A reduced fast ion transport model for the tokamak transport code TRANSP

A reduced fast ion transport model for the tokamak transport code TRANSP

M Podestà, M Gorelenkova and R B White

Princeton Plasma Physics Laboratory, Princeton, NJ 08543, USA

Abstract

Fast ion transport models currently implemented in the tokamak transport code TRANSP (Hawryluk 1980 *Physics of Plasmas Close to Thermonuclear Conditions* (Brussels: CEC)) are not capturing important aspects of the physics associated with *resonant* transport caused by instabilities such as toroidal Alfvén eigenmodes (TAEs). This work describes the implementation of a fast ion transport model consistent with the basic mechanisms of resonant mode–particle interaction. The model is formulated in terms of a *probability distribution function* for the particle's steps in phase space, which is consistent with the Monte Carlo approach used in TRANSP. The proposed model is based on the analysis of the fast ion response to TAE modes through the ORBIT code (White and Chance 1984 *Phys. Fluids* **27** 2455), but it can be generalized to higher frequency modes (e.g. compressional and global Alfvén eigenmodes) and to other numerical codes or theories.

Keywords: fast ion transport, tokamak simulations, Alfvénic instabilities, reduced transport modules

(Some figures may appear in colour only in the online journal)

1. Introduction

Simulations of tokamak discharges have made considerable progress in the past few years. Because of the complexity of the problems that numerical codes aim to solve, two main lines of code development are emerging. On one hand, *first-principles* models target a representation of tokamak physics as close as possible to reality. This is supported by improvements in both experimental measurements and theory, which enable a deeper insight into plasma physics phenomena than was possible only a decade ago. Examples of first-principles models applied to the physics of energetic particles and associated plasma instabilities can be found in [1–6] and references therein. The generality of this approach comes at a cost. Simulations are expensive, in terms both of computing time and of required hardware resources. Use of first-principles codes is thus usually restricted to a limited number of conditions. On the other hand, *reduced* models are often utilized in combination with more general-purpose codes to tackle specific problems in a simplified form. Recent examples include simplified models for computing the relaxed fast ion profile resulting from a given set of Alfvènic instabilities; see [7–9]. Although the accuracy of the results from reduced models is expected to degrade with respect to more comprehensive codes, more insight into the problem under study can be gained by looking at a much larger number of cases. Clearly, the distinction between first-principles and reduced models is sometimes questionable, depending on the amount of simplifications introduced in the models and in their practical implementation. In several aspects, the two approaches are complementary and their combined use can result in improved understanding of tokamak and plasma physics.

This paper reports on the development and initial validation of a new reduced model for fast ion transport by plasma instabilities such as Alfvénic modes. The model will be included in the tokamak transport code TRANSP [10, 11], which is used on several tokamaks either to simulate existing discharges or to develop and predict new plasma scenarios. The general aspects of the new model are discussed in section 2. Section 3 describes the practical implementation and the integration of the new model with the existing TRANSP code and its modules. Initial verification and validation work is then presented in section 4. Section 5 concludes the paper.

2. A description of the new model

The NUBEAM module [12, 13] implemented in TRANSP models fast ion dynamics in tokamaks on the basis of

classical physics. In addition, NUBEAM has four possible implementations for modeling fast ion transport mechanisms different from classical ones. The first two models set the fast ion diffusivity to be proportional to the electron particle diffusivity (with or without corrections for the Ware pinch) through a user-selected multiplier. These models give a fast ion radial flux, Γ_{fi}, that is related to the *electron* density gradient. Clearly, such models do not capture the physics of resonant wave–particle interactions. A third model includes fast ion diffusion and convection coefficients as a function of the radius, $D_{fi}(r)$ and $V_{fi}(r)$, such that $\Gamma_{fi} = -D_{fi}\nabla n_{fi} + n_{fi}V_{fi}$. $n_{fi}(r)$ is the radial fast ion density profile as a function of the normalized minor radius r. The flux is therefore regulated by the spatial fast ion gradient, ∇n_{fi}. Energy- and space-dependent diffusion coefficients can be prescribed in a fourth model for (up to) six classes of particles: passing, barely passing, trapped, deeply trapped, counter-passing and counter-barely passing. The actual diffusivity for each particle during the NUBEAM calculation is obtained as a weighted combination of two of those six coefficients, depending on the orbit localization in phase space. Although some dependence on the specific orbit topology is introduced, the resulting transport is still diffusive in *real space*, i.e. along the radial direction. In general, all of these models do not contain the physics of resonant interaction—and the resulting transport in *phase space*—between instabilities and fast ions. In fact, the resonance condition implies that only narrow regions in phase space may be strongly affected by the modes, whereas adjacent regions are possibly unaffected.

The following features should be included in a new fast ion transport model in TRANSP to mimic the resonant interaction between fast ions and instabilities:

(1) Characterize particles based on their *orbit topology* [14], i.e. in terms of magnetic moment μ, energy E and canonical toroidal angular momentum P_ζ, instead of real-space coordinates such as radius, poloidal/toroidal angles.

(2) Model transport as steps (or kicks) in phase space, for instance kicks in energy associated with the resonant interaction. Radial transport will eventually result from the particle dynamics in phase space, but with no *a priori* assumptions on its nature, such as being diffusive or convective.

(3) Derive transport coefficients from consistent simulations or theory. Whenever possible, experimental data should be used as further constraints.

(4) The model must be suitable for inclusion in the NUBEAM module. In this regard, a Monte Carlo framework seems to offer the best approach.

(5) Calculate variations of E and P_ζ consistently.

The fifth item in the list above is actually the most distinctive feature for transport resulting from resonant wave–particle interactions. On the basis of the guiding center Hamiltonian formulation of the particle's motion in the presence of a mode with toroidal mode number n and frequency $\omega = 2\pi f$, one obtains the relationship [14]

$$\omega P_\zeta - nE = \text{const.} \tag{1}$$

Figure 1. Illustration of the correlated variation of the energy and canonical angular momentum in the presence of a single mode with zero and finite frequency width; see equation (3).

For a single mode, variations in E and P_ζ for particles satisfying equation (1) (given a specific constant on the right-hand side [15]) are thus related through

$$\Delta P_\zeta / \Delta E = n/\omega, \tag{2}$$

which sets a constraint on the allowed trajectories in the (E, P_ζ) space. In reality, if a finite mode frequency width is assumed and if more than one mode are present, ΔE and ΔP_ζ can depart from the ideal (linear) relation of equation (2); see figure 1 and section 2.

One possibility for satisfying the requirements described above would be to develop a module, interfaced with NUBEAM, that models the resonant wave–particle interaction. For example, the model could be a reduced version of the ORBIT code [16]. However, this approach is not very flexible: if one decides to use another code to manage the mode–particle interaction, an entirely new module has to be developed and included in NUBEAM.

Another possibility, that is the one discussed in this paper, is to split the problem into two parts: (i) derive a set of transport coefficients in some given form, and (ii) use those coefficients in NUBEAM for the actual computation of fast ion evolution. By doing this, the NUBEAM part of the problem can be developed independently of the different models (or theories) used to infer the transport coefficients.

The main ingredient of the new model discussed herein is the probability that a particle, whose orbit is characterized by the constants of motion (P_ζ, E, μ), experiences a change over a time δt in energy and canonical angular momentum of magnitude ΔE and ΔP_ζ in the presence of a mode with amplitude A_{mode}. For simplicity, here and in the following it is assumed that μ is conserved. This is a reasonable assumption for low frequency modes with $\omega \ll \omega_{ci}$ (ω_{ci} being the ion cyclotron resonance frequency), such as TAEs[1]. However, the model can be generalized to include μ variations.

In the following sections, the general principles inspiring the new model are first discussed for an ideal case with a single resonance and for a specific class of particles, then generalized to an arbitrary set of modes for the entire fast ion population.

---

[1] It is still debated whether spherical tokamaks such as NSTX represent an exception for magnetic moment conservation. Because the fast ion gyro-radius may be comparable to both the radial width of the modes and the typical scale lengths of equilibrium profiles, conservation of μ is not necessarily satisfied even for low frequency modes.

2.1. The single-mode, single-resonance case

In the presence of a single resonance, the motion of a specific class of particles satisfying equations (1)–(2) is subject to simple constraints in the (E, P_ζ, μ) space [15, 17]. Even for this simple case, the shape of the distribution of steps in energy and canonical angular momentum can be rather complicated and far from a simple bi-Maxwellian distribution, from which diffusion coefficients in E and P_ζ could be readily extracted. However, to illustrate the general ideas of the new model, we first assume that the bi-variate probability density function for ΔE and ΔP_ζ changes of the particle's energy and canonical toroidal momentum can be approximated by

$$p(\Delta E, \Delta P_\zeta | P_\zeta, E, \mu, A_{\text{mode}}) = p_0 \times$$
$$e^{-\frac{1}{2(1-\rho)}\left[\frac{(\Delta E - \Delta E_0)^2}{\sigma_E^2} + \frac{(\Delta P_\zeta - \Delta P_{\zeta 0})^2}{\sigma_{P_\zeta}^2} - 2\rho \frac{(\Delta E - \Delta E_0)(\Delta P_\zeta - \Delta P_{\zeta 0})}{\sigma_E \sigma_{P_\zeta}}\right]} \quad (3)$$

with the normalization factor

$$p_0 = \frac{1}{2\pi \sigma_E \sigma_{P_\zeta} \sqrt{1 - \rho^2}}. \quad (4)$$

(The dependence of all parameters on A_{mode}, e.g. $\sigma_E = \sigma_E(A_{\text{mode}})$, has been omitted to simplify the notation.) Here the variances σ_E and σ_{P_ζ} give the spread of the distribution along the ΔE and ΔP_ζ axes. The *correlation parameter*

$$\rho = \frac{\langle (\Delta E - \Delta E_0)(\Delta P_\zeta - \Delta P_\zeta 0) \rangle}{\sigma_E \sigma_{P_\zeta}} \quad (5)$$

takes into account the coupling between ΔE and ΔP_ζ expressed in equations (1)–(2):

$$\Delta P_\zeta(\Delta E) = \Delta P_{\zeta 0} + \text{sign}(\rho) \times \frac{\sigma_{P_\zeta}}{\sigma_E}(\Delta E - \Delta E_0). \quad (6)$$

The offset (or *convective*) terms ΔE_0 and $\Delta P_{\zeta 0}$ are redundant and $\equiv 0$ for cases in which there is no systematic drift in energy or P_ζ. They are included in the model for generality, e.g. to account for slowing down processes or transport mechanisms other than classical processes.

