UC Irvine

UC Irvine Previously Published Works

Title

A code that simulates fast-ion D α and neutral particle measurements

Permalink

<https://escholarship.org/uc/item/12f712f4>

Journal

Communications in Computational Physics, 10(3)

ISSN

1815-2406

Authors

Heidbrink, WW

Liu, D

Luo, Y

et al.

Publication Date

2011

DOI

10.4208/cicp.190810.080211a

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

A Code that Simulates Fast-Ion D_α and Neutral Particle Measurements

W. W. Heidbrink^{1,*}, D. Liu^{1,3}, Y. Luo^{1,4}, E. Ruskov¹ and B. Geiger²

¹ *Department of Physics and Astronomy, University of California, Irvine, California, CA 92697, USA.*

² *Max-Planck Institute für Plasmaphysik, Garching, Germany.*

³ *Department of Physics, University of Wisconsin-Madison, Madison, WI 53706, USA.*

⁴ *Tri Alpha Energy Corporation, 27211 Burbank, Foothill Ranch, CA 92610, USA.*

Received 19 August 2010; Accepted (in revised version) 8 February 2011

Available online 1 June 2011

Abstract. A code that models signals produced by charge-exchange reactions between fast ions and injected neutral beams in tokamak plasmas is described. With the fast-ion distribution function as input, the code predicts the efflux to a neutral particle analyzer (NPA) diagnostic and the photon radiance of Balmer-alpha light to a fast-ion D_α (FIDA) diagnostic. Reactions with both the primary injected neutrals and with the cloud of secondary "halo" neutrals that surround the beam are treated. Accurate calculation of the fraction of neutrals that occupy excited atomic states (the collisional-radiative transition equations) is an important element of the code. Comparison with TRANSP output and other tests verify the solutions. Judicious selection of grid size and other parameters facilitate efficient solutions. The output of the code has been validated by FIDA measurements on DIII-D but further tests are warranted.

PACS: 52.55.Pi, 52.65.Pp, 52.70.Kz

Key words: Fast ions.

1 Introduction

Supra-thermal populations of energetic ions play an important role in magnetic fusion research. These "fast ions" are created by neutral-beam injection, by RF heating, and in fusion reactions. The distribution function that describes these populations generally is a

*Corresponding author. *Email addresses:* Bill.Heidbrink@uci.edu (W. W. Heidbrink), dliu29@wisc.edu (D. Liu), yluo@trialphaenergy.com (Y. Luo), eruskov@uci.edu (E. Ruskov), bgeiger@ipp.mpg.de (B. Geiger)

complicated function of velocity and configuration-space variables. Measuring the fast-ion distribution function in the harsh magnetic fusion environment is a major diagnostic challenge.

One approach is to exploit charge exchange reactions between energetic deuterium ions and an injected neutral beam. Collection of escaping neutrals is the basis of neutral particle analysis (NPA) [1], a technique that has been applied to tokamak plasmas for nearly five decades [2]. A more recent technique is to analyze the visible photons emitted by hydrogenic fast ions that neutralize in the injected beam [3]. A review of these fast-ion D_α (FIDA) measurements was recently published [4].

Both NPA and FIDA diagnostics provide valuable information about the fast-ion distribution function but also depend sensitively on other plasma parameters and on atomic cross sections. One way to relate the measured signals to theory is to construct a phase-space weight function for each measurement [5]; the signal is the convolution of the fast-ion distribution function with the weight function. As illustrated by the examples in [4], this approach is quite useful for rapid qualitative interpretation of the measurements. It can also be the basis for an inversion algorithm. Although the processes are too complicated for a unique inversion [6], a least-squares minimization scheme that utilizes a weight function can determine which model distribution function agrees best with the data. An example of inference of the distribution function from collective Thomson scattering data was recently published [7].

Alternatively, one can use forward modeling. In this approach, the distribution function is a given quantity supplied by theory. The code described in this paper, dubbed FIDASIM, takes this approach. FIDASIM accepts a theoretical distribution function as input and predicts FIDA and NPA spectra for comparison with the data. The code is designed to compute "active" signals produced by an injected neutral beam. (In reality, collisions with edge neutrals also produce FIDA and NPA signals but the code does not treat these "passive" reactions.) To date, the code has been used to model measurements on the DIII-D and ASDEX-Upgrade conventional tokamaks and on the NSTX and MAST spherical tokamaks. An early version of the code was described in the Appendix of [3]. This paper describes version 3.0 and is organized as follows. Section 2 presents the assumptions and organization of the code. Section 3 describes tests that verify that the code correctly solves the desired equations. Section 4 explains the optimal selection of numerical parameters in terms of physical processes. Section 5 summarizes validation by experiment. Section 6 provides an outlook for further tests and improvements.

2 Model

The code has four main sections (Fig. 1). The first section prepares the data and the second calculates the neutral populations. The third and fourth sections both rely on the first two sections but are independent of each other. One section computes the NPA flux and the other computes the FIDA radiance.

Figure 1: Flow diagram for the FIDASIM code.

2.1 Input data and coordinate mapping

The code begins by collecting the input data. The geometry of the source of injected neutrals is specified first. In some devices (such as NSTX) the detector sightlines intersect several beams, so the code can accommodate multiple beam lines. The code uses the conventions of the NUBEAM module [8] of the TRANSP code [9] to describe the geometry of the viewed neutral beam source (or sources). Each tokamak has its own subroutine called, e.g., BEAM_GEOMETRY_D3D. As in NUBEAM, the neutral beam is described by rectangular source and aperture dimensions and by focal lengths and divergences in both the horizontal and vertical directions. The beam energy, power, and species mix between full-energy, half-energy, and third-energy components are also input parameters.

Next, the code collects information about the detector locations and sightlines. For FIDA, the "detector" location is actually the position of the primary lens (or mirror) of the collection optics, since it is this position that determines the Doppler shift of the emitted radiation. For an NPA, both the sightlines and the solid angles are specified.

Information on the equilibrium is input using the so-called "eqdsk" format produced by the EFIT equilibrium code [10]. For installations that do not use EFIT, a post-processor that is part of the TRANSP distribution can convert TRANSP output files into the desired format.

The code requires profiles of electron density and temperature, ion temperature and toroidal rotation, and impurity density as a function of flux surface. (These quantities are all assumed to be flux functions.) A subroutine exists that converts TRANSP output into the desired format.

The final major piece of input data is the theoretical fast-ion distribution function, which can have a complicated dependence on energy E , pitch $p = v_{\parallel} / v$, and space $\mathbf{r}$. (As in TRANSP, positive p is defined by the direction of the plasma current rather than by the direction of the toroidal field.)

Three distinct coordinate systems are utilized in the initial stages of the code (Fig. 2). The beam and detector geometries are specified in right-handed Cartesian (u, v, z) coordinates with origin the center of the tokamak and z the vertical direction. Plasma parameters are one-dimensional functions of flux coordinates. Because neutrals travel in

Figure 2: Plan view of NSTX. Geometrical neutral beam and detector input to the code is in (u, v, z) coordinates. Neutral beam parameters (upper case labels) follow the TRANSP conventions. The code transforms quantities into (x, y, z) coordinates along the selected beam.

straight lines, a right-handed Cartesian grid (x, y, z) is employed in the main sections of the code. The centerline of the primary beam forms the x axis, with the origin just outside the plasma. After collecting all of the needed input data, the code transforms (u, v) coordinates into (x, y) coordinates and finds the flux coordinate that corresponds to each position in the (x, y, z) mesh. The values of temperature, density, and plasma flow at each cell position are stored in arrays; in addition, several other quantities used in the calculations of collisional-radiative transitions (discussed below) are computed and stored. From the equilibrium data, the magnetic and electric fields in (x, y, z) coordinates at each cell position are computed and stored. Similarly, the (x, y, z) vectors from each cell to the FIDA lens (or lenses) and NPA detectors are placed in arrays.

