

3D full-wave computation of RF modes in magnetised plasmas

Pavel Aleynikov*, Nikolai B. Marushchenko

Max-Planck-Institut für Plasmaphysik, Greifswald, Germany

Abstract

A new 3D full-wave code, named CUWA, is developed to investigate the physics of RF wave propagation in the electron cyclotron frequency range in magnetised plasmas. The code utilises the Finite Difference Time Domain (FDTD) technique and takes advantage of massive parallel computations with Graphics Processing Units (GPU), which allows for a significant acceleration of the computations. As examples of code application we show 3D calculations of the linear transformation of ordinary to extraordinary electron-cyclotron waves and mode coupling in a sheared magnetic field. Thanks to its speed, the GPU-capable code allows for efficient and large parametric scans over a broad range of parameters.

Keywords: FDTD, full-wave, mode conversion, stellarator, OXB,

1. Introduction

RF waves play an important role in plasma physics. In tokamaks and stellarators they are used for heating, current drive and for diagnostic purposes. A large spectrum of problems is connected with Electron Cyclotron (EC) wave propagation in inhomogeneous plasmas. Geometrical optics, also known as the Wentzel-Kramers-Brillouin (WKB) approximation, is a standard tool for studying wave propagation in weakly inhomogeneous plasmas. However a variety of problems exist which fall outside the applicability of the WKB approximation: linear transformation of modes (i.e. mode conversion) [1, 2, 3], mode coupling in strongly sheared magnetic fields [4, 5], cross-polarisation scattering by electro-magnetic fluctuations [6, 7], propagation and reflection of waves in turbulent plasmas [8, 9, 10]. These problems should be considered within a full wave approach.

This paper reports on the development of a new 3D cold plasma full wave code, named CUWA. The code utilises the Finite Difference Time Domain (FDTD) technique [11, 12]. The computation domain is “minimised” around the WKB-trajectory obtained from ray-tracing calculations by means of the convolutional perfectly matched layer (CPML) technique [13]; the background magnetic field is recovered from pre-computed 3D equilibrium data. The code takes advantage of massive parallel computations with Graphics Processing Units (GPUs), which allows for up to $\times 100$ acceleration over a single-CPU. CUWA contains several diagnostic and post-processing tools, including windowed Fast Fourier Transform (wFFT) of the resulting snapshots of the wave

*Corresponding author

Email address: pavel.aleynikov@ipp.mpg.de (Pavel Aleynikov)

fields, which allows for the direct comparison of the full-wave computations with WKB-theory whenever appropriate.

2. Numerical model

2.1. Discretisation

FDTD (also known as Yee’s method) is a standard technique for full-wave simulations in various branches of physics [11, 12]. In fusion plasmas, it has been successfully applied to wave propagation, O-X conversion [14], reflectometry [8] and the scattering of EC waves on plasma turbulence [15].

In this work we mostly follow the FDTD method as given in Chapter 11.3 of [12], with two important modifications: first, we alter the current discretisation scheme and, second, we amend the differential operators of Maxwell’s equations in the boundary layers in accordance with the CPML technique [13].

The system of equations being solved by the code is Maxwell’s equations with the addition of a “cold plasma” response current (\mathbf{J}) equation given by the electron law of motion:

$$\begin{aligned}\frac{\partial}{\partial t}\mathbf{B} &= -\nabla \times \mathbf{E}, \\ \frac{\partial}{\partial t}\mathbf{E} &= c^2\nabla \times \mathbf{B} - \mathbf{J}/\varepsilon_0, \\ \frac{d}{dt}\mathbf{J} + \nu\mathbf{J} &= \varepsilon_0\omega_p^2\mathbf{E} - \mathbf{J} \times \boldsymbol{\omega}_c,\end{aligned}\tag{1}$$

where $\omega_p^2 = n_e e^2 / \varepsilon_0 m_e$ is the electron plasma frequency, ν is the electron collision frequency, $\boldsymbol{\omega}_c = |e|\mathbf{B}_0/m_e$ is the electron cyclotron frequency (vector) corresponding to the background magnetic field \mathbf{B}_0 .

In the FDTD method the field components are discretised on staggered grids in space and time. Note that it is common to discretise the plasma-response current, \mathbf{J} , in such a way that its J_x , J_y and J_z components are co-located in space [12, 14]. This facilitates the current update calculation. However, we find the stability of this scheme to be unsatisfactory, in particular when applied to the CPML region. Therefore we have implemented a slightly more computationally demanding but more stable scheme where \mathbf{J} is discretised in the same way as the wave electric field, \mathbf{E} . This implies that the current update equation involves interpolation of its components [16].

In our implementation the wave magnetic field \mathbf{B} and the plasma response current \mathbf{J} are computed at half-integer timesteps $t = \frac{1}{2}, \frac{3}{2}, \dots, n + \frac{1}{2}$, while the electric field \mathbf{E} is computed at integer timesteps $t = 0, 1, \dots, n + 1$. Such an explicit $\mathbf{E} - \mathbf{J}$ scheme is known to be unstable in some extreme cases and an alternative implementation with implicit stepping therefore exists [17, 18]. However, our explicit $\mathbf{E} - \mathbf{J}$ coupling in combination with the co-location of the corresponding \mathbf{E} and \mathbf{J} components has satisfactory stability in relevant problems, while imposing tolerable overheads in a GPU-specific code. Stabilisation of the explicit $\mathbf{E} - \mathbf{J}$ coupling by modification of the discretisation of the vector product is discussed in Ref. [19].

After the discretisation according to Yee’s method and taking into account CPML alterations (i.e. virtual currents Ψ and stretch coefficients κ - see Section 2.2 for details) the x -components

of the first two equations of the system (1) take the following discrete form:

$$\begin{aligned}
H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} &= H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} \\
&+ C \left(\frac{E_y|_{i,j+\frac{1}{2},k+1}^n - E_y|_{i,j+\frac{1}{2},k}^n}{\kappa_z} - \frac{E_z|_{i,j+1,k+\frac{1}{2}}^n - E_z|_{i,j,k+\frac{1}{2}}^n}{\kappa_y} \right) \\
&- \Psi_{H_{x,y}}|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} + \Psi_{H_{x,z}}|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}}, \\
\end{aligned} \tag{2}$$

$$\begin{aligned}
E_x|_{i+\frac{1}{2},j,k}^{n+1} &= E_x|_{i+\frac{1}{2},j,k}^n \\
&- C \left(\frac{H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - H_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}}}{\kappa_z} - \frac{H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - H_z|_{i+\frac{1}{2},j-\frac{1}{2},k}^{n+\frac{1}{2}}}{\kappa_y} \right) \\
&- J_x|_{i+\frac{1}{2},j,k}^{n+\frac{1}{2}} + \Psi_{E_{x,y}}|_{i+\frac{1}{2},j,k}^{n+1} - \Psi_{E_{x,z}}|_{i+\frac{1}{2},j,k}^{n+1},
\end{aligned}$$

where the electric field has been rescaled as $E \equiv c\varepsilon_0 E$, magnetic field is $H \equiv B/\mu_0$, and the current $J \equiv cJ\Delta t$, with c denoting the speed of light. The time step takes the form of the Courant-Friedrichs-Lewy (CFL) parameter $C \equiv c\frac{\Delta t}{\Delta x}$, where $\Delta x = \Delta y = \Delta z$ is the cell size. In all examples below we use $C = 0.5$. In Eq. (2) and below we adopt index notation, which is typical for the FDTD formulation, i.e. in $H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}}$ the first index x denotes the vector component, the upper index after the vertical bar ($|^{n+\frac{1}{2}}$) denotes the time step and the lower indices after the bar ($|_{i,j+\frac{1}{2},k+\frac{1}{2}}$) denote the location.

