

Verifying raytracing/Fokker–Planck lower-hybrid current drive predictions with self-consistent full-wave/Fokker–Planck simulations

S.J. Frank^{1,†}, J.P. Lee², J.C. Wright¹, I.H. Hutchinson¹ and P.T. Bonoli¹

¹Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

²Department of Nuclear Engineering, Hanyang University, Seoul, South Korea

(Received 18 July 2022; revised 12 October 2022; accepted 14 October 2022)

Raytracing/Fokker–Planck (FP) simulations used to model lower-hybrid current drive (LHCD) often fail to reproduce experimental results, particularly when LHCD is weakly damped. A proposed reason for this discrepancy is the lack of ‘full-wave’ effects, such as diffraction and interference, in raytracing simulations and the breakdown of the raytracing approximation. Previous studies of LHCD using non-Maxwellian full-wave/FP simulations have been performed, but these simulations were not self-consistent and enforced power conservation between the FP and full-wave code using a numerical rescaling factor. Here, we have created a fully self-consistent full-wave/FP model for LHCD that is automatically power conserving. This was accomplished by coupling an overhauled version of the non-Maxwellian TORLH full-wave solver and the CQL3D FP code using the Integrated Plasma Simulator. We performed converged full-wave/FP simulations of Alcator C-Mod discharges and compared them with raytracing. We found that excellent agreement in the power deposition profiles from raytracing and TORLH could be obtained, however, TORLH had somewhat lower current drive efficiency and broader power deposition profiles in some cases. This discrepancy appears to be a result of numerical limitations present in the TORLH model and a small amount of diffractive broadening of the TORLH wave spectrum. Our results suggest full-wave simulation of LHCD is likely not necessary as diffraction and interference represented only a small correction that could not account for the differences between simulations and experiment.

Key words: plasma waves, plasma simulation, fusion plasma

1. Introduction

Lower-hybrid current drive (LHCD) is a radiofrequency (RF) actuator used to efficiently drive current in tokamaks (Fisch 1987). The lower-hybrid current drives via electron Landau damping of slow waves on fast electrons with $v_e \geq 3v_{the}$, in the lower-hybrid (LH) limit, $\Omega_i \ll \omega \ll \Omega_e$, where ω is the wave frequency, $\Omega_s = q_s B/m_s$ is the cyclotron frequency of species s with electric charge q_s and mass m_s in magnetic field B and $v_{the} = \sqrt{2T_e/m_e}$ is the electron thermal velocity for electron temperature T_e (Bonoli 1985). This drives current by direct parallel momentum injection and distortion of the

† Email address for correspondence: frank@psfc.mit.edu

perpendicular distribution such that the parallel resistivity is reduced in the direction of the plasma current (Fisch 1978; Fisch & Boozer 1980).

Despite a long history of experimental demonstrations (Porkolab *et al.* 1984; Bartiromo *et al.* 1986; Moriyama *et al.* 1990; Jacquinot 1991; Ide *et al.* 1992; Bernabei *et al.* 1995; Peysson & Tore Supra Team 2001; Wilson *et al.* 2009; Cesario *et al.* 2010; Wallace *et al.* 2011; Liu *et al.* 2015), LHCD has remained difficult to predict (Bonoli 2014). Difficulty predicting LHCD is attributed to the existence of a ‘spectral gap’, whereby LH waves tend to damp despite being launched with a parallel phase velocity, $v_{\text{ph}\parallel} \sim 6\text{--}8v_{\text{the}0}$, much larger than their linear damping condition of $v_{\text{ph}\parallel} \sim 3v_{\text{the}0}$. A number of mechanisms that upshift the parallel refractive index $N_{\parallel} = k_{\parallel}c/\omega = c/v_{\text{ph}\parallel}$, where k_{\parallel} is the parallel component of the wave vector \mathbf{k} and c is the speed of light in vacuum, closing the spectral gap have been proposed. These mechanisms include: toroidally induced increases in the parallel wavenumber and wave scattering from turbulence (Bonoli & Ott 1981, 1982; Decker *et al.* 2014; Peysson *et al.* 2016; Biswas *et al.* 2020, 2021), parametric wave interactions (Cesario *et al.* 2004; Decker *et al.* 2014; Peysson *et al.* 2016) and diffractive broadening of the wave spectrum (Pereverzev 1992; Wright *et al.* 2009, 2010; Shiraiwa *et al.* 2011a; Wright *et al.* 2014). Which N_{\parallel} upshift mechanism is dominant in a given situation can profoundly affect wave damping and current drive efficiency. Therefore, it is important to understand if and when each of these different mechanisms is important.

Lower-hybrid current drive is typically simulated using raytracing. Raytracing is derived by applying the Wentzel–Kramers–Brillouin (WKB) method, assuming that $kL \gg 1$, where L is the plasma gradient scale length, to the plasma wave equation yielding the ray equations. To calculate non-Maxwellian plasma response and current drive, raytracing simulations are coupled to a Fokker–Planck simulation by formulating a quasilinear diffusion coefficient using the method described in Bonoli & Englade (1986). However, raytracing neglects the effects of diffraction and the $kL \gg 1$ limit breaks down near cutoffs and ray caustics. Solving the wave equation directly using a ‘full-wave’ simulation preserves the effects of diffraction and interference as well as accurately capturing the behaviour of the wave at cutoffs and caustics, and comparison of raytracing with full-wave simulations can allow us to determine if raytracing is accurate despite not including these effects. In this work, we implemented the first fully self-consistent full-wave/FP model that allows us to calculate the LHCD and compare with results obtained with a raytracing/FP model in Alcator C-Mod. Previous LHCD modelling work has included non-Maxwellian effects (Peysson *et al.* 1998; Wright *et al.* 2010, 2014; Shiraiwa *et al.* 2011a; Meneghini 2012), however, these simulations were not performed fully self-consistently. They required *ad hoc* numerical rescaling of the quasilinear diffusion coefficient to ensure the correct RF power was deposited in the FP calculation or employed an approximate method for calculating the non-Maxwellian dielectric tensor. We will describe the construction of our fully self-consistent non-Maxwellian full-wave/FP LHCD model and its application to modelling a set of Alcator C-Mod experiments.

2. The non-Maxwellian TORLH/CQL3D model

In order to perform non-Maxwellian simulations of LHCD using a full-wave simulation code it is necessary to iterate calculations of wave propagation and damping obtained using the plasma Helmholtz equation in the LH limit and a Fokker–Planck equation solver such as CQL3D (Harvey & McCoy 1992; Petrov & Harvey 2016). This iteration solves the system

$$\nabla \times \nabla \times \mathbf{E} = \epsilon_{\perp} \mathbf{E}_{\perp} + i\epsilon_{xy} (\hat{\mathbf{b}} \times \mathbf{E}_{\perp}) + \epsilon_{\parallel}(f) E_{\parallel} \hat{\mathbf{b}} \quad (2.1)$$

$$\frac{Df}{Dt} = C(f) + Q(f, E_{\parallel}), \quad (2.2)$$

where \mathbf{E} is the wave electric field, f is the electron distribution function, $\hat{b} = \mathbf{B}_0/|B_0|$ is the background magnetic field vector, $C(f)$ is the collision operator, $Q(f, E_{\parallel})$ is the divergence of the quasilinear flux and $\epsilon(f)$ is the plasma dielectric tensor with components

$$\epsilon_{\perp} \simeq S = \frac{1}{2}(R + L) \quad (2.3a)$$

$$\epsilon_{xy} \simeq D = \frac{1}{2}(R - L) \quad (2.3b)$$

$$\epsilon_{\parallel}(f) = 1 - \frac{\omega_{pe}^2}{\omega^2} \zeta_e^2 \text{Re}\{Z'(\zeta_e)\} - \frac{\omega_{pe}^2}{\omega^2} + \text{Im}[\chi_{zz}(f)], \quad (2.3c)$$

where S, D, R, L are the cold plasma dielectric components from Stix (1992), $\omega_{ps} = q_s^2 n_s / m_s \epsilon_0$ is the plasma frequency for species s , $\zeta_s = \omega / k_{\parallel} v_{\text{th}s}$ and Z is the plasma dispersion function (Fried & Conte 1961). The non-Maxwellian imaginary correction to $\epsilon_{\parallel}(f)$ here is $\text{Im}[\chi_{zz}(f)]$, discussed further in § 2.2 and Appendix A. TORLH calculates an E -field for a single toroidal mode n_{ϕ} by solving Helmholtz's equation (2.1) discretized with a semi-spectral representation

$$\mathbf{E} = \exp(i(n_{\phi}\phi - \omega t)) \sum_{m=-\infty}^{+\infty} \mathbf{E}^{(m)}(\psi) \exp(im\theta), \quad (2.4)$$

where n_{ϕ} is the toroidal mode number, m is the poloidal mode number and $\mathbf{E}^{(m)}(\psi)$ is represented by a cubic Hermite interpolating polynomial finite element. The TORLH discretization is global along flux surfaces and precisely defines the parallel k vector

$$k_{\parallel} = m(\hat{b} \cdot \nabla\theta) + n_{\phi}(\hat{b} \cdot \nabla\phi) = \frac{m}{n_{\tau}} \frac{B_{\theta}}{|B|} + \frac{n_{\phi}}{R} \frac{B_{\phi}}{|B|}, \quad (2.5)$$

where n_{τ} is a metric coefficient that functions as a generalized minor radius. These properties of the TORLH discretization allow hot plasma effects, like Landau damping, to be resolved without approximations required by fully finite element discretizations (Shiraiwa *et al.* 2011a). Our choice of discretization produces a block tridiagonal system of equations solved using a custom parallel block cyclic reduction solver (Lee & Wright 2014).