Depending on the source of the theoretical fast-ion distribution function, the mapping into the (x, y, z) coordinates can be fairly complicated. To date, distributions produced by the TRANSP NUBEAM [8], ORBIT-RF [11], and CQL3D [12] codes have been used. The desired output is a three-dimensional array of the guiding-center distribution function in the variables energy, pitch, and cell number, $F(E, p, \text{cell})$. The energy and pitch variables have uniform spacing dE and dp . The normalization for F is chosen so that the sum over velocity space is the guiding-center fast-ion density n_f in each cell, i.e.,

$$\sum_i \sum_j F(E_i, p_j) dE dp = n_f. \quad (2.1)$$

Conversion of the NUBEAM output into this form is straightforward, as the NUBEAM distribution is already a function of energy, pitch, and an array of (R, z) positions (R is the major radius); also, apart from a factor of two, the normalization is the same as in Eq. (2.1). Interpolation of the (R, z) positions onto the (x, y, z) grid completes the operation.

Use of results from a Monte Carlo drift-orbit following code such as ORBIT-RF is also straightforward. When running ORBIT-RF, the phase-space coordinates and weights of the orbits are sampled frequently near the time of interest. After reading an output file with this information, FIDASIM selects a phase-space coordinate, then searches for particles that are within the phase-space volume of this coordinate; the desired F is simply the sum of the weights of particles that fall within each phase-space bin.

CQL3D is a Fokker-Planck code. Theoretically, the distribution function f in an axisymmetric tokamak can be described by just three coordinates. CQL3D selects the major radius R_{mid} and pitch angle χ_{mid} at the outer midplane crossing and the speed v as coordinates for f . Although three coordinates suffice for theory, the FIDA and NPA diagnostics require the distribution function throughout the (x, y, z) volume. A particular location is connected to its midplane crossing through the flux surface geometry. In the first comparisons of CQL3D predictions with FIDA data [5, 13], zero banana-width orbits were assumed but, in a recent version of CQL3D, the midplane crossing is shifted by a first order correction for the finite banana width. With or without the finite banana width correction, because of conservation of the first adiabatic invariant μ , the pitch angle changes as the fast ion orbits. The relationship between the pitch angle χ at an arbitrary position

and the pitch angle at the outer midplane is

$$\sin^2 \chi = \left(\frac{B}{B_{mid}} \right) \sin^2 \chi_{mid}, \quad (2.2)$$

where B/B_{mid} is the ratio of the magnetic field at the new position to the field at the outer midplane. Unlike FIDASIM, which includes the Jacobian of the velocity-space coordinate transformation in the definition of F , f in CQL3D does not include the Jacobian, so $\iint f(v, \chi_{mid}) v^2 \sin(\chi_{mid}) dv d\chi_{mid}$ gives the fast-ion density at the midplane. With the further transformation from (v, χ) coordinates into (E, p) coordinates, the desired distribution function $F(E, p)$ that has Jacobian factors included is $F \propto v f$. The values of F in the desired variables are found by selecting a phase space coordinate, calculating the corresponding values of v , R_{mid} and χ_{mid} , and then interpolating over the midplane distribution function.

2.2 Injected and halo neutral densities

The second major section of FIDASIM is devoted to calculation of the injected and halo neutral distributions in real space, velocity space, and energy levels. Since the injected neutrals penetrate the torus in microseconds and the halo forms on the rapid ion-ion collision time ($\sim 10 \mu s$), these neutral populations are time-independent in the code. Note that, although the majority of neutrals occupy the ground state, it is essential to compute the fractional occupation of higher energy levels, as this influences beam attenuation, the probability that charge-exchange events populate higher energy levels and, ultimately, the fraction of neutrals that undergo the detected Balmer alpha radiative transition. In treating the atomic physics of the neutrals, three important simplifications are permissible. First, in some portions of the code, the probability of a reaction $1 - \exp(-x/\lambda)$ is approximated by x/λ if the relevant mean-free path λ is longer than the evaluated step length x . Second, although there are many possible principle quantum numbers n and angular momentum states l available to the neutrals, the strong fine-structure mixing allows the assumption that the population of each quantum state may be grouped as a single population based on the principle quantum number [14]. (Even if this assumption is not strictly valid, the effect is primarily on the polarization and Stark shift of emitted radiation, not the overall intensity. The effect on neutralization probabilities should also be small.) With this assumption, the equations that govern the populations N_j of neutrals in the different energy levels j are written in matrix form as

$$\frac{dN_j}{dt} = \sum_k N_k M_{kj}, \quad (2.3)$$

where M_{kj} is a matrix of rates governing the transitions from levels k to j . Throughout the code, the N_j represent fractional occupational densities, so $\sum_j N_j \leq 1$. (When tracking a neutral, the sum may decrease below unity due to ionization and charge exchange losses.)

Radiative and collisional processes are both included in M . The radiative transition rates are independent of plasma parameters, so these Einstein coefficients [15] are loaded into an array in the initialization stage of the code. In contrast, the collisional rates depend on plasma parameters.

The third major simplification in the treatment of the collisional-radiative transitions is to assume that the speeds of the various species follow the ordering

$$v_e \gg v_f \sim v_n \sim v_i \gg v_I,$$

where the subscripts represent electrons, fast ions, hydrogenic neutrals (both fast and thermal), thermal hydrogenic ions, and impurity ions, respectively. Since the electron distribution function is Maxwellian and the electron thermal speed is much greater than the fastest neutrals, it is expedient to work directly with the reactivities $\langle \sigma v \rangle$ for electron collisions with neutrals. During initialization, the program computes and stores the electron reactivities in each cell. Hydrogenic rates are evaluated using the relative velocity between the ion and the neutral, $|\mathbf{v}_i - \mathbf{v}_n|$, where $\mathbf{v}_i$ and $\mathbf{v}_n$ are the ion and neutral velocities, respectively. For the case of collisions of the fast-ion population with a neutral, $\mathbf{v}_i$ represents the fast-ion velocity. For collisions between the thermal-ion population and the neutrals, it is necessary to average the reactivity over the ion distribution function, which is assumed to be a drifted Maxwellian with temperature T_i and toroidal rotation velocity $\mathbf{v}_{rot}$. With the assumption that the impurity speed is negligible compared to the speed of a hydrogenic neutral, collisions with impurities only depend on the neutral speed v_n . The current version of the code treats fully-stripped carbon as the only impurity species.

Combining the three species, a typical matrix element M_{12} (in this case, the matrix element for excitation from the ground state to the $n=2$ state) is

$$M_{12} = n_e \langle \sigma v \rangle_{12}^{coll,e} + n_d \langle \sigma v \rangle_{12}^{coll,d} + n_C \sigma_{12}^{coll,C} v_n,$$

where n_e , n_d , and n_C are the electron, deuterium, and carbon densities. The deuterium density is not directly measured so it is inferred using quasineutrality: $n_d = n_e - 6n_C$. There is a subtlety associated with the deuterium density, however. In the core of low density plasmas with large fast-ion populations, the fast-ion density n_f can be comparable to the thermal deuterium density. In principle, because the two populations have quite different velocity distributions, the interaction of neutrals with these two populations should be calculated separately. Unfortunately, a correct treatment of neutral collisions with fast ions is quite complicated, requiring calculations similar to those in the final and most time-consuming portion of the code. As an alternative, in the deposition of the injected neutrals, the code ignores the distinction between fast and thermal deuterons, effectively approximating the fast-ion stopping cross section by the thermal-ion cross section.

