UC Irvine

UC Irvine Previously Published Works

Title

Measurement and simulation of passive fast-ion D-alpha emission from the DIII-D tokamak

Permalink

<https://escholarship.org/uc/item/9sr7c8t9>

Journal

Nuclear Fusion, 56(11)

ISSN

0029-5515

Authors

Bolte, NG
Heidbrink, WW
Pace, D
[et al.](#)

Publication Date

2016-09-14

DOI

10.1088/0029-5515/56/11/112023

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

Measurement and simulation of passive fast-ion D-alpha emission from the DIII-D tokamak

Nathan G. Bolte¹, William W. Heidbrink¹, David Pace²,
Michael Van Zeeland² and Xi Chen²

¹ University of California-Irvine, 4129 Frederick Reines Hall, Irvine, CA 92697, USA

² General Atomics, PO Box 85608 San Diego, CA 92186-5608, USA

E-mail: NathanBolte@gmail.com, wwheidbr@uci.edu, pacedc@fusion.gat.com,
vanzeeland@fusion.gat.com and chenxi@fusion.gat.com

Received 19 January 2016, revised 14 August 2016

Accepted for publication 23 August 2016

Published 14 September 2016

Abstract

Spectra of passive fast-ion D-alpha (FIDA) light from beam ions that charge exchange with background neutrals are measured and simulated. The fast ions come from three sources: ions that pass through the diagnostic sightlines on their first full orbit, an axisymmetric confined population, and ions that are expelled into the edge region by instabilities. A passive FIDA simulation (P-FIDASIM) is developed as a forward model for the spectra of the first-orbit fast ions and consists of an experimentally-validated beam deposition model, an ion orbit-following code, a collisional-radiative model, and a synthetic spectrometer. Model validation consists of the simulation of 86 experimental spectra that are obtained using 6 different neutral beam fast-ion sources and 13 different lines of sight. Calibrated spectra are used to estimate the neutral density throughout the cross-section of the tokamak. The resulting 2D neutral density shows the expected increase toward each X-point with average neutral densities of $8 \times 10^9 \text{ cm}^{-3}$ at the plasma boundary and $1 \times 10^{11} \text{ cm}^{-3}$ near the wall. Fast ions that are on passing orbits are expelled by the sawtooth instability more readily than trapped ions. In a sample discharge, approximately 1% of the fast-ion population is ejected into the high neutral density region per sawtooth crash.

Keywords: passive FIDA, edge light, neutral density, FIDASIM, FIDA, sawtooth crash, fast ion losses

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

1. Introduction

Fast ions are produced through neutral beam injection, radio-frequency heating, and fusion reactions. Well-confined fast ions contribute to plasma stability and are required for energy recycling in more reactive devices such as ITER. A large fast-ion population contributes to instabilities and losses in the form of Alfvén eigenmodes and other MHD instabilities and, therefore, the careful study of fast-ion confinement and loss remains a crucial area of research in fusion science.

A sudden loss of fast ions would not only jeopardize plasma performance, but could also—especially in near-reactor conditions—seriously damage the inner wall of the vessel. Accordingly, detection of fast ions near the walls is an important diagnostic challenge.

Two types of fast-ion diagnostics utilize charge exchange (CX) reactions between fast ions and neutrals: neutral particle analysis (NPA) and fast-ion D-alpha (FIDA) spectroscopy. A distinction is made between *active* beam measurements that use CX with *injected* neutrals and *passive* measurements that use *background* neutrals. For active measurements, signals originate near the intersection of the diagnostic sightline with

Figure 1. Lines of sight (LOS) used to measure passive FIDA spectra (shown as line segments). (a) Elevation view showing oblique LOS P1-P11 and vertical LOS V1, V2. Note that the oblique LOS appear curved when projected into 2D. (b) Plan view showing oblique and vertical LOS and beam centerlines (arrows) and naming scheme.

the injected beam. Since the neutral density profile in most tokamaks is usually much higher in the edge region than in the core, passive signals often originate near the plasma edge. For example, passive neutral-particle measurements have been used to detect the expulsion of fast ions from the core for decades [1].

Although more commonly an active diagnostic, passive FIDA signals are also detected. Heidbrink *et al* noticed that diagnostics designed to measure beam emission sometimes detect passive FIDA light in DIII-D and NSTX [2]. The signals are especially pronounced when instabilities expel fast ions to the edge [2, 3]. Dedicated active FIDA systems on MAST [4] and ASDEX Upgrade [5] also pick up large passive FIDA signals under some conditions. These observations of passive FIDA light are the impetus for the work presented here.

The motivation is threefold. First, passive FIDA signals are often an unwanted complication in active FIDA spectroscopy, so it is important to understand this background emission. Second, detection of fast-ion losses is challenging in ITER and other future devices. This paper explores a method to obtain quantitative information about fast-ion losses from spectroscopic data. Third, measurement of the ambient neutral density poses diagnostic challenges in its own right. This paper shows that, in the presence of a known fast-ion source, information about the neutral density profile can be obtained from passive FIDA measurements.

Conceptually, the measured passive FIDA signals originate from three distinct fast-ion populations that are produced by the neutral beams. One population consists of fast ions on their first orbits. Fast ions are born when an injected neutral ionizes. Since the neutral beams are spatially localized, the first orbits are also localized. In measurements with a loss detector, this class of ions was dubbed a ‘light-ion beam probe’ [6]. These ions are created at a particular toroidal location by a selected neutral beam, then execute a banana orbit that traverses the edge region at a different toroidal location near the detector.

For a loss detector, the signal from this toroidally localized source is often much larger than axisymmetric losses [7]. Since the beam deposition and equilibrium fast-ion orbits are well understood, these ions can be used to measure the effect of waves that deflect the orbits [7, 8]. For a passive FIDA diagnostic, this toroidally asymmetric population is detected when fast ions traverse the FIDA sightlines on their first full orbits. This population can be viewed as a known source of fast ions that probe the edge and pedestal regions.

Ions born on confined orbits eventually distribute themselves toroidally. These confined ions are the second population that produce passive FIDA light. This axisymmetric population is described by a Fokker–Planck equation with terms for the source, Coulomb drag and scattering, and charge-exchange losses.

Core fast ions that are expelled into the edge region by instabilities are the third fast-ion population that produces passive FIDA light. These transiently increase the edge confined population; also, ions scattered onto loss orbits can contribute to the signal.

The paper is organized as follows. Section 2 describes the apparatus and experimental conditions. Next, the codes used to model the passive FIDA emission from the three types of fast ions are described (section 3). Analyzed data from the axisymmetric-confined and toroidally-localized populations appear in section 4. Signals produced by ions that are expelled to the edge by the sawtooth instability appear in section 5. Section 6 illustrates inference of the neutral density from passive FIDA measurements. The final section (section 7) presents conclusions and suggestions for future work.

2. Passive FIDA diagnostic and experimental method

The data presented in this article utilize two different groupings of lines of sight (LOS) and two different spectrometers. One group consists of two ‘vertical’ LOS (figure 1) designed

originally to view beam 330° left (figure 1) for use in active FIDA measurements [9]. These vertical LOS have $f/4.4$ lenses mounted under the midplane and light is carried to the spectrometer via $1500\ \mu\text{m}$ fibers. The second group is comprised of eleven ‘oblique’ LOS using $f/1.8$ lenses mounted on top of the tokamak and designed as an upgrade to the FIDA system to view beam 210° right [10]. These eleven LOS each utilize a bundle of three $1500\ \mu\text{m}$ fibers for greater light pickup.

The majority of this work utilizes the original FIDA spectrometer [11]. This spectrometer is a Czerny–Turner with $f/4$ mirrors, an $1800\ \text{grooves mm}^{-1}$ grating, and a $300\ \text{mm}$ focal length. A small, rectangular neutral density filter is placed at the focal plane to heavily reduce the bright, unshifted D_α signal at $656.1\ \text{nm}$ that would otherwise saturate the camera. A charged-coupled device (CCD) camera (VelociCam VC105A from PixelVision) is used to detect the dispersed signal. The CCD chip is a 652×488 pixel ($8 \times 6\ \text{mm}$) array with $1\ \text{ms}$ integration time and is set up to image two LOS at once. The spectral range of this system ($648\text{--}664.5\ \text{nm}$) allows the viewing of both red- and blue-shifted signals.