After the field solve, TORLH performs a post-processing step to obtain the quasilinear diffusion coefficient $D_{\text{ql},\parallel}(E_{\parallel})$, the component of the quasilinear term $Q(f, E_{\parallel})$ dependent on the parallel electric field. The CQL3D quasilinear term may be written (Petrov & Harvey 2016)

$$Q(f, E_{\parallel}) = \left\{ \frac{1}{u^2} \frac{\partial}{\partial u} \left(B_0 \frac{\partial}{\partial u} + C_0 \frac{\partial}{\partial \vartheta} \right) + \frac{1}{u^2 \sin \vartheta} \frac{\partial}{\partial \vartheta} \left(E_0 \frac{\partial}{\partial u} + F_0 \frac{\partial}{\partial \vartheta} \right) \right\} f. \quad (2.6)$$

Here, u is the momentum per rest mass, $p/\gamma m$, ϑ is the velocity-space pitch angle and B_0, C_0, E_0 and F_0 are bounce averaged quasilinear diffusion coefficients. We may relate the local B to the $D_{\text{ql},\parallel}$ from Kennel & Engelmann (1966) using the relativistic form found in

Lerche (1968) and Chapter 17 of Stix (1992)

$$B = u^2 D_{\text{uu}} = u^2 (\cos \vartheta)^2 D_{\text{ql},\parallel}. \quad (2.7)$$

We then perform a normalized zero-orbit-width bounce average of B to obtain B_0

$$B_0 = \lambda \langle B \rangle = v_{\parallel,0} \oint \frac{dl}{v_{\parallel}} B = v_{\parallel,0} \int_0^{2\pi} d\theta \mathcal{J} \frac{B(\theta)}{v_{\parallel}}, \quad (2.8)$$

where the integration element dl is directed along the magnetic field line and $v_{\parallel,0}$ is the parallel velocity at the outboard midplane, and \mathcal{J} is the Jacobian of the magnetic coordinate system. With B_0 we can obtain C_0 , E_0 and F_0

$$C_0 = -\frac{\sin \vartheta}{u \cos \vartheta} B_0, \quad (2.9a)$$

$$E_0 = -\frac{\sin^2 \vartheta}{u \cos \vartheta} B_0, \quad (2.9b)$$

$$F_0 = \frac{\sin^3 \vartheta}{(u \cos \vartheta)^2} B_0. \quad (2.9c)$$

Using the quasilinear diffusion computed by TORLH, CQL3D solves (2.2) to evolve the distribution for some amount of time. The evolved non-Maxwellian distribution function is passed back to a pre-processing routine that calculates a lookup table with values of the imaginary component of ϵ_{\parallel} , $\text{Im}[\chi_{zz}]$ (the only component of the dielectric which meaningfully varies from the Maxwellian) and then TORLH is rerun using the updated non-Maxwellian dielectric. This process is repeated until power deposition profiles stop evolving and the driven current has risen to a steady state value, indicating the system has converged. Convergence is typically obtained after ~ 20 – 50 iterations between TORLH and CQL3D.

To facilitate simulation of multiple current drive scenarios using the TORLH/CQL3D iteration, shown in figure 1, it was necessary to develop an automated framework for performing simulations. We created a set of Python and FORTRAN wrappers for the Integrated Plasma Simulator (IPS) (Elwasif *et al.* 2010). The IPS enabled the TORLH/CQL3D iterations to be performed at large scale with checkpointing and restart capabilities. Using the IPS, we successfully performed large scale TORLH/CQL3D simulations with $\sim 10\,000$ cores on Cori at the National Energy Research Scientific Computing Center (NERSC) that could be iterated for up to 48 h. However, initial simulations using the version of TORLH and the set-up from Wright *et al.* (2014) were not power conserving. In order to obtain physically and numerically self-consistent results, we needed to not only implement the improved TORLH boundary condition and field solver convergence criteria from Frank *et al.* (2022), but we also had to implement significant improvements in the non-Maxwellian components of TORLH including: reformulation of the quasilinear diffusion coefficient, construction of a $\text{Im}[\chi_{zz}]$ lookup table and the lookup table interpolation in TORLH. In the following sections we will describe the key modifications to the non-Maxwellian components of TORLH and develop convergence requirements for them.

2.1. Quasilinear diffusion coefficient formulation

The quasilinear diffusion coefficient, D_{ql} , formulation in TORLH was completely rewritten during the course of this work in order to implement a form that was power

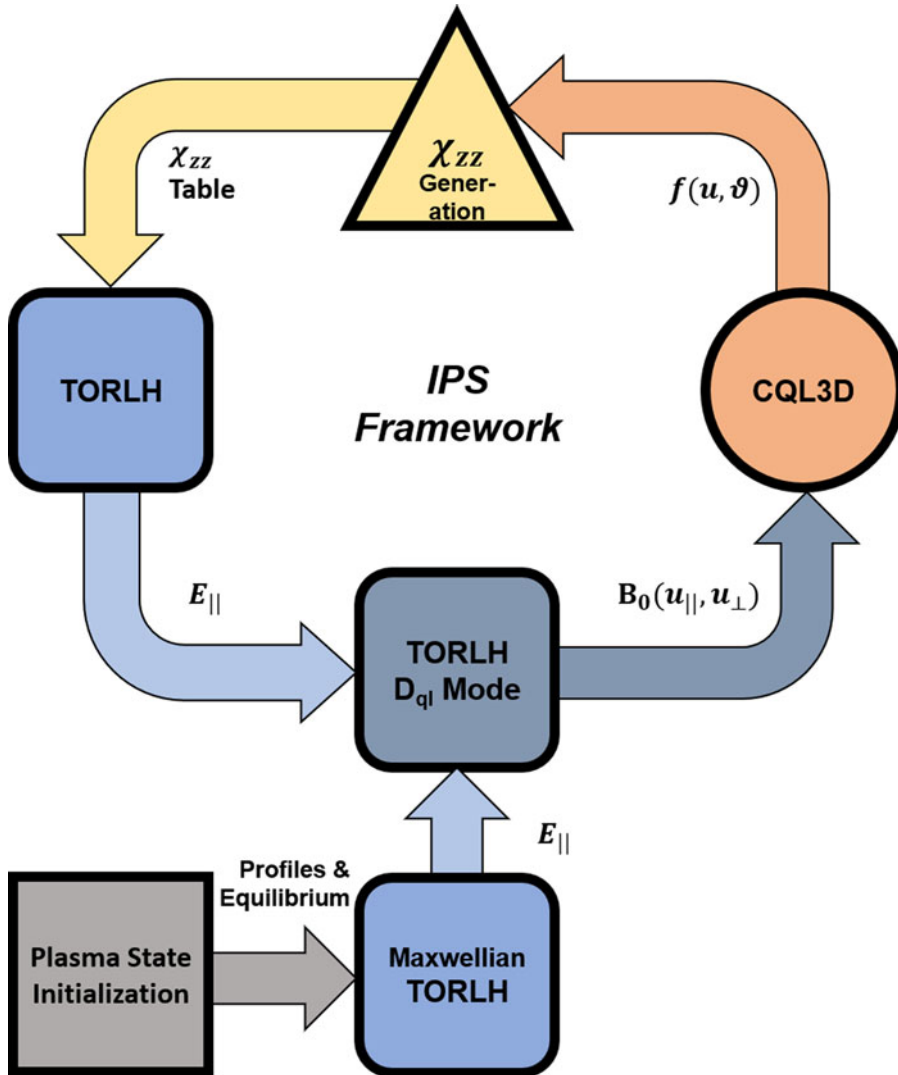


FIGURE 1. Flowchart showing an IPS driven iteration. The IPS creates a ‘plasma state’ containing magnetic equilibrium and plasma profile information then initializes a Maxwellian TORLH simulation from that plasma state. TORLH and CQL3D are then iterated using the IPS framework supplemented by Python and FORTRAN wrappers that manage input files.

conserving. The TORLH diffusion coefficient is derived from the parallel component of the RF quasilinear diffusion coefficient in Kennel & Engelmann (1966)

$$D_{q||} = \frac{\pi e^2}{2m_e^2} \text{Re} \left\{ \int dk_{||1} \int dk_{||2} \frac{v_{||}^2}{c^2} (E_{||1} J_{0,1} \exp(ik_1 \cdot r) N_{||1}) \times \delta(\omega - v_{||} k_{||1}) (E_{||2} J_{0,2} \exp(ik_2 \cdot r) N_{||2}) \right\}. \quad (2.10)$$

This diffusion coefficient parameterizes electron Landau damping by the LH wave. Using (2.7) to rewrite $D_{\text{ql},\parallel}$ in terms of the CQL3D quasilinear diffusion coefficients and applying the TORLH discretization (2.4) yields

$$B(\psi, \theta, u_{\parallel}, u_{\perp}) = \frac{\pi e^2 [E]^2}{2m_e^2 c^2} \text{Re} \left\{ \sum_{n_{\phi}} \sum_{m_1, m_2} \frac{u_{\parallel}^4}{\gamma^2} \left(E_{\parallel}^{(m_2, n_{\phi})} J_{0,2} \exp(im_2 \theta) N_{\parallel}^{(m_2, n_{\phi})} \right)^* \right. \\ \left. \times \delta(\omega - v_{\parallel} k_{\parallel}^{(m_1, n_{\phi})}) \left(E_{\parallel}^{(m_1, n_{\phi})} J_{0,1} \exp(im_1 \theta) N_{\parallel}^{(m_1, n_{\phi})} \right) \right\}, \quad (2.11)$$

where $[E]$ is the electric field normalization in TORLH of 1 V m^{-1} and truncated Bessel function expression $J_{0,n} = J_0(k_{\perp}^{(m_n)} v_{\perp} / \Omega)$ (TORLH does not have a well defined k_{\perp} for a given m, n_{ϕ} , so the k_{\perp} in this expression is obtained from the local slow-wave dispersion relation. This is an approximation is fairly accurate as the fast wave is strongly evanescent in most LHCD scenarios). In this work, we have only used a single toroidal mode due to the computational intensity of TORLH electric field solves, and the n_{ϕ} sum will be suppressed in later equations. To efficiently discretize this problem in velocity space we take advantage of the properties of the δ function. The δ function prescribes that, for a given m number, the quasilinear diffusion coefficient has a uniquely defined v_{\parallel} location in velocity space. Therefore, at a given value of m and u_{\perp} the u_{\parallel} velocity-space location is known. This allows us to rewrite (2.11) in terms of a discretized delta function

$$B(\psi, \theta, m_1 \propto u_{\parallel}, u_{\perp}) = \frac{\pi e^2 [E]^2 c^2}{2\omega m_e^2} \text{Re} \left\{ \sum_{m_1, m_2} \frac{\gamma^2}{\left(N_{\parallel}^{(m_1)} \right)^5} \left(E_{\parallel}^{(m_2)} J_{0,2} \exp(im_2 \theta) N_{\parallel}^{(m_2)} \right)^* \right. \\ \left. \times \frac{1}{\Delta u_{\text{mesh}}} \left(\frac{\partial v_{\parallel}}{\partial u_{\parallel}} \right)^{-1} \left(E_{\parallel}^{(m_1)} J_{0,1} \exp(im_1 \theta) N_{\parallel}^{(m_1)} \right) \right\}. \quad (2.12)$$

Here, we have used a rectangular discrete δ function where Δu_{mesh} is the velocity-space mesh spacing. It is possible to use more complicated form factors for the discrete δ , but these have been found to have a minimal impact on the D_{ql} formulation in the past (Lee *et al.* 2017b). For numerical simplicity, a uniform velocity-space mesh in u_{\perp} and u_{\parallel} is used and the $u_{\parallel} = \gamma v_{\text{ph}\parallel}$ corresponding to a given m number is interpolated to the nearest neighbour point on the u_{\parallel} mesh.

