J. Plasma Phys. (2026), *vol.* 92, E58 © The Author(s), 2026. Published by Cambridge University Press.
This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited. doi:10.1017/S0022377826101639

Ion Cyclotron Resonance Heating wave equation solving using quasimodes

Dirk Van Eester[1], Vincent Maquet[1] and Bernard Reman[1]

[1]Laboratorium voor Plasmafysica – Laboratoire de Physique des Plasmas – ERM/KMS, Belgian 'EUROfusion Consortium' Member, Trilateral Euregio Cluster, Renaissancelaan 30 Avenue de la Renaissance, Brussels 1000, Belgium

Corresponding author: Dirk Van Eester, d.van.eester@fz-juelich.de

(Received 24 October 2025; revision received 23 January 2026; accepted 25 March 2026)

In solving the Ion Cyclotron Resonance Heating wave equation, plane wave Fourier mode $\exp[ikx]$ base functions are often exploited to produce expressions for the dielectric response when including the effect of finite temperature, requiring locally integrating the equation of motion of charged particles by hand. Standard Fourier analysis notation is used: x is the position, k is the wave vector component in the x-direction and i is the imaginary number. In contrast, finite element techniques adopt base functions – typically low - order polynomials – that individually are only non-zero in a small domain. They are a natural and easy approach to capture inhomogeneity effects and readily allow grid refinement to zoom in on regions where this is required. To enable ample realism of kinetic effects while profiting from the rich pool of numerical tools available for solving differential and integro-differential equations relying on finite elements, it is desirable to have a procedure allowing us to profit from both approaches: the detailed physics brought by finite temperature effects – commonly described in terms of Fourier modes – as well as the simplicity from a local polynomial representation. A novel technique is offered to achieve that. It consists of finding the Fourier representation of the localised base functions exploited in finite elements so that the richness of the dielectric response in k-space can be accounted for. The resulting equation is assembled just like the finite element method prescribes but the coefficients of the linear local system are assembled differently. The technique allows us to capture finite temperature corrections in both the parallel and perpendicular directions for a dielectric tensor model of choice. The focus in this paper is on the numerical technique while the expressions for the dielectric response are assumed to be known. A few first examples are briefly discussed: the wave equation solutions for a typical minority heating (H)-D JET plasma are provided for (i) an all-FLR model, (ii) the FLR-0 'tepid' equivalent and (iii) a cold plasma model; FLR refers to the Finite Larmor Radius expansion exploited to account for finite temperature corrections.

Key words: plasma heating, fusion plasma, plasma simulation

1. Introduction

Since it – up to higher-order corrections and depending on the level of realism aimed at – allows us to locally solve the equation of motion of charged particles

undergoing the Lorentz force by hand, it is customary to cast the wave equation in k-space and 'patch' local solutions together to compose the global wave response. A similar 'patchwork' approach is the basic idea of the method presented here: by locally adopting a suitable set of finite element base functions, the necessary computations required to evaluate the dielectric response can be performed by representing the electric field in terms of local base functions in each of the finite elements. Subsequently representing the base functions in terms of their Fourier transform allows us to make use of expressions for the dielectric response cast in k (wave vector)-space, a key asset when including parallel and perpendicular electric field gradient effects. It is stressed that the present paper does not explore new physics; it adopts standardly known models for the dielectric tensor response but pushes their potential further. In particular, it allows us to handle mode conversion in the presence of finite temperature and a poloidal magnetic field since treating mode coupling is at the heart of the technique. Accounting for other physics models (a sketch of possible upgrades to account for inhomogeneity and decorrelation effects is provided in Van Eester (2026)) leaves most of the wave equation solver intact and just requires upgrading the parts actually describing the physics model. The availability of fast computers has triggered renewed interest in upgrading the realism of physics models. Some relevant recent examples are Shiraiwa (2010), Svidzinski *et al.* (2016), Green (2018), Machielsen (2023), Fukuyama (2025), Lamalle (2025), Maquet (2026), Shiraiwa (2026) and Reman (2026).

Restricting to one dimension and to 1 electric field component for the moment to merely show the principle of the proposed technique, the variational form of the wave equation of interest can be written in the general form

$$\int_{-\infty}^{+\infty} \mathrm{d}x \int_{-\infty}^{+\infty} \mathrm{d}k \int_{-\infty}^{+\infty} \mathrm{d}k' F_{k'}^{*} G(k, k', x) E_k \exp[i(k-k')x] = 0, \tag{1.1}$$

where F is the test function and $F_{k'}$ its Fourier component, E is the electric field and E_k its Fourier component and $G(k, k', x)$ contains the vacuum term arising from the curl–curl operator as well as the plasma term. Strictly, $\boldsymbol{k}$ should be a Cartesian wave vector. In practice, one often adopts a convenient local Cartesian frame when evaluating G. The exploited patching implicitly assumes the magnetic field curvature is mild so that its variation can be neglected in a first approximation (upgrades can be found in the literature, see e.g. Lamalle (1997)). The local and global frames are related via a set of rotations. The actual wave equation has 3 electric field components and a test function vector. The above scalar representations are then merely to be generalised to vector–matrix–vector multiplications. The examples provided later in the text result from solving the actual 3-component wave equation. The one-dimensional integrals can also be upgraded to multiple dimensions in an equally straightforward way. This – as well as a more in depth study focussing on the physics rather than on the mathematical/numerical method – is left for later work.

2. Toy example wave equation

We sketch the method by first applying it to an elementary example, solving the wave equation $\mathrm{d}^2E/\mathrm{d}x^2 + K^2(x)E = 0$. $K(x)$ is the local wave vector component. In weak variational form this equation becomes

$$\int \mathrm{d}x \left[-\frac{\mathrm{d}F^*}{\mathrm{d}x}\frac{\mathrm{d}E}{\mathrm{d}x} + F^* K^2(x) E \right] = 0 \tag{2.1}$$

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

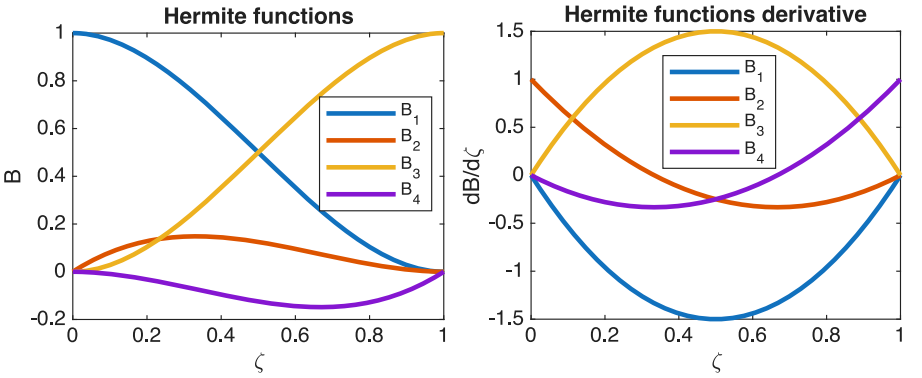


FIGURE 1. The 4 Hermite base functions and their derivatives in a reference unity-width interval.

provided the adopted base functions are sufficiently continuous to allow the surface term to be dropped (which is why base functions that at least ensure the field and its derivative to be continuous across elements are needed for this particular equation; we will come back to this point later in the paper). We subdivide the interval into a set of finite elements $[x_i, x_{i+1}]$. In each of these finite elements E is subsequently expressed in terms of a convenient set of base functions, yielding the local approximation $E_i = \sum_j E_{i,j} B_j$, while the F on which the equation is projected are the various base functions B_j themselves.

