

TORBEAM 2.0, a paraxial beam tracing code for electron-cyclotron beams in fusion plasmas for extended physics applications

E. Poli ^{a,*}, A. Bock ^a, M. Lochbrunner ^a, O. Maj ^a, M. Reich ^a, A. Snicker ^{a,b}, A. Stegmeir ^a, F. Volpe ^c, N. Bertelli ^d, R. Bilato ^a, G.D. Conway ^a, D. Farina ^e, F. Felici ^f, L. Figini ^e, R. Fischer ^a, C. Galperti ^f, T. Happel ^a, Y.R. Lin-Liu ^g, N.B. Marushchenko ^h, U. Mszanowski ^a, F.M. Poli ^c, J. Stober ^a, E. Westerhof ⁱ, R. Zille ^a, A.G. Peeters ^j, G.V. Pereverzev ^{a,1}

^a Max-Planck-Institut für Plasmaphysik, Garching bei München, Germany

^b Department of Applied Physics, Aalto University, Finland

^c Columbia University, New York, NY, USA

^d Princeton Plasma Physics Laboratory, Princeton, NJ, USA

^e Istituto di Fisica del Plasma CNR, Milano, Italy

^f Swiss Plasma Center, EPFL Lausanne, Switzerland

^g Department of Physics and Center for Mathematics and Theoretical Physics, National Central University, Taiwan

^h Max-Planck-Institut für Plasmaphysik, Teilinstitut Greifswald, Germany

ⁱ DIFFER, Nieuwegein, The Netherlands

^j Theoretical Physics V, Department of Physics, University of Bayreuth, Germany

ARTICLE INFO

Article history:

Received 13 October 2017

Received in revised form 7 December 2017

Accepted 22 December 2017

Available online 8 January 2018

Keywords:

Plasma physics

Magnetic confinement

Wave-plasma interactions

Electron cyclotron waves

Paraxial beam tracing

ABSTRACT

The paraxial WKB code TORBEAM (Poli, 2001) is widely used for the description of electron-cyclotron waves in fusion plasmas, retaining diffraction effects through the solution of a set of ordinary differential equations. With respect to its original form, the code has undergone significant transformations and extensions, in terms of both the physical model and the spectrum of applications. The code has been rewritten in Fortran 90 and transformed into a library, which can be called from within different (not necessarily Fortran-based) workflows. The models for both absorption and current drive have been extended, including e.g. fully-relativistic calculation of the absorption coefficient, momentum conservation in electron–electron collisions and the contribution of more than one harmonic to current drive. The code can be run also for reflectometry applications, with relativistic corrections for the electron mass. Formulas that provide the coupling between the reflected beam and the receiver have been developed. Accelerated versions of the code are available, with the reduced physics goal of inferring the location of maximum absorption (including or not the total driven current) for a given setting of the launcher mirrors. Optionally, plasma volumes within given flux surfaces and corresponding values of minimum and maximum magnetic field can be provided externally to speed up the calculation of full driven-current profiles. These can be employed in real-time control algorithms or for fast data analysis.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

The paraxial Wentzel–Kramers–Brillouin (pWKB) method for the propagation of high-frequency waves in plasmas, developed by Pereverzev [1,2], has been implemented numerically in the beam-tracing code TORBEAM [3] and since then used in a variety of studies for present [4–6] and future [7–9] fusion devices, for both heating and current drive applications employing waves in the electron

cyclotron (EC) frequency range and for the analysis and interpretation of short-wavelength diagnostic techniques like reflectometry [10–13]. Applications of the pWKB method to lower-hybrid wave propagation have also been reported [14,15]. TORBEAM has been included in a worldwide benchmark of EC codes presented in [16].

Fig. 1 shows a typical plot of the poloidal projection of the beam propagation in the ASDEX Upgrade tokamak.

The pWKB method reduces the computational effort of solving the vector wave equation in an anisotropic, inhomogeneous and dissipative plasma to the solution of a set of ordinary differential equations (ODEs) describing (a) the centre of the beam (the so-called *reference ray*), (b) the transverse beam structure, expressed by two 3×3 symmetric matrices representing the curvature of

* Corresponding author.

E-mail address: emanuele.poli@ipp.mpg.de (E. Poli).

¹ Deceased.

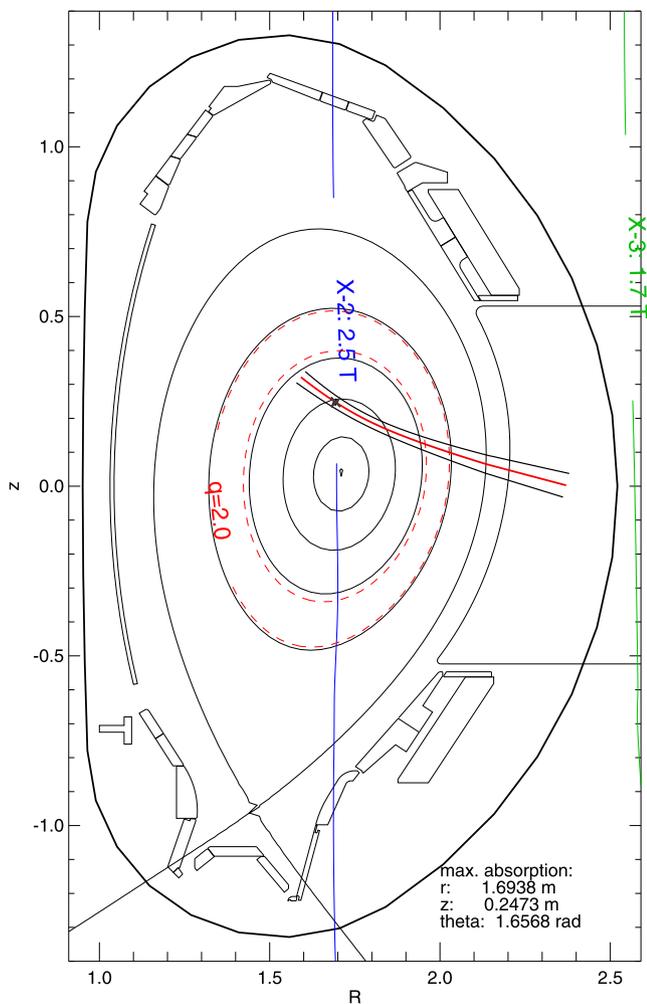


Fig. 1. Propagation of an X-mode, 140 GHz electron cyclotron beam in ASDEX Upgrade as calculated by TORBEAM, visualized as a projection along the cylindrical angle φ . The beam, launched at $(R, Z) = (2.38, 0)$ m, is represented through its reference ray (tracing the maximum of the Gaussian amplitude envelope, red) and the beam width (“peripheral rays”) representing the intersection of the $1/e$ level of the amplitude with a vertical plane. The vertical lines show the position of the second (X-2, blue) and third (X-3, green) cyclotron harmonics. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the beam front and the amplitude profile, assumed to be localized around the reference ray, and (c) the power carried by the beam (which is calculated on the reference ray only). The wave-field amplitude profile transverse to the propagation direction is assumed in the numerical implementation to have a Gaussian shape, although the theory is developed for more generic beams [2]. As a result, the pWKB technique allows the determination of propagation and absorption of EC waves in fusion plasmas, retaining diffraction effects not included in standard ray-tracing codes, at a reduced computational cost even with respect to ray tracing (if more than 3 rays traced). The only numerical “price” that must be paid is the evaluation of second-order derivatives of the involved plasma profiles, but this is in general not critical, because the related overhead is basically limited to the pre-processing of the equilibrium, i.e. the calculation of the finite-difference derivatives on the grid, as explained in Section 3. Due to its numerical simplicity, the method lends itself in particular to applications requiring fast execution times, in particular the inclusion in more comprehensive procedures like transport codes (TORBEAM has been interfaced with the transport solvers ASTRA [17] and TRANSP [18]), integrated

data analysis tools, integrated tokamak modelling [19] (see [20] for a benchmark of TORBEAM against other European EC codes in the frame of the European Integrated Tokamak Modelling framework) or as a part of real-time control algorithms to be executed in parallel to the evolution of experimental plasma discharges [21].

Since the first publication [3], the code has undergone a profound revision. First of all, it has been rewritten in Fortran 90 (the original programming language was Fortran 77) and the source has been split into building blocks, while the modular structure of the original code has been largely conserved. A brief overview of the TORBEAM workflow is given in Section 3. The code has been transformed into a library, which is called by an external process (not necessarily in Fortran), also in charge of the input/output operations. Fortran wrappers for the standalone execution of the programme are available to the user. The main motivation behind this development was the need for an accelerated version to be used in real-time applications [21]. Within this framework, different levels of simplification of the physics scope are now possible, as described in Section 6.