The D_{ql} formulation here is effectively unchanged from that shown in Wright *et al.* (2014), but a key numerical improvement to the D_{ql} formulation has been made. Previously, a sum over all spectral modes, including evanescent wave modes, was performed in TORLH. However, evanescent wave modes are disallowed in the resonant Kennel–Engelmann D_{ql} formulation because waves with $\text{Im}[k_{\parallel}] \neq 0$ are incompatible with the contour integration used to obtain the resonant limit (Kennel & Engelmann 1966). Inclusion of evanescent wave modes in the D_{ql} double mode sum introduces components of D_{ql} at excessively high parallel momenta. This causes interaction with high-energy electrons leading to the formation of sharp structures in the FP solution. These sharp structures in phase space caused numerical instabilities in both TORLH and CQL3D that could in some cases cause the simulations to fail. To avoid this, we check if the wave




FIGURE 2. Plots demonstrating the evanescent cutoff noise reduction method. (a) The CQL3D electron distribution function produced by D_{q1} vs electron energy with (dashed) and without (dotted) the evanescence cutoff. Contours of B_0 vs parallel and perpendicular momentum per rest mass (b) without the cutoff applied and (c) with the cutoff applied.

is evanescent using the following equation derived from the electromagnetic dispersion relation (Bonoli 1985):

$$[(N_{\parallel}^2 - S)(\text{Re}\{\epsilon_{\parallel}(f)\} + S) + D^2]^2 - 4 \text{Re}\{\epsilon_{\parallel}(f)\}S(N_{\parallel}^2 - R)(N_{\parallel}^2 - L) \geq 0. \quad (2.13)$$

For a given mode, if (2.13) is less than zero it is excluded from the double mode sum in (2.12) (doing this calculation efficiently required a rewrite of the quasilinear diffusion coefficient formulation routine in TORLH to modify its parallelization). Ideally, we would solve the dispersion relation with fixed N_{\perp} and ω for N_{\parallel} to determine if the mode had $\text{Im}[k_{\parallel}] \neq 0$, but this is not possible in the TORLH basis. We must settle for calculating when N_{\perp} becomes imaginary with fixed N_{\parallel} and ω . The results of this procedure, highlighted in figure 2, show that it very successfully removes high-energy noise from D_{q1} and CQL3D electron distribution functions.

To self-consistently obtain the desired power, the normalized TORLH electric fields in the D_{q1} formulation must be scaled to the launched power in the experiment P_{target} . To do this, the launched power with the normalized field values in TORLH is obtained from the

integrated Poynting flux at the antenna

$$P_{\text{RF,Poynt}}(\psi_{\text{ant}}) = \frac{\pi l_0^2 [E]^2}{c\mu_0} \int_0^{2\pi} d\theta \frac{R_0}{l_0} n_\tau \sum_{m,m'} \text{Im} \left\{ \exp(i(m-m')\theta) \right. \\ \left. \times \left(E_\eta^{(m)*} R_\zeta^{\text{curl}} E^{(m')} - E_\zeta^{(m)*} R_\eta^{\text{curl}} E^{(m')} \right) \right\}, \quad (2.14)$$

where R^{curl} is the numerical curl operator in TORLH (Brambilla 1996), R_0 is the major radius and $l_0 = c/\omega$ is the TORLH length normalization. In the limit of only Landau damping, as in the simulations here, one can alternatively use the power deposited by Landau damping to scale D_{ql}

$$P_{\text{RF,LD}} = \frac{1}{2} \int dV \text{Re} \{ \mathbf{E}^* \cdot \mathbf{J}_p \} \\ = -\pi\omega\epsilon_0 [E]^2 l_0^3 \int_0^1 d\psi \sum_{m'} \sum_m \text{Im} \left\{ E_{\parallel}^{(m')*}(\psi) \tilde{\chi}_{zz}(m', \psi, \theta) E_{\parallel}^{(m)}(\psi) \right\}, \quad (2.15)$$

where $\tilde{\chi}_{zz}$ is the Fourier transform of the imaginary component of the parallel dielectric

$$\tilde{\chi}_{zz} = \int d\theta \mathcal{J}_n \exp(i(m'-m)\theta) \text{Im}[\chi_{zz}]. \quad (2.16)$$

Using the power calculated in TORLH we may rescale the quasilinear diffusion coefficient

$$B = B \frac{P_{\text{target}}}{P_{\text{RF}}}. \quad (2.17)$$

For linear field response, $P_{\text{RF}} \propto B \propto |E|^2$ making this rescaling physically self-consistent. No further rescaling is necessary, and precise agreement should be obtained between P_{target} and the power calculated in CQL3D using the quasilinear diffusion coefficient.

A key feature of the quasilinear diffusion coefficients calculated by TORLH is their non-positive definiteness. The holes in the diffusion coefficients shown in figure 2 are the result of small non-positive definite components of the D_{ql} that are set to zero after output from TORLH. Non-positive definite quasilinear diffusion is numerically unstable and physically incorrect, but is not unique to TORLH and is a weakness of spectral D_{ql} reconstruction in many full-wave RF solvers (Jaeger *et al.* 2006; Lee *et al.* 2017a,b). The root cause of the non-positive definiteness is from the inconsistency introduced by the use of a resonant, homogeneous, local diffusion coefficient, but a non-local spectral electric field basis and toroidal geometry. This inconsistency leads to improper consideration of parallel magnetic field inhomogeneity ($\mathbf{B} \cdot \nabla \mathbf{B} \neq 0$) during the bounce average (2.8) and also the excessive interaction with evanescent wave modes discussed previously. As this non-positive definiteness occurs throughout the literature (Jaeger *et al.* 2006; Lee *et al.* 2017a,b), a number of techniques to handle it have been developed: the negative components of B_0 may be set to zero (this is done automatically in CQL3D when given a non-positive definite D_{ql}), an advanced positive definite spectral formulation may be used (Kaufman 1972; Lee *et al.* 2017a), or a particle tracking velocity kick formulation of B_0 can be performed (R.W. Harvey, personal communication 2020; Shiraiwa *et al.* 2011a). Advanced spectral formulations or particle tracking, however, are $\mathcal{O}(10^6)$ more computationally expensive than the traditional spectral D_{ql} reconstruction technique described here or the W-dot method from Jaeger *et al.* (2006). Positive definite

D_{ql} formulations also have unfavourable performance scaling with problem size quickly becoming more expensive than the electric field solve. While more sophisticated D_{ql} formulation techniques may be applicable in less computationally expensive simulations of ion cyclotron heating (Lee *et al.* 2017a), their applicability to spectral LHCD calculations is limited. However, preliminary investigations of positive definite calculations have indicated that they successfully remove interaction with evanescent modes, and ensure a positive definite D_{ql} for LH waves (Frank *et al.* 2019).

Previously, 200–400 parallel velocity grid points were used in iterated simulations of TORLH and CQL3D (Wright *et al.* 2014; Yang *et al.* 2014). However, after the rewrite of the D_{ql} formulation performed in this work, the velocity-space convergence requirements of the D_{ql} were reanalysed. To do this, we performed both standalone Maxwellian validation and iterated simulations with CQL3D. For brevity, we will focus on the standalone validation here. For a given distribution function the RF power density \mathcal{P}_{RF} , calculated from power dissipation by integrated Landau damping on a flux surface $\mathbf{J}_p \cdot \mathbf{E}$ (integrated in ψ in (2.15)), may be related the second moment of the quasilinear diffusion term

$$\mathcal{P}_{\text{RF}}(\psi) = \int \int d^3\mathbf{u} (\gamma - 1) m_e c^2 Q(f). \quad (2.18)$$

Using the CQL3D quasilinear diffusion term (2.6), integrating by parts and substituting a relativistic Maxwellian

$$f = n_e \left(\frac{m_e}{2\pi k_B T_e} \right)^{3/2} \exp \left(\frac{-2c^2(\gamma - 1)}{v_{\text{the}}^2} \right) \quad (2.19)$$

yields

$$\mathcal{P}_{\text{RF}} = \frac{4n_e m_e}{\sqrt{\pi} v_{\text{the}}^5} \oint B_0 dl/B \int_{-\infty}^{\infty} du_{\parallel 0} \int_0^{\infty} du_{\perp 0} u_{\perp 0} \frac{B_0}{\gamma^2} \exp \left(\frac{-2c^2(\gamma - 1)}{v_{\text{the}}^2} \right). \quad (2.20)$$

A comparison of the RF power density profiles obtained from (2.18) allows us to evaluate the convergence of the quasilinear diffusion coefficient formulation in velocity space. As the velocity-space grid resolution is increased the agreement between the two \mathcal{P}_{RF} calculations should improve. Convergence of D_{ql} typically occurs when ~ 500 to 1000 parallel velocity grid points are used. An example of such a convergence scan can be found in figure 3. In more strongly damped, high N_{\parallel} , problems the resolution requirement for D_{ql} tends to be larger as power is concentrated in a smaller region of velocity space. When the D_{ql} is under-resolved, power deposition is over-predicted as poorly defined D_{ql} edges cause excess power to be diffused into the distribution function. In these convergence scans we also studied the impact of the evanescent mode cutoff on power deposition profiles. We found adding or removing the evanescent mode cutoff to have no meaningful effect on the power deposition profiles obtained by integrating the quasilinear term. This is perhaps unsurprising as waves which are cut off according to (2.13) have very short evanescence lengths and should not have the opportunity to deposit significant amounts of power in the plasma (the impact of the cutoff was also examined using spot checks on non-Maxwellian simulations. It was found in the simulations performed in § 3 when the last iterate was re-run without the cutoff there was a 2 % increase in the damped power and 6 % increase in the driven current). Usually, a higher resolution than that which is required for convergence of the Maxwellian problem is used for integrated non-Maxwellian simulations as edge effects will be exacerbated when a Landau plateau begins to form. Throughout this work,

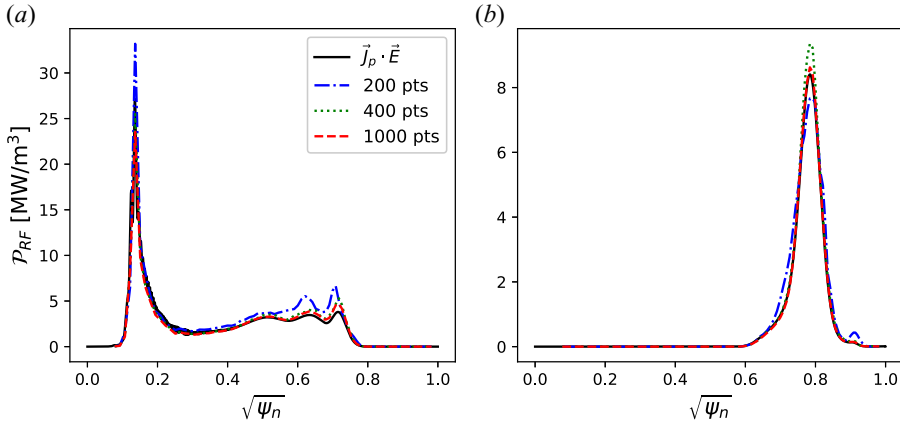


FIGURE 3. Plots of RF power density \mathcal{P}_{RF} vs square-root poloidal flux $\sqrt{\psi_n}$ demonstrating velocity-space convergence of D_{ql} . With increasing numbers of points in parallel velocity space, $\mathcal{N}_{u_{||}}$, convergence improves. Note, for both plots here $\mathcal{N}_{u_{||}} = 2\mathcal{N}_{u_{\perp}}$. In (a) waves are weakly damped in an $N_{||} = -1.6$ simulation, and in (b) waves are strongly single-pass damped in an $N_{||} = -5.0$ simulation.

we use 1000 points in the parallel velocity-space grid as convergence scans showed it consistently provided good convergence for both the non-Maxwellian and Maxwellian problems.