2.1. *Usual finite element treatment*

Strictly speaking, there is no actual need to make the detour via k-space for solving the toy equation. For the present purpose we adopt the Hermite polynomials as base functions in each finite element interval $[x_i, x_{i+1}]$. These functions are $B_1(\zeta) = 1 - 3\zeta^2 + 2\zeta^3$, $B_2(\zeta) = \zeta - 2\zeta^2 + \zeta^3$, $B_3(\zeta) = 3\zeta^2 - 2\zeta^3$, $B_4(\zeta) = -\zeta^2 + \zeta^3$, where ζ is the local variable $\zeta = (x - x_i)/\Delta x$; they are depicted in figure 1. The Hermite polynomials have the advantage that the unknowns of the equation solved have a direct physical meaning at the grid points adopted to discretise the integration domain: they are the E-field and (up to a factor $\Delta x = x_{i+1} - x_i$) its derivative $\mathrm{d}E/\mathrm{d}x$ at the edges of each finite element. Assuming the grid is refined enough to approximate K^2 by its value $\tilde{K}_i^2$ in the middle of the interval $\tilde{x} = (x_i + x_{i+1})/2$ so that it can be moved out of the integration, the local systems in each interval are now of the form

$$\left[-\overline{\overline{B}}_{11} + \tilde{K}_i^2 \overline{\overline{B}}_{00}\right].\boldsymbol{E}_i = 0 \tag{2.2}$$

in which $\overline{\overline{B}}_{mn}$ is the matrix of the integrals of the product of the mth derivative of the base function corresponding to F and the nth derivative of the base function corresponding to E; these can be evaluated by hand once and for all. Assembling all the local systems in a global linear system and adding boundary conditions (e.g. making use of Lagrange multipliers) allows us to solve the wave equation making use of standard matrix inversion. The variation of K^2 can be better accounted for at the price of needing to evaluate products of 3 rather than 2 base functions after expressing K^2 in terms of the same base functions on the adopted grid.

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

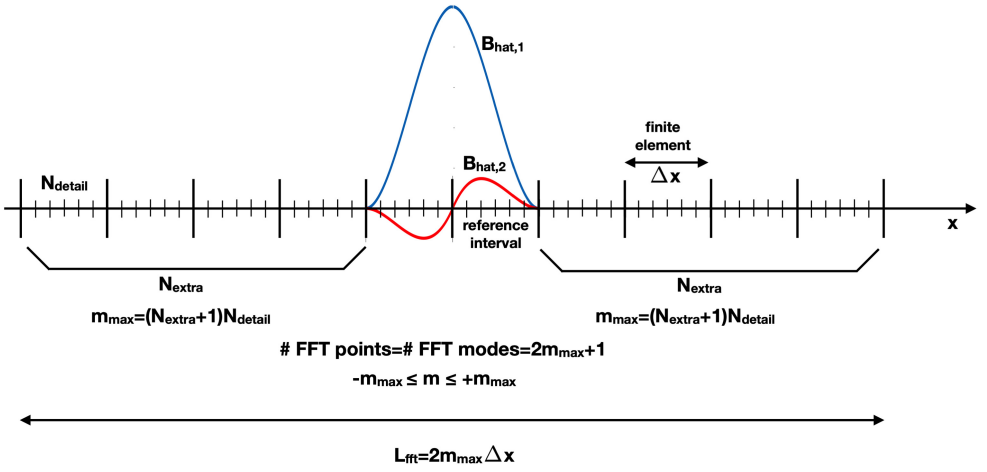


FIGURE 2. Sketch of the reference interval and its neighbour (in which the hat functions have a non-zero value) and the set of neighbouring elements in which a zero padding is imposed when computing the Fourier representation of the Hermite base functions.

2.2. *Quasi-base*

We now introduce a quasi-local base exploiting Fourier analysis: we find the local Fourier representation of the Hermite 'hat' base and then revert back to the original 4 Hermite base functions' 'quasi-Fourier' representation. Expressing F and E in terms of their Fourier transform in the adopted simple equation yields

$$\int \mathrm{d}x \int \mathrm{d}k \int \mathrm{d}k' F^*_{k'}[-k'k + K^2(x)]E_k \exp[i(k-k')x] = 0. \tag{2.3}$$

Determining the Fourier representations of a general function requires integrating over the full domain but the Hermite base functions adopted are only finite on a single interval; as there is a limited number of dominant modes of the Fourier spectrum, one can define a convenient width of the spectrum and sample that by a suitable number of discrete modes k_m. Since the actual Hermite functions have discontinuous jumps across elements for finite values of the base functions, using 'hat' base functions combining the B_j two by two is handy as it avoids the Gibbs phenomenon occurring at the position where the dominant contribution to the required localised integrals is picked up. Combining B_3 in $[x_{i-1}, x_i]$ with B_1 in $[x_i, x_{i+1}]$ yields a basic 'hat' that is zero and has zero derivative at x_{i-1}, but also that is 1 and has zero derivative at x_i and, finally, is zero and has zero derivative at x_{i+1}. Similarly, combining B_4 in $[x_{i-1}, x_i]$ with B_2 in $[x_i, x_{i+1}]$ yields a function of which the derivative has the shape of a 'hat'; it is zero and has zero derivative at x_{i-1}, it takes the value 0 but has a unity ζ-derivative at x_i; finally, it has value zero and has a zero derivative at x_{i+1}. See figure 2 for an impression.

Traditional discrete Fourier transform requires the functions to be periodic while such periodicity is often not present in practice. At first sight, exploiting Fourier analysis thus seems inadequate. But as we seek a representation in a restricted domain and the adopted base functions as well as their derivatives approach zero at the outer edges of the 2-element interval, they are (quasi-)periodic by construction. This is the key to connecting the (global) Fourier treatment with the (local)

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

finite element treatment. The grid distance being Δx and the Fourier base being defined in double that length, the natural spectral resolution Δk for an interval of length L_{fft} sampling the domain on which the hat function is defined is $2\pi/L_{fft} = 2\pi/[2\Delta x]$. In general the thus obtained Δk is too large to capture the spectrum of the waves modelled with sufficient detail so the natural L_{fft} is not a suitable choice. Typical core parameters for the JET machine yield an Ion Cyclotron Resonance Heating (ICRH) fast wave $|k_{fast}|$ around $40\,\mathrm{m}^{-1}$ while for the Bernstein wave and away from the confluence $|k_{Bernstein}| \approx 150\,\mathrm{m}^{-1}$. Hence we would like to cover the region $-k_{\max} \le k \le +k_{\max}$ for $k_{\max} \approx 200\,\mathrm{m}^{-1}$ with a sufficient number of k-modes allowing us to sufficiently faithfully represent local variations (e.g. $\Delta k \approx 1\,\mathrm{m}^{-1}$). But for a finite element grid with spacing $\Delta x = 0.02\,\mathrm{m}$ – needed to capture short wavelength modes when they appear in the core – the natural Δk is already approximately $300\,\mathrm{m}^{-1}$. To have a finer Δk step we need to consider a larger L_{fft} when evaluating the Fourier transform. Adding N_{extra} extra line segments of the same length Δx on each side and padding the functions with zeros there (recall we want to ensure a local representation), yields $\Delta k = 2\pi/[2(N_{extra}+1)\Delta x]$, which can be made as small as desired for any chosen grid distance Δx. The adopted k-spectrum also needs to contain all physically expected modes i.e. we must at least ensure that $k_{\max} = m_{\max}\Delta k$ exceeds some desired value. To guarantee that, a supplementary variable is introduced: N_{detail} is the number of subintervals in which the basic interval of length Δx is subdivided. The length of the considered Fast Fourier Transform (FFT) domain is now $L_{fft} = 2(N_{extra}+1)\Delta x$. In general it is distinct from the width of the actual physical domain. The total number of modes is $N_k = 2m_{\max}+1$, where $m_{\max} = (N_{extra}+1)N_{detail}$ and the modes lie in $-k_{\max} = -m_{\max}\Delta k \le k \le +k_{\max} = m_{\max}\Delta k$. Next, the extended base functions $B_{hat,1or2}$ are evaluated at the N_k sample points and their discrete Fourier transform is computed

$$B_{hat,k_m} = \frac{1}{N_k}\sum_{\mathrm{i}_{detail}=1}^{N_k} \exp[-ik_m x_{FFT,\mathrm{i}_{detail}}] B_{hat}(x_{FFT,\mathrm{i}_{detail}}). \tag{2.4}$$