The original goal of TORBEAM was the calculation of power density (power per unit volume) and current density (current through unit surface) profiles allowing for diffraction effects. This still applies to the code in its present form. With respect to [3], the determination of the absorption coefficient can be performed now also employing fully relativistic expressions for the dielectric-tensor elements. The calculation of the current-drive efficiency has been extended by including the full polarization term in the expression of the quasilinear diffusion coefficient and the momentum-conserving scheme described in [22–24]. A subtle point concerning the implementation of the paraxial WKB method is the algorithm for the calculation of the profiles. Since the power-absorption equation is solved along the reference ray only, a scheme has to be devised for distributing the absorbed power across the involved flux surfaces according to the Gaussian amplitude profile of the beam. Two such schemes are now available, the first one being faster but less accurate near the magnetic axis of the plasma, the second being able to cope also with such cases. A related problem arises when TORBEAM is coupled with a Fokker–Planck code like RELAX [25], which is designed to accept as an input the wave parameters on individual rays, an option to extract this information along so-called extended rays, carrying the information about the energy flow of the diffracting beam, has been implemented [26]. The modifications mentioned above are explained in detail in Sections 4 and 5.

As stated at the beginning, the code is routinely used as a support in the interpretation of reflectometry measurements and to design future diagnostics. The most important modification required by these applications concerns the inclusion of an effective electron mass, following [27], in the cold-plasma dielectric tensor, as required in reactor-grade plasmas [27,28]. Moreover, formulas describing the beam coupling to the receiver in terms of the beam tracing parameters (solutions of the pWKB equations, see Section 2) have been developed and are available as part of the post-processing tools. These extensions of the code, already documented in [11], are briefly reviewed in Section 7.

In its present form, the code assumes toroidal symmetry, i.e. tokamak geometry, although the underlying equations are formulated in arbitrary geometry [2].

The main goal of this paper is to give an overview of the current capabilities of the code, including some selected examples, together with a compact description of the underlying numerical implementation, intended also as a first guidance for the users of the code. For more details on the physical model, the reader is referred to the original publications by Pereverzev [1,2], and to [29–32] for some illustrative examples and applications in critical cases.

2. The paraxial WKB method for the propagation of high-frequency waves in inhomogeneous anisotropic media

The beam tracing method is applied to the solution of the vector wave equation

$$\frac{c^2}{\omega^2} \nabla \times (\nabla \times \mathbf{E}) - \boldsymbol{\epsilon} \cdot \mathbf{E} = 0 \quad (1)$$

in the high-frequency (short wavelength) limit

$$\kappa \equiv \frac{\omega L}{c} = k_0 L \gg 1, \quad (2)$$

where $\boldsymbol{\epsilon}$ is the dielectric tensor, $\omega/2\pi$ is the wave frequency, c is the speed of light, L the inhomogeneity scale of the medium and k_0 the vacuum wave vector. A simplified derivation valid for a Gaussian beam follows from the complex-eikonal *ansatz* for the electric field

$$\mathbf{E}(\mathbf{r}) = A(\mathbf{r})\mathbf{e}(\mathbf{r})e^{ik_0[s(\mathbf{r})+i\phi(\mathbf{r})]}, \quad (3)$$

where the imaginary part of the phase (eikonal) $\phi \geq 0$ is introduced to describe the profile of the field amplitude across the beam cross-section. The wave field (3) is localized in a tubular neighbourhood of size $\mathcal{O}(1/\sqrt{\kappa})$ around the curve defined by $\phi = 0$ (the *reference ray*, see below). Physically, this means that the beam width W satisfies

$$\frac{\lambda}{W} \approx \frac{W}{L} \approx \frac{1}{\sqrt{\kappa}} \quad (4)$$

(λ is the vacuum wavelength). The complex phase of the wave field in Eq. (3) can be Taylor expanded around the reference ray (paraxial expansion) [1,2]:

$$s(\mathbf{r}) = s_0(\mathbf{r}) + \mathcal{N}_\alpha [x_\alpha - q_\alpha] + \frac{1}{2} s_{\alpha\beta} [x_\alpha - q_\alpha][x_\beta - q_\beta], \quad (5)$$

$$\phi(\mathbf{r}) = \frac{1}{2} \phi_{\alpha\beta} [x_\alpha - q_\alpha][x_\beta - q_\beta], \quad (6)$$

where q_α and \mathcal{N}_α are the components of the position vector $\{x_\alpha\} \equiv \mathbf{r}$ and of the vector refractive index $\{N_\alpha\} \equiv \nabla s$, respectively, calculated on the reference ray. Summation over repeated indices is assumed. The second-order coefficients $s_{\alpha\beta}$ describe the change of the wave vector along the wave front and are hence related to its curvature. The coefficients $\phi_{\alpha\beta}$ appear in the real part of the exponential factor of Eq. (3) and are then connected with the width of the amplitude profile, see Eqs. (27) and (28) of [3].

The quantities q_α , \mathcal{N}_α , $s_{\alpha\beta}$ and $\phi_{\alpha\beta}$ are found as the solution of the following set of *ordinary* differential equations (ODEs), known as beam tracing equations [2]:

$$\frac{dq_\alpha}{d\tau} = \frac{\partial H}{\partial N_\alpha}, \quad \frac{d\mathcal{N}_\alpha}{d\tau} = -\frac{\partial H}{\partial x_\alpha}, \quad (7)$$

$$\begin{aligned} \frac{ds_{\alpha\beta}}{d\tau} = & -\frac{\partial^2 H}{\partial x_\alpha \partial x_\beta} - \frac{\partial^2 H}{\partial x_\beta \partial N_\gamma} s_{\alpha\gamma} \\ & - \frac{\partial^2 H}{\partial x_\alpha \partial N_\gamma} s_{\beta\gamma} - \frac{\partial^2 H}{\partial N_\gamma \partial N_\delta} s_{\alpha\gamma} s_{\beta\delta} + \frac{\partial^2 H}{\partial N_\gamma \partial N_\delta} \phi_{\alpha\gamma} \phi_{\beta\delta}, \end{aligned} \quad (8)$$

$$\begin{aligned} \frac{d\phi_{\alpha\beta}}{d\tau} = & -\left(\frac{\partial^2 H}{\partial x_\alpha \partial N_\gamma} + \frac{\partial^2 H}{\partial N_\gamma \partial N_\delta} s_{\alpha\delta} \right) \phi_{\beta\gamma} \\ & - \left(\frac{\partial^2 H}{\partial x_\beta \partial N_\gamma} + \frac{\partial^2 H}{\partial N_\gamma \partial N_\delta} s_{\beta\delta} \right) \phi_{\alpha\gamma}, \end{aligned} \quad (9)$$

where

$$H \equiv \det[-N^2 \mathbf{I} + \mathbf{N}\mathbf{N} + \boldsymbol{\epsilon}^h] = 0 \quad (10)$$

is the dispersion function of geometrical optics. Like in standard ray tracing, the anti-Hermitian part of the dielectric tensor is ordered much smaller (by $1/\kappa$) than the Hermitian part $\boldsymbol{\epsilon}^h$, so that H

depends only on the Hermitian part of the dielectric tensor and hence is real. All the derivatives of H on the right-hand sides of Eqs. (7)–(9) are to be calculated on the *reference ray*, which obeys the Hamiltonian (ray tracing) Eqs. (7).

Since $s_{\alpha\beta}$ and $\phi_{\alpha\beta}$ are symmetric 3×3 matrices, the solution of the beam tracing equations requires the integration of 18 ODEs. According to the paraxial expansion, the equation for the beam power, which descends from the equation for the terms of order $1/\kappa$ in the asymptotic expansion of the wave field, has to be evaluated on the reference ray only. This equation takes the form

$$\frac{dP}{d\tau} = -2\gamma P, \quad (11)$$

where γ is the absorption coefficient, which is defined in terms of the anti-Hermitian part of the dielectric tensor. The whole beam structure, retaining diffraction effects, is thus described through 19 ODEs (to be compared to 7 equations per ray in the case of standard geometrical optics).

3. Code structure

The TORBEAM code as described in [3] was written as a single Fortran 77 file. In the course of its “translation” into Fortran 90, the code was split into several “building blocks” (files), which closely reflect the original structure. This new arrangement of the code is described in this section.

As already mentioned before, the MAIN part of the code has been transformed into a subroutine `torbeam.f90` which represents the library interface. All the arguments of the TORBEAM library are either scalars or one-dimensional arrays, to avoid conflicts or need for re-ordering (row-major vs. column-major). The process calling the library provides as an input the information about the magnetic equilibrium (`eqdata`), the electron density and electron temperature profiles (`prdata`), the dimensions of the respective grids and a set of integer (`intinbeam`) and real (`floatinbeam`) parameters containing the initial conditions for the beam, the main machine parameters, the required integration accuracy and a set of switch parameters for different physics options available in the code, as detailed in the remainder of this paper.

The output varies depending on the application, which is selected through precompiler options. In “real-time” runs, only a reduced set of output values is returned (`rhoresult`), which includes the coordinates of the point of maximum absorption and, if current drive is switched on, some additional ray and plasma parameters at that position, the total driven current and a single-ray estimate of the profile width. The differences between the “real-time” version and the complete version of the code are reviewed in Section 6. For standard heating and current-drive applications, the information returned to the calling process by the library is extended to include the beam trajectory on a poloidal (`t1data`) and toroidal (`t1tdata`) projection, the power deposition and current density profiles (`t2ndata`) and (for diagnostic purposes) the plasma volume as a function of the radial coordinate `volprof`. For reflectometry applications, the output is further extended to include the radial coordinate of the turning point and the values of the density, the refractive index (perpendicular and parallel components) and the magnetic field components at the turning point (`reflout`). No input–output operation needs to be performed from within the library. A complete description of the input and output parameters of the TORBEAM library is provided with the code.