The discretized exact solution of the last equation of the system (1) reads

$$\begin{aligned}
\mathbf{J}|_{i+\frac{1}{2},j,k}^{n+\frac{1}{2}} &= \exp(\mathbf{\Omega}\Delta t)\mathbf{J}|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}} \\
&+ \omega_p^2\Delta t\mathbf{\Omega}^{-1}(\exp(\mathbf{\Omega}\Delta t) - \mathbf{I})\mathbf{E}|_{i+\frac{1}{2},j,k}^n,
\end{aligned} \tag{3}$$

where

$$\mathbf{\Omega} \equiv \begin{bmatrix} -\nu & -\omega_{cz} & \omega_{cy} \\ \omega_{cz} & -\nu & -\omega_{cx} \\ -\omega_{cy} & \omega_{cx} & -\nu \end{bmatrix}.$$

The fully expanded form of these matrix exponentials can be found in Refs. [12, 15]. Note that Eq. (11.36a) in the book [12] has two typos: the signs of the terms proportional to S_1 in the yz and zx matrix elements need to be reversed.

Only the J_x component needs to be updated in the location $(i + \frac{1}{2}, j, k)$, yet the right-hand side of Eq. (3) involves all three components of $\mathbf{J}|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}}$ and $\mathbf{E}|_{i+\frac{1}{2},j,k}^n$ vectors. Therefore their y and z components need to be interpolated from a four neighbouring cells, specifically

$$\begin{aligned}
J_y|_{i+\frac{1}{2},j,k} &= \frac{J_y|_{i,j-\frac{1}{2},k} + J_y|_{i+1,j-\frac{1}{2},k} + J_y|_{i,j+\frac{1}{2},k} + J_y|_{i+1,j+\frac{1}{2},k}}{4}, \\
J_z|_{i+\frac{1}{2},j,k} &= \frac{J_z|_{i,j,k-\frac{1}{2}} + J_z|_{i+1,j,k-\frac{1}{2}} + J_z|_{i,j,k+\frac{1}{2}} + J_z|_{i+1,j,k+\frac{1}{2}}}{4}.
\end{aligned} \tag{4}$$

The interpolation for the electric field components is analogous.

2.2. Perfectly Matched Layer

The minimisation of the computation domain plays a critical role in efficient 3D computations. A physical absorber layer with a non-zero collisionality ($\nu > 0$) is often used in plasma physics computations to truncate the computation domain [14, 15]. With this technique, one would need to amend the computation domain with layers of “many-wavelengths” thickness at each boundary ensuring that any reflected signals are small. This may become problematic in 3D calculations where the domain of interest itself has the size of a few wavelengths. Amending such a domain with six boundary layers may significantly reduce the overall efficiency.

A so-called Convolutional Perfectly Matched Layer (CPML) boundary [13] is instead implemented in our full-wave code to truncate the computation domain. An efficient CPML requires only a few extra Yee cells in each direction - i.e. a fraction of a single “wavelength”. CPML also requires two auxiliary 3D fields (Ψ) for each component of \mathbf{E} and \mathbf{B} , i.e. 12 auxiliary fields in total. However, these auxiliary fields are defined only within the narrow CPML layers, resulting in a negligible overhead in a practical setup.

The conceptual idea behind CPML [13] can be summarised in the following way (further details and a rather complete overview of various PMLs can be found in the book [20]). CPML is an efficient numerical implementation of Complex Frequency Shifted PML [21], in which the differential operators in Eqs. (1) are replaced with the convolution¹:

$$\frac{\partial}{\partial u} \rightarrow \mathcal{L}^{-1} \left\{ \frac{1}{\kappa_u + \frac{\sigma_u}{\alpha_u + i\omega\epsilon_0}} \right\} * \frac{\partial}{\partial u}, \quad (5)$$

where $\mathcal{L}^{-1}\{F(\omega)\}$ denotes the inverse Laplace transform, κ_u and α_u are free parameters accounting for coordinate stretching in the real and complex plains, σ_u is an “effective conductivity”, and $u = x, y, z$. The PML given by Eq. (5) introduces no reflections because it can be viewed as a vacuum with complex-stretched coordinates. In other words, with Eq. (5) the solution of Eq. (1) is *analytically* continued into the complex plane with a subsequent mapping of the complex coordinates back into real space but with a “complex material”.

The implementation of the convolution Eq. (5) would be prohibitively expensive computationally in time domain had it not been for the development of an efficient implementation in Ref. [13]. In the time domain, the inverse Laplace transformed operator from Eq. (5) takes the form of a sum of the unit impulse function $\delta(t)/\kappa_u$ and an exponential decay $\sim e^{-t}h(t)$, where $h(t)$ is the step function. The convolution with the δ -function is trivial (hence the factor $1/\kappa_u$ in Eqs. (2)), while the convolution with an exponential decay $\psi(t) = \int_0^t e^{-a(t-t')} f(t') dt'$ can be accurately found using the following recursive approximation:

$$\psi(t + \Delta t) = e^{-a\Delta t} \psi(t) + \frac{1 - e^{-a\Delta t}}{a} f(t + \Delta t/2), \quad (6)$$

where the next half-step value of the convoluted function $f(t + \Delta t/2)$, namely $\nabla \times \mathbf{B}(t)$ and $\nabla \times \mathbf{E}(t)$, is available in our case thanks to the nature of the leapfrog time-stepping adopted in FDTD algorithm.

¹In the frequency domain this operation takes the form of a coordinate “stretching” by the factor $\kappa_u + \frac{\sigma_u}{\alpha_u + i\omega\epsilon_0}$.