In principle, the set of probability density functions $p(\Delta E, \Delta P_\zeta | E, P_\zeta, \mu, A_{\text{mode}})$ contains all the information required by NUBEAM for calculating the energy and P_ζ steps caused by resonant wave–particle interaction. In practice, if the analytical formulation from equations (3)–(6) had to be adopted, a set of four dimensional (4D) coefficients has to be given as input to NUBEAM:

$$\begin{cases} \sigma_E = \sigma_E(E, P_\zeta, \mu, A_{\text{mode}}) \\ \sigma_{P_\zeta} = \sigma_{P_\zeta}(E, P_\zeta, \mu, A_{\text{mode}}) \\ \Delta E_0 = \Delta E_0(E, P_\zeta, \mu, A_{\text{mode}}) \\ \Delta P_{\zeta 0} = \Delta P_{\zeta 0}(E, P_\zeta, \mu, A_{\text{mode}}) \\ \rho = \rho(E, P_\zeta, \mu, A_{\text{mode}}). \end{cases} \quad (7)$$

Incidentally, a similar term could be added to include variations of the magnetic moment μ in the model, according to a probability density function $\sigma_\mu = \sigma_\mu(E, P_\zeta, \mu, A_{\text{mode}})$.

The dimensionality of the set of transport coefficients in equations (7) can be reduced. ORBIT simulations indicate that σ_E and σ_{P_ζ} have a (roughly) linear dependence on the

normalized mode amplitude A_{mode}; see figure 2. Thus, only the coefficients for a specific value of A_{mode} (for instance $A_{\text{mode}} = 1$) have to be passed to NUBEAM, along with a separate vector of mode amplitude versus time that is used to rescale the coefficients as time evolves. (The validity of this simplification is further discussed in section 4.1 and figure 14.) By doing this, matrices in equations (7) are reduced to three dimensions, i.e. the three variables (E, P_ζ, μ) that identify orbits in phase space.

2.2. Extension to the general, multi-mode case

The expressions introduced in the previous section could be generalized to the case more commonly encountered in experiments of multiple modes (and multiple resonances, even for a single mode) inducing fast ion transport. For example, one could specify a set of coefficients for N modes:

$$\begin{cases} \sigma_{E,i} = \sigma_{E,i}(E, P_\zeta, \mu) \\ \sigma_{P_\zeta,i} = \sigma_{P_\zeta,i}(E, P_\zeta, \mu) \\ \Delta E_{0,i} = \Delta E_{0,i}(E, P_\zeta, \mu) \\ \Delta P_{\zeta 0,i} = \Delta P_{\zeta 0,i}(E, P_\zeta, \mu) \\ \rho_i = \rho_i(E, P_\zeta, \mu) \\ w_i = w_i(E, P_\zeta, \mu) \end{cases} \quad (8)$$

with $i = 1 \dots N$. w_i is the relative weight of the ith probability distribution, such that $\sum_{i=1}^{N} w_i \equiv 1$. The set of probabilities $\{p_i\}$ (see equation (3)) are used as basis functions for modeling the *total* probability $p(\Delta E, \Delta P_\zeta | P_\zeta, E, \mu)$:

$$p(\Delta E, \Delta P_\zeta | P_\zeta, E, \mu) = \sum_{i=1}^{N} w_i p_i(\Delta E, \Delta P_\zeta | P_\zeta, E, \mu). \quad (9)$$

From initial tests based on ORBIT, it seems that $N \lesssim 4$ might be enough for modeling cases with multiple TAE modes, each characterized by different toroidal mode number, frequency and radial mode structure. However, the approach of using the analytical representation for the p_i, complemented by the corresponding weight matrices w_i, seems rather impractical (although feasible), as the number of input files required by TRANSP would rapidly increase for scenarios that require $N > 1$ to model the impact of the modes on the fast ions.

A more straightforward way to define the input required by TRANSP is to use the total $p(\Delta E, \Delta P_\zeta | E, P_\zeta, \mu)$ directly, without any attempt to reconstruct it on the basis of specific expressions such as equations (3)–(8). This requires a single file defining p as a set of two-dimensional (2D) matrices (with variables $\Delta E, \Delta P_\zeta$) as a function of E, P_ζ and μ, with the clear advantage that a large number of scenarios can be effectively modeled (including resonant and stochastic transport) in a relatively simple and general way.

Furthermore, three observations help us to introduce a time dependence in a simplified way:

(1) For some practical cases, the scenario of interest is characterized by modes whose amplitudes vary in time, but whose relative amplitudes do not (at least on average, when time steps $\gtrsim 1$ ms are considered). Therefore,
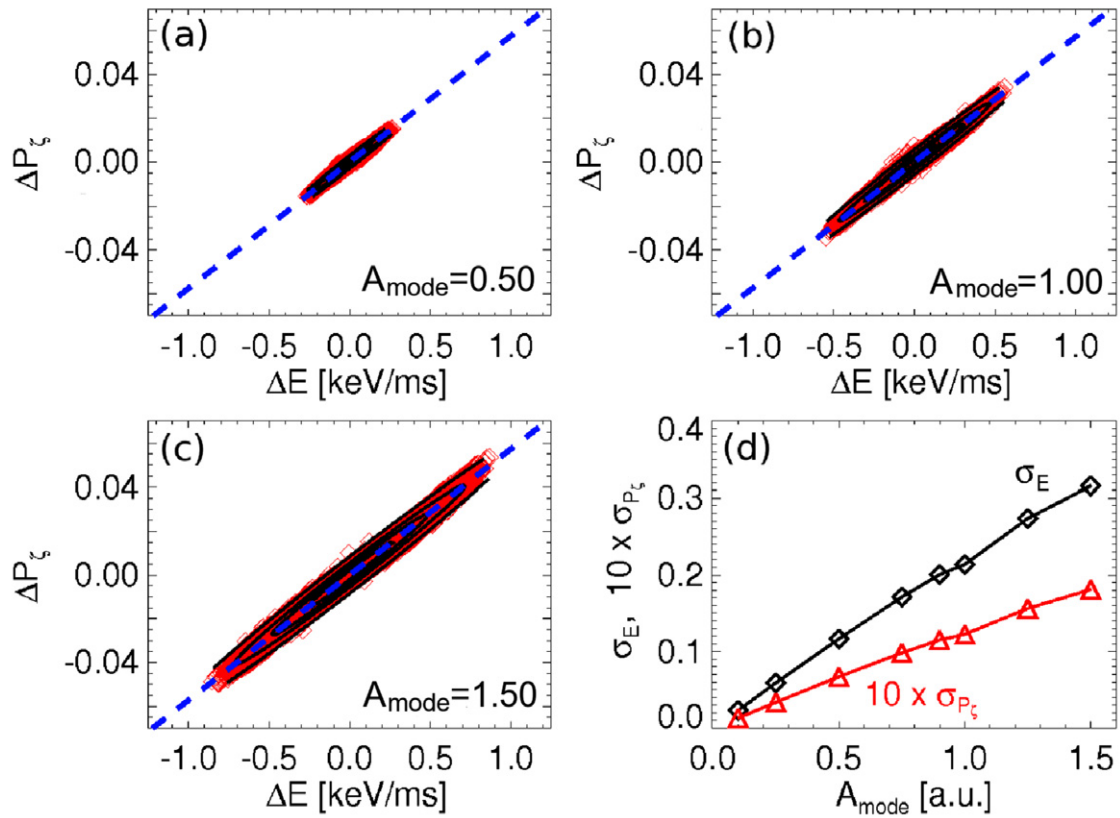
Figure 2. Scan of the mode amplitude (normalized to the experimental value, $A_{mode} = 1$, from the reflectometer's data). ORBIT runs include a single TAE mode. The initial particle distribution consists of 5000 particles, all within the same (E, P_ζ, μ) bin but with different toroidal locations and hence different phases with respect to the mode. The parameters σ_E and σ_{P_ζ} are obtained from a fit of ORBIT data with two bi-variate Maxwellians; see equation (9). (*a*)–(*c*) Contours of the kick probability (solid lines) overlaid on the particle distribution (dots) for different values of the mode amplitude. (*d*) Dependence of the standard deviations σ_E (diamonds, black) and σ_{P_ζ} (triangles, red) on the mode amplitude, showing a roughly linear relationship.

providing the temporal evolution of the *total* mode amplitude, $A_{mode}(t)$, as a 1D vector would be enough to describe the evolution of the entire set of modes.

(2) As noted before, σ_E and σ_{P_ζ} are roughly proportional to the total mode amplitude, e.g. $\sigma_E \propto A_{mode}$ (figure 2(*d*). (Note that if σ_E, σ_{P_ζ} are interpreted as equivalent *diffusion coefficients* such as $D_E \propto \sigma^2/\delta t$, the proportionality gives the expected dependence $D_E \propto A_{mode}^2$ for diffusive transport.)