Once called, the top level of the TORBEAM library performs:

1. The call to the initialization routines, i.e. (a) the conversion of the input as provided by the library arguments into the internal variables of the code and the calculation of the initial conditions for the 19 beam tracing equations (`in.f90`), and

- (b) the computation of the finite-difference derivatives of the plasma profiles required by Eqs. (7)–(9), of the plasma-volume profile and of the flux-surface averages needed for the determination of the driven current (`grid6.f90`);
- The main integration loop, which contains the call to the ODE solver LSODE (`lsode.f90`) [33]. A step in this loop advances the beam tracing equations from τ to $\tau + d\tau$ (this is the central operation performed by the code). When the density on the reference ray reaches a preset value, the integration switches from vacuum to plasma dispersion and the initial conditions¹ are re-calculated in the plasma according to [34] (`interf.f90`). After the integration step, some quantities needed for code diagnostics and for the determination of the driven current are calculated at the new (advanced) point through a call to `coef.f90`. Finally, if the power absorbed during the integration step exceeds a threshold, the current driven is calculated (`currn.f90`);
 - The preparation of the output quantities, mainly the power density and current density profiles (`out.f90`). The algorithms for the calculation of the profiles are described in Section 4.

From this level on, the code has the “cascade” structure inherited from the original version: the ODE integrator `lsode` requires the routine where the beam tracing equations (7)–(11) are specified (`eqs_beam.f90`). This routine in turn requires the ODE coefficients, i.e. the derivatives of H appearing on the right-hand side of the beam tracing equations Eqs. (7)–(9) and the absorption coefficient (`deriv.f90`). The derivatives are evaluated performing a rotation of the reference system to the poloidal plane containing the current integration point (`rot.f90`) and evaluating the electron density and magnetic-field derivatives appearing in the plasma frequency and the cyclotron frequency (`ocop.f90`) and the values of the perpendicular and parallel wave vector (`refr.f90`). The numerical evaluation of the derivatives is achieved through polynomial interpolation (`interp1B.f90`) on the finite-difference grids stored before the start of the integration in a neighbourhood of the current position of the central ray.

A comparison with Sec. 3 of [3] shows that the basic workflow of the code has not changed significantly. For the calculation of the absorption coefficient, the original TORBEAM model has been extended to include the fully-relativistic version of the routine `DAMPBQ` [35] (`westerino.f90`) and the routine `WARMDISP` extracted from the quasi-optical GRAY code [36] (`ecdisp.f90`), which also accounts for the fully-relativistic elements of the dielectric tensor. For current drive, the original `CURBA` routine (`curba.f90`) has been supplemented by the routine `CURGAP` [37] (`TorGA_curgap.f90`) and modifications thereof, as detailed in Section 5 (the momentum-conserving Spitzer function described there is contained in `green_func_ext.f90`).

4. Calculation of power-deposition profiles

The evaluation of the absorption coefficient γ in Eq. (11) is based on the calculation of the imaginary part of the wave vector as obtained by solving the hot-plasma dispersion equation. Originally [3], the solution obtained in the subroutine `DAMPBQ` was limited to the weak-relativistic limit of the hot-plasma dielectric-tensor elements (implying that the relativistic Lorentz factor is expanded to the lowest significant order in the ratio between the particle speed and the speed of light). As stated in Section 1,

¹ The mode of propagation is selected through the corresponding initial conditions for the wave vector, the Hamiltonian function H being provided by Eq. (10) for both modes. Possible cases of imperfect mode purity at launch can be handled through two separate runs of the code, one for each cold-plasma mode.

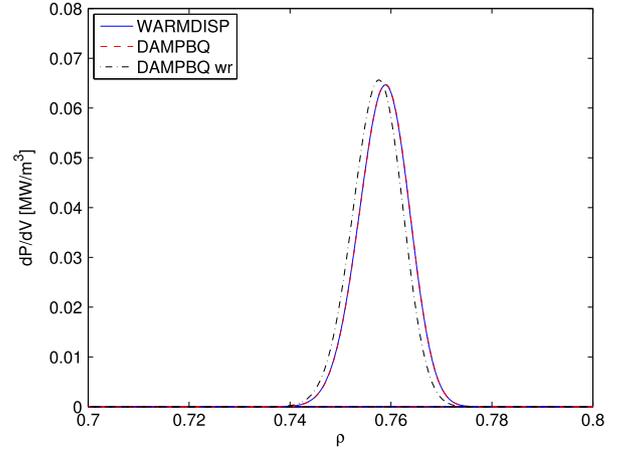


Fig. 2. Power deposition profile obtained with the subroutine `WARMDISP` (blue solid), `DAMPBQ` (red dashed, virtually indistinguishable from the previous one) and with `DAMPBQ` in the weakly-relativistic limit (black dashed-dotted) for a typical ITER case (stabilization of neoclassical tearing modes on the $q = 2$ surface).

the present version of TORBEAM allows for a fully relativistic calculation of the absorption coefficient through the corresponding extension of `DAMPBQ` (from TORAY) and through the subroutine `WARMDISP` (from GRAY). Both routines also account for absorption due to harmonics higher than the second. The results of both routines are nearly identical, as shown in Fig. 2, in which the absorption profiles obtained in the weakly-relativistic approximation (black dash-dotted curves) are compared with fully-relativistic results (practically indistinguishable blue solid and red dashed lines).

While `WARMDISP` returns the perpendicular component of the imaginary part of the wave refractive index N''_{\perp} , `DAMPBQ` returns the projection of N''_{\perp} onto the real refractive index N' [35]. As a consequence, the output of `DAMPBQ` must be divided by the sine of the angle between N' and the confinement magnetic field to be equivalent to the output of `WARMDISP`. The absorption coefficient is calculated as given in Eq. 15) of [3], i.e. $\gamma = (\omega/c)\mathbf{N}'' \cdot \mathbf{V}$, with $\mathbf{V} = \partial H / \partial \mathbf{N}$ (vector parallel to the group velocity).

The spatial distribution of the absorbed power is given as a power-density profile as a function of radial (flux-surface) label, under the usual assumption of an “ergodic” redistribution of the absorbed power along the field lines. As mentioned in Section 1, within the paraxial WKB method the calculation of the power density profile is not as immediate as in the case of ray-based algorithms, in which the power absorption is calculated on each ray and can be assigned straightforwardly to the respective flux surfaces. In the paraxial approach employed in TORBEAM, the quantities directly available after the integration of the beam tracing equations are the power along the reference ray and the Gaussian envelope of the beam. Two algorithms have been developed to compute the power deposition profile starting from this information and are described below. Both are run after the main integration loop has been executed, as described in the previous section.

In the first algorithm, besides the coordinates of the reference ray, the coordinates of an upper and a lower peripheral “ray” are also stored. These are defined as the curves in the poloidal plane corresponding to $1/e$ levels of the electric-field amplitude, determined from the Gaussian amplitude envelope $\phi_{\alpha\beta}$ (Fig. 1). For each point i on the reference ray showing non-zero absorption, the points on the peripheral rays with the same major radius (i.e. same distance from the tokamak symmetry axis) are calculated and the flux surface passing through these points is determined. The radial distance (in terms of the coordinate ρ defined below) between the

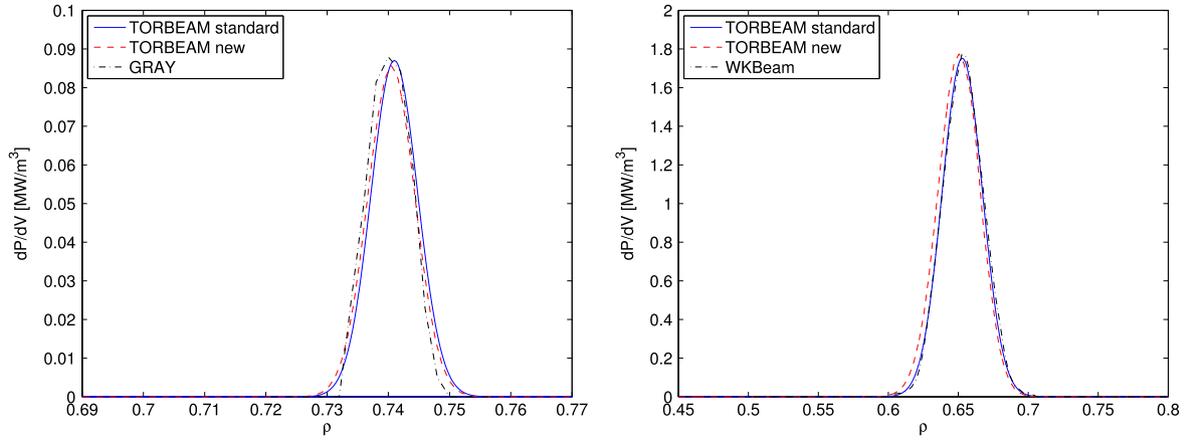


Fig. 3. Power deposition profiles obtained with TORBEAM standard procedure (blue solid) and TORBEAM new procedure (red dashed), for (a) an ITER scenario (left, also GRAY profile shown, black dashed-dotted) and (b) an off-axis ASDEX Upgrade scenario (right, also WKBeam profile shown, black dashed-dotted).