For all species, deexcitation rates are derived from the principle of detailed balance, i.e.,

$$\langle \sigma v \rangle_{u \rightarrow l} = \left(\frac{n_l^2}{n_u^2} \right) \langle \sigma v \rangle_{l \rightarrow u},$$

where u and l represent the upper and lower quantum numbers, respectively.

With the matrix elements M_{kj} defined, the code proceeds to the calculation of the neutral populations associated with the injected beam or beams. This calculation utilizes three subroutines that are employed throughout the remaining portions of the code. One basic subroutine (called TRACK) calculates the trajectory of a neutral through the Cartesian grid, returning the length of the track in each "cell". A second subroutine (called COLRAD) solves the time-dependent collisional-radiative equations (Eq. (2.3)) for the neutral densities in each state N_j ; the number of Balmer-alpha radiative transitions is also computed. A third subroutine (called SPECTRUM) calculates the Stark and Doppler shifts of emitted photons given the local electric and magnetic fields, the velocity of the neutral, and the direction of the photon. (The detector is assumed to measure all emitted polarizations.) Zeeman splitting is negligible. The Doppler-shifted wavelength λ is

$$\lambda = \lambda_0 \sqrt{1 - \left(\frac{v}{c}\right)^2} \left(1 + \frac{v_{\parallel}}{c}\right)^{-1}, \quad (2.4)$$

where λ_0 is the rest wavelength, v is the speed of the neutral, c is the speed of light, and $v_{\parallel}$ is the component of the neutral velocity in the direction of the emitted photon. The electric field in the neutral frame, $\mathbf{E} = \mathbf{E}_{lab} + \mathbf{v} \times \mathbf{B}$, splits the Doppler shifted line through the Stark effect. For a statistical l population, the relative intensities and shifts of the nine spectral lines are given in [16].

Using the known beam geometry, a Monte Carlo procedure launches rays from random positions on the source with a Gaussian distribution of velocities derived from the specified beam divergence. Rays that clear the beam aperture are followed into the plasma. The full, half, and third energy components are assumed to follow the same trajectories but these three populations are treated separately as they progress into the plasma. All neutrals are emitted in the ground state. Through solution of the collisional-radiative equations, the densities and velocities of the full, half, and third neutrals (each as a function of energy level n) are added to the structure that describes each cell. COLRAD also computes the probability of Balmer-alpha emission for each neutral ray and SPECTRUM calculates the spectra produced by each beam component in each cell; this beam-emission light from each cell is accumulated and stored.

As the injected beam attenuates, the code keeps track of charge-exchange events with the deuterium population. These events are the source of halo neutrals. The fast neutrals that produce FIDA light and NPA signals are not counted as halo neutrals, so the charge-exchange source rate is multiplied by $n_d / (n_f + n_d)$ to obtain the halo source rate. (Consequently, the code neglects fast ion collisions with second-generation fast neutrals.) From the known source, the code computes the cloud of halo neutrals that surround each injected beam. The initial velocity is randomly selected based on the local ion temperature and rotation. Because the injected neutrals have large velocities and the charge-exchange cross section is a strong function of relative energy, the true distribution of initial velocities is skewed relative to the thermal-ion distribution, but this effect is neglected in the current version of the code. (According to [14], the deviation associated with this approx-

imation is small.) This neutral is then followed through the cells. If a charge exchange event happens in a cell, the particle is restarted with a new random velocity based on the local ion temperature and rotation. The neutral is followed until it ionizes. To obtain exponential probability distributions [17] for the charge exchange and ionization events, the code computes the ratio of the track length in the cell to the mean-free path, $r = n_d \sigma_{CX} \Delta l$. If $r > \ln(1/\eta)$, where η is a uniform randomly generated number on the interval $[0,1]$, a charge-exchange event occurs in the cell. A similar comparison is performed for ionization. For simplicity, this spatial diffusion is computed assuming that all of the neutrals are in the ground state. (For typical parameters, this is true for 99% of the halo neutrals, so this is an excellent approximation.) The relatively slow timescale of halo formation implies that the energy occupation levels can be approximated by the steady-state collisional-radiative balance, so this is computed after completion of the spatial diffusion calculation. The halo densities (as a function of n) for each cell are then stored.

The code also predicts the spectrum produced by the halo neutrals. Thermal ions that charge exchange with an injected neutral are initially far from collisional-radiative equilibrium but subsequent generations of halo neutrals relax toward equilibrium occupancy levels, so "first-generation" thermal neutrals are treated separately from daughter halo neutrals. The "first-generation" or "direct" charge-exchange light is calculated in the same manner as the FIDA light and is described below. For the daughter halo neutrals, according to Mandl [18], Stark splitting and distortions of the spectrum associated with the energy dependence of the cross section are minor effects, so these effects are ignored. The distribution of light is simply a shifted Maxwellian,

$$g(v_{\parallel}) dv_{\parallel} \propto \exp \left[-\frac{(v_{\parallel} - v_{\parallel,rot})^2}{v_{ti}^2} \right], \quad (2.5)$$

where $v_{\parallel}$ and $v_{\parallel,rot}$ are the components of the emitted light and plasma rotation in the direction of the lens, respectively, and v_{ti} is the ion thermal velocity. The associated Doppler shift relative to the rest wavelength is $\Delta\lambda/\lambda_0 \simeq v_{\parallel}/c$. The rate of emission from a given cell is $n_{Halo*} N_3 A_{32}$, where n_{Halo*} is the halo density after subtraction of first-generation halos, N_3 is the fraction of the halo population in the $n=3$ state, and A_{32} is the Einstein coefficient for the Balmer-alpha transition.

2.3 NPA flux

The NPA flux is found in the third stage of the code. Owing to the large mass difference between electrons and ions and to the small energy exchange in charge-exchange reactions, the angular deflection associated with charge-exchange reactions is $< 1^\circ$ and is ignored throughout the code. Consequently, the pitch p of the escaping neutral is determined entirely by the NPA sightline geometry. The flux of neutrals with energy E_i incident upon a collimated NPA of area A and solid angle $\Delta\Omega$ is

$$\Phi(E_i) = \int F(E_i, p; \mathbf{r}) n_n(\mathbf{r}) \sigma_{cx}(E_{rel}^{i,n}) V_{rel}^{i,n} \frac{\Delta\Omega}{4\pi} e^{-\lambda} A dl (s^{-1}), \quad (2.6)$$

where F is the fast-ion distribution function, n_n is the neutral density, σ_{cx} is the charge-exchange cross section, $V_{rel}^{i,n}$ is the relative velocity between the fast-ion and the neutral, $e^{-\lambda}$ represents reionization of the neutrals, and dl is the differential length along the sightline. Evaluation of Eq. (2.6) is straightforward but there are subtleties. The fast-ion velocity vector is determined by E_i and the detector geometry, so the relative velocity is readily computed, and hence $n_n \sigma_{cx} V_{rel}$ for each of the neutral species. Since the beam divergence is small, the relative velocity for the full, half, and third components is approximated using the nominal velocity vector for the injected neutrals. For the halo neutrals, the reactivity is averaged over a drifted Maxwellian.

The distribution function that is stored in the first stage of the code is the distribution of guiding centers. The detector geometry determines the diagnostic viewing cone but the detected particles have guiding centers that are a gyroradius from the viewing cone. With the velocity $\vec{v}$ known from the viewing geometry, the gyroradius $\vec{\rho}$ is

$$\vec{\rho} = \frac{(\hat{b} \times \vec{v})}{\Omega_{ci}}, \quad (2.7)$$

where $\hat{b}$ is the unit vector in the direction of the magnetic field and Ω_{ci} is the cyclotron frequency. The distribution function used in Eq. (2.6) is evaluated at the guiding center position.