From Eqs. (5) and (6) the CPML auxiliary fields Ψ used in Eq. (2) for the magnetic field update equation become

$$\begin{aligned}\Psi_{H_{x,y}}|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} &= b_y|_{j+\frac{1}{2}} \Psi_{H_{x,y}}|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} \\ &\quad + a_y|_{j+\frac{1}{2}} C(E_z|_{i,j+1,k+\frac{1}{2}}^n - E_z|_{i,j,k+\frac{1}{2}}^n), \\ \Psi_{H_{x,z}}|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} &= b_z|_{k+\frac{1}{2}} \Psi_{H_{x,z}}|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} \\ &\quad + a_z|_{k+\frac{1}{2}} C(E_y|_{i,j+\frac{1}{2},k+1}^n - E_y|_{i,j+\frac{1}{2},k}^n),\end{aligned}\tag{7}$$

where $b_u = e^{-\sigma_u/\kappa_u + \alpha_u}$ and $a_u = \frac{\sigma_u}{\kappa_u(\sigma_u + \kappa_u\sigma_u)}(b_u - 1)$. The best performance is achieved when the coefficients κ_u , α_u , σ_u vary smoothly within the PML when the boundary is approached:

$$\begin{aligned}\sigma_x(x) &= \sigma^{\max}(x/d)^m, \\ \kappa_x(x) &= 1 + (\kappa^{\max} - 1)(x/d)^m, \\ \alpha_x(x) &= \alpha^{\max}(x/d)^{m_1}.\end{aligned}\tag{8}$$

where d is the PML thickness, and x is the distance within the PML. Refs. [11, 13] give a prediction of the optimal value for $\sigma^{\max} \approx 0.7(m+1)C$. Note that the graded coefficients in Eq. (8) needs to be computed at locations corresponding to the definitions of the fields they apply to, i.e. some field components are defined at integer and some at half-integer steps, and the distance x in Eq. (8) needs to be calculated accordingly.

The electric field auxiliary fields $\Psi_{E_{x,y}}$ and $\Psi_{E_{x,z}}$ are updated analogously to Eqs. (7).

2.3. CUWA code structure

The CUWA code uses GPUs to compute the fields update according to Eqs. (2)-(4) and (7). Here we give some details about the general architecture of the code.

The computational setup, including an initialisation of parameters, ray-tracing calculation and reading of the equilibrium data (see Section 2.4) is done on a CPU using the `python` language. Then, the prepared data, namely 3D arrays of plasma frequency $\omega_p(i, j, k)$ and the components of the cyclotron frequency $\omega_{cx}(i, j, k)$, $\omega_{cy}(i, j, k)$ and $\omega_{cz}(i, j, k)$ are copied from RAM to the GPU memory.² Note that this data needs to be defined in the same locations as the current components (see Eq. (3)), i.e. in $(i + \frac{1}{2}, j, k)$, in $(i, j + \frac{1}{2}, k)$ and in $(i, j, k + \frac{1}{2})$. But, in practice, because the background parameters vary slowly, we use a single set of data for updating all three current components or interpolate between the neighbouring nodes using an equation analogous to (4).

The core field-updating routines are written using the CUDA platform [22, 23]. After the initialisation, they are executed in a loop in the following order:

1. Source: $(E_{x,y,z}(t)$ at one of the boundaries)
2. Update $H_{x,y,z}$ using Eq. (2)
3. Copy $J_{x,y,z}$ to $J_{x,y,z}^0$

²Currently CUWA supports only single-GPU calculations, therefore the computation domain is limited by the memory size of the GPU.

4. Update $J_{x,y,z}$ using Eq. (3), where $J_{x,y,z}^{n-\frac{1}{2}} \equiv J_{x,y,z}^0$
5. Update $E_{x,y,z}$ using Eq. (2)

The PML fields (Ψ) are updated using Eq. (7) immediately before the corresponding wave fields. When the iteration stops, the wave fields are copied from the GPU memory to the RAM for post-processing.

The diagnostics implemented in the CUDA GPU code include the calculation of the time-averaged Poynting vector, $\langle \mathbf{E} \times \mathbf{B} \rangle$, and a 2D windowed Fourier Transform (wFFT) of the cross sections of the wave fields. In a homogeneous magnetic field, FFT allows for a direct comparison of the full-wave calculations with the predictions of WKB theory. The 1D sinusoidal window is moved along one of the axis to generate the resulting “dispersion curves”. The batch FFT (containing window-scaled data) is computed on GPUs using the `scikit-cuda` library [24].

2.4. Computation setup

In modern magnetic confinement fusion devices, such as a tokamaks or stellarators, the magnetic field is typically of the order of several tesla. Specifically, in Wendelstein 7-X (W7-X) the magnetic field is 2.5 T, meaning that the vacuum wavelength required for electron cyclotron resonance heating (ECRH) at the second harmonics is approximately 2 mm. This wavelength is much shorter than the size of the plasma, which is of order 1 m. The typical RF beam radius is several cm, which is also much smaller than the plasma size. This disparity between the wave and plasma sizes suggests that the computational domain could be reduced to some small sub-volume of the plasma where the wave is present. The full discretisation of the entire plasma volume with a sub-mm size grid is not feasible.

In order to construct a 3D computational domain which contains the beam and accounts for realistic background plasma parameters (3D magnetic field and density) a preliminary estimate of the beam trajectory in the plasma is needed. This estimate is obtained using the standard ray tracing technique (see, e.g. Ref. [25]). In the WKB approximation Eqs. (1) reduce to the following wave equation [26]

$$D_{\alpha\beta}E_\beta = 0, \quad D_{\alpha\beta} = N_\alpha N_\beta - N^2 \delta_{\alpha\beta} + \varepsilon_{\alpha\beta}, \quad (9)$$

where $\varepsilon_{\alpha\beta}(\omega)$ is the “cold plasma” dielectric tensor, $N_\alpha = k_\alpha c/\omega$ is the refractive index, k_α is the wave vector, and E_β is the polarisation vector. Then, a wave packet can be viewed as a particle with Hamiltonian $\omega(\mathbf{k}, \mathbf{r})$, where the wave-vector \mathbf{k} and the wave packet position \mathbf{r} represent the canonical momentum and spatial coordinate, respectively. Its equations of motion are

$$\begin{aligned} \frac{d\mathbf{k}}{dt} &= -\frac{\partial\omega(\mathbf{k}, \mathbf{r})}{\partial\mathbf{r}}, \\ \frac{d\mathbf{r}}{dt} &= \frac{\partial\omega(\mathbf{k}, \mathbf{r})}{\partial\mathbf{k}}, \end{aligned} \quad (10)$$

where the Hamiltonian is implicitly given by the dispersion relation $\det D_{\alpha\beta} = 0$.

The initial value problem for ray tracing equations (10) is solved using the `scipy.integrate.solve_ivp` interface [27], which provides access to several efficient integration methods, including the Runge-Kutta, the backward differentiation and the Adams methods³. The integration is stopped when the trajectory crosses the last closed flux surface (and the group velocity is directed outward from the plasma).

³A detailed description of some other ray tracing codes can be found in Refs. [28, 29, 30, 31]

CUWA uses pre-computed magnetic equilibrium data, reconstructed using `MConf` library in the same way it is done in TRAVIS code [30] (see Section 4 therein for the details). The `MConf` library supports various formats for the equilibrium data, including 3D VMEC equilibrium data [32] and axisymmetric equilibrium data in EQDSK format [33].