upper and the lower peripheral rays, divided by two, is taken as the half-width w_i of an “infinitesimal”, Gaussian-shaped absorption profile of the form $(dP/dV)_i = C_i \exp[-2(\rho - \rho_i)^2/w_i^2]$, where ρ is the normalized radial coordinate (square root of the normalized flux label) and the factor of two accounts for the fact that the power density is proportional to the electric field squared. The coefficient C_i is fixed in such a way that

$$C_i \int_V \exp[-2(\rho - \rho_i)^2/w_i^2] dV = C_i \int_0^1 \exp[-2(\rho - \rho_i)^2/w_i^2] \frac{dV}{d\rho} d\rho = dP_i, \quad (12)$$

dP_i being the power lost by the beam in the i th integration step and V the plasma volume. The underlying assumptions behind this procedure are that across the beam cross-section the resonance condition is satisfied on a vertical plane and that the absorption-profile width is dominated by the beam width in the poloidal plane, due to the larger curvature of the flux surfaces in the poloidal plane as compared to the curvature in the toroidal plane (the former assumption could be violated in devices for which the toroidal field is not the dominating component, like reversed-field pinches near field reversal). The total absorption profile is calculated as the sum over the infinitesimal profiles, $dP/dV = \sum_i (dP/dV)_i$.

The procedure described above is fast (a complete run of the code including the determination of power and current-drive profiles takes ca. 1 s on a single CPU²) and in most cases accurate, see examples in Fig. 3. Yet, it can fail in distributing correctly the power among flux surfaces when the maximum absorption is localized on the magnetic axis of the tokamak. For these cases, a second algorithm has been devised, based on a numerical integration of the continuity equation for the wave energy density, which in steady-state can be written as

$$\nabla \cdot (|E|^2 \mathbf{V}) = -2\gamma |E|^2, \quad (13)$$

in the region around the beam in which absorption takes place. As shown in the Appendix, from the previous equation one can derive

$$P(\rho) = \frac{2P_0}{\pi} \int_{\Omega(\rho)} \frac{e^{-2\phi(\mathbf{r})}}{|\mathbf{V}(\tau)|W_1(\tau)W_2(\tau)} \left| \frac{1}{P_0} \frac{dP}{d\tau} \right| d^3r, \quad (14)$$

where $\Omega(\rho)$ is the three-dimensional domain enclosed by the flux surface labelled by ρ , ϕ has been defined in Eq. (6), W_1 , W_2 are the principal widths of the elliptic beam cross-section (taken at $1/e$ -level of the electric field amplitude) and P_0 is the injected power.

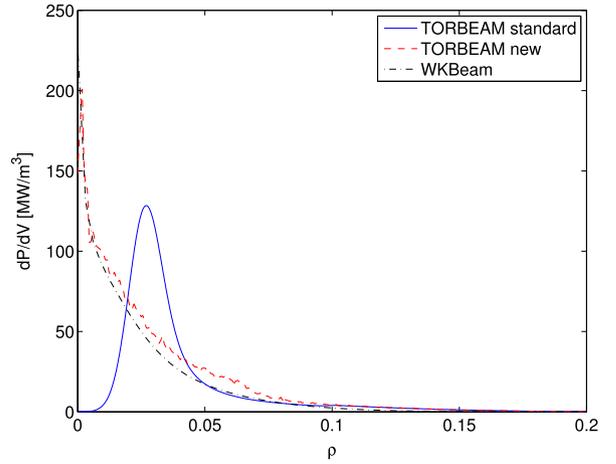


Fig. 4. Power deposition profiles obtained with TORBEAM standard procedure (blue solid), TORBEAM new procedure (red dashed) and WKBeam (black dashed-dotted) for an on-axis ASDEX Upgrade scenario.

The integration domain of Eq. (14) is discretized in a region covering at least four times the transverse beam extension. In this way, basically the whole power carried by the Gaussian beam is taken into account in the discretization. The grid is non-orthogonal, with one axis aligned with the averaged (in the absorption region) beam velocity, the second axis perpendicular to the previous one in the horizontal direction and the third axis in the vertical direction. The physical reason for this choice is again that the resonance region is approximated by a vertical line across the beam extension. Each point on the grid is assigned to a corresponding point on a ρ -grid and the integral in Eq. (14) is substituted by a discrete sum. This new procedure has been compared with the “standard” implementation described before and checked against the quasi-optical code GRAY for an ITER scenario (Fig. 3a) and against the wave-kinetic solver WKBeam code [38] for an off-axis ASDEX Upgrade scenario (Fig. 3b). The agreement of the different profiles is very good.

As a demonstration of the difference between both methods described above for a case with on-axis deposition, Fig. 4 shows a comparison for an ASDEX Upgrade experimental case. Again, the ray-based code WKBeam is used as a benchmark. It can be seen that the standard procedure fails in locating the maximum of the deposition profile on axis. Considering the numerical uncertainties related to the evaluation of the plasma volume and its derivative with respect to ρ so close to the magnetic axis, the agreement between the new procedure implemented in TORBEAM and the

² The performance of the code is discussed in detail in Section 6.

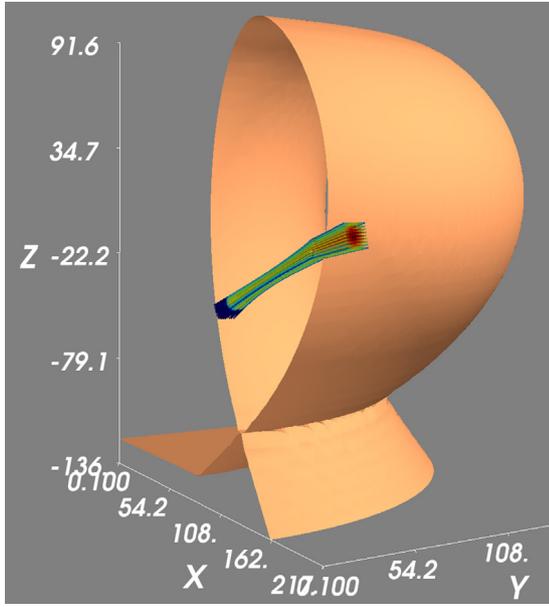


Fig. 5. Visualization of the extended rays calculated for a typical ASDEX Upgrade case.

WKBeam profile is remarkably good. Depending on the injection geometry and on the length of the absorption path, the new method requires ca. one to five seconds for the determination of the absorption profile.

As a further option, TORBEAM can produce the data required on input by the Fokker–Planck code RELAX [25] and thus allows the study of the effects due to the quasi-linear relaxation of the distribution function on the absorption profile. The wave input is provided to RELAX in terms of information (position, parallel refractive index, power fraction and cyclotron frequency) calculated on individual rays, which in the frame of geometrical optics represent the wave energy flow. To provide an analogous information in the frame of the pWKB approach, extended rays describing the energy flow of a diffracting beam can be derived [26,39]. As known from standard geometric optics [40], the action function (phase) (integral of $ds = \mathbf{N} \cdot d\mathbf{r}$ along the rays) solves the Hamilton–Jacobi equation $H(\mathbf{r}, \nabla s) = 0$ and the velocity field \mathbf{V} defined above (tangent to the rays of geometric optics) determines the wave energy transport through Eq. (13) on each ray. When the complex-eikonal ansatz (3) is introduced, assuming $\lambda/W \ll 1$, the relevant Hamilton–Jacobi equation becomes [36,41]

$$H(\mathbf{r}, \nabla s) - \frac{1}{2} \frac{\partial \phi}{\partial x_i} \frac{\partial \phi}{\partial x_j} \frac{\partial^2 H}{\partial N_i \partial N_j}(\mathbf{r}, \nabla s) = 0. \quad (15)$$

As far as the energy flux is concerned, it can be shown [39] that Eq. (13) still holds in the complex-eikonal case, but the function s is now a solution of the modified Hamilton–Jacobi equation (15). In the frame of the paraxial approach implemented in TORBEAM, ∇s can be immediately calculated taking the gradient of s given by Eq. (5) and the equations

$$\frac{d\mathbf{r}}{d\tau} = \frac{\partial H}{\partial \mathbf{N}}(\mathbf{r}, \nabla s) \quad (16)$$

can be advanced together with the beam-tracing equations employing the second-order Heun’s method, see [26] for details. Also the power content of the extended rays is computed, solving Eq. (11) on each of them. Fig. 5 shows the trajectory of the extended rays for the ASDEX Upgrade case of Fig. 3b. For the calculation of 9×9 extended rays, with the integration step halved with respect

to the standard calculation, the execution time is still well below 10 s.