The attenuation of neutrals is computed in the last part of the NPA calculation. The code finds the cell in the detector sightline that is farthest from the detector, then computes the attenuation of neutrals from that point all along the neutral path for several energies. (This calculation is essentially the same as the deposition calculation for the injected neutrals.) With the terms on the right-hand side of Eq. (2.6) known, the predicted spectrum for each detector sightline is found by summing over the viewed cells.

2.4 FIDA radiance

The fourth stage of the code uses a weighted Monte Carlo routine to calculate the FIDA radiance. The fast-ion density n_f and the sum of injected neutral and halo neutral densities $\sum n_n$ have already been calculated as a function of position. The product $n_f \sum n_n$ provides a convenient estimate of the probability of a charge exchange reaction (that neglects the computationally intensive dependence of the reaction rate on the relative velocity), so this product is used to determine how many fast neutrals to launch from each cell. The initial position of the fast neutral within the cell is selected randomly. The initial velocity is found using a Monte Carlo rejection test in the two dimensions that describe the velocity distribution (energy and pitch). The gyroangle is randomly generated, the initial position of the fast ion is shifted by a gyroradius (Eq. (2.7)), and the velocity vector is transformed into (x, y, z) Cartesian coordinates. With the velocity now specified, the actual reaction rate of the fast ion with each of the neutral populations can be computed; the sum of these rates is the weight of this particular fast neutral. (In fact, each individual fast neutral represents a "bundle" of the entire set of neutral n states along the selected

path.) Next, the trajectory of the fast neutral through the cells is computed by TRACK. As the fast neutral travels through each cell, the time-dependent collisional-radiative transitions between states (Eq. (2.3)) is computed by COLRAD, including the number of D_α photons that are emitted. With the velocity of the neutral known, the spectrum of the emitted photons in each cell is computed by SPECTRUM. Finally, the properly weighted spectrum is added to the accumulated spectra in each cell.

The calculation of the direct charge-exchange light from first-generation thermal neutrals is patterned after the FIDA calculation with two differences. Since the halo-light calculation already treats light associated with multiple charge-exchange events, only the injected neutral density is used in the calculation of the charge-exchange probability. The fast-ion distribution function is replaced by a drifted Maxwellian with density n_d .

The output of the FIDA, beam-emission, direct charge-exchange, and halo light are stored in two forms. At the most basic level, the output is the emissivity ϵ_i in the direction of the lens from each cell in units of photons/cm³/s/nm. The code also performs a simple integration over these emissivities ϵ_i along the specified sightlines. For this, TRACK computes the length of the sightline in each cell dl_i and the code sums $\sum_i \epsilon_i dl_i$ to obtain the radiance for each view chord in units of photons/cm²/s/nm.

More sophisticated manipulation of the output of the code is straightforward. The computed radiance treats each view chord as an infinitely-narrow pencil beam. In reality, actual chords have finite transverse extent at the focal plane and even broader spatial extent away from the focal plane. One can easily implement more accurate integration over the stored emissivities to model the actual optics. To model data from an imaging camera acquired with a narrowband filter [19], a post-processor uses TRACK to integrate over the thousands of camera sightlines and sums the spectra over the passband of the filter. Optionally, this post-processor performs similar calculations for the halo and injected-neutral light.

Note that the computed spectra neglect instrumental broadening. For comparison with experiment, the theoretical spectra are convolved with the instrument function.

3 Verification

3.1 Atomic physics

Atomic physics cross sections are important in two places in the code: in the calculation of neutralization probability and in the solution of Eq. (2.3) in COLRAD. The required cross sections and reactivities are available in the literature and in the Atomic Data and Analysis Structure (ADAS) compilation [20,21]. For COLRAD, Eqs. 9 and 10 of [22] give the cross section for proton excitation and impact ionization from the ground state, while [23] contains cross sections for excitation from higher states. Expressions for electron impact ionization as a function of electron temperature T_e and energy level n appear in [24]. Formulas for electron excitation from one energy level to another are in [25]. Impurity cross sections are listed in Eqs. 13-16 of [22]. An alternative compilation of many of the rates

Figure 3: Solution of Eq. (2.3) immediately following neutralization of the fast ion for a typical case (density of $4 \times 10^{13} \text{ cm}^{-3}$, temperature of 4keV). The total time represented by the abscissa is 11.1ns; the fast neutral travels 3cm during this time. The evolution of the $n=3$ state determines the intensity of FIDA light.

appears in [26] but the effect of these differences is small compared to other uncertainties so the current version of the code uses the older compilation by Janev [24]. In COLRAD, rates for energy levels up to $n=7$ are normally employed.

For the initial neutralization probability, cross sections for the charge-exchange reactions between fast ions and neutrals in states $n=1-4$ are given in ADAS [20]. Charge-exchange reactions to states with $n > 4$ are neglected in our calculations because these energy levels are sparsely populated and the cross sections seem uncertain.

The calculations in COLRAD are the most time consuming portion of the code. Fig. 3 shows a typical solution following neutralization of a fast ion. Although only a small fraction of the injected neutrals occupy higher energy levels, the cross section for charge exchange between excited states is several order of magnitudes larger than the cross section for charge exchange from the ground state to an excited state. Consequently, the initial conditions for the collisional-radiative equations (Eq. (2.3)) are far from equilibrium. For example, for the case illustrated in Fig. 3, the steady-state occupation levels for the injected beam are all $\ll 1\%$ for $n > 1$ but, when multiplied by the cross section, the initial occupation fractions for the N_2 , N_3 , and N_4 states are ~ 10 , 3, and 1%, respectively. Rapid decay of the excited states occurs in the first few cells. As a result, it is necessary to solve the full system of equations; a reduced set of equations, such as those proposed in Eqs. 2-7 of [15], are inaccurate when the initial $n=3$ occupation fraction N_3 is higher than the final occupation fraction. A fourth-order Runge-Kutta routine is adopted in the code. As suggested by Hutchinson [15], the code selects the time step based on the smaller of the matrix element timescales,

$$t_{scale} = \min[(M_{11} - M_{22})^{-1}, (M_{11} - M_{33})^{-1}]. \quad (3.1)$$

Usually, an accuracy of $\lesssim 1\%$ in N_3 is obtained with a time step of $t_{scale}/4$ but, if the solution is unphysical (e.g., negative densities or $\sum_i N_i > 1$), the solution is recomputed with half the time step. On the basis of extensive tests, the code ordinarily uses seven energy

levels. Surprisingly, on average, the use of 7 energy levels only takes 7% longer than calculations with 4 energy levels but the accuracy is improved appreciably. (Fewer energy levels underestimates N_3 by $\sim 10\%$.) As Hutchinson notes [15], accuracy is improved by adjusting the coefficient of the highest retained diagonal matrix element so that electrons in the uppermost modeled state are immediately lost. Alternatively, additional energy levels can be retained so that the uppermost states are very sparsely populated; this is the approach adopted in FIDASIM. The number of D_α transitions is proportional to the integral of N_3 over the time in the cell; this is quickly computed with adequate accuracy using an extended trapezoidal rule summation.

COLRAD is called each time a neutral enters a new cell. The collisional-radiative matrix is filled at the beginning of each call. Collision rates with hydrogenic and impurity species must be reevaluated based on the current neutral velocity. Rather than recomputing these rates at each call (particularly the time-consuming integration over the Maxwellian thermal distribution), the code utilizes pre-computed look-up tables of the rates versus neutral energy and ion temperature. (In regions where σv changes rapidly, the look-up table presently uses steps of 0.25 and 1.0keV in temperature and fast-ion energy, respectively.)