The above are evaluated in the reference interval $0 \le x_{fft,\mathrm{i}_{detail}} < L_{fft}$ on a set of equidistant points. We will refer to the obtained set of modes as 'quasimodes' rather than as actual Fourier modes. A key point is that the adopted Fourier coefficients do not actually depend on Δx since Δk has been chosen for a provided Δx so that $\Delta k \Delta x$ is a constant; hence, the Fourier components need only be computed once as they can be interpreted as relating to a fixed set of N_k Fourier angles $0 \le \alpha_{fft} < 2\pi$. Note that the angle used for the FFT has been defined to be $\alpha_{fft} = \pi$ at macroscopic grid point x_i, which is the middle point of the extended interval used for the Fourier transform so that the Fourier modes contain a supplementary phase factor. An alternative would have been to define the FFT interval from $-\pi$ to $+\pi$ and the detailed FFT grid from $-L_{fft}/2$ to $+L_{fft}/2$ so that the reference point x_i corresponds to $x_{detail} = 0$ rather than $x_{detail} = L_{fft}/2$. One of the supplementary advantages of these modes is that we do not require the 'plane wave' base elements to be periodic in the actual physical domain. The role of the quasi-base is merely to locally represent the unknown and its derivatives sufficiently accurately while the value of k has the physical meaning $2\pi/\lambda$, where λ is the local wavelength. Evidently, if the domain is closed so that the solution is periodic, then there is a natural preference for Δk i.e. then the quasimodes are the actual physical modes and the mode number m has a physical meaning.

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

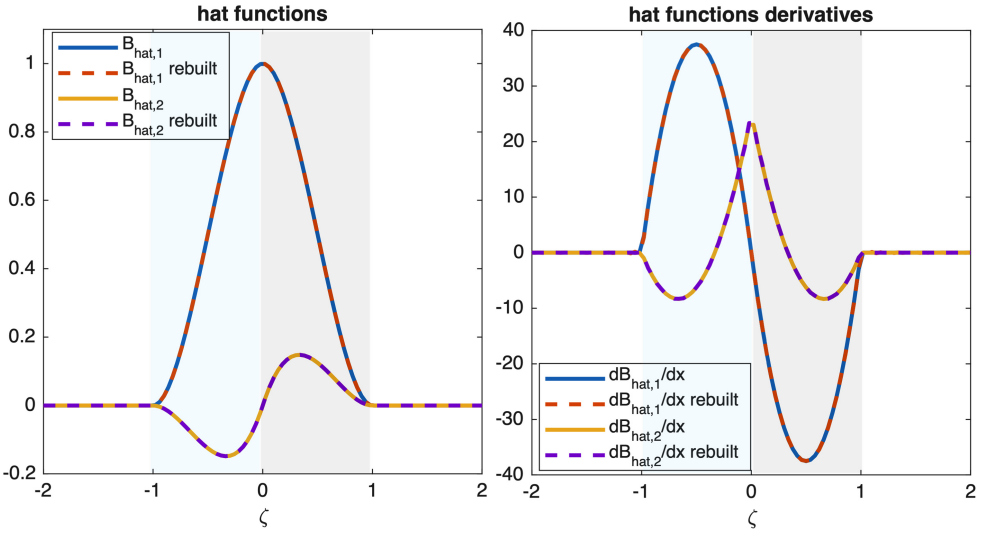


FIGURE 3. Two-element 'hat' functions in a wider interval (clipped at $|\zeta| = 2$) of the original hat functions (full lines) and their reconstruction using the quasi-Fourier base (dashed lines).

Figure 3 shows the hat functions and their derivatives as well as their Fourier base reconstruction. Note the clear break in the derivative of the original second base function at $\zeta = 0$. As the $\exp[ikx]$ have infinitely continuous derivatives, the Fourier representation is smooth at that location but slightly deviates from the actual function. This (so-called Gibbs) phenomenon is unavoidable since the original base functions only guarantee continuity of the base function and its first derivative. By adopting a sufficiently large number of modes the 2 sets of base functions – the original hat functions and the quasi-mode counterparts – can be made to come as close to one another as desired. In practice and for a suitable choice of Δx, the very high-order modes associated with that lie outside the physical spectrum and hence are inconsequential.

The next step is to revert to the original 4 Hermite base functions; in figure 3 they are depicted in the light blue and light grey areas. Since the hat functions $B_{hat,1\&2}$ are defined in double elements $[x_{i-1}, x_{i+1}]$ while we only evaluate B_{1-4} in $[x_i, x_{i+1}]$ (requiring us to shift the $B_{3\&4}$ finite elements to the right) one has

$$\begin{aligned} B_{1,k_m} &= B_{hat,1,k_m}, \\ B_{2,k_m} &= B_{hat,2,k_m}, \\ B_{3,k_m} &= B_{hat,1,k_m} \exp[-im\pi/[N_{extra}+1]], \\ B_{4,k_m} &= B_{hat,2,k_m} \exp[-im\pi/[N_{extra}+1]], \end{aligned} \tag{2.5}$$

because $\exp[ik_m\Delta x(\zeta - 1)] = \exp[ik_m\Delta x\zeta]\exp[-im\Delta k\Delta x]$ and Δk and Δx are linked: in that reference interval B_1 and B_2 are identical to the right halves of the hat functions while B_3 and B_4 are identical to the left halves of the hat functions in the previous interval ($-\Delta x \leq x_{FFT,detail} \leq 0$). Hence the first 2 identities in (2.5) for B_{1,k_m} and B_{2,k_m} do not carry a phase factor while the other 2 identities require a phase shift corresponding to moving the reference position and hence the reference phase in the exponential $\exp[ik_m x]$. Recall that x_i is associated with the angle $\alpha = \pi$

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

rather than zero and each shift by Δx gives rise to a phase shift $m\Delta\alpha$ for quasi-mode k_m. As for the hat functions, the set has to be computed only once, different reference points in the actual physical domain only requiring the evaluation of an extra factor $\exp[-ik_m x_{ref}]$ for each mode. We are now fully equipped to return to the wave equation. The total integral over the physical interval is subdivided in finite elements; we introduce quotes to distinguish between the index corresponding to E and that corresponding to F. Since localised base functions are used, the only finite contributions to the integral are picked up in the interval in which F and E are described by the same set of base functions i.e. $i = i'$; we need – of course – to consider all j for a chosen j'; the wave equation system is a set of projections on the individual $B_{j'}$ in each x-interval. Rather than a continuous k-integral we have a sum over a discrete set of k values. The integral in each finite element $[x_i, x_{i+1}]$ can now either be computed numerically or analytically. We label it as $I(m, m')$. As the reference finite element in the adopted FFT interval has index $N_{extra}+1$ we immediately get

$$\begin{aligned} I(m, m') &= \int_{N_{extra}+1}^{N_{extra}+2} \mathrm{d}\zeta \exp[i(k_m - k_{m'})\Delta x \zeta] \\ &= \frac{\big[\exp[i\Delta k_\zeta (m-m')(N_{extra}+2)] - \exp[i\Delta k_\zeta (m-m')(N_{extra}+1)]\big]}{i\Delta k_\zeta (m-m')}, \end{aligned} \tag{2.6}$$

with $\Delta k_\zeta = \pi/(N_{extra}+1)$; when $m = m'$ the above correctly yields the limit 1. The wave equation system has equations of the form