Toward the end of the campaign, the camera in the Czerny–Turner spectrometer failed, preventing a post-campaign calibration and forcing the use of a second spectrometer. This HoloSpec transmission grating spectrometer uses a Sarnoff CAM1M100 CCD camera to image six simultaneous spectra onto a 1024×1024 pixel chip [10]. To avoid the bright region near $656\ \text{nm}$, this system only images the blue region of D_α emission ($650\text{--}654\ \text{nm}$).

The location of each LOS was determined by back-filling the optical system with light. An intensity calibration was carried out for the HoloSpec spectrometer using a calibrated light source.

The oblique and vertical sightlines are sensitive to different classes of particles. The oblique views are most sensitive to passing particles, while the vertical views are most sensitive to trapped orbits. This is a simple consequence of the Doppler shift. Passing particles have a larger $v_\parallel$, giving them the Doppler shift needed to be seen by the oblique system. Trapped particles have a larger $v_\perp$ and therefore have a more favorable Doppler shift for the vertical system.

The experimental method is designed to produce unambiguous passive FIDA signals with a significant contribution from the toroidally asymmetric population. It is well known that beam prompt losses land within a narrow range of toroidal angles for a given beam and that those angles are a function of magnetic field. A simplified version of P-FIDASIM (section 3.3) is used to scan through LOS and beam options to determine which LOS should be paired with which beam—and at what toroidal field and plasma current—to get the highest signal from first-orbit ions. These calculations consistently predicted higher prompt signals when pairing the oblique views with the 30-left beam. The tokamak is run in the L-mode to avoid edge-localized modes (ELMs) [12], as ELMs eject high-energy particles into the edge region, obscuring the first-orbit signals. ELM bursts can be dealt with in future applications but are avoided altogether in this initial work. Signals would also be swamped by active emission if the LOS were to view an active beam. Thus, any beam that a chosen

Figure 2. Time evolution of (a) plasma current and beam power, (b) central electron temperature from electron cyclotron emission (ECE), and (c) line-average density for central and edge interferometer chords.

LOS crosses is turned off for the duration of the discharge. Figure 2 shows typical traces. Often, only a single modulated beam is the source of fast ions and the plasma current is gradually ramped down (figure 2(a)). Many discharges have sawteeth (figure 2(b)). The electron density is approximately constant (figure 2(c)).

The electron density profile is measured by Thomson scattering [13] and CO_2 interferometers [14]. The electron temperature is measured by Thomson scattering and an electron cyclotron emission (ECE) radiometer [15]. The ion temperature, rotation frequency, and Z_{eff} profiles are inferred from charge exchange recombination spectroscopy measurements of carbon [16]. (Since the plasma-facing components are graphite, carbon is the main impurity.) The safety factor profile is from EFIT equilibrium reconstructions [17] that use magnetics and motional Stark effect [18] data. Plasma profiles are expressed in terms of the normalized toroidal flux, $\rho = \sqrt{\Phi_{\text{tor}}/\Phi_{\text{tor_LCFS}}}$ —where $\Phi_{\text{tor_LCFS}}$ is the flux enclosed by the last closed flux surface (LCFS).

The tokamak (especially the edge-region) is filled with light in the D_α region, which comes from impurity lines, scattering, reflections, bremsstrahlung, and D_α from cold neutrals. These backgrounds need to be removed to reveal the fast-ion component. Beam modulation provides the background signal as a function of time. Normally, the beam providing the prompt signals is modulated at $50\ \text{Hz}$ at 50% duty cycle (figure 2(a)). Spectra are then conditionally averaged by removing the beam-off signals from the beam-on signals, producing the net spectra. This means that spectra are time-averaged based on the *condition* of the current of the neutral beam of interest. Two time-averaged spectra are created: one, averaged over all time points when the beam is turned on and a second spectrum, averaged over all time points when the beam is turned off. A sample set of conditionally-averaged spectra are shown in figure 3(a) and reveals several features. The neutral-density filter at the focal plane of the spectrometer

Figure 3. Conditionally-averaged spectra from discharge 149954 (3100–4984 ms) using oblique LOS P6 while modulating beam 30-left. (a) Average spectrum during beam-on times (black) and during beam-off times (red) with line radiation and neutral-density filtered region labeled. (b) Difference between beam-on and beam-off spectra gives the modulated fast-ion spectrum with error bars (representing variations within the time-averaging window).

Figure 4. Spectrum found by conditionally-averaging edge light based on the timing of a core ECE channel during 26 sawtooth crashes in discharge 149941 (3485–4982 ms, chord P4). The presence of impurity lines is explained by the well-known phenomenon of edge-heating due to the ejection of high temperature particles. The edge heating raises and widens the impurity line signals, thus producing an asymmetry in the time-slice subtraction. The conditional averaging uses 3 ms before each crash and 3 ms after each crash. The curve labeled TRANSP is a FIDASIM spectrum produced using the TRANSP distribution function inside ρ_{inv} for this shot and time (section 3.2). It uses an ad hoc neutral density profile and assumes that the sawtooth crashes preserve the fast-ion velocity. The curve is normalized to the experimental curve and is included to show the plausibility of the sawtooth spectrum, especially the bias toward blue-shifted light resulting from the largely co-going ion population.

as marked in the figure blocks out most light in the region of 6550–6580 Å. Impurity lines from boron, carbon, and oxygen are easily identified as is the heavily-reduced D_α line. The region between the oxygen-V line and the blue-ward side of the neutral-density filter is the region consistent with Doppler shifts expected from co-passing fast ions and the viewing geometry and contains the fast-ion signal. Taking

the difference between the two spectra produces the ‘passive FIDA’ spectrum of figure 3(b).

Spectral quality is improved in the post-processing of signals by removing times when instabilities alter the edge region. The well-known sawtooth instability [19] occurs often in DIII-D L-mode discharges and can cause changes in edge plasma spectra. Therefore, times near crashes

are excluded from the conditional averaging, resulting in larger fast-ion signals, signal baselines closer to zero, and improved line-radiation cancellation. The amounts of time before and after each sawtooth to reject are determined empirically to produce the best final spectra with the smallest magnitude baseline and the greatest reduction of impurity line radiation.

Alternatively, the data may be conditionally averaged to study the effect of the sawtooth instability on the passive FIDA emission. In this case, sawtooth loss spectra are created by conditionally averaging raw spectra against T_e . This is done by averaging all spectra immediately after each sawtooth crash and subtracting from it the average of all spectra immediately before each crash. An example appears in figure 4. Sawtooth crashes eject enough hot plasma that both impurity lines and the bremsstrahlung background rise due to the edge heating (i.e. imperfect background subtraction). Analysis therefore is only done after removing a best estimate of the bremsstrahlung level and by avoiding known impurity-line regions. The regions of the spectrum likely resulting from fast-ion losses are marked in figure 4. Because the viewing geometry is largely tangential, the greater level of blue signal over red is consistent with the mostly co-going fast-ion population of this shot.

3. Passive FIDA simulation

This section describes the models used to simulate the signals produced by each of the three fast-ion populations.