Based on the reference ray trajectory, $\mathbf{R}(s)$, the computational domain can be constructed in various ways depending on the problem. We found that it is often “efficient” to construct a computational box in which the (x, y) plane contains the plasma entry and exit points, while the y basis vector (\mathbf{e}_y) coincides with the ray direction at the entry point, i.e. $\mathbf{e}_y \approx \mathbf{k}c/\omega$.

The wave polarisation, required for the initialisation of the beam, can be determined from Eq. (9). However, a typical CUWA application setup involves launching a wave from the outside of the plasma (i.e. from an external antenna). In this case, direct application of Eq. (9) in the limit of vanishing plasma density is problematic due to degeneracy. Instead, the limiting polarisation of the O- and X-modes entering (or leaving) the plasma can be calculated from the wave equation (9) by expanding it in the small parameter $\hat{\omega}_p^2 = \omega_p^2/\omega^2$ (see, e.g. Ref. [34]). In the Stix reference frame, where $\mathbf{B}_0 = (0, 0, B_0)$ and $\mathbf{N} = (N_\perp, 0, N_\parallel)$, the complex wave amplitude then becomes

$$\begin{aligned} \frac{E_x}{E_z} &= -\frac{N_\parallel}{N_{\perp 0}} \left(1 + \hat{\omega}_p^2 \frac{1 + F_+ - 2\hat{\omega}_c^2/F_+}{2N_{\perp 0}^2(1 - \hat{\omega}_c^2)} \right), \\ \frac{E_y}{E_z} &= i \frac{\hat{\omega}_c}{F_+} \frac{E_x}{E_z}, \end{aligned} \quad (11)$$

for the O-mode polarisation with $E_z \neq 0$ and

$$\begin{aligned} \frac{E_x}{E_y} &= -\frac{i\hat{\omega}_c}{F_- - N_{\perp 0}^2\hat{\omega}_c^2} \left[N_\parallel^2 + \hat{\omega}_p^2 \frac{(\hat{\omega}_c^2 - F_-)(F_- - N_{\perp 0}^2\hat{\omega}_c^2 - N_\parallel^2)}{(1 - \hat{\omega}_c^2)(F_- - N_{\perp 0}^2\hat{\omega}_c^2)} \right], \\ \frac{E_z}{E_y} &= i \frac{\hat{\omega}_c N_\parallel N_{\perp 0}}{F_- - N_{\perp 0}^2\hat{\omega}_c^2} \left[1 + \frac{\hat{\omega}_p^2}{N_{\perp 0}^2(1 - \hat{\omega}_c^2)} \left(\frac{1 + F_-}{2} - \frac{N_\parallel^2 F_-}{F_- - N_{\perp 0}^2\hat{\omega}_c^2} \right) \right], \end{aligned} \quad (12)$$

for the X-mode polarisation with $E_y \neq 0$. Here, $\hat{\omega}_c = \omega_c/\omega$ and

$$F_\pm = \frac{\hat{\omega}_c^2}{2} N_{\perp 0}^2 \left(1 \pm \sqrt{1 + \frac{4N_\parallel^2}{\hat{\omega}_c^2 N_{\perp 0}^4}} \right), \quad (13)$$

with the upper sign (+) corresponding to the O-mode and the lower (-) to the X-mode. All values are taken at the plasma/vacuum boundary, in particular, $N_{\perp 0}^2 = 1 - N_\parallel^2$ and $\hat{\omega}_p^2 \ll 1$. Note also that Eqs. (11)-(12) are derived in the geometrical optics limit, therefore they cannot be applied to cases with a very sharp plasma boundary. A full-wave treatment is required to calculate the limiting polarisation in such case [35].

The spatial distribution of the launched wave field is usually prescribed to have the form of a Gaussian beam [36],

$$\mathbf{E}(\rho, s) = \mathbf{E}_0 \frac{w_0}{w(s)} \exp \left(-\frac{\rho^2}{w^2(s)} - i \left(ks + \frac{k\rho^2}{2R(s)} - \varphi_G(s) \right) \right), \quad (14)$$

where ρ is the radial distance from the centre axis of the beam, s is the axial distance from the beam’s focus (or “waist”, which has radius w_0), $R(s) = s(1 + z_R^2/s^2)$ is the curvature of the

wavefront, $w(s) = w_0\sqrt{1 + s^2/z_R^2}$ is the beam radius, $z_R = \pi w_0^2/\lambda$ is the Rayleigh length, $k = \omega/c$ and $\varphi_G(s) = \arctan(s/z_R)$ is the Gouy phase. The wave polarisation vector \mathbf{E}_0 in the reference frame of the computational box is obtained from Eqs. (11)-(12) after transforming from the Stix reference frame.

Explicitly, the real part of Eq. (14) is prescribed at the first plane after the CPML.

3. Tests and benchmarks against analytical solutions

In this section, we demonstrate CPML performance and show a comparison of our full-wave calculations with analytical solutions in some analytically tractable cases.

3.1. Demonstration of CPML performance

In order to demonstrate CPML performance for the FDTD realisation described above we design a benchmark which represents a generic anisotropic plasma scenario. In this scenario, a divergent Gaussian beam propagating through the plasma escapes from the computation domain through a strongly nonuniform boundary. The entire computational domain is surrounded by a narrow CPML layer. This is shown in Fig. 1, where the wave is launched vertically from the origin $(0, 0)$. The background magnetic field is normal to the plane (x, y) and the plasma density gradient is directed diagonally (bottom right to top left) with the density ranging from 0 to approximately 80% of the O-mode cut-off. The PML parameters are $m_1 = m = 3$, $\alpha^{\max} = 0.5$, $\kappa^{\max} = 5$. It can be seen that at the top PML layer the wave-front incident angle varies from almost normal to quite oblique, ensuring extended benchmarking. The resulting reflection error is estimated from the maximum value of the electric field amplitude recorded in the bottom-right quarter of the domain. Note that the maximum relative error reported in the original work Ref. [13] is $\approx 5 \cdot 10^{-4}$. Similar reflection is observed in our benchmark case (a wave with the cyan contours $[10^{-4}; 10^{-3}]$). We found that the numerical stability in PML can be greatly improved when increasing the α^{\max} parameter at the price that the reflection error is somewhat increased. Note the presence of another wave propagating out through the top right corner (con-




Figure 1: A snapshot of the decimal logarithm of the wave E -field magnitude for a CPML benchmark case. Background magnetic field is normal to the plane (x, y) and plasma density gradient is directed diagonally (bottom right to top left) with the density range from 0 to approximately 80% of the O-mode cut-off (some contours of constant density are shown with dashed curves). The PML parameters are: $m_1 = m = 3$, $\alpha^{\max} = 0.5$, $\kappa^{\max} = 5$.

tours $[10^{-5}; 10^{-4}]$). This spurious wave is associated with an inaccuracy of the Gaussian beam setup at the bottom boundary. Frequently, CPML outperforms the initial condition accuracy (for instance in a typical case where PML lies in vacuum) so that the latter becomes the dominant noise source (i.e. introduces $\sim 10^{-4}$ error in field amplitude). Note that the relative errors are squared when the intensity or mode power are evaluated. The above example is calculated with the resolution of 12 Yee-cells per vacuum wavelength and a CPML size of 12 cells.