(3) Boundaries and the shape of the domain (E, P_ζ, μ) evolve in time during a discharge, for instance because of the evolution of the q-profile. For practical purposes, one can assume that the same domain can be used at different times with appropriate normalizations, e.g. based on the magnetic flux at the boundary for P_ζ, or on the magnetic field on the axis for μ.

The proposed formulation is suitable for a straightforward inclusion of magnetic moment variations, that have been neglected so far. In order to preserve possible correlations between changes of E, P_ζ and μ, the general probability function should then be extended to $p(\Delta E, \Delta P_\zeta, \Delta \mu | E, P_\zeta, \mu)$, i.e. as a six-dimensional matrix. In practice, sampling methods such as the *acceptance–rejection* algorithm can still be used to perform random sampling from the arbitrarily shaped, multi-dimensional probability distribution.

3. Practical implementation of the transport model

3.1. Deriving the transport coefficients

The NSTX [18] reference case used to illustrate how the input for TRANSP can be defined is discussed in [19] and summarized in figure 3. It consists of an NSTX H-mode plasma with bursts of activity of $n = 1$–6 TAEs from ~0.2 to ~0.4 s. The modes have a more stationary character after 0.4 s, when strong low frequency MHD activity is also detected.

TAE modes are analyzed with the NOVA-K code [20] following the procedure outlined in [21]. This provides the mode structure. The perturbation amplitude associated with each mode is obtained by comparing the simulated density response (in arbitrary units) to the modes with that measured through the UCLA reflectometer system [22] installed on NSTX.

The plasma equilibrium at a given time and the mode structures are used in the particle-following code ORBIT to characterize the fast ion population (figure 4) and simulate the fast ion response in terms of transport, including loss, redistribution and energy change. In the following example,
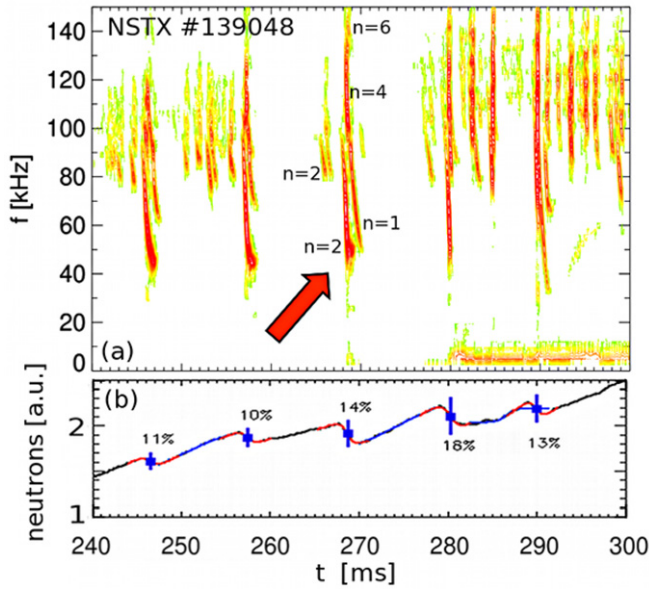
Figure 3. NSTX scenario used to illustrate how input quantities for the new model of fast ion transport in NUBEAM/TRANSP can be derived from an experiment. (*a*) Spectrogram from Mirnov coils located at the plasma edge. (*b*) Neutron rate showing repetitive drops coincident with bursts of TAE activity. The arrow indicates a specific event that will be further discussed later in the paper.

the reference time is chosen at $t \approx 0.27$ s, when the mode amplitude is sufficiently large to obtain good reflectometer data.

Numerically or analytically obtaining the probability function $p(\Delta E, \Delta P_\zeta)$ using expressions for the particle orbit and eigenmode structure is extremely difficult. In particular, modes are localized on flux surfaces, but high energy particle orbits are not, whereby a particle may experience large variations in mode amplitude in a single toroidal transit. The strength and location of resonances can only be determined by special numerical methods [15, 17].

A single ORBIT run with a sufficiently large number of test particles is used to calculate $p(\Delta E, \Delta P_\zeta | E, P_\zeta, \mu)$. Particles are initialized uniformly in the (E, P_ζ, μ) space. Since all particles have equal weight in the standard version of ORBIT, a uniform distribution in phase space is necessary to prevent gradients in the fast ion distribution from affecting the computed $p(\Delta E, \Delta P_\zeta)$. (More efficient schemes may be developed for different codes, for instance to improve the statistics by initializing more particles with adjusted weights in phase-space regions where resonances are active.) Runs simulating ~ 1 ms, corresponding to tens or hundreds of toroidal transit times, are used. The mode amplitude is kept constant ($A_{\mathrm{mode}} = 1$). Mode frequencies are also constant. During the run, each particle is tracked at fixed intervals δt_{sim} and the main parameters (including E, P_ζ, μ, orbit type) are recorded.

ORBIT results are then processed in order to infer $p(\Delta E, \Delta P_\zeta)$. Variations in E and P_ζ are calculated at each step, providing an ensemble of values over the whole phase space:

$$\begin{cases} \Delta E(E, P_\zeta, \mu) \\ \Delta P_\zeta(E, P_\zeta, \mu). \end{cases} \quad (10)$$

This ensemble is then resampled on the basis of a discrete grid in (E, P_ζ, μ), and the probability distribution function is computed as a 2D histogram for each *bin* in the grid. A normalization is applied to ensure that $\sum_{\Delta E, \Delta P_\zeta} p(\Delta E, \Delta P_\zeta) \equiv 1$ for each bin. These steps define the probability matrix that is used as input in NUBEAM/TRANSP.

3.2. Obtaining the mode amplitude scaling factor

Different approaches can be used to compute the mode amplitude scaling factor, $A_{\mathrm{mode}}(t)$. Ideally, experimental measurements of the actual mode amplitude are used. Diagnostics such as reflectometry, electron–cyclotron emission radiometry and beam emission spectroscopy are good examples. In some cases, however, internal measurements of the mode amplitude are limited or not directly available. Additional modeling is required in these cases to infer $A_{\mathrm{mode}}(t)$. In the following, a simple method based on commonly available neutron rate measurements is described. It is assumed that the mode amplitude is known at (at least) one time during the discharge. By varying the mode amplitude in ORBIT simulations with respect to the measured one, one can obtain the relationship between the mode amplitude and expected drop in the neutron rate; see figure 5. This graph, along with the measured neutron rate evolution, can be used to derive the time-dependent $A_{\mathrm{mode}}(t)$, which is then used as input for NUBEAM/TRANSP.

In practice, one calculates the instantaneous relative neutron rate variation, $\Delta R_n(t)/R_n(t)$, from the measured neutron time-trace. At each time, $\Delta R_n(t)/R_n(t)$ is then used to calculate the normalized mode amplitude that corresponds to that change from the relationship shown in figure 5. These steps are illustrated in figure 6. For comparison, the mode amplitude scaling factors from the neutron rate and from the Mirnov coils located at the plasma edge are shown in figure 6(*c*). Both methods succeed in identifying times with TAE bursts. The waveform from the neutron rate appears smoother, mainly because of the filtering of the raw data required to compute time derivatives without large spikes and noise. The consequences of a different input for $A_{\mathrm{mode}}(t)$ are briefly discussed in section 4.2.

3.3. Algorithm implementation in NUBEAM/TRANSP

The fast ion transport model presented in the previous sections has recently been implemented in Fortran and interfaced with the NUBEAM code. A flowchart of the model's implementation is shown in figure 7. The model acts on the fast ion distribution in between two steps of NUBEAM. Initial conditions are read from the so-called *plasma state* structure of TRANSP, which contains information on the machine configuration (e.g. machine size, structures, parameters of NB injection and other auxiliary heating systems) and current plasma and equilibrium parameters (such as the magnetic equilibrium, and plasma profiles).

The first step in the model is to convert the fast ion distribution F_{nb} from real-space variables (R, Z, E, p) to phase-space variables (E, P_ζ, μ). Note that the three variables
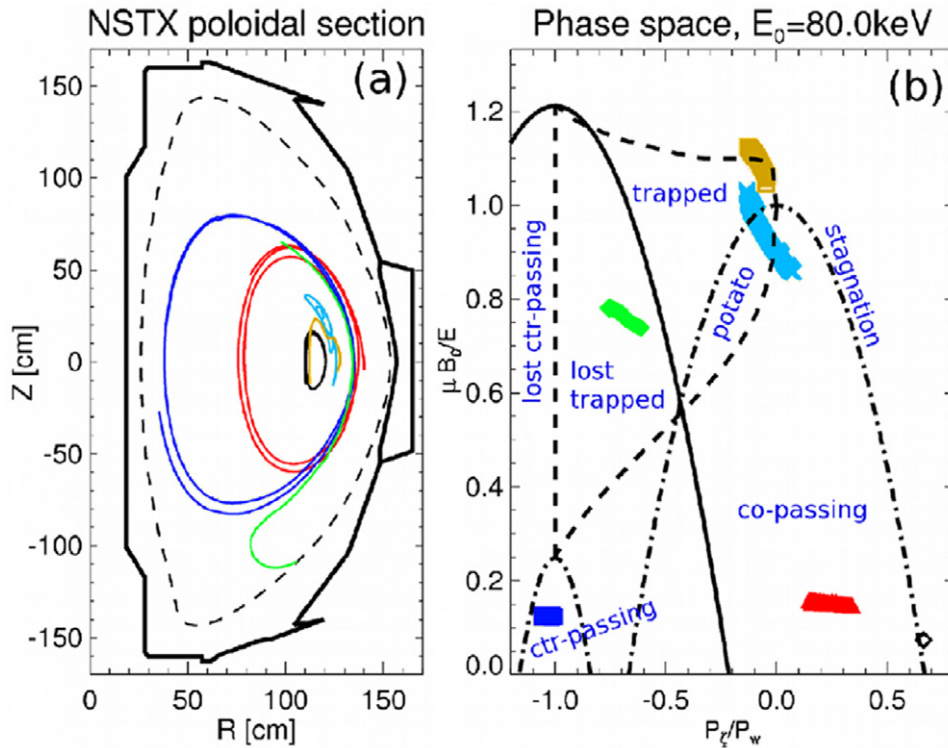
Figure 4. Examples of fast ion orbits in NSTX in the presence of four TAE modes. (*a*) Orbits in real space in an NSTX poloidal cross-section. (*b*) Equivalent phase-space representation.