$$\sum_{j=1}^{4} \sum_{m,m'=1}^{N_k} \big[\mathcal{B}^*_{j',m'} I(m, m') \mathcal{B}_{j,m}\big]\big[-k_{m'}k_m + \tilde{K}_i^2\big] E_{i,j} = 0, \tag{2.7}$$

in which $\mathcal{B}_{j,m} = B_{j,k_m} \exp[-ik_m x_i]$. To avoid overly lengthy notation the sums can compactly be written as vector–matrix–vector multiplications upon storing the base function Fourier mode amplitudes $\mathcal{B}_{j,m}$ in a vector $\boldsymbol{\mathcal{B}}_j$ for each base function B_j. Similarly we store the integrals $I(m, m')$ in a matrix $\overline{\overline{I}}$. Looping over all B_j and B'_j and exploiting their local expression in terms of the Fourier modes, one obtains the 4×4 matrix with elements $\boldsymbol{\mathcal{B}}^*_{j'}.\overline{\overline{I}}.\boldsymbol{\mathcal{B}}_j$. The matrix elements only need to be evaluated once and are used for assembling the global system, the rows corresponding to the base functions composing the electric field and the columns corresponding to the test functions F.

Note that $\mathcal{B}_{j,m}$ only seemingly depends on the index i: when evaluating the Fourier component B_{j,k_m} in the interval $[x_i, x_{i+1}]$, the local variable ζ is introduced so that a factor $\exp[-ik_m x_i]$ appears. But when assembling the linear system, the integral is split into subcontributions over the same set of intervals so a compensating factor appears: both E and E_{k_m} are defined via an integral but the former has an oscillating factor $\exp[ik_m x]$ associated with E_{k_m} while the latter contains the factor $\exp[-ik_m x]$ in its defining integral. As this is the case not only for E but also for the test function F, the actual position of the grid point only matters for the coefficients appearing in the equation (here simply $\tilde{K}_i^2$). The above (2.7) is the quasi-base equivalent of the earlier found system of coupled linear equations (2.2) allowing us to find the $E_{i,j}$. Similar to directly evaluating the elementary integrals of products of base functions,

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

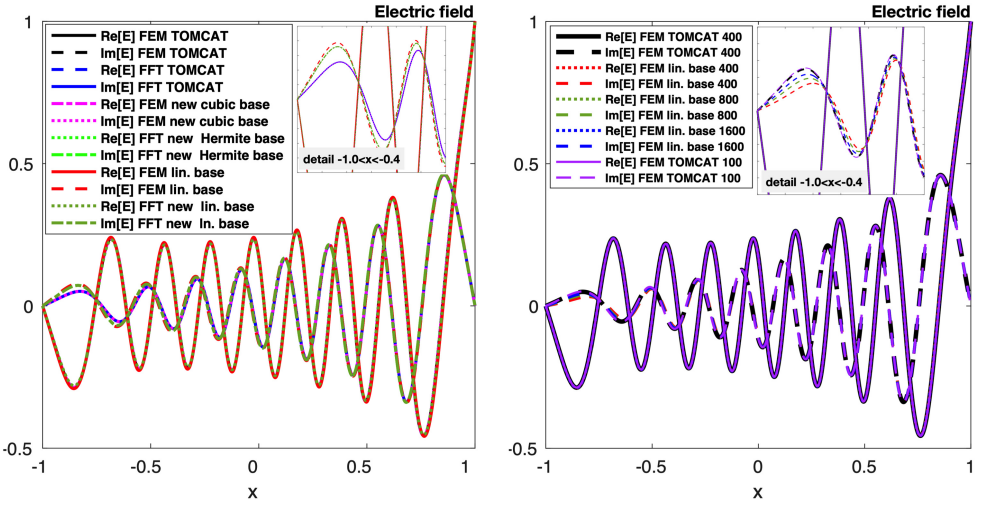


FIGURE 4. Solutions of the toy wave equation. The left plots show the electric field for 3 types of finite element base functions and either using the finite element or the quasimode technique for solving the wave equation. The plot on the right shows the solution obtained using linear base functions for an increasing number of grid points and compares it with the solution obtained adopting higher-order polynomials. The detailed inset plots depict the solutions close to the left edge where the largest difference is observed.

the expressions in the first big straight brackets can be evaluated once and for all: the quasimode algebra being formulated in terms of angles rather than a specific Δx or x_i means that the proposed scheme allows for local x-grid refinement without needing to recompute the quasimode amplitudes of the chosen base functions.

2.3. *Example of the toy wave equation and adopting different base functions*

Figure 4 provides the solutions of an elementary wave equation for a non-uniform K^2 with a mild imaginary part (damping) and solved in various ways; a parabolic profile with $K^2 = (10^3, 50)/m^2$ at $x = 0$ and $K^2 = (-10^2, 0.5)/m^2$ at the edges is adopted and the boundary conditions $E = 1$ at $x = 1m$ and $E = 0$ at $x = -1m$ were imposed.

The key aim of the present paper is to propose a method exploiting standard finite element techniques while allowing to account for detailed expressions of the dielectric response which are usually formulated in k-space. The expressions finally obtained are formally identical to those used for finite element solving except that the coefficients appearing in it are computed differently. A specific set of polynomial base functions – Hermite polynomials – has been adopted so far. Other choices of polynomial base functions can be exploited. The only change that needs to be made is to reformulate the system matrix in terms of these new unknowns by applying the transformation matrix connecting the 2 representations, or – even more directly – to compute the corresponding FFT coefficients $\mathcal{B}_{k_m}$.

The left subplots of figure 4 depict the result for 3 types of finite element base functions when using 100 finite element grid points and corresponding to 3 different types of 'hat' functions: the ones used until now, an upgraded set where continuity of the base elements across the boundaries is ensured up to the third

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

derivative (these seventh-order polynomial base functions were already introduced in Van Eester (1998)) and – finally – a downgraded set where the – piecewise linear – base functions allow us to ensure continuity of E across boundary edges but not of its derivatives. The philosophy for constructing these base functions is to repeat what was done for the Hermite base but extend the adopted philosophy: we seek polynomials that only have zeros as function values as well as derivatives at the individual finite element's edges (i.e. at the finite element grid points) except for 1. This allows us to combine them into 2 neighbouring elements so that they – or one of their derivatives – have the shape of a localised 'hat'. Since the generalised Hermite base functions with up to N continuous derivatives across the element boundaries are of the form $B_i = \sum_{j=0}^{2N-1} B_{i,j}\zeta^j$, their mth derivative takes the form $\sum_{j=m}^{2N-1} B_{i,j}\, j!/(j-m)!\zeta^{j-m}$ and hence – when exploiting this at $\zeta = 0$ and $\zeta = 1$ and for all but one of the derivatives up to N – the set of $2N$ coefficients $B_{i,j}$ for each B_i can readily be found for any desired N.

Although all found solutions depicted in figure 4 are very similar, the ones found adopting linear base functions slightly differ from the others, aside from the fact that the various solutions can hardly be distinguished by the naked eye. Relying on the original finite element method or using the quasimodes makes no noticeable difference either, underlining the potential of the proposed quasimode technique. One may suspect that the reason for the difference between the solutions found using linear base functions and the other solutions comes from the fact that these simple functions fail to correctly pass on the surface contribution i.e. the wave flux $F^*\mathrm{d}E/\mathrm{d}x$ (arising from partial integration to get from the strong to the weak variational form) since the derivative of the piecewise linear 'hat' base functions jumps discontinuously across elements and surface terms have been neglected so far; recall that the Hermite base does allow us to connect the function as well as its derivative and hence the simple wave equation's flux. Computing the surface terms reveals that the volume terms are actually much more important than the surface contributions: solving the equation with or without them hardly makes a difference (one needs to zoom in to see a difference) but, strictly, they need to be included for the specific equation at hand and for the linear hats. The right plots of figure 4 illustrate where the difference actually comes from: these drawings depict a subset of the solutions when gradually increasing the number of grid points. Not surprisingly, it shows that many more grid points are required to reach numerical convergence when adopting lower-order polynomial base functions. Adopting higher-order polynomials allows significantly cutting back on the number of finite element grid points needed.