3.2. Reflection from the cutoff layer

In the simplest 1D benchmark we consider an isotropic linear layer without absorption. The plasma density increases linearly along the y axis, so that the permittivity vanishes at the location y_1 , i.e. $\varepsilon(y) = 1 - y/y_1$. The wave launched from vacuum ($y = 0$) must be reflected from the layer $\varepsilon(y) = 0$. The *exact* analytical solution of the wave equation for the electric field in this case takes the form (see §17 in [35]):

$$\begin{aligned} E(\zeta) &= A\zeta^{\frac{1}{2}} \left[J_{\frac{1}{3}} \left(\frac{2}{3}\zeta^{\frac{3}{2}} \right) + J_{-\frac{1}{3}} \left(\frac{2}{3}\zeta^{\frac{3}{2}} \right) \right] \quad \text{for } \zeta > 0, \\ E(\zeta) &= A(-\zeta)^{\frac{1}{2}} \left[I_{-\frac{1}{3}} \left(\frac{2}{3}(-\zeta)^{\frac{3}{2}} \right) - I_{\frac{1}{3}} \left(\frac{2}{3}(-\zeta)^{\frac{3}{2}} \right) \right] \quad \text{for } \zeta < 0, \end{aligned} \quad (15)$$

where A is a constant, J_n and I_n are ordinary and modified Bessel functions of the first kind, respectively, and the new variable $\zeta \equiv (\omega y_1/c)^{\frac{2}{3}} \varepsilon(y)$.

In the numerical calculation we use 12 cells per wavelength. The length of the domain is 480 cells in y direction⁴. A 10-cell CPML layer is located in the $y < 0$ region. The resulting steady-state electric field distribution, which is achieved after approximately 5000 time-steps, is shown with markers in Fig. 2 (top). Equation 15 is plotted as a solid curve in the same plot. The vertical grey line marks the y_1 location and separates the region where $\zeta, \varepsilon > 0$ from the evanescent layer $\zeta, \varepsilon < 0$. The bottom panel shows the error of the calculation, i.e. the distance between the roots of $E(y) = 0$ obtained from the numerical and analytical solutions. In this setup the error does not exceed the size of the computational cell.

3.3. Interaction of the electromagnetic and plasma waves

Some numerical techniques may experience difficulties when resolving behaviour of a wave near resonances. Such situations also require special treatment when studied analytically. In order to demonstrate that our full-wave code provides well-behaved solutions around resonances we calculate plasma wave excitation by the oblique reflection of a wave from a linear isotropic plasma layer with collisional damping given for relatively small reflection angles. The permittivity in this case is

$$\varepsilon \approx 1 - \frac{y}{y_1} \left(1 + i \frac{\nu}{\omega} \right). \quad (16)$$

This setup is illustrated in Fig. 3, where the calculated electric field is colour-coded and the red curve shows the ray-tracing calculation. In the geometrical optics limit the ray turning point location (indicated with the green line) is given by the relation $y_r = y_1(1 - \sin^2 \theta_0)$, where θ_0 is the wave incidence angle. Therefore, for relatively small θ_0 , such that $|y_1 - y_r| \lesssim \lambda$, a finite evanescent wave, which propagates behind the geometrical reflection point y_r , can excite

⁴Technically this calculation is performed in 2D setup with a very wide Gaussian beam, since our code does not currently support 1D calculations.




Figure 2: Top: Steady-state distribution of the wave electric field (dots) in 1D isotropic plasma with $\varepsilon(y) = 1 - y/y_1$, where $y_1 = 0.0595\text{m}$ (marked with the vertical grey line). The wave frequency is $\omega/2\pi = 140\text{ GHz}$. Analytical solution, Eq. (15), is plotted with the solid curve. Bottom: The distance between the roots of $E(y) = 0$ obtained from numerical and analytical solutions.

plasma (Langmuir) oscillations near the resonance location y_1 (indicated with the dashed line in Fig. 3). Note that electrostatic charge separation near the resonance is due to the electric field component along the density gradient, therefore no singularity arises if $E_y = 0$. Consequently, the considered wave is linearly polarised in the $x - y$ plane.




Figure 3: A snapshot of the electric field amplitude (colour-coded) and the ray tracing trajectory (red curve) for the test case with the permittivity given by Eq. (16). The green line represents y_r and the dashed line represents y_1 .

The problem of wave reflection and generation of Langmuir oscillations near the singularity point, $\varepsilon(y_1) = 0$, can be treated analytically in the limit of rather large incidence angles (see §20 in [35] and references therein). This is the limit where the distance between y_1 and y_r is much greater than the wavelength. Figure 4 shows the squared electric field component E_y^2 around the resonance region computed with our full-wave code (yellow) compared with the approximate analytical solution (green) for $\nu/\omega = 3.4 \times 10^{-4}$ and two incident angles: $\sin \theta_0 = 0.34$, i.e. $\theta_0 \approx 20^\circ$ (a) and $\sin \theta_0 = 0.2$, i.e. $\theta_0 \approx 11.5^\circ$ (b). A grey vertical line marks the geometrical optics reflection point y_r . The analytical solution predicts an exponential (in $\sin^3 \theta_0$) scaling for $|E_y| \sim e^{-\sin^3 \theta_0}/\nu$, which is in agreement with our calculations. In this test the resolution had to be increased to 24 cells per vacuum wavelength in order to resolve the narrow resonance peak. Because the energy flow to the resonance is exponentially slow it takes a very long time for the peak to saturate at the chosen collisional frequency, namely approximately 50000 time-steps

required in the cases of Fig. 4.




Figure 4: Squared electric field component E_y^2 around the resonance region computed with CUWA code (yellow) compared with the approximate analytical solution from [35] (green) for $\nu/\omega = 3.4 \times 10^{-4}$ and two incident angles: $\sin \theta_0 = 0.34$ (a) and $\sin \theta_0 = 0.2$ (b). Grey vertical line marks the geometrical optics reflection point y_r .

4. O-X mode conversion

One of the premier applications of our code is the calculation of the linear transformation of the ordinary (O) to the extraordinary (X) mode in 3D plasmas.