Finally, unlike older versions of TORLH and TORIC (Wright *et al.* 2010, 2014; Lee *et al.* 2017b), the B_0 calculated in TORLH here was passed to CQL3D directly. Exact flux surface matching between the two codes was enforced and a number of intermediate velocity-space interpolations which were present in previous work have been removed. This change markedly improved agreement between the \mathcal{P}_{RF} profiles in CQL3D calculated using the quasilinear diffusion and TORLH calculated using $\vec{J}_p \cdot \vec{E}$ in integrated simulations.

2.2. The $\text{Im}[\chi_{zz}]$ lookup table construction

In order to perform non-Maxwellian simulations of LHCD using TORLH it is necessary to calculate the non-Maxwellian dielectric response. In the case of LH waves this process is straightforward as only the imaginary parallel component of the dielectric $\text{Im}[\epsilon_{||}] = \text{Im}[\chi_{zz,e}]$, which governs Landau damping, undergoes substantial changes (for plateau distributions resulting from Landau damping in the LH limit). However, calculating $\text{Im}[\chi_{zz,e}]$ accurately using realistic fully relativistic distribution functions from CQL3D would be too slow to perform during the dielectric construction in TORLH simulations. Instead, a look up table, similar to those used for complicated dielectrics in AORSA and TORIC (Berry *et al.* 2016; Bertelli *et al.* 2017), is constructed which provides values of $\text{Im}[\chi_{zz,e}]$ on a fixed parallel refractive index $N_{||}$, radial location ψ and flux surface angle θ grid. Then, $\text{Im}[\chi_{zz,e}]$ may be rapidly interpolated from the lookup table to avoid computational and memory bottlenecks during the dielectric construction step in TORLH.

The derivation of $\text{Im}[\chi_{zz,e}]$ in terms of $N_{||}$, ψ and θ for a CQL3D distribution function is performed in Wright *et al.* (2010) and reviewed in Appendix A. However, an overlooked aspect of implementing this $\text{Im}[\chi_{zz}]$ formulation was the method by which the derivatives of the CQL3D distribution were taken. Originally, derivatives of the CQL3D distribution were taken using an uncentred first-order upwind difference. However, verification tests

of TORLH performed during this work revealed, when a $\text{Im}[\chi_{zz,e}]$ produced from a Maxwellian distribution with these derivatives was used in TORLH, poor agreement was obtained with the same simulations using analytic expressions for the Maxwellian $\text{Im}[\chi_{zz,e}]$.

Initially, discrepancies in the Maxwellian regression tests were attributed to the relativistic correction to $\text{Im}[\chi_{zz}]$ s derived numerically from CQL3D distribution functions (the canonical analytic Maxwellian $\text{Im}[\chi_{zz,e}]$ in TORLH using the Z function is non-relativistic), but iterated simulations of TORLH and CQL3D also demonstrated poor agreement between TORLH and CQL3D damped powers. The root cause of the disagreement was found to be a small problem with the derivative method used. In rapidly varying functions, like a plasma distribution function at large u_{\parallel} , it is insufficient to use un-centred differences. Re-centring of the grid at the derivatives' locations is important to obtaining accurate results. This is evident in CQL3D where great care is taken to properly recentre differences (R.W. Harvey, personal communication 2020). To demonstrate this we introduce analytic Maxwellian

$$\frac{\partial F(w)}{\partial w} = -2Cwe^{w^{-2}} \tag{2.21a}$$

$$F(w) = Ce^{w^{-2}}, \tag{2.21b}$$

and a one-dimensional Landau plateau distribution function obtained by solving the Fokker–Planck equation with quasilinear RF diffusion for the Trubnikov collision operator (Trubnikov 1965; Fisch 1978; Karney & Fisch 1979)

$$\frac{\partial F(w)}{\partial w} = -\frac{2F(w)}{w^2(2D(w) + 1/w^3)} \tag{2.22a}$$

$$F(w) = C \exp \left[-\int^w \frac{2w \, dw}{1 + 2w^3 D(w)} \right], \tag{2.22b}$$

for

$$D(w) = \begin{cases} D_0 & w_1 \leq w \leq w_2 \\ 0 & \text{elsewhere.} \end{cases} \tag{2.23}$$

For this choice of a $D(w)$, (2.22) has closed form solution (not previously noted in the literature)

$$F(w) = \begin{cases} C_1 \exp(-w^2) & w < w_1 \\ C_2 \left[\frac{(1 + D_0^{1/3} w)^2}{1 - (2D_0)^{1/3} w + (2D_0)^{2/3} w^2} \right]^{1/(6(2D_0)^{2/3})} & \\ \quad \times \exp \left[-\frac{2\sqrt{3}}{6(2D_0)^{2/3}} \arctan \left(\frac{-1 + 2D_0^{1/3} w}{\sqrt{3}} \right) \right] & w_1 \leq w \leq w_2 \\ C_3 \exp(-w^2) & w_2 < w, \end{cases} \tag{2.24}$$

where C_1 is an arbitrary normalization constant and constants

$$C_2 = C_1 \left[\frac{1 - (2D_0)^{1/3}w_1 + (2D_0)^{2/3}w_1^2}{(1 + D_0^{1/3}w_1)^2} \right]^{1/6(2D_0)^{2/3}} \times \exp \left[-\frac{w_1^2}{2} + \frac{2\sqrt{3}}{6(2D_0)^{2/3}} \arctan \left(\frac{-1 + 2D_0^{1/3}w_1}{\sqrt{3}} \right) \right] \quad (2.25a)$$

$$C_3 = C_2 \left[\frac{(1 + D_0^{1/3}w_2)^2}{1 - (2D_0)^{1/3}w_2 + (2D_0)^{2/3}w_2^2} \right]^{1/6(2D_0)^{2/3}} \times \exp \left[\frac{w_2^2}{2} - \frac{2\sqrt{3}}{6(2D_0)^{2/3}} \arctan \left(\frac{-1 + 2D_0^{1/3}w_2}{\sqrt{3}} \right) \right] \quad (2.25b)$$

Both of these equations (2.21b) and (2.24) have well-defined analytic derivatives (2.21a) and (2.22), except at points on the edges of $D(w)$ where the (2.24) derivatives are undefined. If we ensure our numerical grid does not intersect with the edges of $D(w)$, we can use these equations to precisely evaluate the effectiveness of finite differencing methods. We show a comparison of finite difference methods: first-order uncentred upwind, first-order upwind with grid re-centring and second-order central differences in figure 4. Without re-centring, first-order differences perform extremely poorly, explaining the difficulties experienced with the original $\text{Im}[\chi_{zz,e}]$ calculation. However, when re-centring is applied, first-order upwind differences offer comparable performance to second-order central differences, but unlike higher-order differences, cannot produce the wrong derivative sign. Recentred first-order upwind differences are used in CQL3D, and iterated simulation tests produced the most robust agreement between TORLH and CQL3D when first-order upwind differences were also used in the $\text{Im}[\chi_{zz,e}]$ calculation (even when compared with third- and fourth-order accurate differencing schemes). Presumably, this was the result of increased internal self-consistency. The correction to the derivatives in the $\text{Im}[\chi_{zz,e}]$ calculation immediately improved the agreement between the damped power calculations in TORLH and CQL3D and this correction was key to enabling the work in the following sections.

3. Quasi-steady state simulation

The first tokamak discharge modelled here is Alcator C-Mod shot #1060728011. In this discharge LHCD was produced with the Alcator C-Mod LH1 antenna at 4.6 GHz with 60° phasing corresponding to a launched $N_{\parallel} = -1.6$ (Shiraiwa *et al.* 2011b). Approximately 800 kW of LHCD power was coupled to a plasma with $T_{e0} = 2.3$ keV, $\bar{n}_e = 5 \times 10^{19} \text{ m}^{-3}$ producing a near non-inductive plasma operating at $I_p = 540$ kA with a loop voltage of -0.2 V. This discharge has been modelled extensively in the past using both raytracing (Schmidt 2011) and full-wave simulations (Wright *et al.* 2009, 2010, 2014; Meneghini 2012). Furthermore, this shot is very similar to the shot #1101104011 that has also been subject to much analysis in the literature (Mumgaard 2015; Biswas *et al.* 2020; Baek *et al.* 2021; Biswas *et al.* 2021).