3. General expression

The obtained result can now immediately be generalised: in the interval $[x_i, x_{i+1}]$ and for a specific local base function $B_{j'}$ the corresponding wave equation reads

$$\sum_{j=1}^{4}\sum_{m,m'=1}^{N_k}\left[\mathcal{B}^*_{j',m'} I(m,m')\mathcal{B}_{j,m}\right] G(k_m, k_{m'}, \tilde{x}_i) E_{i,j} = 0, \tag{3.1}$$

in which the expression inside the biggest square brackets – as mentioned in the previous section – just needs to be evaluated once and can be stored and merely reread for later computer runs.

Note that the m and m' in the exponential in $I(m,m')$ appear as a difference so that the integral fades when m and m' go further apart. It was shown in

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

Van Eester (2024) that one does not need to retain all couplings between all k_m and $k_{m'}$ modes – as done in the Jaeger AORSA code Jaeger (2001) – as it suffices to just retain the close couplings, contributions from k-modes further apart having the tendency to be filtered out when integrated over a larger region. Aside from the once-and-for-all obtention FFT of the base functions in a reference interval allowing us to dramatically simplify the algebra and push down the required computational time, this decoupling allows for a further reduction of the CPU time needed to assemble and solve the system. Intended optimisation strategies will briefly be discussed in § 5.

4. First examples of the method

The present paper aims at presenting a numerical and mathematical method allowing us to maximally exploit existing dielectric tensor models while also allowing us to include finite temperature as well as parallel gradient effects. It does not bring fundamentally new physics ingredients but provides a tool to include effects such as mode conversion to the ion Bernstein or ion-cyclotron wave using standard models such as described in Ichimaru (1973), Stix (1997), Swanson (2012) and Van Eester (2021, 2024). As the accent is on the method rather than the physics, only a few preliminary examples are provided; more in-depth exploitation using a wider set of dielectric tensors and ICRH scenarios is reserved for a later publication concentrating on the physics instead. The parameters adopted are loosely inspired on those of JET's most exploited and robust wave heating scheme: H minority heating in a D majority plasma. The toroidal magnetic field on axis is 3.45 T and the antenna frequency is 51 MHz, which places the H fundamental cyclotron layer and the second harmonic cyclotron layer in the centre of the plasma. Poloidal field effects are qualitatively included by imposing a linear growth of $B_{o,pol}/B_{o,tor}$ from the magnetic axis where it is zero, to a maximal value of 0.3 at the edge, crudely representative for JET discharges. The wave equation is solved along a straight, slanted path but away from the equatorial plane so that both the toroidal and poloidal magnetic field components vary. A core ($\rho = 0$) and edge ($\rho = a_p$) density of respectively $N_{e,0} = 8 \times 10^{19}\ \mathrm{m}^{-3}$ and $N_{e,s} = 10^{19}\ \mathrm{m}^{-3}$ was assumed; ρ is the magnetic flux surface labelling variable. Inside the last closed flux surface the density has a profile factor of $\alpha_N = 1.6$ (i.e. $N_e = [N_{e,0} - N_{e,s}][1 - (\rho/a_p)^2]^{\alpha_N} + N_{e,s}$) while in the edge it decays with an exponential decay length of 0.02 m. Core and edge temperatures are 4 keV and 70 eV with a profile factor of 2.5 and an exponential decay length of 0.1 m. The major radius is $R_o = 2.97$ m and the minor radius is $a_p = 0.95$ m. Beyond the last closed flux surface at $\rho = a_p$, the plasma decays towards vacuum. For simplicity the magnetic surfaces were assumed to have a circular cross-section so that all profile computations can simply be done by hand. The integration path runs from $(x = R - R_o, Z) = (-0.95, 0.15)$ m to $(0.95, 0.25)$ m where R is the major radius and Z is the horizontal position with respect to the equatorial plane. As we are solving a one-dimensional wave equation, an assumption needs to be made on the wave dependence in the other two dimensions. Toroidal curvature is included in the model ($k_{tor} = n_{tor}/R$); consistent with the dominant excitation for dipole phasing we chose the toroidal mode number to be $n_{tor} = 27$. The x (or R) dependence of the waves will be determined from solving the wave equation; we assume the $\boldsymbol{k}$ component in the remaining direction to be $k_Z = 5\ \mathrm{m}^{-1}$, mimicking slightly oblique incidence. We adopt $N_{extra} = 240$ and $N_{detail} = 10$. Figure 5 depicts the 3 electric field components for 3 types of dielectric tensor descriptions.

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

(i) The simplest model used is the cold plasma model (Stix 1997; Swanson 2012). In that model, the dielectric tensor has no dependence on the temperature so there is no thermal spreading in the parallel nor the perpendicular direction.

(ii) The 'tepid' model is a simplified version of the hot plasma model: it includes finite temperature corrections in the parallel direction but only retains zero-order finite Larmor radius terms in the perpendicular one (see e.g. Van Eester (2024)).

(iii) The most sophisticated one adopts the full hot dielectric tensor description, the original method of which was presented in Van Eester (1998) (itself inspired by Kennel & Engelmann (1966)), which was applied in truncated finite Larmor radius expansion form at that time; it was later upgraded to include all finite Larmor radius effects in e.g. Van Eester (2021). This underlying model expresses the wave equation in variational form and thereby allows us to ensure a positive definite power balance for all (long or short wavelength) modes admitted by the plasma when the plasma is in thermodynamical equilibrium. The main difference between earlier and the present exploitation is that – the dielectric response being cast in $\boldsymbol{k}$ space but the full response in any given (but single) direction now being accounted for – parallel gradient effects can now also be described, and extension to multi-dimensional application is a technical rather than a difficult matter. Opposite to models that can only describe the perpendicular dynamics when they need to restrain the description to a single toroidal mode number, the advantage of the present model is that it accounts for parallel as well as perpendicular gradients, be it with the level of sophistication of the dielectric response model used. In the present paper the usual 'adiabatic switch on' of the electric field almost always used in ICRH wave modelling is exploited so that the dynamics of the Larmor gyration is accounted for but only the end point of the integral is considered when performing the orbit integral needed to obtain the perturbed distribution function i.e. the dielectric response is local.

Depending on the choice of preferred reference frame the wave equation takes a slightly different form. A typical choice is adopting cylindrical components (E_R, E_Z, E_φ), where E_φ is the toroidal electric field component, and E_R and E_Z are the components in a toroidal cut. The equation can also be expressed in terms of field-aligned variables $(E_{\perp,1}, E_{\perp,2}, E_{//})$, where // and $\perp$ refer to the direction of the static magnetic field. When opting for the latter, various choices are possible for defining the 2 perpendicular directions. The dielectric tensor is most naturally expressed in field-aligned variables. Choosing the first option requires rotating the dielectric tensor to line up with the cylindrical frame; when choosing the latter the curl–curl term needs to be rotated. The former has the drawback that the independent variables contain a part of the perpendicular and the parallel electric field component. Since $E_{//}$ is typically orders of magnitude smaller than $E_\perp$ this choice enhances the risk that the parallel electric field is computed with less accuracy than needed to truthfully capture the effect of Landau damping, which is crucial for electron damping. For that reason, and in spite of the 2 possible formulations mathematically being equivalent, it is preferred here to opt for a field-aligned representation. In what follows $E_{\mathrm{x}} \equiv E_{\perp,1}$, $E_{\mathrm{y}} \equiv E_{\perp,2}$ and $E_{\mathrm{z}} \equiv E_{//}$. The coordinate used is $x = R - R_o$, the horizontal difference between the major radius and the major radius of the magnetic axis (note the difference in notation between x and x).