Electron Cyclotron Resonance Heating (ECRH) is the main plasma heating mechanism in the Wendelstein 7-X (W7-X) stellarator [37]. The X- and O-modes are used in a wide range of operation scenarios. Whereas X2 (frequency corresponding to the 2nd electron cyclotron harmonic) has a cut-off at the density $1.2 \times 10^{20} \text{m}^{-3}$, the O2-mode can be used in plasmas with densities up to $2 \times 10^{20} \text{m}^{-3}$ (the cut-off density is $2.4 \times 10^{20} \text{m}^{-3}$). Possible operation at densities above the cut-off would involve double mode-conversion from O- to slow X- and to Bernstein-mode, i.e. an OXB-scenario [38]. The efficiency of such heating schemes is defined by the O-X conversion efficiency. The physics of O-X conversion is outside the applicability of the routinely used geometrical optics approximation and should be considered within a full-wave approach.

Linear O-X transformation happens when the O-mode cutoff and the X-mode “turning surface” are close to each other. At the intersection of these two surfaces WKB approximation predicts a perfect conversion [1, 39, 2] wave, i.e. a wave with an optimal refractive index, $N_{||}^{opt}$, for given plasma parameters. In reality, however, 3D shaped cutoff surfaces and a finite spectral width of the incident O-mode may reduce the conversion efficiency considerably.

A number of analytical considerations of O-X conversion in 3D plasmas exist [40, 41, 42], which, however, all rely on an idealised 3D geometry at the conversion point. As we will see in the example below in realistic 3D stellarator equilibria the evanescent layer can have a rather complex shape which cannot be properly described in a reduced analytical representation. In contrast, the 3D full wave treatment demonstrated here allows for a direct evaluation of the conversion efficiency for arbitrarily shaped beams and plasmas.

Figure 5 shows an example of a 2D full wave calculation of the O-X conversion process. The O-mode at 140 GHz, launched from vacuum, propagates into the plasma and is partly reflected off the corresponding cut-off layer, where a part of the O-mode is converted into a slow X-mode, which propagates toward the upper-hybrid resonance (dashed line). The magnetic field of approximately 2.5 T is directed along the positive x axis. The density is smoothly ramped up along the y axis the density length scale $L \equiv n_e/\nabla n_e = 1/15$ m at the conversion layer.




Figure 5: An example of the O-X mode conversion calculation. Left plot: snapshot of the wave electric field (colour-coded), ray tracing (solid red). The antenna is located at the origin (0,0). The O-mode cutoff and upper hybrid resonance layers are indicated with the white and the white dashed lines, respectively. The intensity (time averaged normalised Poynting vector amplitude) of the wave at the $y = 0$ plane is illustrated with the white curve. Right plot: windowed Fourier transform of the wave field, i.e the wave refractive index $N_y = N_{\perp}$ (colour-coded) and the refractive index predicted with a ray-tracing (solid).

Despite the fact that WKB-optimal wave launched, the conversion efficiency of the finite width Gaussian beam is much lower than unity. The beam radius and the radius of the curvature at the launching position are 1.2 cm and 1 m, respectively. The reflected part of the O-mode has a “hole” around the reference ray due to the non-uniform O-X conversion. The intensity (time averaged normalised Poynting vector amplitude) of the wave at the $y = 0$ plane is illustrated with the white curve. The overall conversion efficiency is about 60% in this case. The evolution of the perpendicular component of the refractive index is computed using a windowed Fourier transform and is shown in Fig. 5 (right). The incoming O-mode, the reflected O-mode and the slow X-mode (which settles near the upper hybrid resonance layer) can be clearly distinguished in this diagram, which appears to be in a good agreement with expectations from WKB theory (i.e. dispersion curves), validating the numerical solution. Note that the slow X-mode is eventually dissipated numerically without leaving the computation domain. This dissipation is of purely numerical nature. It happens when the wavelength of the mode becomes shorter than 10 computational cells, i.e. $\Delta x/\lambda > 0.1$.

Since W7-X plasmas are strongly shaped, simulations of the O-X conversion in W7-X require 3D evaluation. Figure 6 shows an example of such a calculation. The wave is launched from the bottom of the domain in the y direction, the slow X- mode is contained within the computation domain, and the reflected O-mode power distribution is shown on the plane $x = \text{Const}$). The reflected O-mode power distribution does not have a simple structure (around the central reference ray), as in the 2D case above, but has an asymmetrically distributed power density. This is because of the rather complex structure of the evanescent layer shown in Fig. 7.

To estimate the power conversion efficiency, we compute surface integrals of the time-averaged Poynting vector (i.e. the intensity) over every boundary of the computational domain:




Figure 6: 3D calculation of the scenario similar to the one described in Fig. 5. The decimal logarithm of the time-averaged Poynting vector is colour-coded. The grey tube represents the corresponding WKB trajectory for the reference ray.

$$P_i = \frac{1}{\mu_0} \iint (\langle \mathbf{E} \times \mathbf{B} \rangle \cdot \mathbf{n}) dS, \quad (17)$$

where the index “i” denotes a particular boundary, n represents its normal vector and $\langle \dots \rangle$ denotes time averaging.

In the above case, the O-X conversion efficiency is 85%, calculated as the difference between the intensities of the input (P_{in}) and reflected (P_{ref}) waves. This case represents an optimised scenario in which the W7-X launching antenna position was varied to find the highest conversion within a given equilibrium and density profile. The normalised density gradient near the cutoff is approximately 15m^{-1} . Note that it is not possible to realise the launcher position required for this case in W7-X with the existing ECRH system, and the typical conversion efficiency for realistic configurations and launcher positions is around 50%. Figure 8 shows how the conversion efficiency depends on the launcher azimuth and altitude angles for one of the W7-X launchers [43, 44]. A systematic study of the O-X conversion efficiency in W7-X plasmas was reported in [45].

5. Mode coupling in a sheared magnetic field

Another example of CUWA applications is the calculation of cross-mode coupling at vanishing density [46]. In particular, this effect is observed in LHD, where ECRH beams propagate in strongly sheared magnetic field outside of the plasma [4, 47]. The O- and X-mode’s polarisations are then not well separated and energy transfer occurs between the modes. A number of theoretical approaches for the mode purity study were previously reported, see, for example, Refs. [5, 48]. Full wave calculations permit direct evaluation of mode coupling effects in a realistic magnetic fields [47].

Figure 9 (right) shows an example of a full wave calculation of an LHD-like scenario, i.e. beam propagation in a strongly sheared magnetic field. Note that a realistic LHD setup involves beam propagation in a 2 m duct from the last mirror toward the plasma, whereas Figure 9 only




Figure 7: Structure of the evanescent layer on the O-mode cutoff surface for $N_{\parallel} \approx N_{\parallel}^{opt}$. The thickness of the evanescent layer is colour-coded (the intersection with the X-mode turning surface is indicated with the white contours).