Here, we return to shot #1060728011 using an overhauled TORLH and improved methodology to re-evaluate previous results. As in Frank *et al.* (2022), we have used

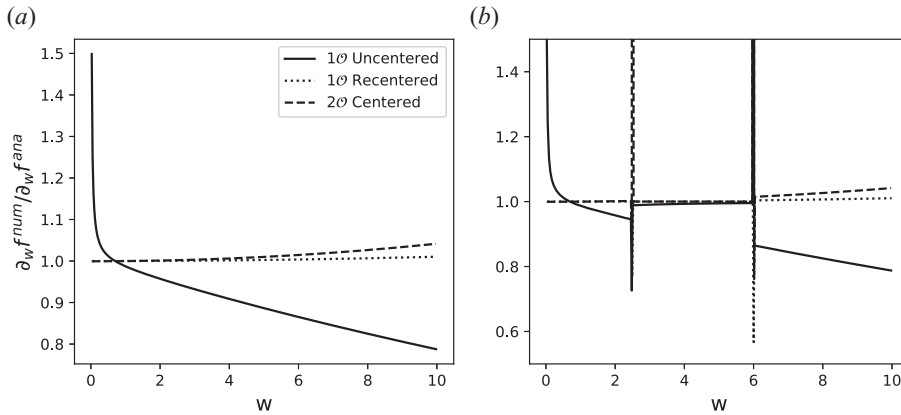


FIGURE 4. The ratio of the analytic and numerical derivatives taken using a first-order upwind forward difference (solid), a re-centred first-order upwind forward difference (dotted) and a second-order central difference (dashed), for (a) the Maxwellian distribution from (2.21) and (b) the non-Maxwellian distribution function from (2.24).

the IPS (Elwasif *et al.* 2010) to create matched raytracing simulations with plasma state files from the full-wave simulation to evaluate the importance of full-wave effects. The GENRAY raytracing simulations here used the same source codes as those in Frank *et al.* (2022) in which the dispersion relation is based on the TORLH dielectric and includes hot plasma corrections to the ϵ_{\parallel} component of the dielectric tensor. The GENRAY simulation of this shot used 400 rays, all with $N_{\parallel} = -1.6$, spread evenly over four grills placed at $\pm 30^{\circ}$ and $\pm 10^{\circ}$ from the outboard midplane. The raytracing simulation results were coupled to a CQL3D simulation that used 60 flux surfaces equi-spaced from $\sqrt{\phi_n} = 0.05$ to 0.95, where ϕ_n is the normalized toroidal flux. The distribution function in the CQL3D simulation was evolved for ~ 125 ms over 35 timesteps. Timestep size was progressively increased from $1\mu\text{s}$ to 10 ms during the simulation for numerical stability (while the time advance of the FP equation in CQL3D is implicit, the updates to the quasilinear diffusion coefficient when CQL3D is coupled with GENRAY are explicit and can become numerically unstable if long timesteps are used initially). In the GENRAY/CQL3D simulation in this case 780 kW of the 800 kW launched power was damped in the plasma.

The TORLH simulations of this discharge were performed using 2047 poloidal modes and 4800 radial finite elements and were run on 255 nodes and 8160 cores on the Cori supercomputer at NERSC. The TORLH resolution settings here follow the requirements detailed in Frank *et al.* (2022), however, higher finite element resolution was needed to combat spectral pollution, which could occur when strong radial discontinuities in the dielectric function for a fixed poloidal mode m appeared due to Landau plateau formation. The boundary condition was matched to the raytracing simulations using the improved boundary condition in Frank *et al.* (2022), and an $N_{\parallel} = -1.6$ was launched. The CQL3D simulations coupled to full-wave simulations used 96 flux surfaces equi-spaced over $\sqrt{\phi_n} = 0.05$ to 0.95. The larger number of flux surfaces in the FP calculation relative to the raytracing/FP simulations was another measure taken to improve numerical stability of the non-Maxwellian TORLH simulations. The TORLH/CQL3D simulations used a ramped timestepping scheme similar to the raytracing/CQL3D simulations to improve numerical stability as TORLH and CQL3D are iterated explicitly within the IPS, as described in § 2. A total of 40 iterations between TORLH and CQL3D were performed here for a total FP time of ~ 90 ms. The TORLH and CQL3D simulation was well converged, TORLH

power matched CQL3D power with discrepancies of $\lesssim 10\%$ without the use of numerical rescaling factors (the TORLH power target was 800 kW and the total damped power in CQL3D was 719 kW).

Comparisons of the \mathcal{P}_{RF} profiles from TORLH/CQL3D and GENRAY/CQL3D simulations, shown in [figure 5](#), demonstrate that excellent power deposition agreement was obtained in the non-Maxwellian simulations, much like in the Maxwellian simulations of this discharge performed in Frank *et al.* (2022). Power deposition profiles at the final iteration calculated by CQL3D from integration of the quasilinear diffusion term for both GENRAY and TORLH simulations closely agreed with one another and agreed with the integrated $J_p \cdot E$ calculated by the TORLH field solver. Despite excellent \mathcal{P}_{RF} agreement, TORLH/CQL3D simulations had lower current drive efficiency than GENRAY/CQL3D simulations. Current profiles in both simulations closely matched the power deposition profiles, but the TORLH/CQL3D simulations drove 288 kA of current with 719 kW of input power compared with the GENRAY/CQL3D simulations which drove 455 kA of current with 800 kW of input power. The shortfall in the TORLH/CQL3D current drive efficiency, which is only $\sim 70\%$ of the GENRAY/CQL3D value, appears to be the result of holes in the quasilinear diffusion coefficient introduced by the non-positive definite nature of the coefficient, discussed in § 2.1. These holes limit the power which can be diffused to higher-energy electrons. The discontinuities from the holes in the quasilinear diffusion coefficient cause a sharp drop in the distribution about the hole. This seems to largely be a two-dimensional effect as the holes do not substantially affect the RF power deposition or current drive profiles, and the power deposition profile is known to be driven primarily by the one-dimensional distribution function behaviour (Bonoli & Englade 1986; Shiraiwa *et al.* 2011a; Meneghini 2012). However, the quasilinear diffusion holes substantially reduce the net driven current which is sensitive to changes to the two-dimensional distribution (Fisch & Boozer 1980). This is demonstrated in [figure 6](#), where it is shown that, despite having very similar current drive profiles, GENRAY/CQL3D simulations drive more current than TORLH/CQL3D simulations, and power is more effectively diffused to higher-energy electrons in GENRAY/CQL3D simulations. We found processing the quasilinear diffusion coefficient produced by TORLH with a biharmonic in-painting algorithm (Chui & Mhaskar 2010; Lasiecka, Damelin & Hoang 2018) to remove the holes from non-positive definiteness increased the current drive efficiency substantially. However, in-painting also spoiled the power calculation's self-consistency and led to an unstable iteration between TORLH and CQL3D that could not be used.

Unlike previous non-Maxwellian simulations of this discharge in Wright *et al.* (2014) there was not as drastic a drop in the non-Maxwellian core electric fields at the first and the last iterations of the TORLH/CQL3D simulation, as shown in [figure 7](#). The drop observed by Wright *et al.* (2014) was due to the physically inconsistent nature of their calculations and an error in the χ_{zz} table interpolation that led to wave damping being over-predicted. Despite having increased field amplitudes in our simulation, constructive interference does not seem to cause substantial differences in the power deposition profiles vs GENRAY. This has a simple explanation; the damped power's dependence on the slope of the distribution, $\propto \exp(-\zeta^2)$ for a Maxwellian distribution or $\propto \partial f / \partial v_{\parallel}$ for a non-Maxwellian distribution, is a much stronger driver of power deposition location in ψ than the damped power's dependence on field intensity $\propto |E_{\parallel}|^2$. Thus, the power deposition profile should only be weakly affected by interference in most cases. Our result here demonstrates that the GENRAY simulations of weakly damped LHCD, which have been found to differ from experiment (Schmidt 2011; Wright *et al.* 2014; Mumgaard 2015), accurately reproduce the power deposition profiles calculated with full-wave simulations. This




FIGURE 5. The RF power density \mathcal{P}_{RF} vs square-root poloidal flux $\sqrt{\psi_N}$ at the final iteration of a TORLH/CQL3D (blue) and a GENRAY/CQL3D (orange) simulation of Alcator C-Mod shot #1060728011. Lines corresponding to the RF power density from integration of the CQL3D quasilinear diffusion term are solid and lines corresponding to RF power density obtained from integration of the TORLH $J_p \cdot E$ are dashed.

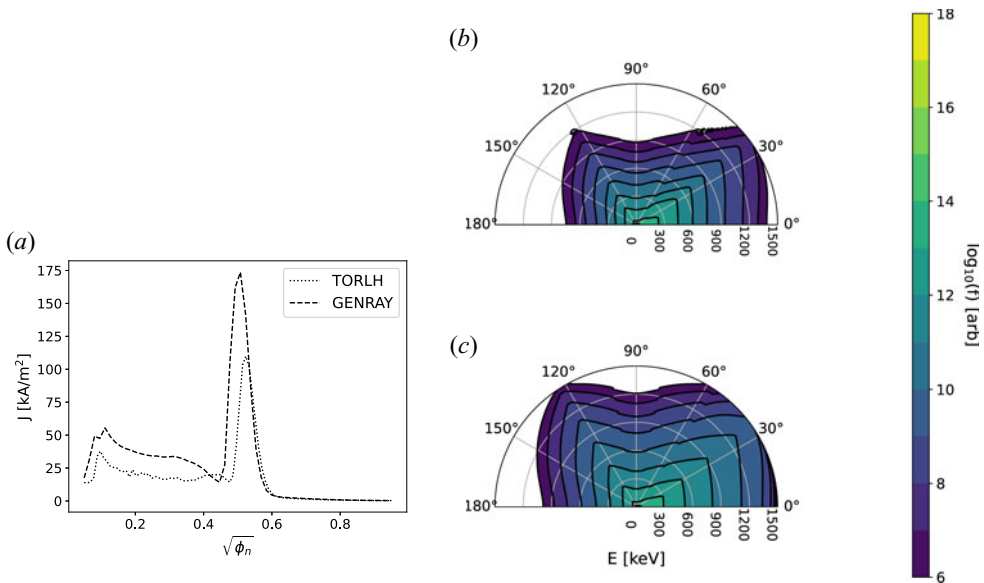


FIGURE 6. (a) The RF driven current density J vs normalized square-root toroidal flux $\sqrt{\phi_n}$ in a TORLH/CQL3D simulation (dotted) and a GENRAY/CQL3D simulation (dashed). Contours of the electron distribution function f at $\sqrt{\phi_n} \sim 0.5$ vs E in keV, and velocity-space pitch angle ϑ , for (b) a TORLH/CQL3D simulation and (c) a GENRAY/CQL3D simulation. Both panels (b,c) use the same contour levels to more effectively show the differences in RF diffusion of power to electrons with high energies.




FIGURE 7. The logarithmic contours of the normalized TORLH electric field (launched field in TORLH is normalized to 1 V m^{-1}) in simulations of Alcator C-Mod shot #1060728011 for Maxwellian damping (a), and non-Maxwellian damping (b).

indicates that the discrepancy between raytracing/FP results and experimentally measured current profiles is not due to full-wave effects such as diffraction or interference. Much like raytracing/FP, full-wave/FP simulations cannot accurately replicate the experimentally measured ohmic-like LHCD profiles which are peaked on axis.