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

Since a standard implicit assumption made is that the unit vectors on which the equation is projected should be as coordinate-independent as possible, we opt for one perpendicular unit vector being as close as possible to pointing in the major radius direction (i.e. $|E_{\perp,1}.\boldsymbol{e}_R|$ being as small as possible, and $E_{\perp,2}$ following from $\boldsymbol{E}_{\perp,2} = \boldsymbol{E}_{//} \times \boldsymbol{E}_{\perp,1}$) rather than being aligned with $\nabla\rho$, where ρ is the magnetic surface labelling parameter. This choice is inspired on what is done in the AORSA code (Jaeger 2001) and avoids the $E_{\perp,1\&2}$ components flipping discontinuously across the magnetic axis.

Independent of the choice of variables, the following boundary conditions are imposed: $E_y = E_z = 0$ at the high field side edge of the integration interval and $E_y = 1$ and $E_z = 0$ at the low field side. Provided the metallic wall is aligned with the magnetic surfaces, these conditions are consistent with metallic boundary conditions except that the 'antenna' is mimicked by a non-zero E_y, reminiscent of the polarising effect of a perfect Faraday screen aligned with the magnetic field lines, shunting out the parallel electric field there.

The top plots in figure 5 depict a fast wave launched from the right (low field side) into a cold plasma; to bend resonances into quasi-resonances a small but finite imaginary part ($\nu/\omega = 10^{-3}$) was added to the driver frequency $\omega \to \omega + i\nu$. As the dominant contribution to the flux is coming from the E_y component and its derivative, and since almost no damping occurs in the cold plasma, the amplitude of that component is crudely constant between the antenna and the core. The associated E_x component shrinks as the density is increasing. Once the wave reaches the plasma centre, the electric field magnitude drops significantly over a short distance: the wave is partly damped, partly converts into another wave, partly tunnels through that central region to reach the high field side (left) and is partly reflected from the high density fast wave core cutoff. On the high field side the transmitted wave amplitude is modest; the wave undergoes a total reflection at the metal wall on the far left. Near the ion–ion hybrid layer in the core the fast wave excites a short wavelength mode which transports part of the wave energy back towards the low field side. Due to the presence of a finite poloidal field the total static magnetic field exhibits a gradient along its field lines and hence the electric field has a finite parallel gradient. As a consequence the fast wave partly mode converts into a forward ion-cyclotron wave: for fixed $k_{//}$, the slow wave can only propagate in the small region between Stix's $S = 0$ and $k_o^2 S = k_{//}^2$ but including poloidal field effects allows the ion-cyclotron wave to emerge (see e.g. Perkins (1977)). Approximately 0.1 m away from where it was excited, this short wavelength wave is already fully absorbed.

The middle plots show the electric field when adopting a tepid plasma model. There are a lot of similarities with the previous plots but some aspects differ. Of particular interest is that the amplitude of the mode converted wave is much smaller. This is the case for 2 reasons: (i) the finite temperature accounted for in the parallel direction now adds electron Landau damping to the picture. This damping has a more global character than the damping occurring due to a finite collisionality at the quasiresonance appearing in the cold plasma. (ii) The wave behaviour is similar on the high field side but has a different amplitude while – more markedly – the wave pattern is clearly different on the low field side. Both are a consequence of the fact that the finite temperature impacts on the interaction between the waves (modifying mode conversion, transmission through and reflection from the confluence region). The excited short wavelength mode is still transporting wave energy from the centre outwards (to the right).

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

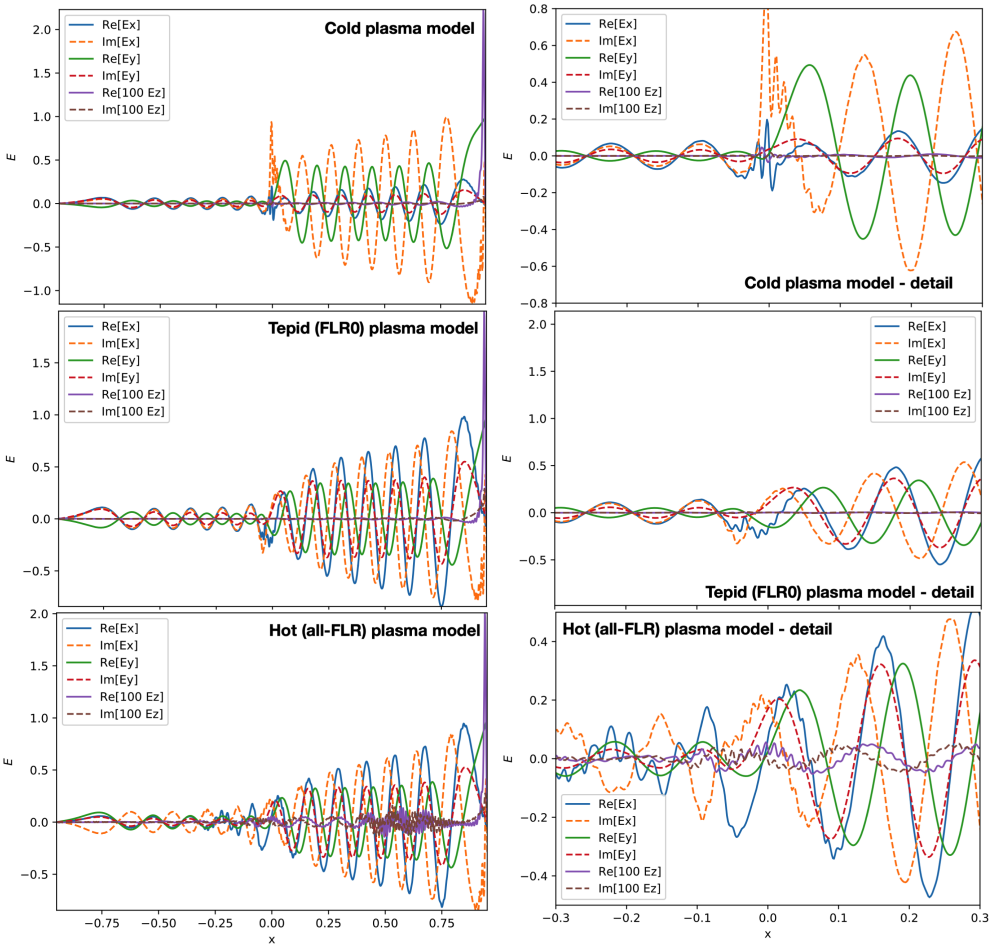


FIGURE 5. Electric field components for a JET (H)-D scenario with 5 % of H: full hot plasma i.e. retaining the full tensor of finite temperature corrections (bottom row), 'tepid' plasma (finite temperature corrections but only to lowest order; middle row) and cold plasma (top row). The left subplots show the 3 complex electric field components plotted in the entire machine while the right subplots zoom in on the ion–ion hybrid region where the fast wave has a mode confluence with a short wavelength mode.

The bottom line of plots in figure 5 shows the electric field when accounting for the full hot dielectric tensor, adding finite temperature effects in the perpendicular direction now as well. The fast wave dynamics at the low field side is very similar to that of the tepid model but 2 differences are to be noted: (i) the short wavelength branch excited in the core is not propagating towards the low but to the high field side now. Rather than the ion-cyclotron wave, the Bernstein wave is excited now. (ii) The E_z component exhibits short wavelength structure superposed on the fast wave field in a large fraction of the low field side domain. Looking back at the corresponding tepid plasma result one notes evidence of the presence of a short wavelength wave in the edge there as well but not deeper inwards. Although a more in-depth analysis is required to reach a more definite conclusion, it is tempting to explain this as short wavelength structure born either at the antenna or near the

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

lower hybrid layer in the middle (tepid) case, and associating it with one of the infinity of short wavelength waves admitted by the hot plasma when adopting an all-Finite Larmor Radius (FLR) model for the bottom case (hot plasma with all finite temperature corrections retained). A physics study to iron out various details is postponed to a more extended paper.