The accuracy of COLRAD was verified in several steps. The matrix elements M agree to within a few percent with the values in Table 1 of [15]. (As mentioned earlier, an exception is the diagonal elements at high n .) The code correctly reproduces the steady-state fractions in Fig. 1 of [15].

3.2 Beam deposition

The calculation of the injected neutral density was compared with TRANSP for an NSTX case. The neutral profiles along the y axis (perpendicular to the neutral beam source) and along the z axis (vertical to the source) axis agree well with the TRANSP simulation results (Fig. 4) and with neutral beam calibration data [27]. The attenuation of the beam is also in reasonable agreement with TRANSP simulations. A summary of one of these comparisons is shown in Fig. 5. Fig. 5(a) shows the neutral density along an NPA sightline from TRANSP and FIDASIM simulations. When halo neutrals are not considered, both codes give very similar neutral density. The figure also shows that the halo neutral density can be comparable to the injected neutral density so that the total neutral density almost doubles when halo neutrals are included in FIDASIM. (The current version of TRANSP redistributes halo neutrals over the entire plasma volume. This approach suffices for power balance calculations but is inaccurate for NPA and FIDA simulations.) Fig. 5(b) shows the attenuation factors for 60keV neutrals along an NPA sightline for the TRANSP and FIDASIM simulations. The small discrepancy between the two curves is probably caused by the use of different cross sections or possibly by the inclusion of multi-step ionization in FIDASIM. Calculations of the attenuation of injected neutrals computed by TRANSP, FIDASIM, and a pencil beam code [28] show similar discrepancies; in general, FIDASIM predicts slightly faster attenuation than the other two codes,

Figure 4: Comparison of injected neutral density profile along the y and z axes from the FIDASIM code (symbols) and from TRANSP (lines) at the beam tangency point ($R=70\text{cm}$) for NSTX Source A.

Figure 5: Comparison between the FIDASIM and TRANSP simulations in NSTX discharge #122631 at $t=0.1\text{s}$. (a) Neutral density, (b) attenuation factor for 60 keV neutrals and (c) differential charge-exchange flux along the NPA sightline.

which may be associated with a more accurate treatment of multistep ionization.

Comparisons of the NPA calculation with TRANSP are shown in Fig. 5(c) and Fig. 6. Fig. 5(c) shows the differential contribution to the 60keV charge-exchange efflux along an NPA sightline for an NSTX case. With halo neutrals neglected, the agreement is satisfactory. Fig. 6 compares spectra after integration over the sightline. At low density,

Figure 6: NPA energy spectra for a sightline with $R_{tan} = 70$ cm from the TRANSP and FIDASIM simulations with and without halos at three different times in NSTX discharge #122631. The density increased from approximately $2 \times 10^{13} \text{ cm}^{-3}$ to $5 \times 10^{13} \text{ cm}^{-3}$ between 0.1 and 0.5 s.

the agreement is satisfactory when halo neutrals are neglected in FIDASIM. At higher densities, the spectral shapes remain similar but FIDASIM predicts smaller signals than TRANSP; this is a consequence of the larger attenuation predicted by FIDASIM.

3.3 Halo neutrals

The halo neutral simulation portion of the code was verified by comparing with a one-dimensional diffusion model. If we assume: (1) a uniform plasma with a circular cross-section neutral-beam injection pattern and (2) no beam attenuation along the neutral-beam centerline, the halo neutrals will diffuse only in the radial direction and the density can be determined from the following simple 1-D diffusion model [29],

$$D \frac{\partial}{\partial r} \left(r \frac{\partial n_h}{\partial r} \right) = n_h n_e \langle \sigma v \rangle_{ei} - \sum_{k=1}^3 n_i n_{b,k}(r) \langle \sigma v \rangle_{cx,k}, \quad (3.2)$$

where n_i , n_e , $n_{b,k}$, and n_h are the ion, electron, k th component of injected neutrals, and halo densities, respectively, $\langle \sigma v \rangle_{ei}$ and $\langle \sigma v \rangle_{cx,k}$ are the electron-impact and charge-exchange reactivities, $D = T_i / m_d \gamma_{cx}$ is the diffusion coefficient and γ_{cx} is estimated by the most probable $n_i \langle \sigma v \rangle_{cx}$ based on the ion temperature. The leftmost term in Eq. (3.2) represents

the halo neutral diffusion. The term $n_h n_e \langle \sigma v \rangle_{ei}$ is the halo decay term due to collisions with electrons and the rightmost term is the halo formation term due to charge exchange reactions between the primary beam neutrals and thermal ions. The halo neutral losses through impurity impact ionization and impurity charge exchange are not included in this simple diffusion model, but they are included in FIDASIM and they have a minor effect (less than 5%) on halo neutral density. Eq. (3.2) is an inhomogeneous modified Bessel equation. It is not easily solved analytically for a spatially varying beam neutral profile but a solution exists for a constant source. To test the halo calculation, the radial direction is divided into many regions in the FIDASIM program and plasma parameters such as the densities are constant in each region. Then the halo neutral density solutions for these regions can be generally expressed as

$$n_h(r)_I = C_1 I_0(\lambda r) + \frac{\gamma_I}{\beta}, \quad (3.3a)$$

$$n_h(r)_{II} = C_2 I_0(\lambda r) + E_2 K_0(\lambda r) + \frac{\gamma_{II}}{\beta}, \dots, \quad (3.3b)$$

$$n_h(r)_{n-1} = C_{n-1} I_0(\lambda r) + E_{n-1} K_0(\lambda r) + \frac{\gamma_{n-1}}{\beta}, \quad (3.3c)$$

$$n_h(r)_n = E_n K_0(\lambda r), \quad (3.3d)$$

where I_0 and K_0 are modified Bessel functions, $\lambda = \sqrt{n_e \langle \sigma v \rangle_{ei} / D}$, γ_j is the source term for each region, and $\beta = n_e \langle \sigma v \rangle_{ei}$ is the loss term. The constant coefficients C_j and E_j are algebraically determined by matching the densities from different regions using continuity of n_H and its first derivative dn_H/dr . The halo neutral density in the whole region can be obtained by combining all of these analytical solutions. Fig. 7 shows the halo neu-

Figure 7: Comparison of halo neutral densities calculated from the 1-D diffusion model (Eq. (3.3)) and the FIDASIM halo diffusion subroutine for plasma profiles with $n_i = n_e = 1.1 \times 10^{14} \text{ cm}^{-3}$, $T_e = 2.5 \text{ keV}$ and $T_i = 1.25 \text{ keV}$. The three dashed steps represent the densities of the full, half, and third energy components of injected neutrals used in both codes.

tral densities calculated from the diffusion model and the Monte-Carlo halo simulation subroutine for identical plasma profiles. The agreement is good.

3.4 FIDA spectra

The SPECTRUM subroutine used to compute the D_α spectra and the weighted Monte Carlo scheme were verified as follows. As part of the initial investigation of the feasibility of FIDA, a simplified model of the expected spectra was developed that ignores atomic physics and assumes that the magnetic field is purely toroidal; Fig. 2 of [3] shows a result calculated by this code. To test the main FIDA simulation loop, we replaced the magnetic field with a toroidal field and modified the cross sections to be independent of velocity. The resulting spectra were consistent with the output of the simple model.