4. Modulation scan simulations

The next set of Alcator C-Mod discharges modelled here were the LHCD modulation experiments performed by Schmidt *et al.* (2011). In these experiments, 400 kW of LHCD was modulated with a 50 % duty cycle and a pulse time of 12.5 ms, in discharges that were primarily ohmic ($V_{\text{loop}} = 1.0 \text{ V}$) at high density $\bar{n} = 9.0 \times 10^{19} \text{ m}^{-3}$, with $T_{e0} \sim 2.2 \text{ keV}$. Three different launcher phasings 60° , 90° and 120° , corresponding to $N_{\parallel} = -1.6$, -2.3 and -3.1 , respectively, were used in these experiments. The GENRAY/CQL3D modelling was performed in the original analysis in Schmidt *et al.* (2011) and Schmidt (2011) and compared with inverted experimental hard X-ray measurements, with mixed results. Qualitative agreement with experiment was obtained when $N_{\parallel} = -2.3$ and -3.1 , but agreement with experiment was very poor when $N_{\parallel} = -1.6$.

We repeated the analysis of these discharges with GENRAY/CQL3D and matched TORLH/CQL3D simulations using the IPS. The GENRAY simulations used 200 rays spread over 4 launch points to mimic the Alcator C-Mod LH1 launcher, the same custom dielectric described in § 3, and used a single fixed N_{\parallel} value (except when otherwise noted). CQL3D in the GENRAY/CQL3D simulations used 48 equi-spaced flux surfaces and took 50, 0.25 ms timesteps for a total FP time of 12.5 ms (equivalent to the LHCD pulse time in the modulation experiments). The TORLH simulations used 2047 poloidal modes and 6000 radial finite elements and were run using 8160 cores and 255 nodes on Cori at NERSC in iterated simulations that lasted $\sim 36 \text{ h}$. The larger finite element number than the simulations in § 3 eliminated spectral pollution which occurred in the edge of the $N_{\parallel} = -2.3$ and -3.1 simulations. The CQL3D settings in the TORLH/CQL3D integrated simulations were the same as those in § 3, but CQL3D was iterated with TORLH only 25 times in order to achieve a FP time of 12.55 ms. All the simulations here used 400 kW of launched power like the experiments. The integrated $\mathbf{J}_p \cdot \mathbf{E}$ and power profiles obtained in




FIGURE 8. The RF power density \mathcal{P}_{RF} vs square-root toroidal flux $\sqrt{\phi}$ calculated by CQL3D coupled with GENRAY (dotted) or TORLH (dashed) for 3 different Alcator C-Mod discharges using different LHCD launch phasings.

CQL3D from integration of the quasilinear diffusion term closely agreed in all cases. Furthermore, the integrated power in all CQL3D simulations was within 2% of the 400 kW target, meaning the degree of self-consistency achieved in these simulations was higher than any previous full-wave FP modelling study (Jaeger *et al.* 2006; Shiraiwa *et al.* 2011a; Wright *et al.* 2014; Lee *et al.* 2017a,b; Bertelli *et al.* 2017).

The RF power deposition profiles obtained from these simulations, shown in figure 8, found once again that GENRAY/CQL3D accurately reproduced the power deposition profiles calculated using TORLH/CQL3D. However, in the $N_{\parallel} = -3.1$ case GENRAY/CQL3D initially predicted a narrower power deposition peak. To obtain a close RF power deposition profile match with TORLH/CQL3D we had to broaden the launched GENRAY spectrum. This difference almost certainly was due to the influence of diffraction in the TORLH simulation, as diffraction is expected to broaden the spectrum symmetrically (Pereverzev 1992). A scan of ΔN_{\parallel} values from 0.1 to 0.6 was performed and a GENRAY launcher spectrum with $\Delta N_{\parallel} = 0.4$ was found to be the best match to the TORLH results. A comparison of the unbroadened and broadened damping is found in figure 9. The spectral broadening needed to account for diffraction is substantially less than the spectral width of the LH1 launcher in C-Mod where $\Delta N_{\parallel} = 1.0$ (Shiraiwa *et al.* 2011b), indicating that the effects of diffractive broadening are relatively small. Furthermore, in the more weakly damped $N_{\parallel} = -2.3$ and -3.1 cases no spectral broadening in GENRAY was needed to obtain good agreement with TORLH. This indicates that the geometric upshift effect (Bonoli & Ott 1981, 1982), included in both simulations, is the dominant gap closure mechanism. Diffraction is only a small correction and relatively unimportant that can impart only modest broadening on the wave spectrum. Importantly, our results agreed well with the GENRAY/CQL3D results in Schmidt *et al.* (2011), and neither the TORLH/CQL3D or GENRAY/CQL3D simulations agreed with experimental measurements in the weak damping, $N_{\parallel} = -1.6$, case where ohmic-like profiles were measured (Schmidt *et al.* 2011).

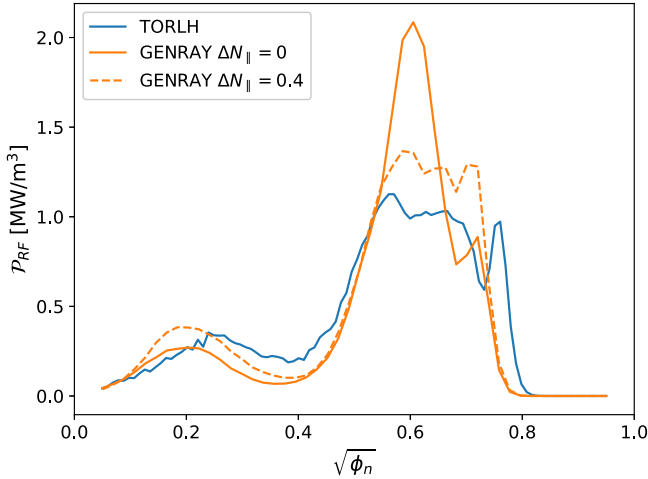


FIGURE 9. The RF power density \mathcal{P}_{RF} vs square-root toroidal flux $\sqrt{\phi}$ calculated by CQL3D coupled with GENRAY (blue) or TORLH (orange) for two different values of GENRAY spectral broadness $\Delta N_{\parallel} = 0$ (solid) and $\Delta N_{\parallel} = 0.4$ (dashed).

Finally, current drive calculations were performed in each case. As in § 3, we observed GENRAY/CQL3D predicted total driven currents fairly close to experiments, but the TORLH/CQL3D simulations systematically under-predicted the current drive as a result of the holes present in the quasilinear diffusion coefficient. Current drive profiles in both the GENRAY/CQL3D and TORLH/CQL3D simulations closely followed the power deposition profiles, however, the TORLH/CQL3D current drive efficiency was again roughly $\sim 70\%$ the value obtained in by the GENRAY/CQL3D simulations.

5. Conclusion and discussion

We performed the first fully self-consistent TORLH/CQL3D simulations of LHCD and demonstrated that full-wave effects, such as diffraction and interference, only weakly affect LH wave power deposition and current drive profiles. Diffractive broadening was quantified and found to be much smaller than the spectral broadening from finite launch spectrum width. In fact, diffraction was only noticeable in situations where the wave was damped in 2–3 passes and in weaker damping the LH spectral gap closure was completely dominated by the geometric upshift effect. Thus, it was determined full-wave effects cannot account for the differences between simulations of LHCD by raytracing/FP models and experiments, and despite breakdowns in the raytracing approximation that can occur at cutoffs and caustics, raytracing/FP simulations robustly reproduced full-wave/FP results. The primary discrepancy between raytracing and full-wave simulations, the reduced current drive efficiency observed in full-wave simulation, may be explained by systemic error introduced from the non-positive definiteness of the full-wave quasilinear diffusion coefficient. This presents an obvious topic for future work, but as positive definite formulations of the quasilinear diffusion coefficient are $> \mathcal{O}(10^6)$ more expensive than the standard formulation of D_{ql} used here (Lee *et al.* 2017a), they were deemed too computationally demanding for implementation in the LHCD problem. Another complication of implementing such a form of the quasilinear diffusion is that it may be necessary to perform a correction for the parallel magnetic field inhomogeneity to $\text{Im}[\chi_{zz}]$ to ensure self-consistency (if the inhomogeneous correction is physically relevant outside of ensuring D_{ql} positive definiteness). However, if positive definite forms of D_{ql} could be

implemented they would have the added benefit of self-consistently including evanescent modes we have removed in the simulations here as they do not model resonant damping with a delta function.

Our TORLH/CQL3D results are the first full-wave LHCD simulations to demonstrate both precise RF power deposition and integrated power agreement between the full-wave electric field and the FP simulation results without the use of numerical power rescaling factors. To do this many of the non-Maxwellian components in the TORLH/CQL3D simulations were rewritten with a focus on improving self-consistency. We implemented noise reduction and fixed errors in the quasilinear diffusion coefficient formulation, improved the $\text{Im}[\chi_{zz}]$ formulation in TORLH to use f derivatives matched to those in CQL3D, implemented a robust iteration framework using the IPS and improved TORLH's performance so that it could be run at higher resolution to suppress spectral pollution. This allowed us to perform the very large, 8000+ cores for 36+ h, simulations here. In total, this study used ~ 100 million total CPU hours over 4 years on NERSC.

In summary, full-wave effect influences on LHCD power and current deposition profiles appear to be small. Our simulations indicate raytracing/FP simulations accurately reproduce LH wave propagation and damping in the core of tokamaks, but in some cases with stronger damping there may be a small correction related to diffraction that will usually be smaller than other corrections, such as finite launcher spectrum width. If the LHCD is moderately damped, and a launcher with a very small spectral width is used or very accurate simulations are required, an effective way to account for diffractive broadening might be the use of a higher-order accurate WKB method such as beamtracing that captures diffraction (Bertelli *et al.* 2012; Poli *et al.* 2018; Hugon, Bizarro & Rodrigues 2020).

Acknowledgements

The authors would like to thank D.B. Batchelor for his help with the IPS.

Editor Cary Forest thanks the referees for their advice in evaluating this article.

Funding

This work was supported by Scientific Discovery Through Advanced Computing (SCIDAC) Contract No. DE-SC0018090 and Department of Energy grant: DE-FG02-91ER54109. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award FES-ERCAP0020035.

Declaration of interests

The authors report no conflict of interest.