5. A note on optimisation

We anticipate several avenues to optimise code and thereby minimise the CPU time required to assemble and solve the linear system corresponding to the equation. Reducing the required CPU time and memory is envisaged using a number of simple ideas.

(i) Jaeger's method (Jaeger 2001) does not involve finite elements nor does it involve a local representation of the field. It yields a system with a full system matrix and with the (global) Fourier components of the electric field as the unknowns. In the set of equations he adopts – but limiting oneself to only one-dimensional application in this paper – the electric field is first written in terms of a discretised form of its Fourier representation $E = \int \mathrm{d}k E_k \exp[ikx] \approx \sum_m E_{k_m} \exp[ik_m x]$ and then the wave equation is imposed at as many grid points as there are Fourier components to be found, allowing us to find a unique solution when imposing boundary conditions via Lagrange multipliers. A first and direct optimisation of the method Jaeger proposed consists of only retaining key couplings, as was demonstrated in Van Eester (2024): rather than evaluating the wave equation itself and doing it at specific (x-grid) points for a given set of (k, k'), that approach locally adopts a polynomial approximation of the dielectric response in terms of x via the local $\zeta = [x - x_i]/[x_{i+1} - x_i]$ i.e. $G(x, k, k') \approx \sum_j G_j(k, k')\zeta^j$ so that integrals of the form $\int \mathrm{d}x F^*_{k'} G(x, k, k') E_k \exp[i(k - k')x] = F^*_{k'}[\int \mathrm{d}x G(x, k, k') \exp[i(k - k')x]] E_k$ can be evaluated by hand after subdividing the physical domain into a suitable set of finite elements. This method partly uses the philosophy of the finite element technique as it splits the domain into finite sections $[x_i, x_{i+1}]$ but still exploits global base functions ($\exp[ik_x x]$) rather than local polynomial ones to write the electric field. The key idea is that the integration over x acts as a filter: it is found that the number of side diagonals (i.e. couplings between modes) to be retained is a small fraction of the total number of unknowns since the integral tends to approach zero the more k is separated from k' because of the gradually more rapidly oscillating factor $\exp[i(k - k')x]$ in the integrand. Going from a point-by-point to an element-by-element approach is essential: all local evaluations of couplings – as was done in AORSA – have similar amplitudes while the integrated response allows us to set apart important and less important couplings.

Dedicated sparse matrix solvers can be exploited. A crude guess of the CPU time and memory gain can immediately be provided from the scaling for simple Gaussian elimination. When the number of unknowns is n and the fraction of important side diagonals is f then matrix inversion when using brute force inversion computational effort scales as n^3 while if only considering the needed side diagonals the time required scales as $n^3 f^2$. If $f = 0.2$ then this reduces the needed effort by a factor 25. Going beyond one dimension increases the absolute but reduces the relative need of side diagonals; in two dimensions the relative gain is a further factor 25.

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

(ii) Aside from the time spent actually solving the system, time gain can also be achieved when assembling the system. In view of the fine structure in k-space that can be incorporated, the assembly will actually be the main place where optimisation is required. Compared with a matrix fill factor 1, only retaining the couplings actually needed reduces the number of evaluations to be made by the same factor $1/f^2$.

(iii) Moreover, the faster variation of the integrands resulting from $\exp[i(k-k')x]$ rather than from $G(x, k, k')$ – which involves Bessel functions that commonly vary fairly slowly as a function of k – the function G can be evaluated on a crude grid subset of k and k' values, and the intermediate values on the actual k-grid can be computed by simple interpolation. This procedure need not be static: monitoring the local k-gradient of the function G, an automatic procedure for grid refinement can be installed to only zoom in on k-regions where more detail is actually needed. Recall moreover that the Fourier components of the polynomial base functions only need to be evaluated once and for all. Although allowing us to include effects that normally require an integro-differential approach and yield a well-filled system matrix, the method proposed here is local and hence the system is sparse both in x- and (k, k')-space.

In spite of the fact that the accent in the present paper is on the numerical method and not on the physics model, it is worthwhile mentioning that there is a basic difference between the dielectric tensor considered by Jaeger and the dielectric response operator used here to assess the wave propagation in the plasma using the all-FLR model: Jaeger exploits the homogeneous plasma tensor while the operator exploited here allows direct interfacing with the Fokker–Planck equation and properly ensures positive definite power absorption for Maxwellian distribution functions i.e. properly accounts for plasma inhomogeneity; for a discussion see Van Eester *et al*. (2026).

6. Conclusions and prospects

A general purpose solution technique inspired by a mix of the finite element and Fourier analysis methods has been proposed to find the solution of the ICRH wave equation accounting for parallel as well as perpendicular gradients. It relies on a representation of localised finite element base functions in terms of a set of 'quasi-modes'. The philosophy is to first find a 'quasi-Fourier-mode' representation of the local base functions by performing a discrete Fourier transform in a reference interval, padding neighbouring elements by zeros to achieve a suitable grid sampling in k-space to enable capturing of the needed physics of the wave–particle interaction. In view of the direct physical interpretation of the unknowns associated with them – the values of the electric field components and their derivatives at the finite element grid points – the Hermite functions have been chosen as reference base functions. It was illustrated that this philosophy can easily be extended to ensure higher-order derivatives of the unknowns are continuous as well. As the technique locally reproduces spatial dependence of the base function but does so in terms of k-modes, it allows accounting for the full k-spectrum of the dielectric response. The integro-differential all-FLR dielectric response of the inhomogeneous plasma in the presence of both parallel and perpendicular gradients is then accounted for by a simple vector–matrix–vector representation. The shape of the obtained linear equation system is identical to that when using the finite element approach, the only

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

difference being that the system coefficients are evaluated in an alternative way. The constructive and destructive interference yielding the coupling of nearby modes and the decoupling of well-separated ones is captured in a matrix that can be computed once and for all.

The method combines the possibility of exploiting all-FLR dielectric tensor descriptions formulated in k-space available in the literature while accounting for parallel gradient effects (arising due to the presence of the finite poloidal magnetic field or toroidal non-axisymmetry). The method allows us to describe short as well as long wavelength modes and offers options for CPU speed-up, combining advantages of various methods earlier proposed in the literature.

To illustrate the potential of the method, a few relevant examples have been provided. They illustrate the fast wave mode conversion to either the Bernstein or the ion-cyclotron wave depending on the relative importance of temperature and poloidal field effects for typical JET H minority heating parameters. Three types of models have been considered to illustrate the method, all of which are already known in the literature: a cold plasma description that allows to highlight the impact of the presence of a $k_{//}$-upshift and giving rise to the emergence of the ion-cyclotron wave, a 'tepid' (FLR-0) plasma description lacking the finite temperature corrected perpendicular dynamics but having the parallel dynamics resulting from the poloidal field and when the temperature is finite, and a full hot plasma (all-FLR) dielectric tensor model allowing us to account both for parallel and perpendicular gradients illustrating how the Bernstein wave emerges in the description.

The focus of the present paper is the proposed numerical technique rather than a physics exploration. Only already existing models for the dielectric response were exploited here. The level of physics realism of a model is the level of richness of the physics model adopted. The examples provided rely on the usual assumption that the electric field is 'adiabatically switched on' so that only the end point contribution along the orbit contributes.

The technique can readily be generalised to multiple dimensions. Together with an in-depth physics study (e.g. adding a power deposition computation via post-processing to assess the particular role of finite temperature corrections and parallel gradients), going to two-dimensional (axisymmetric tokamak) or possibly even three-dimensional (stellarator) applications, and the incorporation of the option to account for non-Maxwellian distributions are intended next steps of the work. Optimisation of CPU time and memory along the lines discussed earlier will be necessary subpart of these efforts.