5. Modelling of electron-cyclotron current drive

The calculation of the current-drive efficiency in TORBEAM is based on routines implementing the adjoint method [42]. In addition to CURBA, which was already present in the original version of the code [3], the routine CURGAP has been added. This routine includes the exact polarization-dependent radio-frequency quasilinear diffusion operator [43] in evaluating the current drive efficiency (this feature is not described in the original publication [37], where also the small gyroradius expansion is explicitly used, while the routine allows for the complete Bessel-function terms). Both CURBA and CURGAP employ the high-speed limit for the determination of the Spitzer function calculated in the frame of the adjoint method (see below). In TORBEAM, the Spitzer function can now be determined, as an extension of CURGAP, including the momentum-conserving scheme described in [22–24], which takes the transfer of momentum to the bulk electrons and their contribution to the parallel electron flow into account. Results including this effect in TORBEAM were reported for the first time in the frame of a study of the EC current-drive efficiency in DEMO-sized machines [8]. Meanwhile, this extended version of CURGAP has been extensively used in ITER calculations [9,44] and is now also the standard option for applications in the ASDEX Upgrade tokamak [45]. A further extension of CURGAP to include consistently the contribution of more than one cyclotron harmonic to the current drive has been implemented and is described in the second part of this section.

Following the notation of [37], the adjoint method requires the determination of the response function χ , through which the driven current density can be calculated as

$$\frac{j_{\parallel}}{B} = -e \left\langle \int d^3v \chi \mathbf{S}_{\text{rf}}(f_M) \right\rangle, \quad (17)$$

where angular brackets denote flux-surface average and \mathbf{S}_{rf} is the radio-frequency-induced quasilinear diffusion operator [46,43]. The response function is written factorizing its dependence on particle energy and pitch angle, $\chi = \text{sgn}(u_{\parallel}) F(u) H(\lambda)$, with $\mathbf{u} = \gamma \mathbf{v}/c = \mathbf{p}/mc$ the normalized momentum ($\gamma = \sqrt{1+u^2}$ is here the relativistic Lorentz factor) and $\lambda = (B_{\text{max}}/B)(u_{\perp}^2/u^2)$. The energy-dependent part is referred to as the Spitzer function. In CURGAP, it is in general given by Eq. (33) of [37]. If momentum conservation is included, this expression is replaced by the relativistic adaptation of the Spitzer function obtained in the Appendix of [47]. Being based on a fifth-degree polynomial in u , whose coefficients can be determined in a straightforward way, the momentum-conserving formulation of the Spitzer function is computationally very efficient. Due to this reason, this option is also adopted in the “real-time” version of the code when current drive is required, see Section 6. To quote an example of the contribution of the various effects mentioned above to the final result, for an ITER standard-H-mode scenario with EC current driven on the $q = 3/2$ surface, TORBEAM returns a total driven current $I_{\text{CD}} = -7.90$ kA/MW employing CURBA, $I_{\text{CD}} = -8.30$ kA/MW with CURGAP (full polarization term) and $I_{\text{CD}} = -8.94$ kA/MW if momentum conservation is included. Here, the minus sign indicates that the current is driven in clockwise direction if the tokamak is viewed from above.

As mentioned above, a further important improvement has been introduced in the calculation of the current drive efficiency in order to account consistently for the contribution of each cyclotron harmonic. The current drive efficiency η , defined as the ratio between the current and power densities, can be written in terms of the quantities introduced before as

$$\eta \propto \frac{\langle \int d^3v \chi \mathbf{S}_{\text{rf}}(f_M) \rangle}{\langle \int d^3v \mathcal{E} \mathbf{S}_{\text{rf}}(f_M) \rangle}, \quad (18)$$

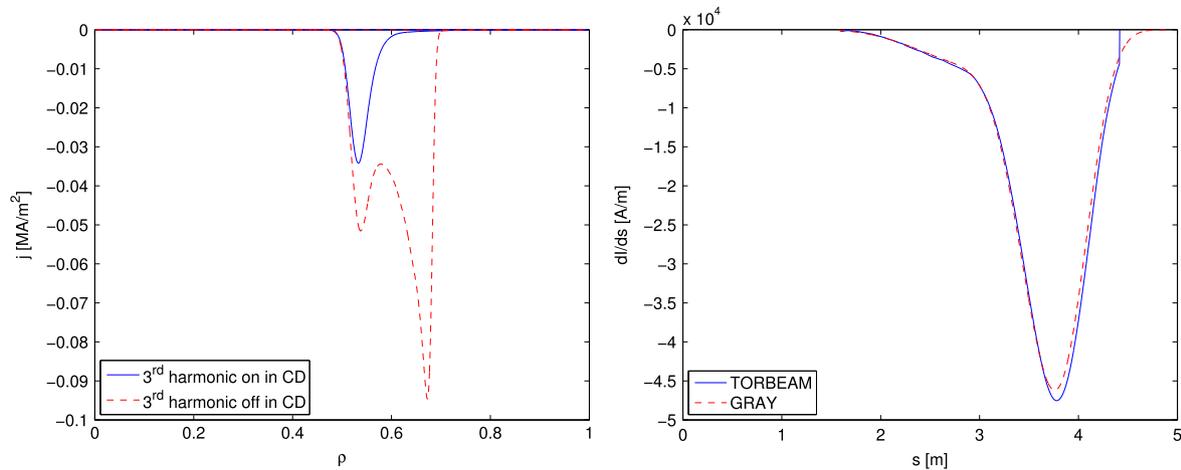


Fig. 6. Left: Current drive profile obtained considering the contribution of second and third harmonic consistently (blue solid line) and assuming that all the power drives a current on the lowest accessible harmonic (red dashed). Right: Comparison of the current driven per unit length along the central ray in TORBEAM and GRAY (s being the arclength along the ray).

where $\varepsilon = mc^2(\gamma - 1)$ is the particle energy. The quasilinear flux S_{rf} contains a sum over the cyclotron harmonics, see e.g. Eqs. (14) and (15) of [36]. If more than one harmonic contributes to the wave-particle interaction, both integrals appearing in Eq. (18) become actually a sum of integrals, one for each harmonic. However, since the lowest harmonic present in the plasma usually dominates the absorption, often in practical implementations (as in the original TORBEAM code) only the lowest harmonic is considered in the calculation of the current drive efficiency (18). On the other hand, in particular for large values of the parallel wave vector, the lowest harmonic may exhibit a rather modest absorption when it becomes accessible in low-field-side launch scenarios, as the corresponding resonance curve lies on the tail of the electron distribution function. Hence, only few electrons are available for absorption. Since high-energy electrons have a low collisionality, however, the numerator of Eq. (18) weights them more strongly than the denominator, and the current drive efficiency evaluated on the lowest harmonic resonance curve only can be fairly large. Thus, attributing all the driven current to the lowest harmonic, while the power absorption is still governed by the successive harmonic, would result in an unrealistically large driven current. This is illustrated rather strikingly by a half-field scenario for ITER (Fig. 6, left), where a spurious peak (red dashed curve) appears in the current density profile when all the power absorbed by the plasma (mainly by third X-mode harmonic) is used to calculate the driven current with η determined retaining only the contribution of the second harmonic as soon as it becomes accessible. Incorporating also the contribution from the third harmonic in both the numerator and the denominator of Eq. (18) eliminates this unphysical effect (blue solid line) and the current drive becomes nearly identical to that calculated by the GRAY code, in which this effect is included (Fig. 6, right). The associated computational effort remains modest, as the contribution of the current-drive routine to the total running time is usually relatively small, see e.g. Table 2.

6. Real-time applications

Future fusion reactors, and more so fusion power plants, need to be equipped with a series of control systems able to react in real time to given plasma signals. Such systems are being tested in present experiments. EC waves are foreseen as actuators in a number of applications, including control, mitigation and suppression of MHD instabilities, in particular sawtooth oscillations and NTMs. In these applications, real-time ray tracing is an option to obtain information on the mirror settings needed in order to reach

the rational surface around which the instability develops [21]. Fast execution of a wave-heating module is also required in real-time simulations of the plasma profiles, or for numerical optimization of tokamak scenarios, as performed by the RAPTOR code [48,49]. More in general, the possibility of accelerating the code execution is highly desirable in all applications requiring a repeated evaluation of the beam trajectory in loops of some sort, as is for instance the case in transport codes. The execution time of the code has been monitored through accurate time profiling to locate the most critical (time-consuming) operations. The numbers quoted below have been obtained running TORBEAM for the medium-size tokamaks ASDEX Upgrade and TCV on Solaris SunOS 5.11 and Linux SLES 11 64-bit machines, which are presently the standard platforms available at the Max-Planck-Institut für Plasmaphysik. The exact execution time is of course machine-dependent, but the relative speed-up is less sensitive to the details of the architecture. Performance improvement can be achieved from running the code on dedicated processors when executing the control algorithms.