Appendix A. The $\text{Im}[\chi_{zz}]$ formulation from CQL3D plasma distributions

To begin our derivation of a non-Maxwellian $\text{Im}[\chi_{zz,e}]$ appropriate for use in TORLH we take the $n = 0$ zz components of the relativistic hot plasma dielectric for a general distribution function from Stix (1992)

$$\chi_{zz,e} = 2\pi \frac{\omega_{pe,0}^2}{\omega \Omega_{e,0}} \int_0^\infty dp_\perp p_\perp \int_{-\infty}^\infty p_\parallel dp_\parallel \left(\frac{\Omega_e}{\omega - k_\parallel v_\parallel} \right) J_0^2(k_\perp v_\perp / \Omega_e) \frac{\partial f_e}{\partial p_\parallel}, \quad (\text{A1})$$

where relativistic momentum $\mathbf{p} = \gamma m_0 \mathbf{v}$ and we have denoted quantities which use the rest mass, m_0 , with subscript 0 and quantities that utilize the relativistic mass γm_0 without a subscript, i.e. $\Omega = \Omega_0/\gamma$. Integrating in p_{\parallel} and taking the imaginary component from the residue of the pole in the resonant denominator yields

$$\text{Im}[\chi_{zz,e}] = -2\pi^2 m_0 \frac{\omega_{pe,0}^2}{k_{\parallel}^2} \int_0^{\infty} dp_{\perp} p_{\perp} J_0^2 \left(\frac{k_{\perp} p_{\perp}}{m_0 \Omega_{e,0}} \right) \frac{\partial p_{\parallel}}{\partial v_{\parallel}} \frac{\partial f_e}{\partial p_{\parallel}} \Big|_{v_{\parallel}=\omega/k_{\parallel}}. \tag{A2}$$

We want to create a lookup table in TORLH that is a function of radial location ψ , flux surface angle θ and N_{\parallel} using derivatives of the CQL3D distribution function at the outboard midplane $F_0(\psi, u, \vartheta)$. We transform (A2) into a more convenient form in terms of CQL3D and TORLH variables with the following relations:

$$u_n = p/m_0c = \sqrt{u_{\parallel n}^2 + u_{\perp n}^2} = \sqrt{\frac{N_{\parallel}^2 u_{\perp}^2 + 1}{N_{\parallel}^2 - 1}} \tag{A3a}$$

$$\gamma = \sqrt{1 + u_n^2} \tag{A3b}$$

$$u_{\parallel n} = \sigma_{\pm} \left(\frac{1 + u_{\perp n}^2}{N_{\parallel}^2 - 1} \right)^{1/2} \tag{A3c}$$

$$u_{\perp n} = p_{\perp}/m_0c \tag{A3d}$$

$$\frac{\partial p_{\parallel}}{\partial v_{\parallel}} = \frac{m_0 \gamma^3}{1 + u_{\perp}^2}, \tag{A3e}$$

where u_n is the normalized momentum per rest mass, and $\sigma_{\pm} = \text{Sign}(\omega/k_{\parallel})$. Using these equations we may rewrite (A2) in terms of our normalized table variables

$$\text{Im}[\chi_{zz}](\psi, \theta, n_{\parallel}) = -\alpha \int_0^{\infty} du_{\perp}^2 J_0^2 \left(\frac{k_{\perp} u_{\perp} c}{\Omega_{e,0}} \right) (1 + u_{\perp}^2)^{1/2} \frac{\partial f_e}{\partial p_{\parallel}} \Big|_{v_{\parallel}=\omega/k_{\parallel}} \tag{A4a}$$

$$\alpha = \pi^2 \frac{\omega_{pe,0}^2}{\omega^2} (m_{e,0}c)^4 \frac{N_{\parallel}}{(N_{\parallel}^2 - 1)^{3/2}}. \tag{A4b}$$

Finally, we must rewrite the derivative of the distribution function $\partial f/\partial p_{\parallel}$ in terms of the distribution function at the outboard midplane that is output by CQL3D this requires that we relate $f(p_{\parallel}, p_{\perp}, \theta, \psi)$ to $F_0(u, \vartheta_0(\vartheta, \theta), \psi)$, where ϑ is the velocity-space pitch angle. These quantities may be related by applying conservation of momentum and magnetic moment yielding equation

$$\frac{\partial f}{\partial p_{\parallel}} = \frac{\cos \vartheta}{m_{e,0}} \left(\frac{\partial F_0}{\partial u} - \frac{\tan \vartheta_0}{u} \frac{\partial F_0}{\partial \vartheta_0} \right). \tag{A5}$$

With this our derivation of χ_{zz} is complete. To construct the lookup table a grid in θ and N_{\parallel} must be specified (the ψ grid is set automatically based on the CQL3D flux surface locations), and an integration grid in u_{\perp} must also be specified. Then, χ_{zz} is calculated from the CQL3D distribution function at each grid point. In order to find $\partial f/\partial p_{\parallel}$ the nearest neighbour point on the CQL3D grid point corresponding to the Landau resonance location at a given u_{\perp} and u_{\parallel} is used (no noticeable improvement in self-consistency was observed with more complicated interpolations).

REFERENCES

- BAEK, S.G., BISWAS, B., WALLACE, G.M., BONOLI, P.T., DING, B.J., LI, M.H., LI, Y.C., WANG, Y.F., WANG, M., WU, C.B., *et al.* 2021 A model investigation of the impact of lower hybrid wave scattering angle on current drive profile in EAST and Alcator C-Mod. *Nucl. Fusion* **61** (10), 106034.
- BARTIROMO, R., HESSE, M., SÖLDNER, F.X., BURHENN, R., FUSSMANN, G., LEUTERER, F., MURMANN, H., ECKHARTT, D., EBERHAGEN, A., GIULIANA, A., *et al.* 1986 Experimental study of non-thermal electron generation by lower hybrid waves in the ASDEX tokamak. *Nucl. Fusion* **26** (8), 1106.
- BERNABEI, S., DAUGHNEY, C., EFTHIMION, P., HOOKE, W., HOSEA, J., JOBES, F., MARTIN, A., MAZZUCATO, E., MESERVEY, E., MOTLEY, R., *et al.* 1995 Lower-hybrid current drive in the PLT tokamak. *Phys. Rev. Lett.* **49** (17), 1255.
- BERRY, L.A., JAEGER, E.F., PHILLIPS, C.K., LAU, C.H., BERTELLI, N. & GREEN, D.L. 2016 A generalized plasma dispersion function for electron damping in tokamak plasmas. *Phys. Plasmas* **23** (10), 102504.
- BERTELLI, N., MAJ, O., POLI, E., HARVEY, R., WRIGHT, J.C., BONOLI, P.T., PHILLIPS, C.K., SMIRNOV, A.P., VALEO, E. & WILSON, J.R. 2012 Paraxial Wentzel–Kramers–Brillouin method applied to the lower hybrid wave propagation. *Phys. Plasmas* **19** (8), 082510.
- BERTELLI, N., VALEO, E.J., GREEN, D.L., GORELENKOVA, M., PHILLIPS, C.K., PODESTÀ, M., LEE, J.P., WRIGHT, J.C. & JAEGER, E.F. 2017 Full-wave simulations of ICRF heating regimes in toroidal plasma with non-Maxwellian distribution functions. *Nucl. Fusion* **57** (5), 056035.
- BISWAS, B., BAEK, S.G., BONOLI, P.T., SHIRAIWA, S., WALLACE, G. & WHITE, A. 2020 Study of turbulence-induced refraction of lower hybrid waves using synthetic scrape-off layer filaments. *Plasma Phys. Control. Fusion* **62** (11), 115006.
- BISWAS, B., SHIRAIWA, S., BAEK, S.-G., BONOLI, P., RAM, A. & WHITE, A.E. 2021 A hybrid full-wave Markov chain approach to calculating radio-frequency wave scattering from scrape-off layer filaments. *J. Plasma Phys.* **87** (5), 905870510.
- BONOLI, P.T. 1985 Linear theory of lower hybrid waves in tokamak plasmas. In *Wave Heating and Current Drive in Plasmas* (ed. V.L. Granatstein & P.L. Colestock). Gordon and Breach Science Publishers.
- BONOLI, P.T. 2014 Review of recent experimental and modeling progress in the lower hybrid range of frequencies at ITER relevant parameters. *Phys. Plasmas* **21** (6), 061508.
- BONOLI, P.T. & ENGLADE, R.C. 1986 Simulation model for lower hybrid current drive. *Phys. Fluids* **29**, 2937.
- BONOLI, P.T. & OTT, E. 1981 Accessibility and energy deposition of lower-hybrid waves in a tokamak with density fluctuations. *Phys. Rev. Lett.* **46** (6), 424.
- BONOLI, P.T. & OTT, E. 1982 Toroidal and scattering effects on lower-hybrid wave propagation. *Phys. Fluids* **25** (2), 359.
- BRAMBILLA, M. 1996 A full wave code for ion cyclotron wave in toroidal plasmas. *Tech. Rep.* IPP 5/66. Max-Planck-Institute for Plasma Physics.
- CESARIO, R., AMICUCCI, L., CASTALDO, M., MARINUCCI, M., PANACCIONE, L., SANTINI, F., TUDISCO, O., APICELLA, M.L., CALABRO, G., CIANFARANI, C., *et al.* 2010 Current drive at plasma densities required for thermonuclear reactors. *Nat. Commun.* **1** (1), 55.
- CESARIO, R., CARDINALI, A., CASTALDO, C., PAOLETTI, F. & MAZON, D. 2004 Modeling of a lower-hybrid current drive by including spectral broadening induced by parametric instability in tokamak plasmas. *Phys. Rev. Lett.* **92** (17), 175002.
- CHUI, C.K. & MHASKAR, H.N. 2010 MRA contextual-recovery extension of smooth functions on manifolds. *Appl. Comput. Harmon. Anal.* **28** (1), 104–113.
- DECKER, J., PEYSSON, Y., ARTAUD, J.-F., NILSSON, E., EKEDAHL, A., GONICHE, M., HILLAIRET, J. & MAZON, D. 2014 Damping of lower hybrid waves in large spectral gap configurations. *Phys. Plasmas* **21** (9), 092504.
- ELWASIF, W.R., BERNHOLDT, D.E., SHET, A.G., FOLEY, S.S., BRAMLEY, R., BATCHELOR, D.B. & BERRY, L.A. 2010 The design and implementation of the SWIM integrated plasma simulator. In *2010 18th Euromicro Conference on Parallel, Distributed and Network-Based Processing*, pp. 419–427. IEEE.

