effect entirely overlooked here. Both classical and neoclassical diffusion involve collisionality (and hence density and temperature), while the latter also depends on the local safety factor (hence depending on resistivity – so temperature – reigning current diffusion). Moreover, the plasma dynamics is not only governed by diffusion–convection: MHD phenomena happen on a much faster time scale than many diffusion phenomena and have a significant impact on the evolution of the discharge. It was illustrated how particle source modulation (pellets being the typical example) has a significant impact on the pedestal gradients but the resulting MHD relaxation (ELMs) is – aside from hints on why gradients are more sensitive to sources in transient phases – well beyond the scope of the here used elementary models. Likewise, much more sophisticated Fokker–Planck codes accounting for effects such as bounce averaging and the deviations of orbits from magnetic surfaces when the particle energy is large should be accounted for to allow a more than qualitative demonstration of the deviations of transient from steady-state solutions. Because several of the plasma constituents are routinely simultaneously directly or indirectly heated and hence have non-Maxwellian distributions, proper modelling equally requires accounting for the dynamics of all populations and e.g. requires solving of a set of coupled Fokker–Planck equations. The collisional flux in the Fokker–Planck equation is (see e.g. Karney 1986)

$$S_c^{a/b} = \frac{q_a^2 q_b^2}{8\pi\epsilon_0^2 m_a} \ln \Lambda^{a/b} \int d\mathbf{v}' \mathbf{U} \cdot \left[\frac{F_{o,a}(\mathbf{v})}{m_b} \frac{\partial F_{o,b}(\mathbf{v}')}{\partial \mathbf{v}'} - \frac{F_{o,b}(\mathbf{v}')}{m_a} \frac{\partial F_{o,a}(\mathbf{v})}{\partial \mathbf{v}} \right], \quad (5.1)$$

where $U = [u^2 1 - \mathbf{u}\mathbf{u}]/u^3$ in which $\mathbf{u} = \mathbf{v} - \mathbf{v}'$ is the relative velocity; the distributions of the test particle ‘a’ ($F_{o,a}$) and that of the population ‘b’ on which it is colliding ($F_{o,b}$) integrate to the densities N_a and N_b . Because electrons and ions occupy different regions in velocity space in view of their vastly different mass (hence u is typically big), their collisional interaction is less efficient than that between ion populations and hence it takes longer to reach a steady state when both electrons and ions are involved. Moreover, it can be seen in the above expression that the charges, masses and densities of the populations matter. Even when a population is not directly heated it will be heated via collisions. By definition, the cross-talk between 2 species is symmetric (what is collisionally lost from species a is gained by species b) but the overall impact needs not to be symmetric i.e. for both distributions being non-Maxwellian transferring power from one species to the other can be more efficient than *vice versa*. A fraction of the power intended to heat fusion fuel ions indirectly via Coulomb collisions will unavoidably not be channelled to these species. Directly affecting the distribution of one or both the fusion fuel ions to maximise the number of particles where the fusion cross-section peaks is to be preferred.

Acknowledgements

The authors wish to thank Dr L. Garzotti for illuminating discussions.

Editor Per Helander Special Collection: Invited Contributions from the 20th European Fusion Theory Conference 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.

Declaration of interests

The authors report no conflict of interest.

Appendix A. Adopted time stepping scheme

Adopting a grid of ρ positions (1-D transport equation) or $(v_{\perp}, v_{\parallel})$ values (2-D Fokker–Planck equation), storing the unknown density N , temperature T or distribution F_o as a vector \mathbf{H} and adopting a finite difference scheme to express the spatial derivatives of the unknown at a given location in terms of the values of the unknown function at that same position as well as at its neighbours (e.g. $\partial h/\partial \zeta(\zeta_i)$ is approximated by $[h(\zeta_{i+1}) - h(\zeta_{i-1})]/[2\Delta\zeta]$ for points away from the edges of the integration domain), all equations studied in the present paper can formally be written in the form

$$\frac{\partial \mathbf{H}}{\partial t} = \mathbf{M} \cdot \mathbf{H} + \mathbf{S}_H, \quad (\text{A1})$$