Figure 5. Scan of the fractional neutron rate drop as a function of the constant mode amplitude from the ORBIT code. $A_{rel} = 1$ corresponds to the nominal experimental values of the mode amplitude for which $p(\Delta E, \Delta P_\zeta)$ is computed.

(E, P_ζ, μ) (and sometimes the sign of the parallel velocity) define a complete orbit in an axisymmetric system, whereas (R, Z, E, p) define a single point on that orbit. Since the time scale for distribution modification is long compared to transit times, variables (E, P_ζ, μ) are the relevant ones. Within the NUBEAM/TRANSP framework, this conversion from single-point (instantaneous) to orbit representation of the particle distribution is valid because of the statistical approach used in the code, i.e. the code evolves the distribution as a statistical ensemble rather than keeping track of each particle separately.

Information on the mode amplitude evolution for the current NUBEAM step is recovered from the input file. Then, the model evolves each particle according to its location in phase space and to the given mode amplitude. After the full Monte Carlo step for (E, P_ζ, μ), all remaining particle parameters are updated and the possibly new position in (R, Z) is recomputed. This leads to the conversion of $F_{nb}(E, P_\zeta, \mu)$ back to the initial representation in terms of (R, Z, E, p) required by NUBEAM/TRANSP. Once the updated F_{nb} is known, information is passed back to NUBEAM/TRANSP to update the plasma state and begin a new time step.

The full procedure is summarized here below:

(1) A TRANSP input file is defined for the step probability distribution function, $p(\Delta E, \Delta P_\zeta | E, P_\zeta, \mu)$, that is used as input to TRANSP. This consists of a five-dimensional (5D) matrix of probability values. Two variables define the $(\Delta E, \Delta P_\zeta)$ grid over which the probability for a given region in phase space (E, P_ζ, μ) is provided. The remaining three variables define the discrete regions (or *bins*) into which the phase space is divided.

(2) The time evolution of the total mode amplitude is specified as a 1D array in a separate input file.

(3) The two input files, along with the appropriate switches in the standard TRANSP parameter list, completely describe the input for a given set of modes with (possibly) time-varying amplitude.

(4) At the beginning of the TRANSP run, the probability matrix is recovered from the input files. Variables in $p(\Delta E, \Delta P_\zeta)$ are updated during the run to preserve the correct orbit topology, based on the q-profile, and magnetic equilibrium [14].
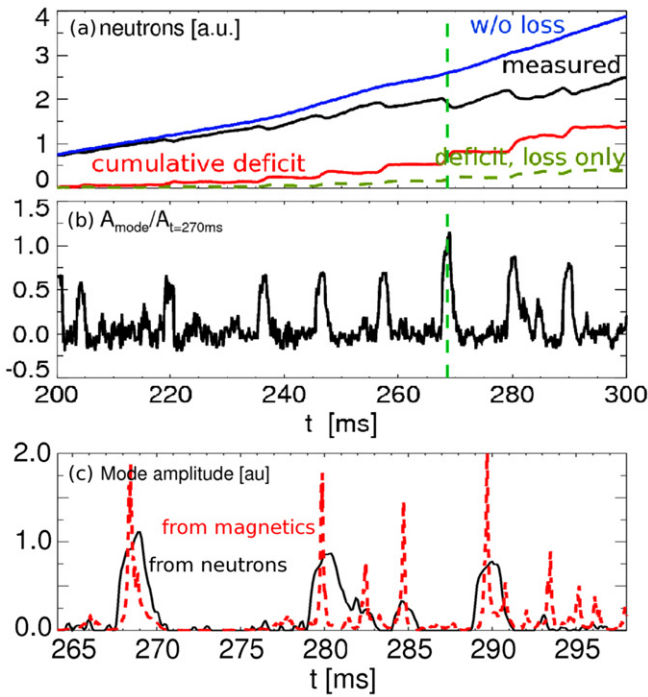
Figure 6. Calculation of the normalized mode amplitude evolution based on the measured neutron rate and using the relationship between the neutron variation and the normalized mode amplitude from ORBIT; see figure 5. (*a*) Measured neutron rate (black) and cumulative neutron deficit (red). Also shown are the reconstructed neutron rate in the absence of modes (blue) and the deficit that would result from fast ion loss only (green), i.e. without taking into account the energy variation induced by the wave–particle interaction with the TAE modes. (*b*) Mode amplitude $A_{mode}(t)$ normalized to the peak mode amplitude at $t = 268.7$ ms, obtained from the measured neutron rate evolution and the graph in figure 5. (*c*) Comparison between $A_{mode}(t)$ from neutrons and from Mirnov coil data over a limited time range.

(5) The next step is to update the trajectory of the particles at each iteration of NUBEAM. Classical mechanisms (slowing down, scattering) are treated in the standard way. Only the effects of resonant modes are discussed here.

(6) For a particle characterized by values (P_ζ, E, μ), one can identify the corresponding *bin* in the probability matrix and proceed to the random sampling of the kicks in E and P_ζ for the next iteration. The steps for the ith particle are indicated as ΔE_i and $\Delta P_{\zeta,i}$. (More details on how a particle's energy and P_ζ may be advanced in practice are given in the next section.)

(7) Now proceed to the calculation of the E and P_ζ variations for a time step of duration δt_{step}. On the basis of the time step used to compute $p(\Delta E, \Delta P_\zeta)$, indicated as t_{step}^{sim}, the step δt_{step} is divided into N_{steps} smaller time intervals of duration t_{step}^{sim} and the total variation of E and P_ζ is calculated as the result of a *correlated random walk* in time. The mode amplitude during the time step may vary according to the input $A_{mode}(t)$. The ΔE and ΔP_ζ steps are calculated from the randomly extracted ΔE_i and $\Delta P_{\zeta,i}$, and E and P_ζ are then updated.

(8) Loop over particles.

(9) Update other variables such as the radius and poloidal angle of the particles.

3.4. Monte Carlo extraction of steps ΔE_i, $\Delta P_{\zeta,i}$

The random extraction of ΔE_i and $\Delta P_{\zeta,i}$ is arguably the most important step in the procedure depicted above. An important feature is that the particle motion is characterized by different time scales; see figure 8. Firstly, particles oscillate in the wavefield over periods of $\simeq 1/f_{mode}$, which is typically much shorter than the time step used in simulations with NUBEAM/TRANSP. For many particles with similar parameters but different phase with respect to the modes, this fast motion results in a spreading of energy and P_ζ around the initial values. Secondly, the average particle energy and P_ζ drift over time scales of several toroidal transit times. This second time scale is the relevant one for the process under study. What makes these constraints important for calculating the particle trajectory in the (P_ζ, E, μ) space is that the particle motion is not purely periodic and can be skewed toward positive as well as negative values of ΔE and ΔP_ζ. In other words, over time scales of hundreds of μs a specific particle can (i) fluctuate around its instantaneous energy value, but (ii) slowly drift away from the initial energy; see figure 8. When the concept is extended from a single particle to a statistically significant ensemble of particles, a net gain or loss of energy can emerge.

Moreover, each $p(\Delta E, \Delta P_\zeta)$ for a given bin (P_ζ, E, μ) may, in general, contain both resonant and non-resonant particles because of the discrete grid used to cover the (P_ζ, E, μ) space. The algorithm used to evolve F_{nb} must be able to keep track of the 'class' to which a particle belongs, or else, for example, a non-resonant particle will have a finite probability of experiencing large energy and P_ζ excursions, which would lead to wrong results. To illustrate that, suppose the $p(\Delta E, \Delta P_\zeta)$ for a specific bin is composed of two distinct components, that are hereafter indicated by their typical energy steps only. The first group has negligible $\Delta E_1 \sim 0$ and is classified as *non-resonant* particles. The second group has large $\Delta E_2 \sim 5$ keV and is dubbed the *resonant* particles. If the average step ΔE is sampled randomly at each sub-step, particles in the first group may receive an occasional kick $\Delta E_2 \gg \Delta E_1$. As a result of the unbiased sampling, the final distribution will be much broader than it should be, even if the extra kick does not happen frequently for particles in the first group. The main consequence of developing a scheme to evolve F_{nb} in time is that the motion has to be described as a *correlated random walk*, rather than a simple random walk with every particle's step totally uncorrelated with the previous one. This is implemented in the new model in a semi-empirical way. The following explanation is given for energy variations only, then extended to the $\Delta E, \Delta P_\zeta$ case under consideration by implicitly invoking the constraint in equation (2).