4 Numerics

Numerical input parameters affect both the accuracy and computational expense of a calculation. Uncertainties in plasma parameters (particularly electron density) introduce uncertainties in the predicted radiance or efflux of 20% or more (Appendix A of [30]), so extremely fine grids are wasteful and unnecessary. Moreover, uncertainties in atomic cross sections also introduce considerable uncertainties in the predictions. This section shows examples of the sensitivity of the predicted radiance to numerical input parameters and provides recommendations for these choices in terms of the relevant physical processes. As in [30], we again study the case of an MHD-quiescent DIII-D plasma (shot 122060 at $t = 2.05\text{s}$) with good agreement between theory and experiment for the FIDA spectrum. The neutral-beam injection energy is 80keV in this plasma. The previous study explored the dependence of the predictions on uncertainties in plasma parameters. Here, we explore the dependence on numerical input parameters. Not surprisingly, the total number of grid cells and the number of beam Monte Carlo (MC) particles affect the computational time the most.

The original version of the code was written in the Interactive Data Language (IDL). Recently, the fourth portion of the code, the Monte Carlo calculation of the FIDA spectra, was converted into Fortran 90. The Fortran version of the code is an order of magnitude faster than the IDL version. For the IDL version, on a fairly modern Dual-Core AMD Opteron Processor running at 2.8GHz, the wall clock time spent per neutral is $\sim 10\text{ms}$. For a typical FIDA simulation with 10^7 reneutrals this translates to about 28 hours. About 35% of that time is spent solving collisional-radiative equations in COLRAD; calculation of the neutralization probability is also time-consuming. For a standard DIII-D FIDA grid size choice of $\{nx, ny, nz\} = \{31, 21, 21\}$, the preparatory steps performed in the first two stages of the code takes about 4 additional hours. The halo neutral density calculation and the mapping of the beam distribution function to the (x, y, z) grid take ~ 1 hour each.

4.1 Number of Monte Carlo reneutrals launched

There is a linear dependence between the number of launched neutralized fast ions (or "reneutrals") and the computational time in the fourth section of the code. Fig. 8(a) compares the spectra at the location of peak D_α emissivity ($R = 187.5\text{cm}$) for several simulations with varying number of launched reneutrals. To eliminate the influence of the random number seed choice, the same seed value was used in all five simulations. Other simulations with random seed and 10^7 reneutrals have shown that the MC noise is $< 5\%$, for spectra with $E_\lambda < 30\text{keV}$. (Each wavelength is associated with an equivalent energy along the line-of-sight known as E_λ ; since the Doppler shift only measures one component of the velocity, E_λ is the minimum energy of the reneutrals that produce a particular Doppler shift [30].)

Numerical MC noise is proportional to $1/\sqrt{N}$, where N is the number of particles in the simulation. The computationally most expensive simulation in our study used 9×10^7

Figure 8: (a) Comparison of FIDA spectra from simulations with 10^6 to 9×10^7 Monte Carlo particles. (The reduction in FIDA light for small Doppler shifts is an artifact caused by truncating the fast-ion distribution function at $E_{\min} = 10\text{keV}$.) (b) FIDA spectra normalized to the spectrum computed with 90 million particles. The dashed vertical lines relate Doppler shifts to equivalent energies along the line-of-sight E_λ .

Figure 9: Comparison of FIDA profiles for simulations with 10^6 to 9×10^7 Monte Carlo particles. The spectra are integrated over Doppler shifts larger than (a) $E_\lambda \geq 10\text{keV}$ and (b) $E_\lambda \geq 30\text{keV}$.

reneutrals. Considering the resulting spectra as the most accurate, it is instructive to compare the ratios of the spectra from simulations with less particles to the spectra from this particular simulation (Fig. 8(b)). As expected, the simulation with 4×10^7 reneutrals is essentially indistinguishable from that with 9×10^7 reneutrals. The agreement gets progressively worse as the number of MC particles in the simulation decreases, especially for wavelengths with high red and blue shifts, because only a small fraction of the beam distribution function has energies close to 80keV; moreover, the likelihood of an energetic $n=3$ neutral moving directly along the FIDA line-of-sight is quite small.

FIDA spectra from several radial positions are often integrated over a specific wavelength range to obtain spatial profiles. Fig. 9 shows the calculated profiles for different numbers of MC particles. If the spectral integration is from $E_\lambda = 10\text{keV}$ and above, good spatial profiles are obtained with just 10^6 reneutrals but more particles are needed if the lower limit of spectral integration is $E_\lambda = 30\text{keV}$. Therefore, two factors determine the choice of the number of reneutrals that the code needs to launch: the integrating spectral band of the diagnostics, and the energy range where the beam interaction of interest takes place. For a FIDA diagnostic that starts integrating close to the thermal ion energies, simulations with 10^6 particles are sufficient. However, if the integration starts at half the beam injection energy, ~ 3 times more particles are needed. Resolving fine spectral details close to the injection energy requires an additional order of magnitude more particles.

4.2 Simulation volume

The computational grid must enclose most of the interacting beam ions and neutral particles. Fig. 10 illustrates the geometry of the beam line that injects 2.5MW of 80keV deuterium neutrals into the DIII-D plasma under consideration. Ten vertical FIDA sightlines collect radiation along the beam x -axis. The fast-ion density calculated by TRANSP and the halo and beam neutral densities calculated by FIDASIM are shown. The halo and beam neutral densities are comparable in magnitude, while the fast-ion density is three orders of magnitude higher because in tokamaks the fast ions are very well confined. The dashed rectangle with x - y dimensions of 120cm by 60cm represents the standard horizontal grid size in our simulations. The neutrals contour plots shown in Fig. 10 are from simulations using x - y grid size of $160 \times 120\text{cm}$. The cloud of halo neutrals is "blown" in the direction of plasma rotation. Clearly, the $120 \times 60\text{cm}$ horizontal grid size truncates some halo neutrals but the effect on the FIDA profiles is insignificant. Further reduction of the grid in the x -direction to 90cm truncates a sizable fraction of the halo neutrals but the effect on the calculated spectra is still small because D_α light due to charge exchange with these neutrals rarely reaches the vertically positioned FIDA detectors.

For vertically-viewing FIDA chords, the spatial extent in the vertical direction is most important. The neutral beams have low divergence and, apart from steady attenuation along the x -axis, the injected neutral profile in the y - z plane does not change appreciably. (For example, the injected neutral profile at $x = 70\text{cm}$ is very similar to the profile at

Figure 10: (a) Midplane ($z=0$) cross section of a quadrant of the DIII-D tokamak, showing the intersection of the vertical FIDA sightlines (triangles) and the (x,y) coordinate system associated with the active neutral beam. (b) Midplane halo neutral density. The distribution is asymmetrical due to the plasma flow. (c) Beam-ion density. (d) Injected neutral density.

$x=2$ cm that is shown in Fig. 11(a), apart from a reduction by a factor of four.) However, the situation with the halo neutrals is dramatically different. When the injected neutrals enter the plasma, the beam halo is limited to the region of high injected neutral density (Fig. 11(b)). Further down the x -axis, the plasma density increases and, as more injected neutrals charge exchange with thermal deuterium ions, the beam halo profiles spread in both transverse directions (Fig. 11(c)).

Figure 11: Vertical (y,z) planes of (a) injected neutral density near the plasma edge, (b) halo neutral density near the plasma edge, and (c) halo neutral density in the plasma interior.