3.1. Modulated confined population

Since the modulation period is shorter than the slowing-down time, which is ~ 100 ms in the core and ~ 30 ms at $\rho = 0.925$, it is convenient to separate the axisymmetric fast-ion distribution function F into a slowly-varying component $\bar{F}$ and a modulated component $\tilde{F}$, $F = \bar{F} + \tilde{F}$. For the analysis in section 4, the background are removed through timeslice subtraction, as if the signal was produced by an active beam. This eliminates unwanted contaminants like impurity lines from the spectra but it also effectively removes passive FIDA light produced by the time-average distribution $\bar{F}$. Only the modulated distribution $\tilde{F}$ contributes to the measured signal. Modeling of this signal follows the procedure described in [20]. The distribution function is computed by the TRANSP NUBEAM code [21] at intervals of 2 ms using 10^6 Monte Carlo markers for the same time period as averaged experimentally. The distribution functions are input to the FIDASIM [22] synthetic diagnostic code. FIDASIM normally simulates active FIDA signals produced by the injected full, half, and third-energy beam components and the surrounding beam halo. To simulate passive signals from the confined population, the injected neutrals are eliminated and the halo neutrals are replaced with a neutral density population in local thermodynamic equilibrium that is distributed spatially according to the profile of figure 19. (For the chords analyzed here, the passive signals are 24 and 65 times smaller than a typical active signal.) Next, the predicted

Figure 5. Conditionally averaged FIDA signals (*) for the oblique channels that intersect the midplane at (a) $R = 205$ and (b) 214 cm. The predicted contributions of the confined axisymmetric population (dashed line), the toroidally-asymmetric population (dotted line), and the best-fit sum of these contributions (solid line) are also shown. The signals are averaged between 3.2–4.5 s, integrated between 650–655 nm, and the mean is subtracted for both theory and experiment.

spectra are conditionally averaged and integrated over wavelength in the same manner as the experimental data.

Figure 5 shows the result of this analysis for one of the discharges in this study. Although theoretically expected signals from each of the three fast-ion populations are computed separately, in experiment, their effects are additive. The discharge of figure 5 does not have sawteeth, so the episodic contributions to the signal from instability bursts are negligible. With a modulated source beam, the signal from the toroidally-asymmetric population is expected to appear as a square wave, with a slight ~ 0.1 ms delay relative to the modulated source. (The delay is caused by the orbital transit time between ionization and the FIDA sightlines. The expected square wave response is observed on loss-detector signals [7, 8]). The contribution from the confined axisymmetric population is computed by TRANSP and FIDASIM, as described above. For both of the analyzed chords, the best fit to the measured signal is obtained by assuming that 60% of the signal is produced by the square-wave response to the toroidally-asymmetric population and 40% is produced by the modulated confined population. The agreement between the measured and modeled signals is excellent. This analysis successfully explains the observed time evolution of the passive FIDA signals (figure 5).

Figure 6. Fast ion distribution output from TRANSP for shot 149941 just inside the sawtooth inversion radius of $\rho = 0.43$.

3.2. Population expelled by sawtooth crashes

In this section, rough TRANSP-FIDASIM calculation is used to confirm that the experimental data shown in figure 4 are in fact FIDA signals from sawtooth-crash-ejected fast ions. First, the TRANSP NUBEAM code [21] is run for the same shot and conditions to determine the fast-ion distribution function. The fast ions *just inside* the sawtooth inversion radius ρ_{inv} are summed to obtain an estimate of the portion of the fast ion distribution likely to be expelled to the plasma edge, $F(E, p)$, where E is the fast-ion energy and $p = v_{\parallel}/v$ is the pitch (see figure 6). Next, FIDASIM is used to calculate weight functions [23], $W(E, p; \lambda)$, as a function of wavelength λ for the views in question. Usually, FIDASIM employs beam-injected neutrals however, for the calculation presented here, FIDASIM is modified to use a cold edge neutral population. The simulated spectra for the desired chords are the integral of the product of F and W , i.e. $S = \iint F W dE dp$, where the integral is over pitch and energy. One oblique chord, P4, is used and the calculated spectrum is compared with the measured spectrum in figure 4. (A similar calculation is also performed for a vertical chord, V1.) Using the fast-ion distribution function just inside ρ_{inv} gives a first-order estimate of sawtooth spectra which essentially assumes that sawtooth crashes eject a uniform sampling of the fast-ion population into the edge region. This is an over-simplification, in that sawtooth crashes are known to preferentially affect passing particles [24–26]. In addition, particles redistributed outwardly undergo changes in energy and pitch as a result of the process and therefore are not simply samples of the confined ion distribution. In addition, the TRANSP spectrum shown in figure 4 uses an improvised edge neutral density spatial profile. However, despite these known differences, the TRANSP/FIDASIM spectrum has a striking similarity to the experimental spectrum. In particular, the bias towards blue-shifted light is apparent.

The TRANSP code includes an optional Kadomtsev model [27] of the sawtooth crash. In section 5.2, this option is used to estimate the fraction of fast ions that are expelled into the high neutral-density region near the plasma edge.

3.3. Toroidally asymmetric population

A new code dubbed P-FIDASIM is developed to simulate signals produced by the toroidally asymmetric population on their first orbits. P-FIDASIM has six main parts. A beam attenuation model determines the initial locations and velocities of ions born from the ionization of injected neutrals. An ion orbit-tracking and capturing module calculates the ion orbits and accumulates the portions of orbits that pass through the chosen lines of sight. A neutral density model calculates a normalized neutral density from plasma profiles. An interpolation module finds all the pertinent values at each captured ion location. The ‘atomic loop’ portion of the code assembles several sub-modules that include a collisional-radiative (CR) model to find the photon emission rates and a module to find the Doppler and Stark wavelength shifts for each LOS. Finally, the ‘spectra’ module compiles and organizes the emission information into spectral intensity versus wavelength plots. A schematic of the simulation is shown in figure 7.

The beam model calculates the ion birth locations and velocities for the chosen neutral beam. The injected neutral gas is mostly ionized by charge exchange with deuterium ions, but also by electron impact as well as charge exchange with non-fully stripped carbon impurities. The beam model takes, as inputs, electron and carbon density profiles, the electron temperature profile, the equilibrium fields, the beam energy, and the number of desired ions. The deuterium ion density is inferred from the electron and carbon densities using quasi-neutrality. Many lines, or ‘rays’, are randomly chosen as potential pathways for neutrals. Together, all of the rays reproduce the known divergence, width, and height of the beam. The beam attenuation profile is then calculated along each ray using the atomic data and analysis structure (ADAS) beam-stopping routines³. Random selection is again used to choose ionization locations along each ray. Beam attenuation rates along each ray are then used to give ion birth rates. Combining this list of birth rates with the desired number of

³ www.adas.ac.uk

Figure 7. Schematic of the simulation of passive FIDA light produced by first-orbit ions.

ions gives a final list of ion birth locations and velocities. Any trajectories that intersect beam collimators are removed from the final output. Any neutrals that do not ionize are collected at the opposite wall and are given in the model's output as 'shinethrough'. The beam model is experimentally validated by comparison to camera images [28].

The ion capture module calculates full ion orbits and determines if and where each ion enters the volume of the chosen lines-of-sight. The main inputs for this code are the equilibrium fields, the ion birth locations and velocities from the beam model, and the LOS geometry. Ion orbits are found by solving the Lorentz force equation using a predictor-corrector Adams–Bashforth–Moulton differential equation solver⁴ and are calculated until either the ions hit the wall or travel 40 m (empirically found by increasing the maximum number of steps per orbit until the spectra from successive simulation runs were within photon-noise levels of each other). The time it takes to calculate the ion orbits makes this the most intensive set of calculations in the simulation. The output of this module is a list of ion locations and velocities detected within the volume of each LOS.

The neutral density model determines the normalized spatial dependence of neutral density resulting from the attenuation of neutrals launched uniformly from the walls. Neutral attenuation is calculated based on the dominant terms of electron-impact ionization and charge exchange with deuterium. While other, better, codes exist that calculate realistic neutral densities in 2D or 3D (e.g. [29]), these codes require neutral gas pressure and D_α/D_β data that were not available during these studies. The employed model is less complete and far simpler by comparison, however, it utilizes data we know with confidence: the 1D plasma profiles. Neutrals do not respond to magnetic and electric fields and there exist localized

neutral sources and sinks and therefore neutral density is not generally a flux function. However, the neutral profile does attenuate rapidly in the $\nabla\rho$ direction so, for simplicity, neutral density is assumed to be a flux function at this point in the simulation (prior to inversion with experimental data—section 6). The 1D neutral density model begins by 'launching' a neutral atom inwardly from the wall with a nominal energy where it encounters the experimentally-determined electron and main-ion profiles. The motion of the neutral is along $\nabla\rho$ —perpendicular to the magnetic field lines—and therefore the profiles are unchanged perpendicular to the neutral's line of motion, thus justifying the 1D nature of the model. Initial velocities are chosen randomly in 3D at 3 eV (a typical temperature after dissociation from D_2) and only the component along the 1D line is considered. Ions are also launched with random thermal velocities and the probability of CX between the ions and the neutral, the probability of ionization from electrons, and the neutral's motion through the 1D cells along the line of motion are all monitored. Charge exchange is interpreted as a new neutral being 'born' which is then given the velocity of the incoming ion. Ionization means the neutral is lost and a new neutral is launched from the wall to start the process over. If neither of these events take place before the time it would take the neutral to pass to the next cell, then the neutral is moved to the next 1D cell where the tracking process begins again. The neutral density is then proportional to the time spent in each cell and is normalized by cell size, number of Monte Carlo particles, and an arbitrary scaling factor (to keep the maximum value close to 1). Neutrals continue to be sourced inwardly from the wall until a dependence on ρ is converged upon. The 1D model is used repeatedly to find the density for the entire (R,z) cross-section of the tokamak.