Figure 8: Contours of the O-X conversion efficiency for a given plasma parameters over varying beam aiming angles (for a fixed location of the launcher).

shows a calculation of the last 15 cm of a beam propagation. A very high-purity O-mode (see Eq. (11)) is launched from the bottom of the computational domain toward the plasma, and the density of the plasma rises above the X-mode “cutoff” so that the spurious X-mode is reflected. In this particular example the reflected wave amplitude is only 3% of the wave amplitude of the incoming O-mode. Note that when the realistic LHD magnetic field is replaced with a unidirectional magnetic field the X-mode is not generated. This is demonstrated in Fig. 9 (left). Our preliminary analysis of W7-X relevant plasma shows negligible cross-polarisation power exchange. This is mainly because of a significantly lower magnetic shear in W7-X, as compared to LHD.

6. Summary

In this paper we report on the development of the 3D full wave GPU-capable code CUWA. The code is equipped with CPML “boundary conditions”, which permit efficient truncation of the computational domain and windowed Fourier Transform for a direct comparison of the full wave




Figure 9: O-mode propagation in vanishing density for a unidirectional magnetic field (left) and for an LHD-like sheared magnetic field (right). Power transfer to the spurious X-mode (reflected from the plasma) is observed in the sheared case (right). The solid curves correspond to a ray-tracing calculations of the O and X modes. Dashed curves are the contours of constant density.

solution with the predictions of the WKB approximation. The code uses ray-tracing calculations to construct a computational domain with realistic plasma parameters.

The accuracy and stability of the employed numerical techniques are demonstrated in benchmarks against analytically tractable examples.

Also, we demonstrate CUWA application to a problem of linear O-X mode conversion in 3D plasmas and to a problem of mode-coupling in a strongly sheared magnetic field.

A great advantage of the GPU-capable code is its speed. The 3D computation shown in Fig. 6 takes only a few minutes (depending on the GPU device), while analogous 2D computations are typically no longer than several seconds. This allows for efficient and large parametric scans over a broad range of parameters, and makes it possible to solve otherwise demanding optimisation problems or problems involving statistical averaging over a large number of calculations, i.e. study of the effects of turbulence on waves propagation.

7. Acknowledgements

The authors are grateful for helpful comments from Per Helander and would like to acknowledge useful discussions with Torsten Stange and Yuriy Turkin.

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

- [1] J. Preinhaelter and V. Kopecký. Penetration of high-frequency waves into a weakly inhomogeneous magnetized plasma at oblique incidence and their transformation to Bernstein modes. *Journal of Plasma Physics*, 10(01):1, aug 1973.
- [2] E. Mjølhus. Coupling to z mode near critical angle. *Journal of Plasma Physics*, 31(01):7, feb 1984.

- [3] H. P. Laqua, V. Erckmann, H. J. Hartfuß, and H. Laqua. Resonant and nonresonant electron cyclotron heating at densities above the plasma cutoff by o-x-b mode conversion at the w7-as stellarator. *Physical Review Letters*, 78(18):3467–3470, may 1997.
- [4] T Notake, S Kubo, T Shimozuma, H Idei, Y Yoshimura, S Inagaki, K Ohkubo, S Kobayashi, Y Mizuno, S Ito, Y Takita, T Watari, K Narihara, T Morisaki, I Yamada, Y Nagayama, K Tanaka, S Sakakibara, R Kumazawa, T Seki, K Saito, T Mutoh, A Shimizu, A Komori, and the LHD experimental group. Optimization of incident wave polarization for ECRH in LHD. *Plasma Physics and Controlled Fusion*, 47(3):531–544, feb 2005.
- [5] S. E. Segre and V. Zanza. Evolution of polarization of an electromagnetic wave propagating in magnetized plasma: Comparison between two alternative formalisms. *Physics of Plasmas*, 18(9):092107, sep 2011.
- [6] T Lehner, J. M Rax, and X. L Zou. Linear mode conversion by magnetic fluctuations in inhomogeneous magnetized plasmas. *Europhysics Letters (EPL)*, 8(8):759–764, apr 1989.
- [7] X. L. Zou, L. Colas, M. Paume, J. M. Chareau, L. Laurent, P. Devynck, and D. Gresillon. Internal magnetic turbulence measurement in plasma by cross polarization scattering. *Physical Review Letters*, 75(6):1090–1093, aug 1995.
- [8] J H Irby, S Home, I H Hutchinson, and P C Stek. 2d full-wave simulation of ordinary mode reflectometry. *Plasma Physics and Controlled Fusion*, 35(5):601–618, may 1993.
- [9] Abhay K. Ram, Kyriakos Hizanidis, and Yannis Kominis. Scattering of radio frequency waves by blobs in tokamak plasmas. *Physics of Plasmas*, 20(5):056110, may 2013.
- [10] A Köhn, L Guidi, E Holzhauer, O Maj, E Poli, A Snicker, and H Weber. Microwave beam broadening due to turbulent plasma density fluctuations within the limit of the born approximation and beyond. *Plasma Physics and Controlled Fusion*, 60(7):075006, may 2018.
- [11] Allen Taflove and Susan C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. ARTECH HOUSE INC, 2005.
- [12] Umran S. Inan and Robert A. Marshall. *Numerical Electromagnetics: The FDTD Method*. Cambridge University Press; 1 edition, 2011.
- [13] J Alan Roden and Stephen D Gedney. Convolutional pml (cpml): An efficient fdtd implementation of the cfs-pml for arbitrary media. *Microwave and optical technology letters*, 27(5):334–338, 2000.
- [14] A Köhn, Á Cappa, E Holzhauer, F Castejón, Á Fernández, and U Stroth. Full-wave calculation of the o-x-b mode conversion of gaussian beams in a cylindrical plasma. *Plasma Physics and Controlled Fusion*, 50(8):085018, jun 2008.
- [15] T R N Williams, A Köhn, M R O Brien, and R G L Vann. Propagation in 3d of microwaves through density perturbations. *Plasma Physics and Controlled Fusion*, 56(7):075010, may 2014.
- [16] Yaxin Yu and J.J. Simpson. An e-j collocated 3-d FDTD model of electromagnetic wave propagation in magnetized cold plasma. *IEEE Transactions on Antennas and Propagation*, 58(2):469–478, feb 2010.
- [17] David N. Smithe. Finite-difference time-domain simulation of fusion plasmas at radiofrequency time scales. *Physics of Plasmas*, 14(5):056104, may 2007.
- [18] Jean-Pierre Bérenger. An implicit FDTD scheme for the propagation of VLF-LF radio waves in the earth-ionosphere waveguide. *Comptes Rendus Physique*, 15(5):393–402, may 2014.
- [19] Filipe da Silva, Martin Campos Pinto, Bruno Després, and Stéphane Heuraux. Stable explicit coupling of the yee scheme with a linear current model in fluctuating magnetized plasmas. *Journal of Computational Physics*, 295:24–45, aug 2015.
- [20] Jean-Pierre Berenger. *Perfectly Matched Layer (PML) for Computational Electromagnetics*. Morgan and Claypool Publishers, 2007.
- [21] Mustafa Kuzuoglu and Raj Mittra. Frequency dependence of the constitutive parameters of causal perfectly matched anisotropic absorbers. *IEEE Microwave and Guided wave letters*, 6(12):447–449, 1996.
- [22] Shane Cook. *CUDA Programming*. Elsevier LTD, Oxford, 2012.
- [23] Andreas Klöckner, Nicolas Pinto, Yunsup Lee, Bryan Catanzaro, Paul Ivanov, and Ahmed Fasih. PyCUDA and PyOpenCL: A scripting-based approach to GPU run-time code generation. *Parallel Computing*, 38(3):157–174, mar 2012.
- [24] Lev E. Givon, Thomas Unterthiner, N. Benjamin Erichson, David Wei Chiang, Eric Larson, Luke Pfister, Sander Dieleman, Gregory R. Lee, Stefan van der Walt, Bryant Menn, Teodor Mihai Moldovan, Frédéric Bastien, Xing Shi, Jan Schlüter, Brian Thomas, Chris Capdevila, Alex Rubinsteyn, Michael M. Forbes, Jacob Frelinger, Tim Klein, Bruce Merry, Nate Merrill, Lars Pastewka, Li Yong Liu, S. Clarkson, Michael Rader, Steve Taylor, Arnaud Bergeron, Nikul H. Ukani, Feng Wang, Wing-Kit Lee, and Yiyin Zhou. scikit-cuda 0.5.2: a Python interface to GPU-powered libraries, November 2018. <http://dx.doi.org/10.5281/zenodo.1479820>.
- [25] Thomas H. Stix. *Waves in Plasmas*. American Institute of Physics, 1992.
- [26] V.D. Shafranov. Electromagnetic waves in a plasma. *Reviews of Plasma Physics vol. 3, ed M.A. Leontovich (New York: Consultants Bureau)*, 1967.
- [27] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001. [Online;