Blue-shifted FIDA profiles from a set of four simulations with 10^7 MC particles, where a single vertical dimension was varied, are shown in Fig. 12. Our baseline simulation uses values of the grid half-width of 30 and 40cm in the y - and z -directions, respectively. This standard vertical grid is illustrated in Fig. 11(c). Doubling the size of the grid in the y -direction and thus including essentially all beam halo neutrals in that direction does not change the FIDA profile. When the y -width is reduced to 20cm, the profile uniformly decreases by $\sim 5\%$ (not shown), which is within the MC noise level. In contrast, extending the grid in the z -direction results in a $\sim 10\%$ increase in the predicted profile. If the vertical grid dimension is reduced to a z -width of 20cm, the predicted values decrease by $\sim 25\%$.

To understand the impact of these choices on the predicted profile, consider an extreme case with *no* halo-neutral contribution. (The FIDASIM code has an option to turn off the halo-neutral calculation.) As shown in Fig. 13, turning off the halo produces profiles and spectra that are $\sim 50\%$ lower than the baseline values. Neglect of the halo not only affects the magnitude of the predicted signal but also affects the shape of the spatial profile and of the spectrum. The profile variations are principally due to variations in the ratio of injected: halo neutral densities along different sightlines. The spectral variations are ultimately associated with the strong energy dependence of the charge-exchange cross sections. Halo neutrals have lower velocities than injected neutrals, so the relative velocity of fast ions with different velocities (and, hence, Doppler shifts) is different for collisions with halo neutrals than for injected neutrals. Although the "no-halo" case shown in Fig. 13 is unrealistic, it is useful as an upper-bound on the magnitude of effects associated with truncation of the box-size. The variations in signal shown in Fig. 11 depend on the fraction of the halo population that is included in the simulation. These results indicate that the viewing geometry of the FIDA diagnostic ultimately determines the extent to which the halo neutral's volume must be included in the simulation. For FIDA chords that view horizontally, expanding the grid along the y -axis is more important than expanding it along the z -axis.

Ultimately, the required simulation volume depends on the spatial extent of the halo, which scales as the geometric mean of the mean-free paths for ionization and charge-exchange evaluated at the ion temperature [31]. These mean-free paths are both inversely proportional to the density, with only weak dependencies on T_e and T_i (for typical temperatures), so the required simulation volume scales approximately as $1/n_e$.

4.3 Cell size

Once the simulation volume in the (x, y, z) space is defined, the size of each cell needs to be determined. The calculated signals are line integrals over three-dimensional emissivity profiles. Fig. 14 shows a two-dimensional example. If the grid is too coarse, pixelation is evident in the resulting profile. The severity of this effect depends primarily on the gradient of the emissivity profile in directions perpendicular to the sightline; pixelation is also more severe when a sightline is oriented along one of the (x, y, z) axes.

Physically, the emissivity profiles depend on numerous quantities including the beam

Figure 12: Dependence of the calculated FIDA radiance on the volume of the simulation. The spectra are integrated over wavelengths above $E_\lambda = 10\text{keV}$.

Figure 13: (a) Ratio of the blue-shifted FIDA radiance for simulations without halo neutrals and with halo neutrals vs. major radius. The spectra are integrated over wavelengths above $E_\lambda = 10\text{keV}$. (b) Ratio of the FIDA spectral radiance for simulations without halo neutrals and with halo neutrals for the FIDA channel at $R = 187.5\text{cm}$. Two values of E_λ are indicated.

Figure 14: (a) FIDA emissivity in an (x,z) plane, with the standard (33×21) grid boundaries overlaid. The red lines illustrate the range of sightline angles graphed in the lower figure. (b) Calculated radiance vs. sightline angle for different choices of grid size in the (x,z) plane. The flat portions of the curves are caused by pixelation.

injection geometry and the plasma parameters; gradients in any of these quantities impact transverse emissivity gradients. The FIDA and NPA emissivities are approximately proportional to the product of fast-ion and neutral densities, $n_f \sum n_n$; these typically change

over scale lengths of several centimeters. The injected neutral light depends on the beam attenuation and can change appreciably in a few centimeters for densities $\gtrsim 5 \times 10^{13} \text{cm}^{-3}$. Generally speaking, owing to the line integration, the cell size in the approximate direction of the sightlines can be twice as large as in transverse directions.

Finer grids increase the computational expense. Our baseline simulations for the nominally vertical sightlines shown in Fig. 10(a) employ x - y - z grid cells of $4 \times 3 \times 4 \text{cm}$. To study the effect of coarser or finer grids, we use the standard $120 \times 60 \times 80 \text{cm}$ computational domain and 10^7 particles. Refining the grid size in the vertical direction to $4 \times 3 \times 0.8 \text{cm}$ or, in the horizontal direction, to $2.5 \times 1.2 \times 4 \text{cm}$ increases the number of grid cells n_g by a factor of 5 and 3.6, respectively. The corresponding computational times are 63% and 22% higher. The effect of these changes on the calculated spatial profiles integrated above $E_\lambda = 10 \text{keV}$ is negligible. Coarsening of the x - y - z grid size to $6 \times 4 \times 5.5 \text{cm}$ leads to a computational time savings of $\sim 25\%$. For this case, the FIDA spatial profile resembles the baseline but deviates slightly more than expected from the MC noise level. While further coarsening of the grid to $12 \times 10 \times 14 \text{cm}$ leads to a computational savings of 35%, the profile is deformed in shape and has decreased values between 10% and 40% at different radial chord locations.

To avoid pixelation in simulations of two-dimensional imaging data with thousands of sightlines, cell sizes need to be 1-2cm. On the other hand, because the bandpass filters employed in imaging diagnostics average over wavelength, coarse velocity-space grids are permissible.

5 Validation

The calculation of the injected neutrals was compared with experimental measurements of the beam-emission light in a DIII-D experiment [32]. After passing through a bandpass filter, two-dimensional images of the light were measured with a CCD camera. The code predictions are in good agreement with measurements of the vertical extent of the beam and of the beam penetration as a function of density.

In a recent DIII-D experiment [33], halo light was measured for sightlines that are outside the footprint of the injected neutral beam. Good agreement with code predictions is observed.

The first detailed quantitative comparison of the predicted FIDA spectrum with experiment was reported in [34]. In MHD-quiescent DIII-D plasmas, code predictions based on the fast-ion distribution function predicted by NUBEAM have the same spectral shape as experiment and the intensity of the FIDA signal agrees to within 25%. A later DIII-D experiment [35] also found agreement to within about 25% between the spectral shape, radial profile, and absolute intensity for quiet plasmas with ion temperatures below $\sim 3 \text{keV}$.

A two-dimensional measurement of the profile of FIDA light was made with a bandpass filter and imaging CCD camera on DIII-D [19]. The profile shape and dependence

of the signal on beam-injection angle agreed well with code predictions.

FIDASIM predicts the radiance from the injected beam, the halos, and the fast ions. In addition, a post-processor that uses the output of the second stage of the code calculates the expected visible bremsstrahlung radiance. Comparison of the relative intensity of these four features is a useful check that is independent of any experimental errors in the intensity calibration. To date, the most detailed comparison of this type was performed on ASDEX-Upgrade [36] and shows good agreement for all four spectral contributions.

6 Outlook

With plasma profiles and a fast-ion distribution function as input, FIDASIM predicts the flux measured by NPAs and the radiance measured by a D_α spectrometer. The code has a rather complete physics model that treats the atomic physics of these processes with scarcely any approximations. One possible area of improvement is a post-processor that replaces the approximation of infinitesimal sightlines with an accurate treatment of the collection optics. A more challenging upgrade is needed to treat plasmas where the fast-ion density is comparable to the thermal-ion density.

The present code is computationally intensive. Parallelization of the MC routines or further optimization of the most time consuming subroutines is desirable. A complementary reduced model that is sufficiently fast to make predictions between discharges (for example) is needed.