The interpolation module gathers the various profiles and interpolates their values onto the specific ion locations determined in the ion-capture module. Several profiles are

⁴ <http://cow.physics.wisc.edu/~craigm/idl/down/ddeabm.txt>

Figure 8. Atomic population densities during a charge exchange event, showing the validity of the local approximation.

interpolated onto *each* ion location: ion temperature, impurity density, electron density, electron temperature, and neutral density. The deuteron density at each ion location is found by quasi-neutrality: $n_D = n_e - 6n_c$, where n_c is the carbon density. The electron collision rate, $n_e \langle \sigma v \rangle$, is also found at each ion location. The fractional population of each atomic level (up to $n = 4$) is determined for all neutrals assuming collisional-radiative equilibrium using local plasma parameters. Lastly, the equilibrium magnetic field is found at each ion location for later use in calculating the Stark effect.

The ‘atomic loop’ module carries out the atomic physics calculations by running a collisional-radiative model. Using the charge-exchange probability between promptly-lost fast ions and cold edge-neutrals, it gives the photon emission rates and determines the Doppler and Stark shifts. The probability that edge neutrals and fast ions charge exchange is a function of the relative population of the quantum states of the neutrals, the relative velocity between the ions and the neutrals, and the quantum states that the new, fast neutrals are born into. Because edge neutrals are ~ 10 eV while fast ions are ~ 80 keV, a great simplification can be made by letting the relative velocity be the fast-ion velocity, i.e. $\vec{v}_{\text{rel}} = \vec{v}_{\text{fi}} - \vec{v}_{\text{n}} \approx \vec{v}_{\text{fi}}$. This assumption is not valid in the core region and while the core is not ignored by the simulation, the neutrals there are implicitly assumed to be cold. This assumption is more valid in discharges with higher core electron density as this reduces the core neutrals and forces the edge light to dominate. The background neutrals are limited to the first four quantum states as are the fast neutrals resulting from CX as the occupation levels for $n > 4$ are very low in a collisional-radiative equilibrium. The collisional radiative (CR) model considers intermediate atomic levels up to $n = 7$ and outputs the photon emission rate for the Balmer-alpha $n = 3 \rightarrow 2$ relaxation. The Doppler shift is then calculated by taking the dot product of the fast neutral’s velocity (i.e. that of the fast ion at the time of CX) and the unit vector of the LOS. The ‘local approximation’—that $\vec{x}_{\text{fast_ion}} \approx \vec{x}_{\text{cx_neutral}}$ —is justified as the time between CX and emission is on the order of several nanoseconds meaning

the distance traveled before emission is $\lesssim 2$ cm while the LOS diameter for the oblique views is ~ 10 cm (i.e. the neutral radiates before leaving the LOS). Excitation and radiation subsequent to CX is quickly returned to equilibrium levels as can be seen in figure 8. The figure shows the atomic population densities during a CX event. At a density of $5 \times 10^{19} \text{ m}^{-3}$, the populations return to equilibrium over a distance of 5 cm (at 80 keV). This local approximation reduces the spatial accuracy of the simulation but is a great simplification over tracking each neutral until it radiates. The Stark effect further shifts the wavelength of light based on the electric field experienced by the fast neutral.

The ‘spectra’ section of the code takes the raw, unorganized spectral data from the atomic loop module and bins the photon counts together into wavelength bins while separating them according to which LOS they are in.

The appropriate units of spectral radiance are obtained by multiplying by the length of the LOS, dividing by the volume of the LOS, and dividing by the isotropic solid angle, 4π sr. Finally, spectra are scaled to realistic values by multiplying by the ratio of experimental beam current to simulated beam current.

Spectra are further modified to more realistically represent experimental spectra by the application of instrumental broadening and response functions. Instrumental broadening is the well-known tendency of physical optics to broaden spectral features. The instrumental response function is a normalized spectral response to white light that provides a correction for variations in responsivity as a function of wavelength, including the effect of the neutral density filter.

P-FIDASIM is benchmarked against the well-established FIDASIM code [22] to verify that P-FIDASIM is correctly solving the appropriate mathematical problem. Significant differences in purpose and methodologies between the two codes require unphysical but overlapping conditions. The parameters of the benchmark are as follows: all plasma profiles are uniform for $\rho \leq 1.0$ and zero for $\rho > 1.0$: electron density is $5 \times 10^{12} \text{ cm}^{-3}$, impurity density is zero ($Z_{\text{eff}} = 1$), and

Figure 9. Comparison of simulation results for FIDASIM and P-FIDASIM. **Upper:** sample spectra for oblique LOS P9b. **Lower:** P-FIDASIM results normalized by FIDASIM results for all six outermost LOS. The average (RMS) difference is 13.6%.

$T_e = T_i = 100 \text{ eV}$. The equilibrium fields come from DIII-D's magnetic equilibrium fitting code, EFIT, for discharge 152817 at 3000 ms. Creation of the fast-ion population requires more care as FIDASIM requires the input of a toroidally-symmetric fast-ion distribution function while P-FIDASIM requires the 3D orbits of individual ions. Equivalent inputs are created for both simulations in the following manner: beam 30-left is used as a starting point for a fast-ion population. Only full-energy, 80 keV ions born in the edge region ($R \geq 2.0 \text{ m}$) and near the midplane ($|Z| \leq 3.0 \text{ cm}$) are considered. The orbits are calculated for this small subset of ions and are then 'smeared out' toroidally by replicating the orbits with ever-increasing toroidal angles. The result is a small but toroidally-symmetric fast-ion population. The orbits themselves are entered into P-FIDASIM and are binned on a 2D grid for FIDASIM. A uniform density of charge-exchange neutrals of $1 \times 10^{10} \text{ cm}^{-3}$ is defined in the volume of the 210-right beam as it intersects the 'oblique' sightlines. The beam is given a mere 50 eV of kinetic energy and zero temperature.

The results of the benchmark are summarized in figure 9. While the physics of these spectra is not relevant (this is a contrived case for benchmarking only), features can still be understood. The predominant blue-shift of the spectra is consistent with a large population of co-passing ions from beam injection being viewed by largely tangential lines of sight. The far blue-shifted peak is due to passing particles having

a mostly toroidal velocity and therefore being well-aligned to spend more time in the LOS. The peak at the red-shifted end of the spectra is due to the physically downward drift of trapped particles when their gyro-orbits are tangential to the LOS and therefore more time is spent in that LOS. The total average difference between the simulations is 13.6% for the six LOS used. Considering the different purposes and methodologies of these two simulations, this error seems low enough to verify that the passive simulation is indeed solving the correct set of equations.

The success of P-FIDASIM depends heavily on knowing the plasma profiles in the edge region as it is ions born in the edge that are poorly confined and therefore enter the higher neutral density of the edge. However, plasma profiles are not always measured at DIII-D beyond the LCFS ($\rho = 1.0$). For some of this work, edge profiles were determined by extrapolating core profiles into the edge region. Thus, sensitivity studies are performed to determine the effects of systematic changes in edge profiles resulting from extrapolation. Edge data, when available, are fit well to an exponential and thus can be parametrized. The average exponential parameter is then used in cases where edge data are not available.