- FISCH, N.J. 1978 Confining a tokamak plasma with rf-driven currents. *Phys. Rev. Lett.* **42** (6), 410.
- FISCH, N.J. 1987 Theory of current drive in plasmas. *Rev. Mod. Phys.* **59**, 175–234.
- FISCH, N.J. & BOOZER, A.H. 1980 Creating an asymmetric plasma resistivity with waves. *Phys. Rev. Lett.* **45** (9), 720.
- FRANK, S., BONOLI, P.T., BATCHELOR, D.B., WRIGHT, J.C. & HUTCHINSON, I.H. 2019 Coupling of lower-hybrid full wave and 3D Fokker–Planck codes in weak damping scenarios. In *APS Division of Plasma Physics Meeting Abstracts*, p. UP10.069. APS.
- FRANK, S.J., WRIGHT, J.C., HUTCHINSON, I.H. & BONOLI, P.T. 2022 An assessment of full-wave effects on Maxwellian lower-hybrid wave damping. *Plasma Phys. Control. Fusion* **64** (10), 105023.
- FRIED, B.D. & CONTE, S.C. 1961 *The Plasma Dispersion Function*. Academic.
- HARVEY, R.W. & MCCOY, M.G. 1992 The CQL3D code. In *Proceedings of the IAEA TCM on Advances in Sim. and Modeling of Thermonuclear Plasmas Montreal*, pp. 489–526. IOP Publishing.
- HUGON, H., BIZARRO, J.P.S. & RODRIGUES, P. 2020 An improved asymptotic matching technique to trace the wave amplitude of rays across singularities: application to lower-hybrid wave propagation in tokamaks. *Phys. Plasmas* **27** (8), 082508.
- IDE, S., IMAI, T., USHIGUSA, K., NAITO, O., IKEDA, Y., NEMOTO, M. & SATO, M. 1992 Investigation of the wave spectrum gap in the JT-60 LHCD plasma. *Nucl. Fusion* **32** (2), 282.
- JACQUINOT, J. 1991 Heating, current drive and confinement regimes with the JET ICRH and LHCD systems. *Plasma Phys. Control. Fusion* **33** (13), 1657.
- JAEGER, E.F., BERRY, L.A., AHERN, S.D., BARRETT, R.F., BATCHELOR, D.B., CARTER, M.D., D'AZEVEDO, E.F., MOORE, R.D., HARVEY, R.W., MYRA, J.R., *et al.* 2006 Self-consistent full-wave and Fokker–Planck calculations for ion cyclotron heating in non-Maxwellian plasmas. *Phys. Plasmas* **13** (5), 056101.
- KARNEY, C.F.F. & FISCH, N.J. 1979 Numerical studies of current generation by radio-frequency traveling waves. *Phys. Fluids* **22** (9), 1817–1824.
- KAUFMAN, A.N. 1972 Quasilinear diffusion of an axisymmetric toroidal plasma. *Phys. Fluids* **15** (6), 1063–1069.
- KENNEL, C.F. & ENGELMANN, F. 1966 Velocity space diffusion from weak plasma turbulence in a magnetic field. *Phys. Fluids* **9** (12), 2377–2388.
- LASIECKA, I., DAMELIN, S.B. & HOANG, N.S. 2018 On surface completion and image inpainting by biharmonic functions: numerical aspects. *Int'l J. Math. Sci.* **2018**, 3950312.
- LEE, J., SMITHE, D., WRIGHT, J. & BONOLI, P. 2017a A positive-definite form of bounce-averaged quasilinear velocity diffusion for the parallel inhomogeneity in a tokamak. *Plasma Phys. Control. Fusion* **60** (2), 025007.
- LEE, J.P. & WRIGHT, J.C. 2014 A block-tridiagonal solver with two level parallelization for finite element-spectral codes. *Comput. Phys. Commun.* **185**, 2598.
- LEE, J., WRIGHT, J., BERTELLI, N., JAEGER, E.F., VALEO, E., HARVEY, R. & BONOLI, P. 2017b Quasilinear diffusion coefficients in a finite Larmor radius expansion for ion cyclotron heated plasmas. *Phys. Plasmas* **24** (5), 052502.
- LERCHE, I. 1968 Quasilinear theory of resonant diffusion in a magneto-active, relativistic plasma. *Phys. Fluids* **11** (8), 1720–1727.
- LIU, F.K., DING, B.J., LI, J.G., WAN, B.N., SHAN, J.F., WANG, M., LIU, L., ZHAO, L.M., LI, M.H., LI, Y.C., *et al.* 2015 First results of LHCD experiments with 4.6 GHz system toward steady-state plasma in EAST. *Nucl. Fusion* **55** (12), 123022.
- MENEGHINI, O. 2012 Full-wave modeling of lower hybrid waves on Alcator C-Mod. PhD thesis, Massachusetts Institute of Technology.
- MORIYAMA, S., NAKAMURA, Y., NAGAO, A., NAKAMURA, K., HIRAKI, N. & ITOH, S. 1990 Ultra-long pulse operation using lower hybrid waves on the superconducting high field tokamak TRIAM-1M. *Nucl. Fusion* **30** (1), 47.
- MUMGAARD, R.T. 2015 Lower hybrid current drive on Alcator C-Mod: measurements with an upgraded MSE diagnostic and comparisons to simulation. PhD thesis, Massachusetts Institute of Technology.
- PEREVERZEV, G.V. 1992 Use of the multidimensional WKB method to describe propagation of lower hybrid waves in tokamak plasma. *Nucl. Fusion* **32** (7), 1091–1106.

- PETROV, YU.V. & HARVEY, R.W. 2016 A fully-neoclassical finite-orbit-width version of the CQL3D Fokker–Planck code. *Plasma Phys. Control. Fusion* **58** (11), 115001.
- PEYSSON, Y., DECKER, J., NILSSON, E., ARTAUD, J.-F., EKEDAHN, A., GONICHE, M., HILLAIRET, J., DING, B., LI, M., BONOLI, P.T., *et al.* 2016 Advances in modeling of lower hybrid current drive. *Plasma Phys. Control. Fusion* **58** (4), 044008.
- PEYSSON, Y., SÉBELIN, E., LITAUDON, X., MOREAU, D., MIELLOU, J.-C., SHOUCRI, M.M. & SHKAROFSKY, I.P. 1998 Full wave modelling of lower hybrid current drive in tokamaks. *Nucl. Fusion* **38** (6), 939–944.
- PEYSSON, Y. & TORE SUPRA TEAM 2001 High power lower hybrid current drive experiments in the Tore Supra tokamak. *Nucl. Fusion* **41** (11), 1703.
- POLI, E., BOCK, A., LOCHBRUNNER, M., MAJ, O., REICH, M., SNICKER, A., STEGMEIR, A., VOLPE, F., BERTELLI, N., BILATO, R., *et al.* 2018 Torbeam 2.0, a paraxial beam tracing code for electron-cyclotron beams in fusion plasmas for extended physics applications. *Comput. Phys. Commun.* **225**, 36–46.
- PORKOLAB, M., SCHUSS, J.J., LLOYD, B., TAKASE, Y., TEXTER, S., BONOLI, P., FIORE, C., GANDY, R., GWINN, D., LIPSCHULTZ, B., *et al.* 1984 Observation of lower-hybrid current drive at high densities in the Alcator C tokamak. *Phys. Rev. Lett.* **53** (5), 450–453.
- SCHMIDT, A. 2011 Measurements and modeling of lower hybrid driven fast electrons on Alcator C-Mod. PhD thesis, Massachusetts Institute of Technology.
- SCHMIDT, A., BONOLI, P.T., MENEGHINI, O., PARKER, R.R., PORKOLAB, M., SHIRAIWA, S., WALLACE, G., WRIGHT, J.C., HARVEY, R.W. & WILSON, J.R. 2011 Investigation of lower hybrid physics through power modulation experiments on Alcator C-Mod. *Phys. Plasmas* **18** (5), 056122.
- SHIRAIWA, S., KO, J., PARKER, R., SCHMIDT, A.E., SCOTT, S., GREENWALD, M., HUBBARD, A.E., HUGHES, J., MA, Y., PODPALY, Y., *et al.* 2011a Full wave effects on lower hybrid wave spectrum and driven current profile in tokamak plasmas. *Phys. Plasmas* **18**, 080705.
- SHIRAIWA, S., MENEGHINI, O., PARKER, R.R., WALLACE, G., WILSON, J., FAUST, I., LAU, C., MUMGAARD, R., SCOTT, S., WUKITCH, S., *et al.* 2011b Design, and initial experiment results of a novel LH launcher on Alcator C-Mod. *Nucl. Fusion* **51** (10), 103024.
- STIX, T.H. 1992 *Waves in Plasmas*, 2nd edn. Springer.
- TRUBNIKOV, B.A. 1965 Particle interactions in a fully ionized plasma. *Rev. Plasma Phys.* **1**, 105.
- WALLACE, G.M., HUBBARD, A.E., BONOLI, P.T., FAUST, I.C., HARVEY, R.W., HUGHES, J.W., LABOMBARD, B.L., MENEGHINI, O., PARKER, R.R., SCHMIDT, A.E., *et al.* 2011 Lower hybrid current drive at high density in Alcator C-Mod. *Nucl. Fusion* **51** (8), 083032.
- WILSON, J.R., PARKER, R., BITTER, M., BONOLI, P.T., FIORE, C., HARVEY, R.W., HILL, K., HUBBARD, A.E., HUGHES, J.W., INCE-CUSHMAN, A., *et al.* 2009 Lower hybrid heating and current drive on the Alcator C-Mod tokamak. *Nucl. Fusion* **49** (11), 115015.
- WRIGHT, J.C., BADER, A., BERRY, L.A., BONOLI, P.T., HARVEY, R.W., JAEGER, E.F., LEE, J.P., SCHMIDT, A., D'AZEVEDO, E., FAUST, I., *et al.* 2014 Time dependent evolution of RF-generated non-thermal particle distributions in fusion plasmas. *Plasma Phys. Control. Fusion* **56** (4), 045007.
- WRIGHT, J.C., BONOLI, P.T., SCHMIDT, A.E., PHILIPS, C.K., VALEO, E.J., HARVEY, R.W. & BRAMBILLA, M.A. 2009 An assessment of full wave effects on the propagation and absorption of lower hybrid waves. *Phys. Plasmas* **16** (7), 072502.
- WRIGHT, J.C., LEE, J.P., VALEO, E., BONOLI, P., PHILIPS, C.K., JAEGER, E.F. & HARVEY, R.W. 2010 Challenges in self-consistent full-wave simulations of lower hybrid waves. *IEEE Trans. Plasma Sci.* **38** (9), 2136.
- YANG, C., BONOLI, P.T., WRIGHT, J.C., DING, B.J., PARKER, R., SHIRAIWA, S. & LI, M.H. 2014 Modelling of the EAST lower-hybrid current drive experiment using GENRAY/CQL3D and TORLH/CQL3D. *Plasma Phys. Control. Fusion* **56** (12), 125003.