Although many aspects of the code have been successfully validated by experiment, additional comparisons are desirable. The FIDA predictions have been validated in conventional tokamaks but have not yet been confirmed in a spherical tokamak. Validation of the NPA model is also a future task.

Acknowledgments

This work has benefited from the contributions of a large number of scientists. Keith Burrell, Bill Davis, Rainer Fischer, Manuel García-Muñoz, Doug McCune, and Mike Van Zee-land gave valuable advice. Contributors to the code itself include Brian Grierson, Clive Michael, Chris Muscatello, Novi Pablant, Mario Podestá, and Wayne Solomon. We are also indebted to our experimental collaborators on DIII-D and NSTX. The originating developer of ADAS is the JET Joint Undertaking. This work was funded by the U.S. Department of Energy under SC-G903402 and DE-FC02-04ER54698 and DE-FG02-06ER54867.

References

- [1] I. H. Hutchinson, *Principles of Plasma Diagnostics*, Cambridge University Press, New York, 1987.
- [2] L. A. Artsimovich, V. V. Afrostimov and I. P. Gladkovskij et al., in *Plasma Physics and Controlled Nuclear Fusion Research 1965*, Volume 2, pp. 595, Vienna, 1966, IAEA.

- [3] W. W. Heidbrink, K. H. Burrell, Y. Luo, N. A. Pablant and E. Ruskov, Hydrogenic fast-ion diagnostic using Balmer-alpha light, *Plasma. Phys. Control. Fusion.*, 46 (2004), 1855–1875.
- [4] W. W. Heidbrink, Fast-ion D_α measurements of the fast-ion distribution (invited), *Rev. Sci. Instrum.*, 81 (2010), 10D727.
- [5] W. W. Heidbrink, Y. Luo and K. H. Burrell et al., Measurements of fast-ion acceleration at cyclotron harmonics using Balmer-alpha spectroscopy, *Plasma. Phys. Control. Fusion.*, 49 (2007), 1457–1475.
- [6] J. Egedal and H. Bindslev, Reconstruction of gyrotropic phase-space distributions from 1D projections, *Phys. Plasma.*, 11 (2004), 2191–2198.
- [7] M. Salewski, F. Meo and M. Stejner et al., Comparison of fast ion collective Thomson scattering measurements at ASDEX upgrade with numerical simulations, *Nucl. Fusion.*, 50 (2010), 035012.
- [8] A. Pankin, D. Mccune, R. Andre, G. Bateman and A. Kritiz, The tokamak Monte Carlo fast ion module NUBEAM in the national transport code collaboration library, *Comput. Phys. Commun.*, 159 (2004), 157–184.
- [9] R. V. Budny, A standard DT supershot simulation, *Nucl. Fusion.*, 34 (1994), 1247,
- [10] L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellman and W. Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nucl. Fusion.*, 25 (1985), 1611–1622.
- [11] V. S. Chan, S. C. Chiu and Y. A. Omelchenko, Radio-frequency-driven radial current and plasma rotation in a tokamak, *Phys. Plasma.*, 9 (2002), 501–510.
- [12] R. W. Harvey and M. G. Mccoy, The CQL3D Fokker-Planck code, in *Proceedings of IAEA Technical Committee Meeting on Advances in Simulation and Modeling of Thermonuclear Plasmas*, Montreal, 1992, IAEA.
- [13] D. Liu, W. W. Heidbrink and M. Podesta et al., Profiles of fast ions that are accelerated by high harmonic fast waves in the national spherical torus experiment, *Plasma. Phys. Control. Fusion.*, 52 (2010), 025006.
- [14] T. A. Gianakon, R. J. Fonck, J. D. Callen, R. D. Durst and J. S. Kim, Effects of edge plasma turbulence on radial correlation length measurements with BES, *Rev. Sci. Instrum.*, 63 (1992), 4931–4933.
- [15] I. H. Hutchinson, Excited-state populations in neutral beam emission, *Plasma. Phys. Control. Fusion.*, 44 (2002), 71–84.
- [16] G. W. Series, *The Spectrum of Atomic Hydrogen*, Oxford University Press, New York, 1987.
- [17] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, Cambridge.
- [18] W. Mandl, Development of active Balmer-alpha spectroscopy at JET, Technical Report JET-IR(92)05, JET, 1992.
- [19] M. A. Van Zeeland, W. W. Heidbrink and J. Yu, Fast ion D_α imaging in the DIII-D tokamak, *Plasma. Phys. Control. Fusion.*, 51 (2009), 055001.
- [20] <http://adas.phys.ath.ac.uk/>.
- [21] H. Anderson, M. Von Hellermann and R. Hoekstra et al., Neutral beam stopping and emission in fusion plasmas I: deuterium beams, *Plasma. Phys. Control. Fusion.*, 42 (2000), 781–806.
- [22] R. K. Janev, C. D. Boley and D. E. Post, Penetration of energetic neutral beams into fusion plasmas, *Nucl. Fusion.*, 29 (1989), 2125.
- [23] J. G. Lodge, I. C. Percival and D. Richards, Semi-empirical cross sections for excitation of hydrogen atoms by protons, *J. Phys. B.*, 9 (1976), 239–254.

- [24] R. K. Janev, W. D. Langer, Jr. D. E. Post and Jr. K. Evans, *Elementary Processes in Hydrogen-Helium Plasmas*, Springer-Verlag, Berlin, 1987.
- [25] L. Vriens and A. H. M. Smeets, Cross-section and rate formulas for electron-impact ionization, excitation, deexcitation, and total depopulation of excited atoms, *Phys. Rev. A.*, 22 (1980), 940–951.
- [26] R. Janev, D. Reiter and U. Samm, Technical Report 4105, Juel Report, 2004, <http://www.eir-ene.de>.
- [27] R. E. Bell, Carbon ion plume emission produced by charge exchange with neutral beams on NSTX, Technical Report PPPL-4167, Princeton Plasma Physics Laboratory, 2006.
- [28] D. F. Finkenthal, The Measurement of Absolute Helium Ion Density Profiles on the DIII-D Tokamak Using Charge Exchange Recombination Spectroscopy, PhD thesis, University of California, Berkeley, 1994.
- [29] V. Tang, Experimental and Numerical Characterization of Ion-cyclotron Heated Protons on the Alcator C-Mod Tokamak, PhD thesis, Massachusetts Institute of Technology, 2006.
- [30] Y. Luo, W. W. Heidbrink, K. H. Burrell, D. H. Kaplan and P. Gohil, Measurement of the D_α spectrum produced by fast ions in DIII-D, *Rev. Sci. Instrum.*, 78 (2007), 033505.
- [31] R. J. Goldston and P. H. Rutherford, *Introduction to Plasma Physics*, Institute of Physics Publishing, Bristol, 1995.
- [32] M. A. Van Zeeland, J. H. Yu and N. H. Brooks et al., Active and passive spectroscopic imaging in the DIII-D tokamak, *Plasma. Phys. Control. Fusion.*, 52 (2010), 045006.
- [33] B. Grierson et al., *Nucl. Fusion*, (2011), submitted.
- [34] Y. Luo, W. W. Heidbrink, E. Ruskov, K. H. Burrell and W. M. Solomon, Fast-ion D_α measurements and simulations in quiet plasmas, *Phys. Plasmas.*, 14 (2007), 112503.
- [35] W. W. Heidbrink, M. Murakami and J. M. Park et al., Beam-ion confinement for different injection geometries, *Plasma. Phys. Control. Fusion.*, 51 (2009), 125001.
- [36] B. Geiger et al., Fast-ion D_α measurements at ASDEX upgrade, *Plasma. Phys. Control. Fusion.*, 53 (2011), 065010.