To evolve the particles, the energy variation after a time δt_{step} has to be reconstructed starting from the available $p(\Delta E)$, which is computed for $t_{step}^{sim} \ll \delta t_{step}$. Begin by associating a step ΔE, randomly extracted from $p(\Delta E)$, with a particle (figure 9). Assume that $\pm \Delta E$ represents the maximum energy
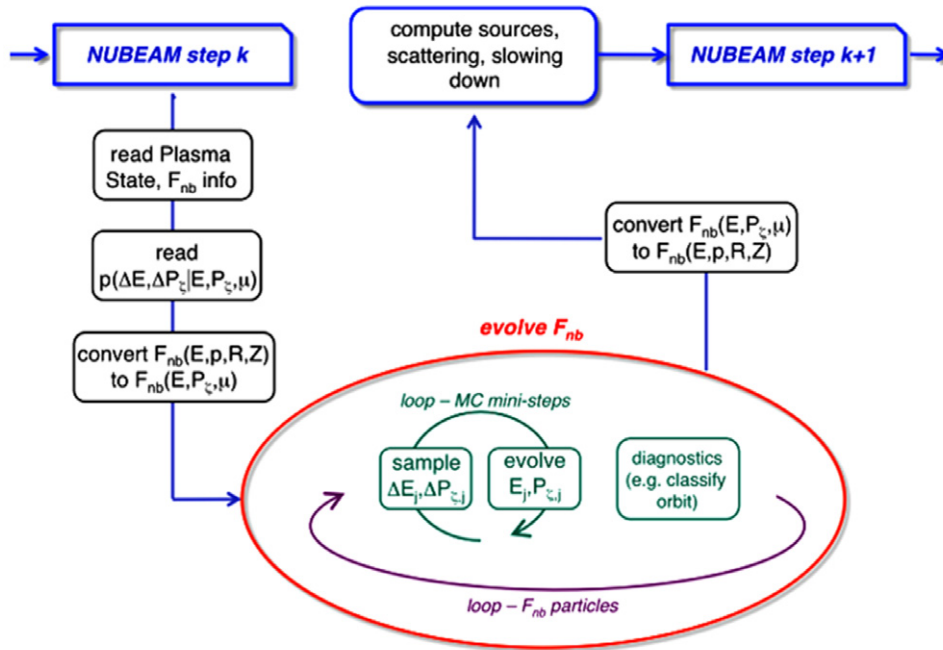
**Figure 7.** Flowchart of the model as implemented in NUBEAM.

Figure 8. Example of E, P_ζ temporal evolution for particles with different initial values of (E, P_ζ). Particles in red and cyan are lost after a few microseconds.

variation that that particle can experience under the effects of the oscillating wavefield. The particle's energy will vary in time as

$$E(t) \sim E_0 + \Delta E \sin[\phi(t) + \phi_0] + \delta E(t), \quad (11)$$

where $\phi(t)$ is the phase between the particle and wave(s), ϕ_0 is a random initial phase and $\delta E(t)$ accounts for possible (secular) drifts caused by resonances. For δE sufficiently small over a step of duration $t_{\text{step}}^{\text{sim}}$, the probability that the particle will have

an energy variation x at a certain time is thus

$$p(x) \approx \frac{1}{\pi \sqrt{\Delta E^2 - x^2}}, \qquad -\Delta E \leqslant x \leqslant \Delta E, \quad (12)$$

which is further simplified as

$$p(\Delta E) \approx \frac{\delta(\Delta E^+) + \delta(\Delta E^-)}{2} \quad (13)$$

with $\Delta E^+ \doteq +\Delta E$ and $\Delta E^- \doteq -\Delta E$. This particle is now assumed to be representative of an ensemble of particles with
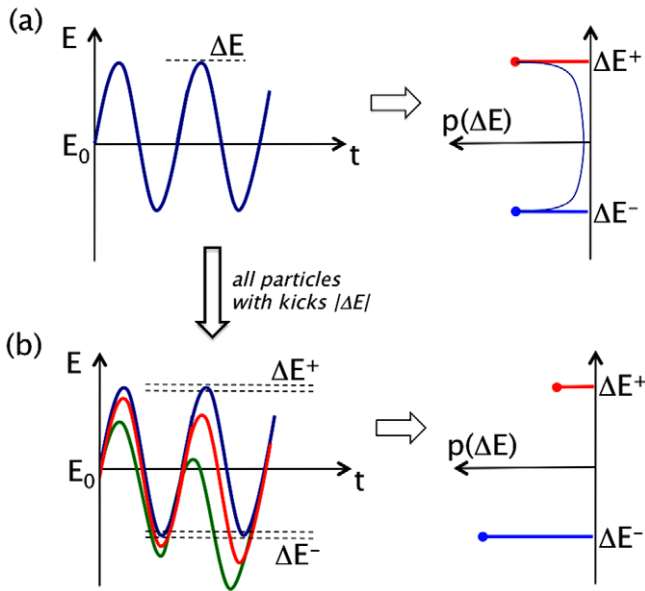
Figure 9. Qualitative derivation of the scheme for evolving the energy of the particles and P_ζ. (*a*) For a single particle oscillating in the wavefield, the probability function for experiencing a step ΔE, shown by the line in the right panel of (*a*), is approximated by two delta functions at $\pm \Delta E$. (*b*) When all particles experiencing a kick $|\Delta E|$ and a drift in energy are taken into account, an asymmetric distribution appears.

characteristic step $|\Delta E|$; see figure 9(*b*). In general, $p(\Delta E)$ may not be symmetric for the ensemble if $\delta E(t) \neq 0$. Estimate the average energy change, $\overline{\Delta E}$, after N_{steps} of duration $t_{\text{step}}^{\text{sim}}$ for the ensemble characterized by this $|\Delta E|$:

$$\overline{\Delta E} \approx \sum_{j=1}^{N_{\text{steps}}^+} \Delta E^+ + \sum_{k=1}^{N_{\text{steps}}^-} \Delta E^-, \qquad (14)$$

where N_{steps}^+ (N_{steps}^-) is the number of steps with positive (negative) energy variation.

Now define the probability for either positive or negative variations, $p_+ \doteq p(+\Delta E)$ and $p_- \doteq p(-\Delta E)$. For N_{steps} sufficiently large, the probability that the particle will experience a positive energy kick at each step is given by $p_+/(p_+ + p_-)$ (and similarly for negative kicks); therefore equation 14 can be rewritten as

$$\underbrace{[\sum_{j=1}^{N_{\text{steps}}^+} (+1)}_{N_{\text{steps}} \frac{p_+}{p_+ + p_-}} + \underbrace{\sum_{k=1}^{N_{\text{steps}}^-} (-1)]}_{N_{\text{steps}} \frac{p_-}{p_+ + p_-}} \times \Delta E \rightarrow \underbrace{\sum_{j=1}^{N_{\text{steps}}} S_{r,k}}_{N_{\text{steps}} \frac{p_+ - p_-}{p_+ + p_-}} \Delta E, \qquad (15)$$

where $S_{r,k}$ is the sign of a random number uniformly extracted from $[-p_-/(p_+ + p_-), p_+/(p_+ + p_-)]$ at each step. Finally, since the step ordering in equation (14) is probabilistic, the overall $\overline{\Delta E}$ from equations (14)–(15) can be cast as

$$\overline{\Delta E} \approx \sum_{k=1}^{N_{\text{steps}}} S_{r,k} \Delta E \, A_{\text{mode},k} \qquad (16)$$

with the substitution $\Delta E \rightarrow \Delta E \, A_{\text{mode},k}$ at each step k to account for the proportionality $\sigma_E \propto A_{\text{mode}}$; cf figure 2.

On the basis of this semi-empirical derivation, the overall ΔE and ΔP_ζ over a time δt_{step} are calculated as the result of multiple steps, each of duration $t_{\text{step}}^{\text{sim}}$ that is *sufficiently shorter* than δt_{step}. For example, the tests presented hereafter have $t_{\text{step}}^{\text{sim}} \ll \delta t_{\text{step}}/10$. The kicks ΔE and ΔP_ζ are calculated for the first step only. The only exception is when a particle moves from one bin to another in (P_ζ, E, μ) before the full step δt_{step} has been covered, for example because it experiences a large energy kick. In this case, new values of ΔE and ΔP_ζ are sampled on the basis of the new phase-space coordinates.

4. Verification of the reduced model

4.1. Verification against ORBIT simulations

In this section, the algorithm that implements the new transport model is verified against full ORBIT simulations for the scenario introduced in figure 3. Because a realistic modeling of fast ion dynamics in phase space lies at the core of the new model, the ability of the code to identify and evolve fast ion orbits is first tested. For a given fast ion distribution (e.g. from NUBEAM/TRANSP), orbits are classified and their trajectory reconstructed in both real space and phase space [14]; see figures 10–11.

Figure 10(*a*) shows a co-passing, confined fast ion that is pushed into a trapped confined region by the interaction with TAE modes. The change in orbit type is clear from its projection in the (R, Z) poloidal cross-section of NSTX. The evolution of the orbit in the corresponding (P_ζ, μ) plane (figure 10(*b*)) shows that, as the particle loses energy and reaches a more negative Ψ (i.e., larger minor radius), it first moves on a potato orbit and finally it becomes a trapped particle.

Similarly, figure 11 shows the loss process for a co-passing, confined particle at larger pitch that crosses the loss boundary and hits the wall. It is interesting to note that the final orbit resides in a region of phase space where co-passing lost and counter-passing confined orbits coexist (figure 11(*b*)). By classifying the orbit of the particles at each Monte Carlo time step, the code can still resolve the correct orbit evolution.