As mentioned previously, the typical run time of the code in its complete version (including the calculation of power and current density profiles) is of the order of one second. For implementation in present real-time workflows for NTM control, an execution time of around 10 ms should be targeted (see discussion in [21]). In its complete version with standard settings, most of the CPU time is spent by TORBEAM in the calculation of the output profiles. For a highly resolved case (5000 points in ρ), more than 90% of the time is spent in the output routine. This number drops to around 55% for 500-points profiles. The most time-demanding operation at this level is the iterated interpolation of the plasma-volume profile, cf. Eq. (12). For those applications not relying explicitly on the exact form of the absorption profile, the call to the output routine OUT-PROF in out.f90 can be skipped altogether. By this, the execution time drops below 100 ms, which is, however, still insufficient for the targeted real-time applications. Once the calculation of the absorption profile has been dropped, the remaining most time-consuming operation (about 30 ms), namely the calculation of the plasma volume as a function of the radial coordinate ρ , becomes unnecessary and can be skipped as well. A further speed-up is obtained by reducing the number of equations to be integrated to seven, i.e. dropping Eqs. (8), (9) and solving only for Eqs. (7), (11). This step, which turns TORBEAM into a ray-tracing code computing a single ray, makes also the preparation of second-derivatives finite-difference arrays unnecessary. Finally, the position of maximum absorption returned by the code is computed as the position reached at the last integration step for which the power decrement

Table 1

Time profiling for a run of the “real-time” version of TORBEAM (ASDEX Upgrade parameters, no current drive). The total time needed for this run was 8.1 ms (first line). It can be seen that most of the time required for advancing the beam (1sode, 3.4 ms) is actually spent in the calculation of the coefficients of the beam tracing equations (deriv, 3.1 ms), rather than in the integrator itself. The seven steps needed from the antenna to the plasma (dervvac) do not give a measurable contribution (< 0.05 ms). The routines beginning with “int . . .” perform interpolation of various plasma quantities.

torbeam	was called	1 times,	total	8.1 ms
grid	was called	1 times,	total	1.4 ms
intpsi	was called	674 times,	total	0.3 ms
dervvac	was called	7 times,	total	0.0 ms
lsode	was called	580 times,	total	3.4 ms
coef	was called	213 times,	total	2.0 ms
intpb	was called	572 times,	total	0.2 ms
intdab	was called	572 times,	total	2.6 ms
intpop	was called	554 times,	total	0.0 ms
deriv	was called	359 times,	total	3.1 ms
intpte	was called	859 times,	total	0.1 ms

dP is larger than at the previous step. The code can be stopped either at that position or (if the current drive is also computed) when dP falls below a preset threshold. The steps described above bring the execution time in the required ballpark of 10–20 ms (the calculation of the total driven current employing the momentum-conserving scheme described in Section 5 requires an additional CPU effort below 2–3 ms). An additional reduction of the execution time has been achieved by avoiding the determination of the plasma coefficients in `coef.f90` (Section 3) after each integration step but rather extrapolating them from the previous values for a pre-set number of time steps, after which the actual coefficients are computed again. The corresponding speed-up is a few ($\lesssim 5$) ms. Finally, the dimensions of the grids on which the magnetic equilibrium and the kinetic profiles are defined have been limited to 150 points.

Table 1 shows a sample of the TORBEAM profiling obtained for the off-axis-heating case on ASDEX Upgrade already considered in the profile benchmark shown in Fig. 3b (run on a single Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80 GHz processor).

In this run, 17% of the time was spent in the preliminary calculations (preparation of finite-difference arrays, GRID), 42% in the integration loop (LSODE), which includes as the main part the calculation of the right-hand side of the beam tracing equations (DERIV), 25% in the calculation of the coefficients needed after the integration step (COEF). These three steps account hence for ca. 84% of the execution time. It is interesting to notice that 20%–25% of the time needed for the calculation of the coefficients of the beam tracing equations in DERIV is employed for the determination of the absorption coefficient γ (in real-time applications, the weakly-relativistic version of DAMPBQ is used, which turns out to have the shortest execution time). The remaining 75%–80% of the CPU time spent in DERIV is primarily employed in the interpolation of the plasma profiles, in particular the magnetic field components (INTDAB). An example for the implementation of the real-time version of TORBEAM in an actual control scheme is described in [21]. A second example is the recent inclusion of TORBEAM in the real-time control system SCD of the TCV tokamak [50]. Since this system is entirely programmed in Simulink[®], it was necessary to include TORBEAM in a Simulink S-function block. This was achieved by writing an S-function wrapper that calls the Fortran library. Various separate instances of TORBEAM are run (one for each EC launcher) on different processors on a dedicated node of the control system. The execution times for a TCV equilibrium with grid size of 28×65 in the poloidal plane and radial grid of 21 points for density and temperature profiles are of the order $\lesssim 5$ ms for typical cases.

A different accelerated version of TORBEAM has been developed to speed up applications in which the code is used in

Table 2

Time profiling for the accelerated version of TORBEAM including deposition profiles (total CPU time 75.3 ms).

torbeam	was called	1 times,	total	75.3 ms
grid	was called	1 times,	total	2.7 ms
intpsi	was called	1957 times,	total	1.9 ms
dervvac	was called	11 times,	total	0.0 ms
lsode	was called	308 times,	total	20.6 ms
coef	was called	308 times,	total	4.4 ms
intpb	was called	1307 times,	total	0.8 ms
intdab	was called	1308 times,	total	13.9 ms
intpop	was called	1307 times,	total	0.9 ms
deriv	was called	999 times,	total	19.2 ms
intpte	was called	1237 times,	total	0.0 ms
intpdvrtprof	was called	135280 times,	total	16.7 ms
currn	was called	10 times,	total	1.0 ms
intpdv	was called	5 times,	total	0.0 ms
outprof	was called	1 times,	total	40.5 ms

integrated data analysis of kinetic profiles and equilibrium reconstruction [51,52]. In this case, an equilibrium reconstruction code solving the Grad-Shafranov equation (GSE) is coupled with the current diffusion equation solving the temporal evolution of the current profile between the time points the GSE is solved. The neo-classical current diffusion depends on the kinetic profiles as well as on the driven current from neutral beams and microwaves. In these applications, the calling process is supposed to provide the code some of the geometric quantities usually calculated internally, so that again the calculation of the plasma volumes in GRID can be omitted. Also the maximum and minimum values of the magnetic field on a given flux surface are now linearly interpolated on a user-supplied grid and not internally. Moreover, the calculation of the profiles is performed on a radial grid of 500 points, which corresponds to a radial resolution of $\Delta\rho = 2 \times 10^{-3}$, so that also the time needed to compute the deposition profiles remains below ca. 50 ms. The targeted execution time of this version for a typical ASDEX Upgrade should remain below 100 ms. An example of the time profiling for an ASDEX Upgrade application is shown in Table 2.

7. Reflectometry applications

The reflectometry diagnostic is based on the injection of waves in the EC frequency range under cutoff condition, with the aim of extracting information about the plasma properties from the measurement of the reflected beam [53]. The integration of ray tracing (and even more beam tracing) equations becomes challenging close to turning points, in particular for injection nearly perpendicular to the cutoff layer, as in this case all the three components of the group velocity are very small at the turning point. Nevertheless, the advancement of the equations is usually still possible, provided the input profiles (electron density in particular) are smooth enough,³ and delivers important information for the analysis of the measurement [54,55]. For a discussion of the applicability of the pWKB method to reflectometry see also [31]. From the point of view of the numerical implementation, apart from the additional output parameters mentioned in Section 3, the main extension introduced in the code for reflectometry studies is the renormalization of the electron mass needed to account for relativistic effects for propagation close to the cutoff [27]. This is obtained by rescaling the electron mass appearing in both the plasma frequency and the cyclotron frequency according to (see [27])

$$m_{\text{eff}} = m \left(1 + \frac{5T_e}{mc^2} \right)^{1/2}. \quad (19)$$

³ To further facilitate the numerical integration, in reflectometry runs the integration step in TORBEAM is reduced with increasing density. Tools to extend and smooth density and temperature profiles are provided with the code.

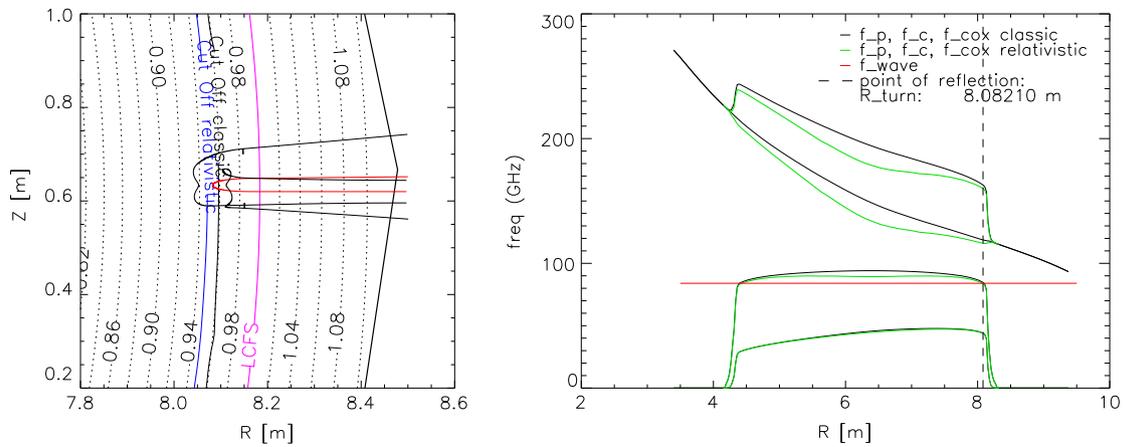


Fig. 7. Left: Propagation of a reflectometry beam (injected from the low-field side with a frequency of 84.9 GHz, O-mode), showing also the “classical” (cold plasma) and relativistic corrected position of the cutoff. The reference ray is the red central line. Right: Profiles of cyclotron frequency (f_c), plasma frequency (f_p) and cutoff frequencies for X and O-mode. The injection frequency is represented by the solid horizontal red line, while the vertical dotted line shows the position of the turning point.