Jaeger adopts 'Cartesian' coordinates R and Z – and hence the associated Fourier modes $\exp[i(k_R R + k_Z Z + n\varphi)]$ – in a toroidal cut $\varphi = \text{constant}$ in a tokamak and equally adopts $\boldsymbol{E}$-field components (E_R, E_Z, E_φ) along the associated coordinate lines. Although appealing due to its simplicity, a drawback of this choice is that neither magnetic surfaces nor the metallic walls surrounding the plasma coincide with edges of the adopted finite element volumes. The approach presented here does not involve the actual (global) $\boldsymbol{E}$-field's Fourier components but just mimics the Fourier transform to locally represent the electric field and its gradients, while yielding a system of equations that formally is identical to that obtained using finite elements (be it that the coefficients are assembled in a different way). This allows a bigger freedom in the choice of the elementary volumes. When going beyond one-dimensional application, an appealing choice is to exploit a field-aligned reference and adopt (deformed) cubes or prisms, one direction of the sides of which is in the parallel direction, while the other 2 are in the plane perpendicular to it

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

but 'as close as possible' to the unit vectors $\boldsymbol{e}_R$ and $\boldsymbol{e}_Z$. Such a choice was already pioneered by Jaeger. It ensures the magnetic axis is a regular point. Cubes more easily allow defining locally independent coordinates to immediately extend the philosophy adopted in the present paper to multi-dimensional application, but prevent easy fitting all around curved surfaces such as metallic walls. Triangle-shaped faces of a prism more elegantly and easily allow us to approximate any closed shape on a surface. Barycentric coordinates (one of which is dependent on the other 2, and each of the 3 being identically 0 on only one of the triangle sides and 1 at the opposite corner) are typically exploited, and associated base functions that allow us to form two-dimensional 'hat' functions inside the triangle can be used; depending on the exact choice of base functions surface terms will need to be included or can be omitted. Finding the most optimal choice is left for later work.

Although illustrated here for the specific problem of ICRH, the method presented here is likely useful for other applications.

Acknowledgements

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

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

Declaration of interests

The authors report no conflicts of interest.

REFERENCES

FUKUYAMA, A. & KHAN, S.A. 2025 Kinetic full wave analyses in inhomogeneous plasmas using integral form of dielectric tensor. In *RF Power in Plasmas Conference (Hohenkammer, Germany, 2025)*. Avialable at: https://www.ipp.mpg.de/5533835/.

GREEN, D.L., BERRY, L.A., SIMPSON, A.B. & YOUNKIN, T.R. 2018 Kinetic-j: a computational kernel for solving the linearized Vlasov equation applied to calculations of the kinetic, configuration space plasma current for time harmonic wave electric fields. *Comput. Phys. Commun.* **225**, 122.

ICHIMARU, S. 1973 *Basic Principles of Plasma Physics. A Statistical Approach*. W.A. Benjamin Inc.

JAEGER, E.F., BERRY, L.A., D'AZEVEDO, E., BATCHELOR, D.B. & CARTER, M.D. 2001 All-orders spectral calculation of radio-frequency heating in two-dimensional toroidal plasmas. *Phys. Plasmas* **8**, 1573.

KENNEL, C.F. & ENGELMANN, F. 1966 Velocity space diffusion from weak plasma turbulence in a magnetic field. *Phys. Fluids* **9**, 2377.

LAMALLE, P.U. 1997 On the radiofrequency response of tokamak plasmas. *Plasma Phys. Control. Fusion* **39**, 1409.

LAMALLE, P.U., REMAN, B.C.G., SLABY, CHR, VAN EESTER, D., LOUCHE, F., GEUZAINE, CHR, ZALESKI, J. & MORAL SANCHEZ, E. 2025 Integral dielectric kernel approach to modelling RF heating in toroidal plasmas. In *RF Power in Plasmas Conference (Hohenkammer, Germany, 2025)*. Avialable at: https://www.ipp.mpg.de/5533820/.

MACHIELSEN, M. 2023 Exact expression for the hot plasma conductivity kernel in configuration space. *Fundam. Plasma Phys.* **3**, 100008.

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press

MAQUET, V., REMAN, B.C.G., VAN EESTER, D., ADRIAENS, A. & RAGONA, R. 2026 First Implementation of a Fast 2D Wave-Solver Using the Budé Method with NGSolve: initial validation without poloidal fields, with further investigation incorporating them. In *25th Topical Conference on Radio-Frequency Power in Plasmas (RFPPC2025)*, *EPJ Web of Conferences*, vol. 346, p. 01020.

PERKINS, F.W. 1977 Heating tokamaks via the ion-cyclotron and ion-ion hybrid resonances. *Nucl. Fusion* **17**, 1197.

REMAN, B.C.G., LAMALLE, P.U., ZALESKI, J., SLABY, C., VAN EESTER, D., GEUZAINE, C., LOUCHE, F., LERCHE, E.A., MORAL SANCHEZ, E. & MAQUET, V. 2026 Integral dielectric kernel implementation to model RF heating in toroidal plasmas. In *25th Topical Conference on Radio-Frequency Power in Plasmas (RFPPC2025), EPJ Web of Conferences*, vol. 346, p. 01019.

SHIRAIWA, S., BERTELLI, N., KIM, E.H., WRIGHT, J., VERSTRAETEN, M. & LAMALLE, P.U. 2026 Finite element modeling of RF waves in fusion plasmas: progress in past decades and future role. In *25th Topical Conference on Radio-Frequency Power in Plasmas (RFPPC2025)*, *EPJ Web of Conferences*, vol. 346, p. 01001.

SHIRAIWA, S., MENEGHINI, O., PARKER, R., BONOLI, P., GARRETT, M., KAUFMAN, M.C., WRIGHT, J.C. & WUKITCH, S. 2010 Plasma wave simulation based on a versatile FEM solver. *Phys. Plasmas* **17**, 056119.

STIX, T.H. 1997 *Waves in Plasmas*. American Institute of Physics.

SVIDZINSKI, V.A., KIM, J.S., SPENCER, J.A., ZHAO, L., GALKIN, S.A. & EVSTATIEV, E.G. 2016 Hot plasma dielectric response to radio-frequency fields in inhomogeneous magnetic field. *Phys. Plasmas* **23**, 112101.

SWANSON, D.G. 2012 *Plasma Waves*. Academic Press.

VAN EESTER, D. 2026 Towards solving the ICRH wave and Fokker–Planck equations self-consistently. submitted to *Nucl. Fusion*. arXiv 2601.02130.

VAN EESTER, D. & KOCH, R. 1998 A variational principle for studying fast-wave mode conversion. *Plasma Phys. Control. Fusion* **40**, 1949.

VAN EESTER, D. & LERCHE, E.A. 2021 Solving the all-FLR ICRH integro-differential wave equation as a high-order differential equation for studying combined ICRH-NBI heating. *Nucl. Fusion* **61**, 016024.

VAN EESTER, D. & LERCHE, E.A. 2024 A mixed Fourier-variational approach to solve differential or integro-differential wave equations for magnetised plasmas. *Plasma Phys. Control. Fusion* **66**, 045002.

VAN EESTER, D., MAQUET, V., REMAN, B.C.G., CARPIAUX, H., VALENTIN, B., LERCHE, E.A. & LAMALLE, P.U. 2026 Almost-off-the-shelf tools for ICRH modelling. In *25th Topical Conference on Radio-Frequency Power in Plasmas (RFPPC2025)*, *EPJ Web of Conferences*, vol. 346, p. 01016.

https://doi.org/10.1017/S0022377826101639 Published online by Cambridge University Press