The black curves in figure 10 show typical plasma profiles used for simulations. The red curves show the variation in the extrapolated regions used for this sensitivity study and table 1 summarizes the net effect of the profile changes. Note that a

Figure 10. Edge profile changes used for sensitivity studies. Profiles for $\rho \leq 1.0$ are fits to data for discharge 152817 at 3001 ms. Profiles in the region $\rho > 1.0$ are fits to data of the form $F = F_0 \rho^{-\alpha}$ where F is the profile in question and F_0 is that profile's value at $\rho = 1.0$. These represent extrapolations when edge-data is not available. α_{original} is the normal value used for this discharge while α_{new} is used in the edge-profile sensitivity study.

Table 1. Effects of the profile changes shown in figure 10.

Profile	Profile change	Change for LOS:					
		Innermost			Outermost		
		P1	P3	P5	P7	P9	P11
T_e	+361	−5.2	−5.2	−6.5	+5.5	+2.8	−2.4
T_i	−45	$\sim 10^{-6}$	$\sim 10^{-6}$	$\sim 10^{-6}$	$\sim 10^{-6}$	$\sim 10^{-6}$	$\sim 10^{-6}$
n_e	+245	+125	+116	+178	+46	−10	+18
n_c	+226	+11	+14	+18	n/a	n/a	n/a

Note: Numbers are the average percentage change with the sign indicating a rise or decline of the parameter.

361% increase in the edge value of the electron temperature only changes the spectra by ~ 2 –6% and that a 45% decrease in T_i has no real effect on the spectra. By far, the greatest effect comes from changes in the electron density. Increasing n_e in the edge increases electron-impact ionization of the incoming beam—causing the deposition of more ions in the edge that later enter the LOS—and also affects the stimulated excitation and relaxation of the CX fast neutrals. Increasing the edge carbon density also increases signals by increasing the beam

deposition in that region. For cases where edge profile measurements are not available or when time variations are large, the large changes in edge profiles used in this sensitivity study (while arbitrary in number) show that even wildly-varying extrapolations only affect the final spectra by a factor of ~ 2 .

As an example of how changing edge profiles changes spectral shape, figure 11 shows how the electron density profile change affects the spectra. Notice that changes occur in the spectral shape as well as intensity.

Figure 11. Changes in spectra associated with the changes to the electron density profile shown in figure 10.

4. Comparison of P-FIDASIM simulations with beam modulation data

This section compares a database of signals that are produced by the toroidally-asymmetric and modulated-confined populations with simulations of expected signals from the toroidally-asymmetric population. Ideally, the simulated signals would also include the contribution of the modulated-confined population but that contribution is neglected. As will be shown, the agreement between theory and experiment is still good. The error introduced by neglect of the modulated-confined population is modest for two reasons.

- (i) Due to the unavailability of an absolute intensity calibration (section 2), the comparison is of relative signal levels. For the case that was analyzed in detail (figure 5), the toroidally-asymmetric population contributes 60% of the signal, while the modulated-confined population contributes 40% of the signal. The data are analyzed by subtracting the ‘beam-off’ signal from the ‘beam-on’ signal. Because the modulated-confined response is delayed relative to the beam timing (figure 5), this further reduces the modulated-confined contribution to the averaged signal. Moreover, as long as the relative contributions of the toroidally-asymmetric and modulated-confined populations are similar for the different discharges in the database, a *relative* comparison of experimental and theoretical signal levels is largely unaffected.

Figure 12. Distribution function calculated by TRANSP for the modulated-confined population near the crossing of the oblique P8 sightline with the midplane. The net ‘beam-on minus beam-off’ distribution function is obtained from 11 cycles between 3.2–4.5 s for the discharge shown in figure 5. The triangles indicate velocities that are heavily populated by the toroidally-asymmetric population.

- (ii) The expected spectra produced by the two populations are similar. The toroidally-asymmetric distribution function has three narrow peaks associated with the full, half, and third energy components of the beams; because the beam divergence is small, the pitch distribution of each peak is also narrow where the orbit intersects a FIDA sightline.

Figure 13. Comparison between PFIDASIM simulation and experiment. Experimental spectra are smoothed using a moving average with a width of 2 Å. Note that missing regions in the simulated spectra can be considered zero values as no photons were detected. Spectra were normalized by their maximum values. Note that plasma parameters are held largely constant over the time-averaging window, with the exception that cases (c) and (d) have decreasing plasma current throughout the window.

This differs greatly from the slowing-down distribution of the time-average confined population $\bar{F}$; however, the averaging technique eliminates the contribution of $\bar{F}$ and only measures the contribution of the modulated population $\tilde{F}$. As shown in figure 12, the $\tilde{F}$ population has a distribution function that peaks at the same locations as the toroidally-asymmetric first-orbit distribution. Since the FIDA spectral features are determined by the fast-ion velocity vectors, the FIDA spectra produced by the two populations are similar.

Figure 13 shows representative comparisons between the P-FIDASIM simulations and experiments. The lack of calibration for the spectrometer camera forces these plots to be in arbitrary units. However, it can be seen that several features of the experimental signals are reproduced by P-FIDASIM. For example, the four cases recreate the appropriate Doppler shifts: showing a balance of blue and red shift for a vertical LOS—sensitive to trapped ions—and a predominately blue shifted signal for oblique LOS—sensitive to co-passing ions. Figure 13(a) reproduces the doublet of the main peak as well as the matching amplitude across the neutral-density filter region. Figure 13(c) matches the spectral shape very well and contains the experimentally-observed shoulder to the left of the main peak. A poor experimental signal-to-noise-ratio (SNR) can be seen in figure 13(d), however, the spectral shapes are still similar. Low SNR is a common challenge throughout the dataset and particularly for the vertical LOS due to their lower optical throughput.

To investigate the wavelength dependencies systematically, both experimental and simulated spectra are averaged

in 10 pixel-wide bins (~ 5 Å) to have a reduced set of data points and to focus the comparison on more general features instead of fine structure. Comparison consists of calculating the Spearman rank correlation coefficient, R , where $R = 0$ means a completely random dataset while $R = 1$ indicates a perfect positive correlation between the two variables. Various parameters are considered to group datasets together to test their correlation as a function of the changing parameter.

Several experimental parameters are varied to produce a dataset of 43 experiments. Each of these 43 cases uses two LOS and therefore provides 86 experimental spectra, which are simulated by P-FIDASIM. The parameters changed to produce these different cases are: fast-ion source beam, LOS, plasma current, plasma current slope, and toroidal magnetic field. In post-experimental data analysis, several other parameters can be considered as well like signal integration time, peak signal level, plasma shape, and information that only comes from the simulation like ion location at the time of emission.

Figure 14 shows the correlation between simulation and experiment. For convenience, simulated spectra are scaled by a single factor for each LOS. This factor is chosen as the average factor required to minimize the χ^2 difference between experiment and simulation for each LOS. Many subsets of data (for a given density range or LOS for example) show a correlation greater than that of the entire set. One example subset is given in figure 14(a). Notice the linear relationship, which suggests that, while the simulation may have systematic differences from experiment, it does reproduce the wavelength dependence well. All datasets together have a correlation of 0.57

Figure 14. Calculated versus measured modulated signal. Both experimental and simulated spectra are averaged in 10 pixel-wide bins (~ 5 Å). Both axes are in arbitrary radiance units. (a) Sample correlation scatter plot showing good correlation when utilizing sightlines P6 & P8. This represents 6 spectra from 3 discharges. (b) The correlation for all 43 discharges (86 spectra) representing many experimental variations in density, temperature, magnetic field, elongation, plasma current, etc.

with subsets reaching 0.80. The parameters considered (and their ranges) are: peak experimental signal (1.3–226.2 a.u.), average experimental signal (0.2–28.5 a.u.), ion source (all but 150-right and 210-left), time-averaged I_p (0.2–1.2 MA), signal integration time (0.1–2.0 s), LOS (P2, P4, P6, P8, P10, V1–2), average vertical position of ions (–0.3–0.43 m), I_p slope (+, –, 0), n_e (0.6 – $2.4 \times 10^{19} \text{ m}^{-3}$) (outbound of the magnetic axis), B_T (–2.1–2.1 T), plasma shape (circle versus oval), and 1 versus 2 X-points. The results show that higher experimental signal levels give better correlations with the simulation. This makes sense because the SNR is higher. Also related to SNR is the signal integration time, for which a better correlation is seen at higher integration times. There is a better correlation for more positive plasma current slopes. There seems to be a ‘preference’ for more triangular and elongated plasmas, which may be from pushing the plasma core closer to the higher neutral-density regions near the X-points. The simulation is more accurate at lower electron density as is seen by a 25% difference in correlation in examining discharges ranging from 0.6 to $2.4 \times 10^{19} \text{ m}^{-3}$. This may again be an issue of SNR as lower electron density means higher neutral density and therefore more FIDA radiation.