This introduces a dependence of the plasma frequency and the cyclotron frequency, which appear in the Hamiltonian function H and its derivatives on the right-hand side of Eqs. (7)–(9), on the electron temperature. Consequently, first and second-order derivatives of the temperature profile need to be calculated, which are otherwise absent since the propagation of the beam is computed in usual heating and current drive applications according to the cold-plasma dispersion relation. These corrections were added for a design study of the ITER low-field-side reflectometer [11]. A typical example of beam propagation and cutoff profiles is shown in Fig. 7.

In addition to the quantities at the turning point returned by the code in reflectometry applications, see Section 3, the information about the beam parameters provided by the pWKB method allows one to determine also the coupling between the beam and the receiver antenna, which can be useful in assessing the measured signal and guide the design of future reflectometers, as described in [11]. The coupling efficiency T_a is given [56,57] by the squared modulus of the integral (taken on the antenna plane) of the product of the beam field, given in terms of Eqs. (3), (5), (6) and the antenna radiation pattern, which simply follows from the vacuum solution for Gaussian beams, Eqs. (1), (13), and (14) of [11]. Taking the centre of the antenna opening to be at $(x_0, 0, 0)$ and integrating over the plane $x_0 = \text{const.}$, the result is given in terms of the beam tracing variables introduced in Section 2 as follows (Eq. (15) of [11]):

$$T_a = \frac{16}{W_0^2 W_1 W_2 \det[\sigma + \Sigma]} \times \exp(\phi_{\alpha\beta} q_\alpha q_\beta) \exp(\text{Re}[(\sigma + \Sigma)_{\alpha\beta}^{-1} b_\alpha b_\beta]), \quad (20)$$

with $\sigma_{\alpha\beta} = 2\delta_{\alpha\beta}/W_0^2$ (W_0 being the radius of the antenna pattern and $\delta_{\alpha\beta}$ the Kronecker δ), $\Sigma_{\alpha\beta} = \phi_{\alpha\beta} + i s_{\alpha\beta}$ and $b_\alpha = \Sigma_{\alpha\beta} q_\beta - i N_\alpha$, where now the indices α and β run only over the (y, z) plane. A routine that calculates T_a as part of the post-processing and visualization is also supplied with the code.

To conclude this section, it is remarked that TORBEAM is applied routinely also to Doppler reflectometry, which is a special application where the reflectometry probing beam is incident obliquely to the cutoff layer (see Fig. 8) and is used to measure the turbulence propagation velocity $u_\perp = 2f_D/k_\perp$ from the Doppler frequency shift f_D in the backscattered signal [10]. Here, the ray and beam-tracing equations appear to give rather robust estimations of both the ray-turning point, i.e. the region of maximum backscatter) and the component beam wavenumbers k_\perp and k_\parallel , which are required to convert the Doppler shift to a velocity. The beam envelope behaviour may also provide relevant information on the spatial and wavenumber resolution [58,55].

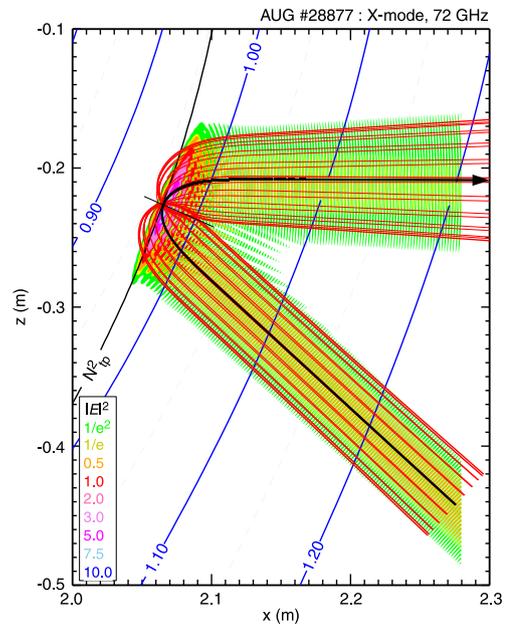


Fig. 8. Full-wave $|E|^2$ contours overlaid with TORBEAM envelope (red) and central ray (black) for an X-mode, 72 GHz, Doppler reflectometer simulation of AUG shot 28877. Launch antenna is in lower right-hand corner. Flux surfaces (ρ_{pol}) in blue. N^2 (turning point) layer in black. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

8. Conclusions

The paraxial WKB method is a powerful tool to calculate the propagation of a high-frequency beam in anisotropic inhomogeneous media, retaining diffraction effects with low computational costs. Its numerical implementation in the code TORBEAM allows fast determination of power absorption and current density profiles in tokamaks, and is useful also in diagnostics applications like reflectometry. With respect to the original publication [3], significant advances have been made, which are reviewed in this paper. The functionalities of the code have been extended, in particular to include accelerated versions suitable for real-time usage in control loops running in parallel with an actual plasma discharge or for integrated data analysis. From the point of view of the physics model, the calculation of absorption and current drive includes now state-of-the-art capabilities like the fully relativistic

treatment of the dielectric tensor in the calculation of the absorption coefficient and momentum conservation in the calculation of current drive. The possibility of accounting for the separate contribution of different harmonics in the calculation of the driven current has been introduced. A theory-based algorithm for the calculation of the absorption profile, applicable also for on-axis deposition, has been implemented and tested. This is important, for instance, for advanced tokamak scenarios relying on counter-current drive on axis. The application of the code to reflectometry studies has been discussed.

Further improvements and extensions of TORBEAM can be considered. From the point of view of the underlying physical model, for instance, although the cold-plasma approximation should be sufficient in most of the cases of interest, including warm-plasma dielectric-tensor elements in the calculation of the Hamiltonian function H would enable the “anomalous” dispersion near the EC resonance [59] to be accounted for. From the numerical point of view, further steps to streamline the calculation and achieve even faster performances in the accelerated versions could be undertaken. Very desirable, although more speculative, since the corresponding theory has not been developed to date, would be the possibility to account for the beam distortion due to asymmetric absorption or the inclusion of scattering effects (as retained e.g. in the WKBeam code [38]) within the pWKB method.

The code revisions are managed through an svn repository. This paper describes revision 606. The code is distributed by the main author upon request, after a software agreement with the author's home institution (Max-Planck-Institut für Plasmaphysik), excluding the commercial exploitation of the code, has been signed.

Acknowledgments

This work is supported in part by DOE, under contract DE-AC02-09CH11466. The authors acknowledge the help of M. Rampp at the Max Planck Computing and Data Facility for providing the interpolation routine used in the new calculation of the absorption profile. We are indebted to F. Sommer for developing the graphical user interface under which TORBEAM presently runs on the ASDEX Upgrade infrastructure (and which has been used to obtain Fig. 1). Many thanks also to O. Meneghini for incorporating TORBEAM in the OMFIT frame at DIII-D and to E. Fable for the collaboration during the coupling of the TORBEAM to the ASTRA transport solver.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix. Power deposition of paraxial beams

Here we sketch the derivation of Eq. (14). The starting point is the steady-state continuity equation for the wave energy density U [46]

$$\nabla \cdot (U \mathbf{v}_g) = -2\tilde{\gamma}U, \quad (21)$$

where \mathbf{v}_g is the group velocity and $\tilde{\gamma}$ is related to the absorption coefficient γ introduced in Eq. (11), see Eq. (31). The energy density is related to the electric-field amplitude through [46]

$$U = \frac{\omega |\partial_\omega H|}{16\pi} |E|^2 \quad (22)$$

Since

$$\mathbf{v}_g = -\frac{c}{\omega} \frac{\partial H / \partial \mathbf{N}}{\partial H / \partial \omega} = -\frac{c}{\omega \partial_\omega H} \mathbf{V} \quad (23)$$

one has $U \mathbf{v}_g = c|E|^2 \mathbf{V} / 16\pi$ (for $\partial_\omega H < 0$). Considering a region $\Omega(\psi)$ of the plasma inside the flux surface labelled by ψ , with

boundary $\partial\Omega(\psi)$, the deposited power is the energy that crosses $\partial\Omega(\psi)$ per unit time,

$$\begin{aligned} P(\psi) &= - \int_{\partial\Omega(\psi)} U \mathbf{v}_g \cdot d\mathbf{S} = - \int_{\Omega(\psi)} \nabla \cdot (U \mathbf{v}_g) d^3r \\ &= 2 \int_{\Omega(\psi)} \tilde{\gamma} U d^3r, \end{aligned} \quad (24)$$

where the second step follows from Gauss's theorem and the third one from Eq. (21).