A comparison of the evolutions of the entire distribution from ORBIT and from the reduced model is shown in figure 12. Histograms for E and P_ζ variations are shown for co-passing, trapped and all particles in the original distribution. Considering the simplicity of the model, the agreement between the two codes is satisfactory over at least three orders of magnitude, as indicated by the high values of the coefficient of correlation, $\rho \geqslant 0.97$, between results from ORBIT and from the new model. Possibly, the only exception is in the energy and P_ζ variations for trapped particles, for which the relative error between the model and ORBIT results approaches 1 for the bulk of the distribution around $\Delta E = 0$, $\Delta P_\zeta = 0$. This is attributed to the relatively poor statistics for this phase-space region, although uncertainties in the reconstructed probability $p(\Delta E, \Delta P_\zeta | E, P_\zeta, \mu)$ cannot be completely ruled out.
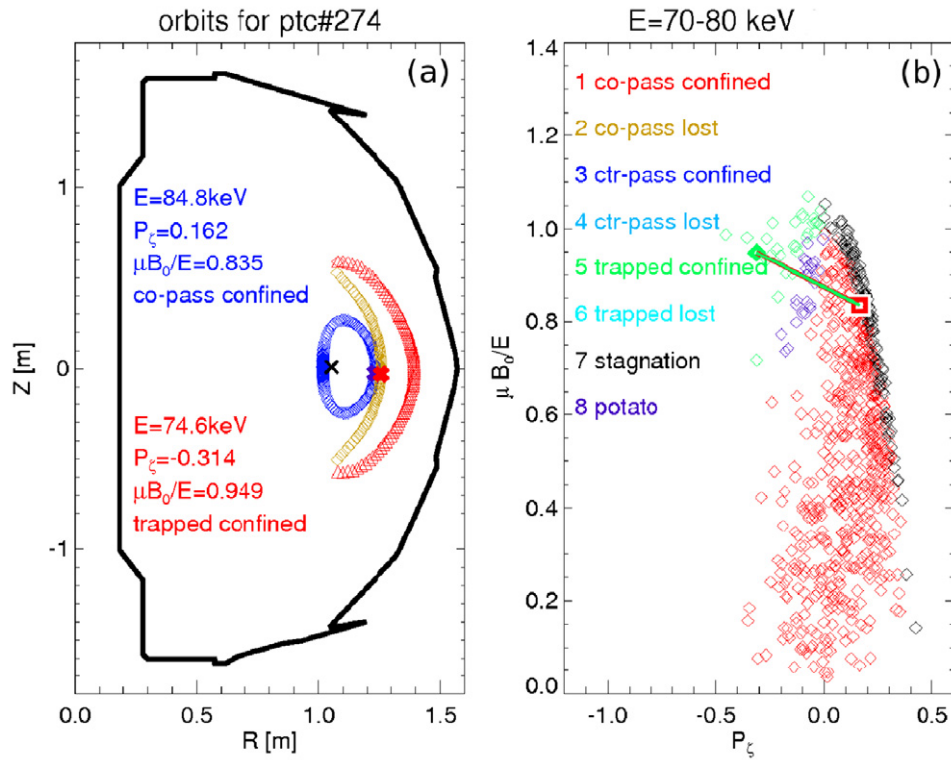
Figure 10. Illustration of orbit topology change induced by TAE modes in (*a*) real space and (*b*) phase space. The particle's parameters are reported in (*a*). Dark (light) colours for the orbit indicate co-going (counter-going) motion. The particle's trajectory in phase space is shown in (*b*) as a colored line connecting the initial orbit (red square) to the final orbit (green diamond). Final locations of other particles with similar initial energy are shown by the symbols.
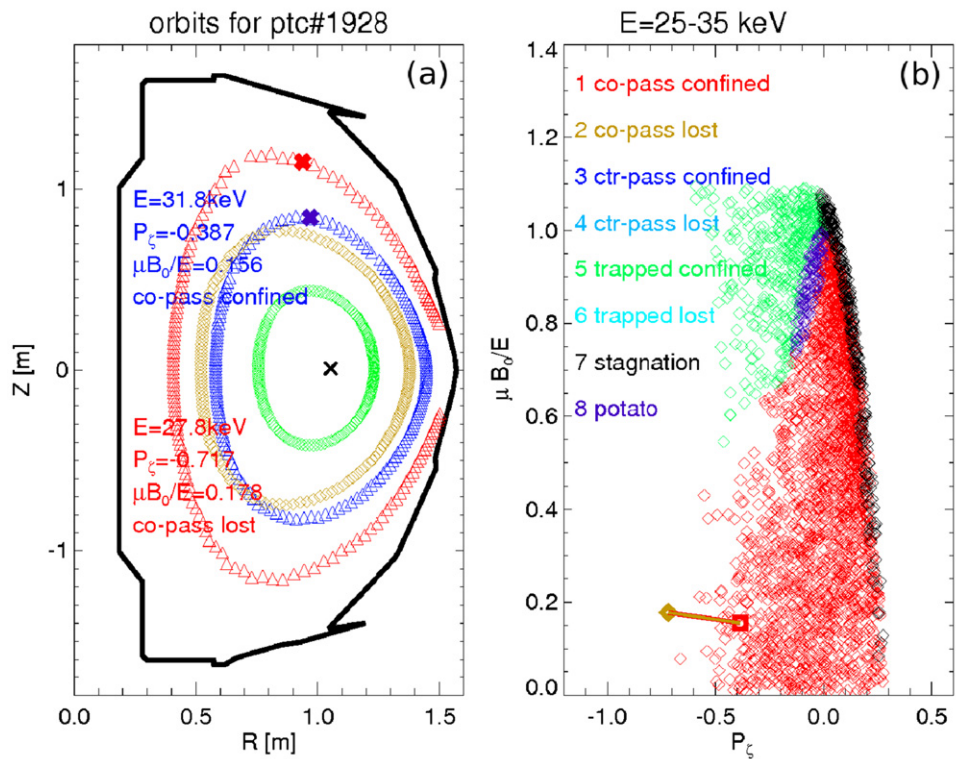

Figure 11. As figure 10, but for a lost co-passing particle. Note that, in addition to the particle's initial (blue) and final (red) orbits, the code correctly identified other orbits for confined counter-passing particles (yellow and green orbits) in the same phase-space region.
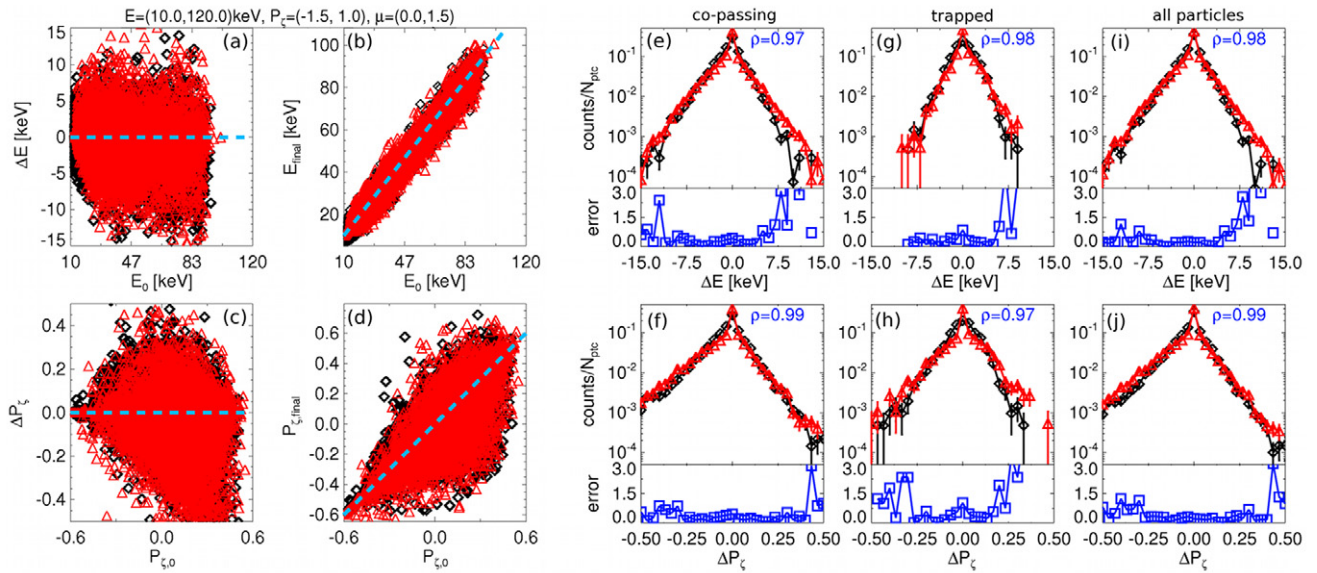
Figure 12. Test of the Fortran algorithm for evolving F_{nb} over 5 ms according to a given $p(\Delta E, \Delta P_\zeta | P_\zeta, E, \mu)$ and $A_{mode}(t)$. Results from the Fortran code (red triangles) are compared to full simulations performed with ORBIT (black diamonds). Through ORBIT, $p(\Delta E, \Delta P_\zeta)$ is calculated using steps of $t_{step}^{sim} = 25\ \mu s$. Scatter plots of energy and P_ζ variations are shown in (a)–(d). The corresponding distributions from the model and from ORBIT are compared in (e)–(j). For each case, ρ indicates the coefficient of correlation between the model and ORBIT results. The bottom panels show the relative error between the model and ORBIT results as a function of ΔE and ΔP_ζ. Test particles are selected with values of E, P_ζ and μ corresponding to ((e), (f)) co-passing, ((g), (h)) trapped and ((i), (j)) all particles with $E \geqslant 10\ keV$.