5. Passive FIDA bursts from sawteeth

An example of passive FIDA signal produced by sawteeth was already given in figure 4. As shown in that figure, the measured conditionally-averaged spectrum has a spectral shape in qualitative agreement with the spectrum predicted by the simple calculation described in section 3.2. This section shows that the passive FIDA signals have the expected dependence on sightline geometry and that passing ions are more likely to be expelled into the edge region than trapped ions. In addition, a method to quantify the magnitude of the transport is given.

5.1. Dependence on sightline geometry

The spectrum shown in figure 4 is typical of the data for oblique views. Figure 15 compares the amount of blue-shifted light with the amount of red-shifted light for 38 distinct experimental conditions. For the oblique views, the amount of blue-shifted light following the sawtooth crash is usually much larger than the intensity of red-shifted light. This is not the case for vertical views, however. For vertical views, the intensity is similar on both sides of the cold D_α line. The difference occurs because oblique views preferentially detect passing ions, while vertical views preferentially detect trapped ions. Since the passing beam ions are predominately co-going (figure 6), the oblique views detect far more blue-shifted light than red-shifted light in these plasmas.

To confirm that the observed variation is theoretically expected, figure 15 includes the rough estimate using TRANSP and FIDASIM that was shown in figure 4, as well as a similar estimate for a vertical view. The predicted ratio of blue-shifted to red-shifted light is large for the oblique view but near unity for the vertical view, in qualitative agreement with the data.

As expected, the sawtooth-induced signal is larger in discharges with larger sawteeth. In figure 16, the amplitude of the sawtooth is quantified by the reduction in central electron temperature $\Delta T_e/T_e$ and by the sawtooth inversion radius, ρ_{inv} . (The inversion radius, ρ_{inv} , is the normalized minor radius where the ECE signal is unchanged at the sawtooth crash.) To account for variations in neutral density, the sawtooth-induced signal is normalized by the beam-modulation signal. (This normalization is justified in section 5.2.) The sawtooth signal increases with both $\Delta T_e/T_e$ and ρ_{inv} (figure 16). For the blue-shifted portion of the spectrum (figures 16(a) and (b)), the signal is twice as large for the oblique views that are sensitive to passing ions as for the vertical views that are sensitive to trapped ions (average value of 2.4 versus 1.2). For the red-shifted data (figures 16(c) and (d)), the average signals are much smaller: 0.7 for the oblique views and 0.9 for the vertical views. This is expected, as trapped ions cause the majority of the red-shifted signal.

The enhanced *jump* in *edge* signal reported here is consistent with previous active FIDA measurements with the same instruments that showed a larger *drop* in *core* signal for the oblique chords than for the vertical chords at the sawtooth crash [25]. Evidently, larger losses of passing ions in the core appear as larger increases of passing ions in the edge. The stronger effect of sawtooth crashes on passing particles is

Figure 15. A comparison of the blue portion of the spectrum with the red portion for the entire data set of conditionally-averaged sawtooth spectra. Each value is the mean spectral intensity in ranges equidistant from the cold line (6590–6610 Å and 6512–6532 Å). The strong bias towards blue-shifted light for the oblique views is consistent with the use of co-going beams and the greater sensitivity to passing particles. The crude sawtooth spectral model using TRANSP and FIDASIM (section 3.2) also shows a blue bias for the oblique chord. The vertical views are more balanced between red and blue, which is consistent with their greater sensitivity to trapped particles. This spectral balance is further validated by the TRANSP/FIDASIM data point for a vertical chord (which comes from a spectrum analogous to the TRANSP curve in figure 4).

Figure 16. Sawtooth signal relative to beam-modulation signal as functions of sawtooth ((a), (c)) amplitude and ((b), (d)) inversion radius. The data are integrated over ((a), (b)) the blue region of the sawtooth spectrum from 6510–6550 Å or ((c), (d)) the red region from 6590–6610 Å. The ordinate is normalized by the corresponding beam-modulation average from the blue-shifted region.

theoretically predicted [30]. Sawtooth crashes often involve magnetic reconnection and alteration of flux surfaces. Passing particles follow field lines closely and so are affected more by reconnection than trapped particles that traverse many flux surfaces, effectively averaging over the local perturbations. The increased sensitivity of passing ions to the sawtooth instability has also been measured at ASDEX-Upgrade [26] and at TEXTOR [31].

5.2. Quantification of fast-ion transport

Since FIDA light is produced by fast ions, it should be possible to relate measured signals to the fast-ion distribution function. In practice, four ingredients are needed to make this connection. First, one needs an accurate intensity calibration. Second, one needs to know the neutral density. Third, one needs an accurate forward model that relates the distribution function to FIDA spectra. Fourth, one needs sufficient FIDA chords to infer the distribution function.

Progress towards inference of the distribution function from *active* FIDA data has been significant [32]. In the case of active FIDA, the ability to calculate accurately the injected neutral density is well established, as is the forward model embedded in FIDASIM. The relationship between FIDA views and the distribution function is also thoroughly understood [33]. In contrast, for passive FIDA measurements, the neutral density in the edge and pedestal regions have relatively large uncertainties.

An empirical estimate of the magnitude of transport to the edge region is available, however. A readily measured experimental quantity is the ratio of passive FIDA signal at a sawtooth crash relative to the signal produced by a modulated beam, S_{ST}/S_{mod} . The distribution function produced by a modulated beam is another known quantity, consisting of toroidally-asymmetric and axisymmetric-confined populations, so the methods of sections 3.1 and 3.3 can be used to calculate the modulated fast-ion density n_{mod}^{sim} in the edge region. Assuming a simple proportionality between the signals produced by the modulated population and the population expelled by sawteeth, the density of fast ions transported to the edge region by a sawtooth crash is

$$n_{ST} = n_{mod}^{sim} \frac{S_{ST}}{S_{mod}}. \quad (1)$$

The validity of this estimate hinges on two assumptions that may or may not be satisfied in practice. The most important (and questionable) assumption is that the velocity distribution of the expelled fast ions resembles the velocity distribution of the modulated population. The FIDA weight function depends sensitively on energy and pitch so, if the distribution functions differ, the magnitude of the relative signals is altered. As with inference of the distribution function from active FIDA data, comparison over multiple wavelengths with multiple sightlines provides more detail about the true distribution function.

The second assumption concerns the neutral density profile. A FIDA signal depends upon the product of fast-ion and neutral densities $n_f n_n$, integrated over the sightline. Hence,

the calculated fast-ion density n_{mod}^{sim} depends on the assumed neutral profile. It is quite unlikely that the expelled fast ions have the same spatial profile as the modulated fast-ion populations. Nevertheless, since the neutral-density profile is heavily weighted toward the edge, the inferred density n_{ST} can be viewed as an average edge fast-ion density, weighted by the neutral-density profile.