- accessed [today].
- [28] A.P. Smirnov and R.W. Harvey. Calculations of the current drive in diii-d with the genray ray tracing code. *Bull. Am. Phys. Soc.*, 1995.
 - [29] Y Peysson, J Decker, and L Morini. A versatile ray-tracing code for studying rf wave propagation in toroidal magnetized plasmas. *Plasma Physics and Controlled Fusion*, 54(4):045003, mar 2012.
 - [30] N.B. Marushchenko, Y. Turkin, and H. Maassberg. Ray-tracing code TRAVIS for ECR heating, EC current drive and ECE diagnostic. *Computer Physics Communications*, 185(1):165–176, jan 2014.
 - [31] Pavel Aleynikov and Boris Breizman. Stability analysis of runaway-driven waves in a tokamak. *Nuclear Fusion*, 55(4):043014, mar 2015.
 - [32] S. P. Hirshman. Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Physics of Fluids*, 26(12):3553, 1983.
 - [33] *EQDSK file format*, https://w3.pppl.gov/ntcc/TORAY/G_EQDSK.pdf.
 - [34] F. M. A. Smits. Elliptical polarisation for oblique ec-wave launch, 1993. Proc. of 8th Joint Workshop on Electron Cyclotron Emission and Electron Cyclotron Resonance Heating (1993) Vol 2, p. 324.
 - [35] V.L. Ginzburg. *Propagation of Electromagnetic Waves in Plasma*. Pergamon Press, Oxford, 1964.
 - [36] Paul F. Goldsmith. *Quasioptical Systems: Gaussian Beam Quasioptical Propagation and Applications*. Wiley-IEEE Press, 1998.
 - [37] V. Erckmann, P. Brand, H. Braune, G. Dammertz, G. Gantenbein, W. Kasperek, H. P. Laqua, H. Maassberg, N. B. Marushchenko, G. Michel, M. Thumm, Y. Turkin, M. Weissgerber, A. Weller, W7-X ECRH Team at IPP Greifswald, W7-X ECRH Team at FZK Karlsruhe, and W7-X ECRH Team at IPF Stuttgart. Electron cyclotron heating for w7-x: Physics and technology. *Fusion Science and Technology*, 52(2):291–312, 2007.
 - [38] Heinrich Peter Laqua. Electron Bernstein wave heating and diagnostic. *Plasma Physics and Controlled Fusion*, 49(4):R1–R42, mar 2007.
 - [39] Harold Weitzner and D. B. Batchelor. Conversion between cold plasma modes in an inhomogeneous plasma. *Physics of Fluids*, 22(7):1355, 1979.
 - [40] A G Shalashov and E D Gospodchikov. On o-x mode conversion near the cut-off surfaces in 3d sheared magnetic field. *Plasma Physics and Controlled Fusion*, 52(11):115001, sep 2010.
 - [41] A Popov. On o-x mode conversion in a 3d inhomogeneous plasma. *Plasma Physics and Controlled Fusion*, 53(6):065016, apr 2011.
 - [42] Harold Weitzner. O-x mode conversion in a non-symmetric torus at electron cyclotron frequencies. *Physics of Plasmas*, 24(2):022502, feb 2017.
 - [43] Pavel Aleynikov and Nikolai B. Marushchenko. 3d full-wave modelling and mode conversion in realistic w7-x plasmas. *EPJ Web of Conferences*, 149:03007, 2017.
 - [44] Pavel Aleynikov and Nikolai B. Marushchenko. 3d full-wave modelling and mode conversion in realistic w7-x plasmas. *20th Joint Workshop on Electron Cyclotron Emission (ECE) and Electron Cyclotron Resonance Heating (ECRH), May 14-17, 2018, Greifswald, Germany*, 2017.
 - [45] Pavel Aleynikov and Nikolai B. Marushchenko. Ecrh and mode conversion in overdense w7-x plasmas. *27th IAEA Fusion Energy Conference, Gujarat, India*, 2018.
 - [46] K. Nagasaki, A. Ejiri, T. Mizuuchi, T. Obiki, H. Okada, F. Sano, H. Zushi, S. Besshou, and K. Kondo. Effects of magnetic shear on electron cyclotron resonance heating in heliotron/torsatron configurations. *Physics of Plasmas*, 6(2):556–564, feb 1999.
 - [47] T. Ii Tsujimura, S. Kubo, H. Takahashi, R. Makino, R. Seki, Y. Yoshimura, H. Igami, T. Shimozuma, K. Ida, C. Suzuki, M. Emoto, M. Yokoyama, T. Kobayashi, C. Moon, K. Nagaoka, M. Osakabe, S. Kobayashi, S. Ito, Y. Mizuno, K. Okada, A. Ejiri, and T. Mutoh. Development and application of a ray-tracing code integrating with 3d equilibrium mapping in LHD ECH experiments. *Nuclear Fusion*, 55(12):123019, nov 2015.
 - [48] I Vakulchuk, Aleynikov P., and Marushchenko N.B. The wave-mode purity in ecrh: advanced 3d ray-tracing modeling for w7-x. *43rd EPS Conference on Plasma Physics*, 2016.