Figure 13. Probability distributions for the final energy step ΔE from full ORBIT simulations (black diamonds) and from the model (red triangles). The mode amplitude $A_{mode}(t)$ is scaled by a factor of 0.25 up to 1.5. Error bars are simply estimated as proportional to the square root of the number of counts in each energy bin.
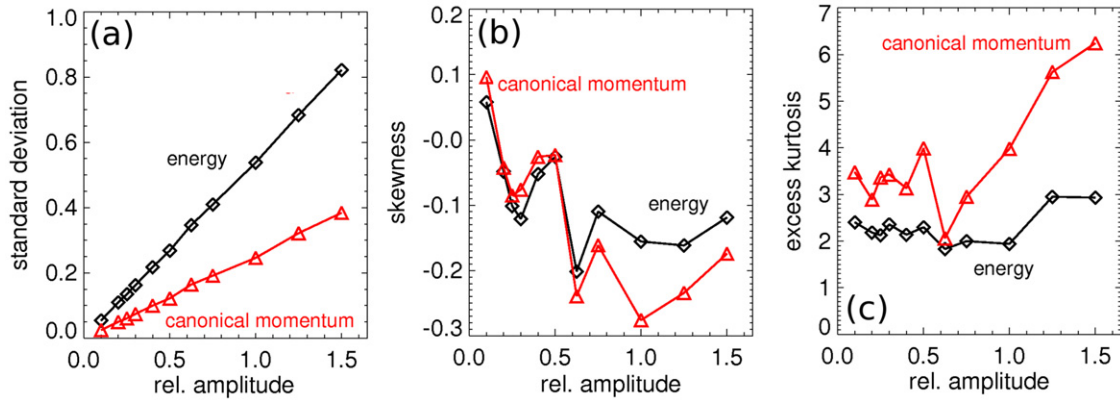
Figure 14. Variations of the shape of the $p(\Delta E)$ and $p(\Delta P_\zeta)$ probability distributions as a function of the mode amplitude scaling factor are not directly observed in the standard deviation (*a*), but do appear in both the skewness (*b*) and the kurtosis (*c*) as they change with the scaling factor.

Larger discrepancies between the Fortran algorithm and full ORBIT runs are observed when the overall mode amplitude evolution is scaled to smaller or larger values than those used for computing $p(\Delta E, \Delta P_\zeta)$; see figure 13. The cause for the increased discrepancy is illustrated in figure 14, where the second, third and fourth moments of the distributions of ΔE and ΔP_ζ are shown as a function of the amplitude scaling factor (the nominal mode amplitude is 1) in the presence of three TAE modes. As seen in the figure, the standard deviation does indeed increase linearly with the mode amplitude, similarly to what was already shown for a single mode (see figure 2(*d*)). However, both the skewness and the kurtosis vary with the mode amplitude, indicating that the shape of the probability distribution from ORBIT does vary with A_{mode}. In this case, the assumption of simple linear scaling of ΔE and ΔP_ζ with A_{mode} is not adequate for accurately reproducing the actual dependence $p = p(A_{\text{mode}})$. Improvements to the model to include more accurate scaling of the probability shape are under consideration. Nonetheless, this exercise provides an indication of the uncertainties in reconstructed fast ion evolution associated with computing $p(\Delta E, \Delta P_\zeta)$ at a single time for the reduced model, and then applying the same probability distribution to a broader time window.

The computed fast ion evolution in real space for this NSTX scenario is shown in figure 15. TAEs cause a redistribution to the outer minor radii, with a drop in the number of fast ions near the plasma center ($\sqrt{\Psi} \lesssim 0.2$) of O(10%). The corresponding increase in fast ion population at outer radii appears small, but it should be considered that particles are diluted into regions with much larger volume. The relative variation shown in figure 15(*b*) is more suitable for quantifying the fast ion density increase at $\sqrt{\Psi} \gtrsim 0.4$. For this case, losses remain limited to $\ll$1%, or O(10) lost particles from the initial population of 34 000 particles from NUBEAM.

The redistribution of fast ions is expected to affect other quantities, such as the amount of NB-driven current, here simply approximated by $I_{\text{NB}} \propto \sum_k p_k \sqrt{E_k}$, where the sum runs over all particles with $E \geqslant 10$ keV at each minor radius $\sqrt{\Psi}$. The results are shown in figure 15(*c*). Redistribution of core fast ions causes a drop of I_{NB} near the magnetic axis, and a slight increase for $\sqrt{\Psi} \gtrsim 0.6$.
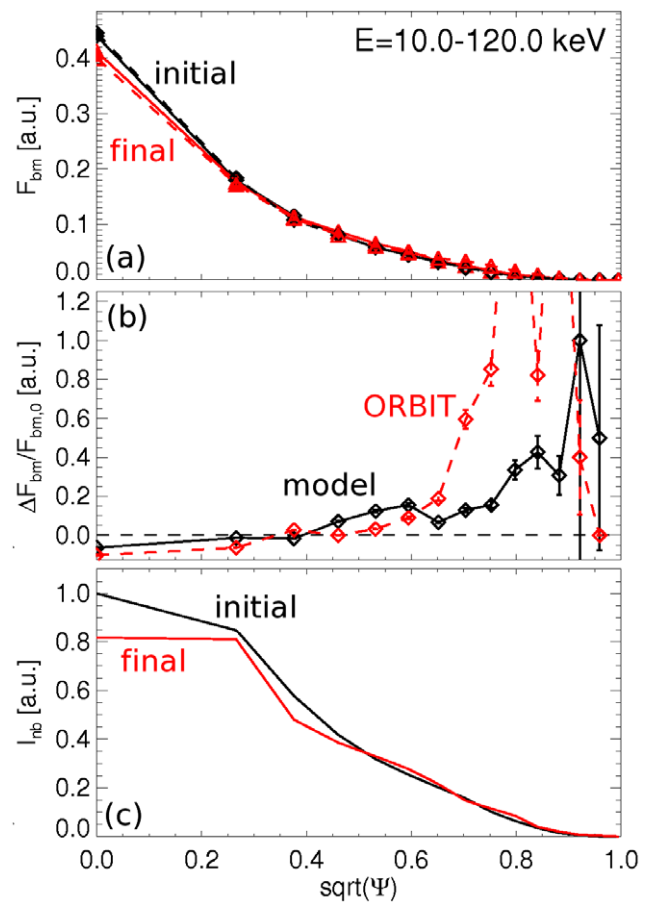

Figure 15. Computed redistribution of F_{nb} over a 5 ms step with a burst of TAE activity. Shown are (*a*) the profile variation and (*b*) the relative profile variation as a function of the normalized minor radius. Solid (dashed) lines in (*a*) refer to the Fortran algorithm (ORBIT) simulation. Statistical uncertainties are assumed to be equal to the square root of the counts. Panel (*c*) shows the estimated variation in the NB-driven current profile.

4.2. Initial tests within the NUBEAM/TRANSP framework

Before its final implementation in the NUBEAM module of TRANSP, the reduced model is tested with a stand-alone version of NUBEAM. Iterations between the two codes are used to test all of the steps illustrated in the conceptual

Figure 16. (*a*) Mode amplitude scaling factor derived from the measured neutron rate; see figure 6. (*b*) Measured (hatched region) and simulated neutron rates. Rates are normalized to their value at $t = 266$–267 ms (gray shaded region). Symbols refer to different time steps δt_{step} for the macroscopic NUBEAM iterations. (*c*) Same as (*b*), but forcing the reduced model to resample ΔE and ΔP_ζ after a different number of time steps (parameter N in the figure), each of which has duration $t_{\text{step}}^{\text{sim}} = 25\,\mu$s.

flowchart of figure 7. Plasma profiles such as the density and temperature ones are assumed to be constant in time. The magnetic configuration is also considered as fixed. Without modifying this background scenario, the code evolves the fast ion population and other quantities related to it, for instance the neutron rate and NB-driven current profile.

The goal of these initial tests is to verify that the new model is capable of evolving fast ions over time periods of the order of (or longer than) typical collisional and slowing down times, i.e. $\gtrsim 10$–20 ms for NSTX plasmas. Considering the scheme in figure 7, adjustable parameters for this simulation are the total duration of the simulation and the length of each NUBEAM step $k \rightarrow k + 1$ required to cover that time range. The latter also defines how frequently the fast ion population is updated by the new model.

The results of simulations with varying step size are shown in figure 16, where measured and reconstructed neutron rates, R_n, are compared for a specific NSTX discharge. The mode amplitude scaling factor is derived from the neutron rate as discussed in figure 6. The measured neutron rate is divided by the central deuterium density at each time to account for


Figure 17. Dependence of the simulated neutron rate on the mode amplitude. (*a*) Mode amplitude scaling factor rescaled to $k_{\text{scale}} = 0.75$, 1 and 1.25 times its nominal value. (*b*) Measured and simulated neutron rate, including the case with no modes for reference. The simulations have $\delta t_{\text{step}} = 2$ ms, $N = 80$. Solid lines represent linear fits of the simulated neutron rate, showing the same rate of increase in regions where the mode activity is negligibly small.

the fact that plasma profiles are kept constant in time in the simulation. Neutron rates are then normalized to their value before the TAE bursts begin, $t = 266$–267 ms.