This method of absolute edge density determination is applied to sawtooth crashes for shot 149941, where the average number of fast ions lost per sawtooth is estimated for times between 3485 and 4982 ms using chord P4. Three main components go into this calculation: the sawtooth edge light, the modulated edge light, and the line-averaged first-orbit density along the chord (from P-FIDASIM). The sawtooth light is taken as the integral of the ‘suspected FIDA signal’ in figure 4 after removing the Bremsstrahlung offset. The modulated light is taken as the integral of the modulated spectrum for this discharge. The ratio of signals is $S_{ST}/S_{mod} = 6.1$. The fast-ion density of the toroidally-asymmetric population is estimated by P-FIDASIM. (The contribution of the modulated-confined population is ignored in this rough estimate because, when averaged over ‘beam-on’ and ‘beam-off’ periods, the net axisymmetric signal is small for chord P4.) The estimated line-averaged modulated density is $n_{mod}^{sim} = 7.5 \times 10^{15} \text{ m}^{-3}$ so, using equation (1), the sawtooth-induced edge density is $n_{ST} = 4.5 \times 10^{16} \text{ m}^{-3}$.

This empirical density estimate is reasonable. The TRANSP code includes an optional Kadomtsev model of the sawtooth crash [27]. A representative sawtooth crash in discharge 149941 is simulated using this model. The number of confined fast ions drops 1.7% due to the sawtooth crash. To compare with the empirically-derived density n_{ST} , assume that the measured fast ions are expelled into a volume $V_{loss} = (2\pi R)(\frac{\pi}{4}\kappa a)\Delta r_n$, assuming a 45° poloidal angle of losses, where R is the major radius (2.18 m), κ is the average elongation (1.16), a is the minor radius of the last closed flux surface (LCFS) (0.64 m), and Δr_n is a characteristic neutral density penetration depth (0.21 m). This gives a loss volume of 1.7 m^3 and the number of lost ions, $N_{ST} = n_{ST} V_{loss}$, as 7.6×10^{16} . The total number of confined fast ions calculated by TRANSP is 6.3×10^{18} , meaning that the sawtooth crashes eject on average 1.2% of the fast-ions, in rough agreement with the TRANSP Kadomtsev calculation.

6. Neutral density measurement

This section shows how, in the absence of an independent measurement of the neutral density, passive FIDA data can be used to estimate the neutral density profile. The inferred profile is in reasonable agreement with published L-mode neutral density profiles for DIII-D.

The specific data used to find the neutral density are as follows: discharge number 152817 uses beam 30-left as a modulated fast-ion source at 50% duty cycle. Beams 150-right and 330-right are active at 8% and 50% duty cycle to maintain the equilibrium and, therefore, these times are excluded from the conditional averaging, while the remaining 5 beams are

Figure 17. Calibrated, experimental spectra of passive FIDA from first-orbit ions used to find the neutral density.

turned off. Calibrated spectra (figure 17) are collected by the HoloSpec spectrometer for LOS P1, P3, P5, P7, P9, and P11. Spectra are time-averaged from 1800 ms to 4600 ms. The profiles and equilibrium are taken to be those at 3001 ms. Over the course of the discharge, the electron temperature and density are virtually constant, while the plasma current gradually ramps down. Sawtooth crashes occur throughout.

For simplicity, only the toroidally-asymmetric contribution to the modulated passive FIDA signals is considered. As discussed in section 3.3, P-FIDASIM self-consistently calculates a 2D normalized neutral density based on the plasma profiles. Discharge 152817 has double X-points and therefore a z -dependence is assumed for the neutral density in addition to the ρ dependence, the form and scaling of which will be determined by the following inversion process.

The inversion is an optimization that solves the over-determined least-squares problem that finds the neutral-density scaling (or source) terms, α , as a function of z , that minimize the difference between the theoretical (T) and experimental (E) spectra. The 2D grid of normalized neutral density is then multiplied by this scaling factor to obtain a neutral density that is scaled to realistic values and has a dependence on ρ and z (in practice, R and z).

The inversion has three terms: the unknown neutral density scaling factor α as a function of z , the experimental spectra E as a function of wavelength and LOS number, and the simulated spectra T as a function of wavelength, LOS number, and z . Ignoring the LOS designation, all LOS data are lumped together as one large dataset. The key equation to invert for line-integrated signals then becomes

$$E = T\alpha \quad (2)$$

and inverting gives the optimal neutral density scaling factors:

$$\alpha = T^{-1}E. \quad (3)$$

The inversion is performed using a generalization of a non-negative least-squares solver⁵. The code allows for arbitrary

restrictions to be placed on α ; however, non-negative was the only restriction used in this work. Multiplying the normalized neutral density by α (according to their matching z values) gives the final values of $n_n(R, z)$.

The inversion method was tested and optimized using manually-created test data to determine the optimal number of vertical z -bins that best reproduces an input ‘known’ neutral density within minimal acceptable error limits. The optimization showed the inversion to be more stable when assuming the neutral density rises symmetrically away from the midplane (i.e. that $\alpha(z) = \alpha(|z|)$). Physically, the double X-point of this equilibrium (discharge 152817) suggests increasing neutral density away from the midplane and computationally this makes sense when considering that the lines of sight have similar pathlengths through the upper and lower regions. The inversion, therefore, cannot distinguish between light emitted in the upper regions and light emitted in the lower regions. The inversion method is found to reproduce the manually-created input neutral density with less than 5% error when using four vertical bins. Figure 18 shows the input α for the inversion test along with the output.

Figure 17 shows the experimental spectra used to find the DIII-D neutral density. Figure 19(a) shows the α found by inversion. Note that it rises quickly away from the midplane as hypothesized. The contour plot in figure 19(b) is the final neutral density obtained. The scaled neutral density has realistic values and the z -dependence is made apparent by the low densities closer to the midplane. The figure also shows the average of the contour as a function of ρ , emphasizing the exponential rise of neutral density from the core to the wall. The average neutral density at the LCFS is $8.1 \times 10^9 \text{ cm}^{-3}$ and the average value at the edge ($\rho = 1.3$) is $1.1 \times 10^{11} \text{ cm}^{-3}$. To emphasize the role of the neutral density’s spatial distribution in shaping the fast-ion spectra, figure 20 shows the simulated spectra before and after the vertical weighting of the 1D neutral density profiles. The complicated interplay between the neutral density spatial distribution and the spectral shape and amplitude helps explain how only 6 LOS can be used to infer the neutral density profile.

⁵ www-astro.physics.ox.ac.uk/~mxc/software/bvls.pro

Figure 18. Neutral density scaling factor resulting from a test of the inversion process is less than 5% different from the input value.

Figure 19. (a) Alpha neutral density scaling factor from simulation inversion for discharge 152817. (b) Predicted neutral density from combining simulated neutral density and α . (c) Average of contour plot.

Figure 21 is a comparison between these measurements and previous DEGAS simulations of neutral flux into the core along the LCFS [34] and serves as an order-of-magnitude validation of the calculated neutral density profile. Comparisons [35] between P-FIDASIM and neutral density simulations by Carreras *et al* [36] also show reasonable consistency. Although the plasma conditions for the published profiles differ from the plasma conditions here, the profiles of [34, 36] do come from L-mode discharges and the L-mode case in figure 21 has comparable electron density and temperature: n_e within 9% and T_e within 15%. The P-FIDASIM results have similar spatial dependencies (rising toward the wall and X-points) to

these older works and are within an order of magnitude in neutral density and flux into the core.

7. Summary, outlook, and implications

The DIII-D passive FIDA signals originate from three distinct fast-ion populations. A toroidally-asymmetric population is caused by newly-born ions that traverse the FIDA sightlines on their first orbits. A second population consists of axisymmetric confined fast ions in the edge region. A third population is expelled into the edge region by sawteeth.

Figure 20. Simulated spectra before and after neutral density scaling. Note that, because α is a function of z , the spectra change *shape* after scaling. The ‘original’ simulation is normalized to the peak of the ‘scaled’ simulation.

Figure 21. Data points: neutral deuterium flux into the core for L-mode (circle) and ELMy H-mode plasmas (solid square) as calculated for DIII-D by Groth using the DEGAS code (figure 3 of [34]). Hollow points have X-point fueling while solid points have neutral leakage from divertor. Solid vertical line is single X-point location. Blue curve: the neutral deuterium flux into the core at the LCFS for an L-mode discharge as calculated by P-FIDASIM for discharge 152817 is the neutral density at the LCFS multiplied by the average neutral velocity normal to the flux surface. Dashed vertical lines are double X-point locations. Compare only the right half of the plot as the older work has one X-point while the new work has two and consider only solid circles (L-modes without active X-point fueling). (Reprinted from Groth *et al* 2005 *J. Nucl. Mater.* **337** 425, copyright (2005) with permission from Elsevier.)