where \mathbf{S}_H is the ‘source’ term, independent of \mathbf{H} . Both depend on t as well as on position (transport equation) or parallel and perpendicular velocity (Fokker–Planck equation). The individual rows of the set of linear equations with system matrix \mathbf{M} are the finite difference equivalents of the equation expressed at each of the grid points, supplemented by a suitable set of boundary conditions. Numerically accounting for the time evolution can be done using the Crank–Nicolson scheme (Crank & Nicolson 1947). This replaces the partial time derivative of the sought unknown \mathbf{H} by $[\mathbf{H}_{j+1/2} - \mathbf{H}_{j-1/2}]/[\Delta t]$ on the left-hand side while that on the right-hand side is written as $[\mathbf{H}_{j+1/2} + \mathbf{H}_{j-1/2}]/2$ for a series of time steps $t = j\Delta t$. The system matrix \mathbf{M} as well as the source term \mathbf{S}_H are evaluated at $t = j\Delta t$. The system can then be solved by the time stepping scheme

$$\mathbf{H}_{j+1/2} = \left[\frac{2}{\Delta t} \mathbf{1} - \mathbf{M} \right]^{-1} \cdot \left[\left[\frac{2}{\Delta t} \mathbf{1} + \mathbf{M} \right] \cdot \mathbf{H}_{j-1/2} + 2\mathbf{S}_H \right], \quad (\text{A2})$$

where the matrices are the finite difference equivalent operators acting on the vector of unknown \mathbf{H} values. The advantage of the Crank–Nicolson scheme is that fairly large time steps can be taken.

The necessary boundary conditions are imposed relying on Lagrange multipliers; see e.g. Babuska (1973).

REFERENCES

- ANGIONI, C. 2021 Impurity transport in tokamak plasmas, theory, modelling and comparison with experiments. *Plasma Phys. Control. Fusion* **63**, 073001.
- BABUSKA, I. 1973 The finite element method with Lagrangian multipliers. *Numer. Math.* **20**, 179–192.
- BERTELLI, N., JAEGER, E.F., HOSEA, J.C., PHILLIPS, C.K., BERRY, L., BONOLI, P.T., GERHARDT, S.P., GREEN, D., LEBLANC, B., PERKINS, R.J., *et al.* 2016 Full wave simulations of fast wave efficiency and power losses in the scrape-off layer of tokamak plasmas in mid/high harmonic and minority heating regimes. *Nucl. Fusion* **56**, 016019.
- BOURDELLE, C. 2015 Turbulent Transport in Tokamak Plasmas: bridging theory and experiment. *Habilitation à diriger des recherches, Aix-Marseille Université*, <https://theses.hal.science/tel-01113299/document>.
- BRAMBILLA, M. & BILATO, R. 2009 Advances in numerical simulations of ion cyclotron heating of non-Maxwellian plasmas. *Nucl. Fusion* **49**, 085004.
- BRESLAU, J., GORELENKOVA, M., POLI, F., SACHDEV, J., PANKIN, A., PERUMPILLY, G., YUAN, X. & GLANT, L. 2018 TRANSP is a 1.5D equilibrium and transport solver for interpretation and prediction of tokamak discharges. *TRANSP. Computer Software. USDOE Office of Science (SC), Fusion Energy Sciences (FES) (SC-24)*. doi:10.11578/dc.201806274.
- COSTER, D., BASIUK, V., PEREVERZEV, G. & KALUPIN, D. 2010 The European transport solver. *IEEE Trans. Plasma Sci.* **38** (9), 2085–2092.
- CRANK, J. & NICOLSON, P. 1947 A practical method for numerical evaluation of solutions of partial differential equations of the heat conduction type. *Proc. Camb. Phil. Soc.* **43** (1), 50–67.
- FEDERICI, G., MAVIGLIA, F. & HOLDEN, J. 2022 DEMO: outcome of the pre-conceptual design phase. Special Issue of *Fusion Engineering and Design* on the European Programme towards DEMO; various papers.
- GAFFEY, J.D. 1976 Energetic ion distribution resulting from neutral beam injection in tokamaks. *J. Plasma Phys.* **16** (2), 149–169.
- GARZOTTI, L., CHALLIS, C., DUMONT, R., FRIGIONE, D., GRAVES, J., LERCHE, E., MAILLOUX, J., MANTSINEN, M., RIMINI, F., CASSON, F., *et al.* 2019 Scenario development for D-T operation at JET. *Nucl. Fusion* **59**, 076037.
- HELLSTEN, T., JOHNSON, T., CARLSSON, J. & ERIKSSON, L.-G. 2004 Effects of finite drift orbit width and RF-induced spatial transport on plasma heated by ICRH. *Nucl. Fusion* **44** (8), 892.
- HOELZL, M., HUIJSMANS, G.T.A., PAMELA, S.J.P., BÉCOULET, M., NARDON, E., ARTOLA, F.J., NKONGA, B., ATANASIU, C.V., BANDARU, V., BHOLE, A., *et al.* 2021 The JOREK non-linear extended MHD code and applications to large-scale instabilities and their control in magnetically confined fusion plasmas. *Nucl. Fusion* **61**, 065001.
- JAEGER, E.F., BERRY, L.A., D'AZEVEDO, E., BATCHELOR, D.B. & CARTER, M.D. 2001 All-orders spectral calculation of radio-frequency heating in two-dimensional toroidal plasmas. *Phys. Plasmas* **8** (5), 1573–1583.
- JUCKER, M., GRAVES, J.P., COOPER, W.A., MELLET, N., JOHNSON, T. & BRUNNER, S. 2011 Integrated modeling for ion cyclotron resonant heating in toroidal systems. *Comput. Phys. Commun.* **182** (4), 912–925.
- KARNEY, C.F.F. 1986 Fokker-Planck and quasilinear codes. *Comput. Phys. Rep.* **4**, 183.
- KAZAKOV, Y., VAN EESTER, D. & ONGENA, J. 2015 Plasma heating in present-day and future fusion machines. In *12th Carolus Magnus Summer School on Plasma and Fusion Energy Physics*, ed. K. Crombé. Schriften des Forschungszentrums Jülich Reihe Energie und Umwelt/Energy and Environment, vol. 298, pp. 290–297. Forschungszentrum Jülich GmbH Zentralbibliothek.
- LENNHOLM, M., MCKEAN, R., MOONEY, R., TVALASHVILI, G., ARTASERSE, G., BARUZZO, M., BELONOHY, E., CALABRO, G., CARVALHO, I.S., CHALLIS, C.D., *et al.* 2021 Statistical assessment of ELM triggering by pellets on JET. *Nucl. Fusion* **61**, 036035.
- LERCHE, E.A., LENNHOLM, M., CARVALHO, I.S., DUMORTIER, P., DURODIÉ, F., VAN EESTER, D., GRAVES, J., JACQUET, P., MURARI, A. & JET CONTRIBUTORS 2017 Sawtooth pacing with on-axis ICRH modulation in JET-ILW. *Nucl. Fusion* **57**, 036027.