Figure 16(*b*) demonstrates that the new model is indeed reproducing the correct neutron rate (or, rather, its correct temporal evolution). Exceptions to this are for the runs in which iterations between the model and NUBEAM happen every 0.5 ms and 4 ms. The reason for this discrepancy is twofold, as shown in figure 16(*c*). When δt_{steps} between two NUBEAM iterations is $\ll 1$ ms, the steps ΔE and ΔP_ζ are resampled too frequently in the model. The requirement of a *correlated* random walk for the fast ion evolution is violated, resulting in a net increase in the total transport (see section 3.3). This is confirmed by a test run with the standard $\delta t_{\text{step}} = 2$ ms but forcing the resampling of ΔE and ΔP_ζ every $N = 20$ steps, instead of the default $N = 80$. Transport is artificially enhanced in this case. The second reason for the discrepancy appears when δt_{steps} approaches the reciprocal of the NB injection rate. Since the new model is applied here at the beginning of each step, this means that newly injected particles are not redistributed, resulting in reduced transport. The three cases with $\delta t_{\text{steps}} = 4$ ms in figure 16(*c*) confirm this interpretation. Transport is underpredicted (i.e., the neutron rate is overestimated) with respect to the measured one for $\delta t_{\text{steps}} = 1$–2 ms, unless N is reduced to <50, thus causing a fictitious increase in the transport as explained above. These two examples indicate that the value of δt_{steps} must be chosen in the NUBEAM run to be much larger than the duration of the microsteps in the model, $t_{\text{step}}^{\text{sim}}$, but sufficiently shorter than the collisional and NB injection time scales.
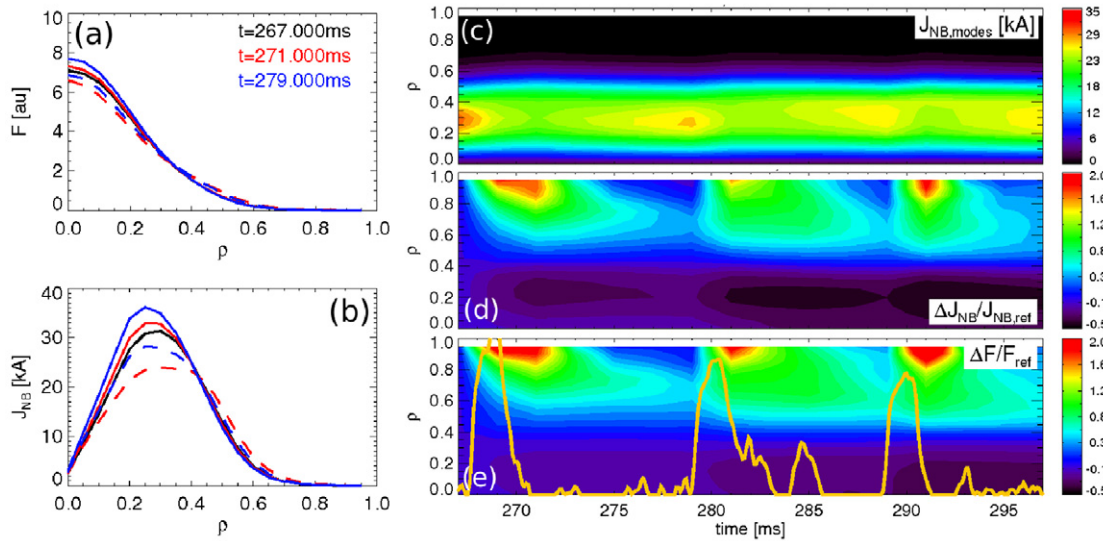
Figure 18. Simulated effects of bursting TAE modes on fast ion and NB-driven current profiles. The radial variable ρ corresponds to the normalized toroidal flux. Simulations have $\delta t_{step} = 2$ ms, $N = 80$. (*a*) Radial fast ion profile before, just after and 10 ms after the first TAE burst occurs. (*b*) Same as (*a*), but for the NB-driven current profile. Solid (dashed) lines in (*a*) and (*b*) refer to simulations without (with) enhanced transport from TAEs. The bottom panels show the temporal evolution of the (*c*) NB-driven current $J_{nb}(\rho)$, (*d*) relative variation of $J_{nb}(\rho)$ normalized to the reference case with no modes and (*e*) relative variation of the fast ion profile $F_{nb}(\rho)$ normalized to the case with no modes. The solid line in (*c*) shows $A_{mode}(t)$.

Once the correct step size is adopted, it is useful to assess the sensitivity of the simulation results to the input parameters, and especially to the mode amplitude scaling factor, A_{mode}. Figure 17 illustrates the computed neutron rate for mode amplitudes varying by $\pm 25\%$ from the nominal value. The case with $A_{mode} = 0$ is also shown for reference. It can be seen that the time evolution of the neutron rate is quite sensitive to A_{mode}. This is because R_n at a given time depends in part on the previous history of the fast ion population. Inaccurate reconstructions of the fast ion transport at earlier times cumulate, at least on time scales comparable to the slowing down time, leading to increased discrepancy for the wrong choice of A_{mode}. In practice, the neutron rate also evolves at a rate set by the NB injection rate at times with negligible mode activity. Any deviation between runs with different values of A_{mode} is attributed to previous drops in R_n, which scales with A_{mode}.

The sensitivity of the reconstructed R_n to the input A_{mode} has two additional implications. Firstly, it is a confirmation (although indirect) that the implementation of the new model within NUBEAM/TRANSP does reflect the fast ion dynamics as obtained from the original full simulations with the ORBIT code. Secondly, it may be used to infer with some confidence transport levels associated with other modes that are not included in the computation of $p(\Delta E, \Delta P_\zeta)$, but may be present in the original experiment. Another implication of the tight dependence $R_n = R_n(A_{mode})$ is that it may enable simulations in which the measured neutron rate is supplied as input, and A_{mode} is determined during the run to match it. The possible implementation of this R_n-*feedback* scheme directly inside the model is left as future work.

Time evolutions of the fast ion and NB-driven current profiles are finally considered as the last example in this section; see figure 18. Like for the neutron rate, significant

drops are observed in both F_{nb} and J_{nb} when bursts of TAEs occur. It is important to note that, as for the neutron rate, the effects of each burst are not limited in time to the duration of the burst, but propagate over a much longer time span, of the order of the slowing down time.

5. Conclusions

A new fast ion transport model has been developed for the tokamak transport code TRANSP. The model is based on a Monte Carlo approach, for mimicking the effects of instabilities on the fast ion population. A probability distribution matrix, along with a mode amplitude scaling factor, is used to define the fast ion response to the modes as a function of time and fast ion phase-space coordinates. Preliminary verification of the new model against the guiding center code ORBIT has been performed. Initial tests with a stand-alone version of the NUBEAM module, which evolves the fast ion distribution in TRANSP, are successful. The implementation of the model in the NUBEAM module of TRANSP is under way. It will be followed by extensive verification and validation work to assess the potential of the new model as well as its limitations.

Acknowledgment

References

[1] Spong D A, Carreras B A and Hedrick C L 1992 *Phys. Fluids* B **4** 3316
[2] Briguglio S, Vlad G, Zonca F and Kar C 1995 *Phys. Plasmas* **2** 3711

[3] Park W, Belova E V, Fu G Y, Tang X Z, Strauss H R and Sugiyama L E 1999 *Phys. Plasmas* **6** 1796

[4] Lang J, Chen Y, Parker S E and Fu G-Y 2009 *Phys. Plasmas* **16** 102101

[5] Todo Y, Shinohara K, Takechi M and Ishikawa M 2005 *Phys. Plasmas* **12** 012503

[6] Schneller M, Lauber Ph, Bilato R, Garcia-Munoz M, Brudgam M, Gunter S and the ASDEX Upgrade Team 2013 *Nucl. Fusion* **53** 123003

[7] Ghantous K, Gorelenkov N N, Berk H L, Heidbrink W W and Van Zeeland M A 2012 *Phys. Plasmas* **19** 092511

[8] Heidbrink W W, Van Zeeland M A, Austin M E, Bass E M, Ghantous K, Gorelenkov N N, Grierson B A, Spong D A and Tobias B J 2013 *Nucl. Fusion* **53** 093006

[9] Waltz R E and Bass E M 2013 Prediction of the fusion alpha density profile in ITER from local marginal stability to Alfvèn eigenmodes *13th IAEA-TM EP Meeting (Beijing, China, 2013)*

[10] Hawryluk R J 1980 An empirical approach to tokamak transport *Physics of Plasmas Close to Thermonuclear Conditions* (Brussels: CEC) vol 1, pp 19–46

[11] For more details on the TRANSP code, please visit the TRANSP webpage at http://w3.pppl.gov/p̄share/help/transp.htm

[12] Pankin A, McCune D, Andre R, Bateman G and Kritz A 2004 The tokamak Monte Carlo fast ion module NUBEAM in the National Transport Code Collaboration library *Comput. Phys. Commun.* **159** 157

[13] Goldston R J, McCune D C, Towner H H, Davis S L, Hawryluk R J and Schmidt G L 1981 *J. Comput. Phys.* **43** 61

[14] White R B 2006 *The Theory of Toroidally Confined Plasmas* 2nd edn (London: Imperial College Press)

[15] White R B 2012 Modification of particle distribution by MHD instabilities I *Commun. Nonlinear Sci. Numer. Simul.* **17** 2200

[16] White R B and Chance M S 1984 Hamiltonian guiding center drift orbit calculation for plasmas of arbitrary cross section *Phys. Fluids* **27** 2455

[17] White R B 2011 Modification of particle distributions by magnetohydrodynamic instabilities II *Plasma Phys. Control. Fusion* **53** 085018

[18] Ono M *et al* and the NSTX Team 2000 *Nucl. Fusion* **40** 557

[19] Fredrickson E D *et al* 2013 *Nucl. Fusion* **53** 013006

[20] Cheng C Z 1992 *Phys. Rep.* **211** 1

[21] Fredrickson E D *et al* 2009 *Phys. Plasmas* **16** 122505

[22] Crocker N A *et al* 2008 *Phys. Plasmas* **15** 102502