In the absence of sawteeth, the observed time evolution of the signal agrees well with the predicted evolution produced by a combination of the toroidally-asymmetric and axisymmetric-confined sources (figure 5). Expected changes in the toroidally-asymmetric population account for the measured dependencies of the spectra on parameters such as sightline or modulated source (figure 14).

Passive FIDA signals produced by sawteeth have the expected spectral shape (figure 4). The signals increase with

increasing sawtooth amplitude (figure 16). The transport of passing fast ions to the edge region is larger than the transport of trapped ions (figure 15).

FIDA signals are proportional to the product of neutral and fast-ion densities integrated over the sightline, $\int n_t n_n d\ell$. (There is also a velocity-space dependence.) This implies that, if one of these two densities is known, the other can be inferred. This paper explores both possibilities. In section 6, the neutral density is inferred from absolutely-calibrated passive FIDA measurements, using the toroidally-asymmetric fast-ion population as a known source. The inferred neutral-density profile is reasonable in magnitude and shape.

Section 5.2 proceeds the opposite way. The passive FIDA signal produced by sawtooth bursts is compared to the signal produced by beam modulation. By assuming that a similar neutral-density profile and fast-ion velocity distribution produce both signals, one infers that $\sim 1\%$ of the fast ions are expelled into the edge region by the sawtooth crash, in reasonable agreement with a TRANSP estimate.

The goal of this paper is to illustrate the potential of passive FIDA measurements for diagnosis of the neutral density and fast-ion transport to the edge region. Many improvements to the work reported here are possible. The neutral-density profile of figure 19 is based on measurements by 6 channels. An obvious improvement is to diagnose the edge region with more chords with multiple viewing angles. As with any tomographic inversion, the more chords, the better the inversion.

More chords with different angles would also benefit the edge fast-ion measurement. With multiple views, the techniques that are being developed to infer the fast-ion distribution function from active FIDA data could be applied to passive FIDA data. Knowledge of the neutral density is also crucial. This could come from passive FIDA measurements with a known source, as described in section 6, but independent determination of the neutral-density profile would be even better.

Since passive FIDA signals contribute to the background, this work has obvious implications for active FIDA diagnostics. Existing FIDA diagnostics treat the background one of three ways: beam modulation, toroidally-displaced reference views, or modeling of the background emission. The

work here has implications for all three. In the case of beam modulation, one can calculate the contribution of the modulated confined population to the signal, as in section 3.1. The toroidally-asymmetric contribution to the background can also be estimated with a code such as P-FIDASIM (section 3.3). This contribution varies with toroidal angle, so it is particularly important to consider when using reference views for background subtraction. Finally, if neither beam modulation nor reference views are employed for background subtraction, calculations of the passive backgrounds from both populations should be employed.

This work also has implications for diagnosis of fast-ion transport into the edge region. ITER plans spectroscopic instruments that will measure light from both helium and deuterium. Instabilities that transport a significant fraction of the confined fast-ion population into the edge region are likely to produce measurable signals. If the neutral-density profile is known, these data can provide quantitative information on the edge alpha and deuterium densities.

Acknowledgments

We thank the DIII-D team for their support and an anonymous referee for stressing the importance of the confined axisymmetric population. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Awards DE-FC02-04ER54698 and DE-FG03-94ER54271. DIII-D data shown in this paper can be obtained in digital format by following the links at http://fusion.gat.com/global/D3D_DMP.

References

- [1] Heidbrink W.W. and Sadler G.J. 1994 *Nucl. Fusion* **34** 535
- [2] Heidbrink W.W., McKee G.R., Smith D.R. and Bortolon A. 2011 *Plasma Phys. Control. Fusion* **53** 085007
- [3] Heidbrink W.W. et al 2011 *Plasma Phys. Control. Fusion* **53** 085028
- [4] Michael C.A. et al 2013 *Plasma Phys. Control. Fusion* **55** 095007
- [5] Geiger B. et al 2013 *Rev. Sci. Instrum.* **84** 113502
- [6] Chen X. et al 2014 *Rev. Sci. Instrum.* **85** 11E701
- [7] Chen X. et al 2013 *Phys. Rev. Lett.* **110** 065004
- [8] Chen X. et al 2013 *Nucl. Fusion* **53** 123019
- [9] Luo Y., Burrell K.H. and Heidbrink W.W. 2004 *Rev. Sci. Instrum.* **75** 3468
- [10] Muscatello C.M., Heidbrink W.W., Taussig D. and Burrell K.H. 2010 *Rev. Sci. Instrum.* **81** 10D316
- [11] Luo Y., Heidbrink W.W., Burrell K.H., Kaplan D.H. and Gohil P. 2007 *Rev. Sci. Instrum.* **78** 033505
- [12] Keilhacker M. et al 1984 *Plasma Phys. Control. Fusion* **26** 49
- [13] Carlstrom T.N. et al 1992 *Rev. Sci. Instrum.* **63** 4901
- [14] Carlstrom T.N., Ahlgren D.R. and Crosbie J. 1988 *Rev. Sci. Instrum.* **59** 1063
- [15] Austin M.E. and Lohr J. 2003 *Rev. Sci. Instrum.* **74** 1457
- [16] Gohil P., Burrell K.H., Groebner R.J. and Seraydarian R.P. 1990 *Rev. Sci. Instrum.* **61** 2949
- [17] Lao L.L., St. John H., Stambaugh R.D., Kellman A.G. and Pfeiffer W. 1985 *Nucl. Fusion* **25** 1611
- [18] Rice B.W., Nilson D.G. and Wroblewski D. 1995 *Rev. Sci. Instrum.* **66** 373
- [19] Von Goeler S., Stodiek W. and Sauthoff N. 1974 *Phys. Rev. Lett.* **33** 1201
- [20] Heidbrink W.W. et al 2016 *Nucl. Fusion* **56** 112011
- [21] Pankin A., McCune D., Andre R., Bateman G. and Kritza A. 2004 *Comput. Phys. Commun.* **159** 157
- [22] Heidbrink W.W., Liu D., Luo Y., Ruskov E. and Geiger B. 2011 *Commun. Comput. Phys.* **10** 716
- [23] Geiger B. 2013 Fast-ion transport studies using FIDA spectroscopy at the ASDEX Upgrade tokamak *PhD Thesis* Ludwig-Maximilians-Universität München
- [24] Nielsen S.K. et al 2011 *Nucl. Fusion* **51** 063014
- [25] Muscatello C.M. et al 2012 *Plasma Phys. Control. Fusion* **54** 025006
- [26] Geiger B. et al 2014 *Nucl. Fusion* **54** 022005
- [27] Kadomtsev B.B. 1975 *Fiz. Plazmy* **1** 710
- [28] Van Zeeland M.A. et al 2010 *Plasma Phys. Control. Fusion* **52** 045006
- [29] Stotler D. and Karney C. 1994 *Contrib. Plasma Phys.* **34** 392
- [30] Kolesnichenko Y.I., Lutsenko V.V., White R.B. and Yakovenko Y.V. 2000 *Nucl. Fusion* **40** 1325
- [31] Nielsen S.K. et al 2010 *Plasma Phys. Control. Fusion* **52** 092001
- [32] Salewski M. et al 2014 *Nucl. Fusion* **54** 023005
- [33] Salewski M. et al 2014 *Plasma Phys. Control. Fusion* **56** 105005
- [34] Groth M. et al 2005 *J. Nucl. Mater.* **337** 425
- [35] Bolte N.G. 2015 Measurement and simulation of deuterium balmer-alpha emission from first-orbit fast ions and the application to neutral density and general fast-ion loss detection in the DIII-D tokamak *PhD Thesis* UC Irvine
- [36] Carreras B.A. et al 1998 *Phys. Plasmas* **5** 2623