- MAGGI, C., *et al.* 2023 Special nuclear fusion issue of papers presenting results from the JET tritium and deuterium/tritium campaign. *Nucl. Fusion* **63**, 112001–112015.
- MAILLOUX, J., *et al.* 2022 Overview of JET results for optimising ITER operation. *Nucl. Fusion* **62**, 042026.
- MANTICA, P., ANGIONI, C., BAIOCCHI, B., BARUZZO, M., BEURSKENS, M.N.A., BIZARRO, J.P.S., BUDNY, R.V., BURATTI, P., CASATI, A., CHALLIS, C., *et al.* 2011 Ion heat transport studies in JET. *Plasma Phys. Control. Fusion* **53**, 124033.
- MANTICA, P., BONANOMI, N., MARIANI, A., CARVALHO, P., DELABIE, E., GARCIA, J., HAWKES, N., JOHNSON, T., KEELING, D., SERTOLI, M., *et al.* 2021 The role of electron-scale turbulence in the JET tokamak: experiments and modelling. *Nucl. Fusion* **61**, 096014.
- ROMANELLI, M., COELHO, R., COSTER, D. & FERREIRA, J. 2020 Code integration, data verification, and models validation using the ITER Integrated Modeling and Analysis System (IMAS) in EUROfusion. *Fusion Sci. Technol.* **76** (5), 1–7.
- STIX, T.H. 1992 *Waves in Plasmas*. AIP.
- VAN EESTER, D. 1994 Self-consistent modelling of ion cyclotron resonance heating. *Plasma Phys. Control. Fusion* **36**, 1327–1354.
- VAN EESTER, D. & KOCH, R. 1998 A variational principle for studying fast-wave mode conversion. *Plasma Phys. Control. Fusion* **40**, 1949–1975.