To proceed further, we write the electric-field solution, Eq. (3), introducing the relative amplitude $a(\tau)$

$$\mathbf{E} = A_0 a(\tau) \mathbf{e}(\tau) e^{ik_0(s(\mathbf{r}) + i\phi(\mathbf{r}))} \quad (25)$$

defined such that $a(0) = 1$. Hence

$$|E|^2 = |A_0|^2 |a(\tau)|^2 e^{-2k_0\phi}. \quad (26)$$

In the pWKB method, ϕ is a non-negative definite quadratic form, see Eq. (6), describing the elliptic cross section of the wave beam with principal widths W_1 and W_2 . The $1/e$ half-width of the wave intensity ($\propto |E|^2$) is hence given by $w_{1,2} = W_{1,2}/\sqrt{2}$. As a consequence, the integral of $|E|^2$ across a constant- τ surface is

$$\int |E|^2 dS = \pi |A_0|^2 |a(\tau)|^2 w_1(\tau) w_2(\tau). \quad (27)$$

The power crossing this surface is hence

$$P(\tau) = \frac{c}{16} |\mathbf{V}(\tau)| |A_0|^2 |a(\tau)|^2 w_1(\tau) w_2(\tau). \quad (28)$$

Using the previous equations, the relation between Eqs. (21) and (11) can be established. Integration of Eq. (21) inside a “flux tube” following the beam from the antenna plane to the position labelled by τ yields (with $d^3r = |V(\tau)| d\tau dS$)

$$\begin{aligned} P(\tau) - P_0 &= -2 \int \tilde{\gamma} U d^3r \\ &= -\frac{\omega |A_0|^2}{8} \int_0^\tau \tilde{\gamma} |\partial_\omega H| |\mathbf{V}| |a|^2 w_1 w_2 d\tau'. \end{aligned} \quad (29)$$

Deriving the previous expression and substituting Eq. (28), one finds

$$\frac{dP}{d\tau} = -2 \frac{\omega |\partial_\omega H|}{c} \tilde{\gamma} P(\tau) \quad (30)$$

from which Eq. (11) is recovered with

$$\gamma = \frac{\omega |\partial_\omega H|}{c} \tilde{\gamma} = \frac{V}{v_g} \tilde{\gamma}. \quad (31)$$

Finally, with $\tilde{\gamma}$ obtained from Eq. (30) and U from Eq. (22), and using Eqs. (26) and (28), Eq. (24) can be straightforwardly cast into the form (14) given in the main text of this paper.

References

- [1] G.V. Pereverzev, Rev. Plasma Phys. 19 (1996) 1.
- [2] G.V. Pereverzev, Phys. Plasmas 5 (1998) 3529.
- [3] E. Poli, A.G. Peeters, G.V. Pereverzev, Comput. Phys. Comm. 136 (2001) 90.
- [4] R.C. Wolf, et al., Nucl. Fusion 41 (2001) 1259.
- [5] S. Günter, et al., Nucl. Fusion 45 (2005) S98.
- [6] R. Bilato, et al., Nucl. Fusion 49 (2009) 075020.
- [7] M.A. Henderson, et al., Fusion Eng. Des. 82 (2007) 454.
- [8] E. Poli, G. Tardini, H. Zohm, et al., Nucl. Fusion 53 (2013) 013011.
- [9] L. Figini, et al., Plasma Phys. Control. Fusion 57 (2015) 054015.
- [10] G.D. Conway, E. Poli, T. Happel, Plasma Fusion Res. 5 (2010) S2005.
- [11] A. Stegmeier, G.D. Conway, E. Poli, E. Strumberger, Fusion Eng. Des. 86 (2011) 2928.
- [12] D.A. Thomas, et al., Nucl. Fusion 56 (2016) 026013.
- [13] J.C. Hillesheim, et al., Phys. Rev. Lett. 116 (2016) 065002.
- [14] G.V. Pereverzev, Nucl. Fusion 32 (1992) 1091.

- [15] N. Bertelli, et al., *Phys. Plasmas* 19 (2012) 082510.
- [16] R. Prater, et al., *Nucl. Fusion* 48 (2008) 035006.
- [17] G.V. Pereverzev, P.N. Yushmanov, ASTRA: Automated System for Transport Analysis in Tokamaks, IPP Report 5/98, 2002.
- [18] R. Hawryluk, in: B. Coppi, G. Leotta, D. Pfirsch, R. Pozzoli, E. Sindoni (Eds.), *Physics of Plasmas Close to Thermonuclear Conditions, Vol. 1*, Pergamon Press, Oxford, 1981, pp. 19–46.
- [19] G.L. Falchetto, et al., *Nucl. Fusion* 54 (2014) 043018.
- [20] L. Figini, et al., *EPJ Web Conf.* 32 (2012) 01011.
- [21] M. Reich, et al., *Fusion Eng. Des.* 100 (2015) 73.
- [22] N.B. Marushchenko, H. Maassberg, Yu Turkin, *Nucl. Fusion* 48 (2008) 054002.
- [23] N.B. Marushchenko, H. Maassberg, Yu Turkin, *Nucl. Fusion* 49 (2009) 129801.
- [24] N.B. Marushchenko, et al., *Phys. Plasmas* 18 (2011) 032501.
- [25] E. Westerhof, RELAX, a computer code for the study of collisional and wave driven relaxation of the electron distribution function in toroidal geometry, Rijnhuizen Report RR-92-211 CA, 1992.
- [26] O. Maj, E. Poli, E. Westerhof, *J. Phys. Conf. Ser.* 401 (2012) 012013.
- [27] E. Mazzucato, *Phys. Fluids B* 4 (1992) 3460.
- [28] H. Bindslev, *Plasma Phys. Control. Fusion* 35 (1993) 1093.
- [29] E. Poli, G.V. Pereverzev, A.G. Peeters, *Phys. Plasmas* 6 (1999) 5.
- [30] E. Poli, G.V. Pereverzev, A.G. Peeters, M. Bornatici, *Fusion Eng. Des.* 53 (2001) 9.
- [31] O. Maj, G.V. Pereverzev, E. Poli, *Phys. Plasmas* 16 (2009) 062105.
- [32] O. Maj, A.A. Balakin, E. Poli, *Plasma Phys. Control. Fusion* 52 (2010) 085006.
- [33] A.C. Hindmarsh, *Acm-Signum Newslett.* 15 (1980) 10.
- [34] E. Poli, A.G. Peeters, G.V. Pereverzev, *Phys. Plasmas* 8 (2001) 4325.
- [35] E. Westerhof, Implementation of TORAY at JET, Rijnhuizen Report RR-89-183, 1989.
- [36] D. Farina, *Fusion Sci. Technol.* 52 (2007) 154.
- [37] Y.R. Lin-Liu, V.S. Chan, R. Prater, *Phys. Plasmas* 10 (2003) 4064.
- [38] H. Weber, O. Maj, E. Poli, *EPJ Web Conf.* 87 (2015) 01002.
- [39] O. Maj, A. Mariani, E. Poli, D. Farina, *Phys. Plasmas* 20 (2013) 042122. <http://dx.doi.org/10.1063/1.4802935>. <http://link.aip.org/link/?PHP/20/042122&aem ail=author>.
- [40] Yu A. Kravtsov, Yu I. Orlov, *Geometrical Optics of Inhomogeneous Media*, Springer Verlag, Berlin, 1990.
- [41] E. Mazzucato, *Phys. Fluids B* 1 (1989) 1855.
- [42] T.M. Antonsen, K.R. Chu, *Phys. Fluids* 25 (1982) 1295.
- [43] R.W. Harvey, M.G. McCoy, The CQL3D Fokker-Plank code, GA Report GA-A20978, 1992.
- [44] E. Poli, C. Angioni, F.J. Casson, D. Farina, L. Figini, T.P. Goodman, O. Maj, O. Sauter, H. Weber, H. Zohm, G. Saibene, M.A. Henderson, *Nucl. Fusion* 55 (2015) 013023.
- [45] H. Zohm, et al., *Nucl. Fusion* 55 (2015) 104010.
- [46] M. Brambilla, *Kinetic Theory of Plasma Waves*, Clarendon Press, Oxford, 1998.
- [47] M. Romé, et al., *Plasma Phys. Control. Fusion* 40 (1998) 511.
- [48] F. Felici, et al., *Nucl. Fusion* 51 (2011) 083052.
- [49] F. Felici, O. Sauter, *Plasma Phys. Control. Fusion* 54 (2012) 025002.
- [50] H. Anand, et al., *Nucl. Fusion* 57 (2017) 056005.
- [51] R. Fischer, et al., *Fusion Sci. Technol.* 58 (2010) 675.
- [52] R. Fischer, et al., *Fusion Sci. Technol.* 69 (2016) 526.
- [53] I.H. Hutchinson, *Principles of Plasma Diagnostics*, Cambridge University Press, 1987.
- [54] T. Happel, et al., *Phys. Plasmas* 22 (2015) 032503.
- [55] G.D. Conway, et al., *Plasma Phys. Control. Fusion* (2018) in preparation.
- [56] H. Kogelnik, Coupling and Conversion Coefficients for Optical Modes, in: *Microwave Research Institute Symposia*, vol. 14, Polytechnic Press, New York, 1964.
- [57] W.B. Joyce, B. DeLoach, *Appl. Opt.* 23 (1984) 4188.
- [58] G.D. Conway, et al., *Proc. 12th Intl. Reflectometry Workshop - IRW12*, 2015. http://Www.Fz-Juelich.de/Conferences/IRW12/EN/Publications/_Node.Html.
- [59] E. Westerhof, *Plasma Phys. Control. Fusion* 39 (1997) 1015.