

UKAEA-CCFE-PR(23)128

P. Ollus, S. Allan, J.R. Harrison, A. Jackson, T. Kurki-Suonio, K.G. McClements, C.A. Michael, D. Moulton, B.S. Patel, M. Robson, A. Snicker, J. Varje, C. Vincent, the MAST-U team, the EUROfusion WPTE team

Validating the simulation of beam-ion charge exchange in MAST Upgrade

Enquiries about copyright and reproduction should in the first instance be addressed to the UKAEA Publications Officer, Culham Science Centre, Building K1/O/83 Abingdon, Oxfordshire, OX14 3DB, UK. The United Kingdom Atomic Energy Authority is the copyright holder.

The contents of this document and all other UKAEA Preprints, Reports and Conference Papers are available to view online free at scientific-publications.ukaea.uk/

Validating the simulation of beam-ion charge exchange in MAST Upgrade

P. Ollus, S. Allan, J.R. Harrison, A. Jackson, T. Kurki-Suonio, K.G. McClements, C.A. Michael, D. Moulton, B.S. Patel, M. Robson, A. Snicker, J. Varje, C. Vincent, the MAST-U team, the EUROfusion WPTE team

Validating the simulation of beam-ion charge exchange in MAST Upgrade

P. Ollus^{a,*}, S. Allan^b, J. R. Harrison^b, A. R. Jackson^{b,c}, T. Kurki-Suonio^a, K. G. McClements^b, C. A. Michael^d, D. Moulton^b, B. S. Patel^b, M. Robson^b, A. Snicker^a, J. Varje^e, C. Vincent^b and the MAST-U Team¹

^aDepartment of Applied Physics, Aalto University, P.O. Box 11100, 00076 AALTO, Finland

^bUnited Kingdom Atomic Energy Authority, Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK

^cDepartment of Physics, University of Durham, South Road, Durham, DH1 3LE, United Kingdom of Great Britain and Northern Ireland

^dDepartment of Physics and Astronomy, University of California - Los Angeles, Los Angeles, CA 90095-7099, United States of America

^eTokamak Energy Ltd., 173 Brook Dr, Milton Park, Abingdon, UK

*Corresponding author

¹See the author list of J. R. Harrison et al. 2019 Nucl. Fusion, 59 112011

E-mail: patrik.ollus@aalto.fi (P. Ollus)

June 2023

Abstract. Predictions of the impact of charge-exchange (CX) reactions on beam ions in the MAST Upgrade spherical tokamak have been compared to measurements carried out with a fission chamber (neutron fluxes) and a Fast Ion Deuterium-Alpha (FIDA) diagnostic. A simple model was developed to reconstruct the outer-midplane neutral density based on Thomson scattering data for the electron density and temperature, and on measurements of deuterium-alpha emission from edge neutrals. The main computational tools used in this study were the ASCOT orbit-following code and the FIDASIM code for producing synthetic FIDA signals.

The neutral density reconstructed using the simple model is found to qualitatively agree with SOLPS-ITER modelling and to yield a synthetic passive FIDA signal that is consistent with measurement. It was observed that when CX losses of beam ions are accounted for, predictions of neutron emission rates are more consistent with the measurements. It was also necessary to account for CX losses of beam ions in simulations to reproduce the measured passive FIDA signal. The results suggest that the neutral density reconstruction was a good approximation, that CX with edge neutrals causes significant beam-ion losses in MAST Upgrade, and that the ASCOT fast-ion CX model can be used to accurately predict the redistribution and loss of beam ions due to CX.

Keywords: fast ions, neutral density, neutron rate, passive FIDA, MAST Upgrade,

ASCOT, FIDASIM

1. Introduction

Charge exchange (CX) with edge neutrals caused significant losses of neutral-beam ions from the plasma in the Mega Amp Spherical Tokamak (MAST) [1]. CX losses decrease the effective heating power and current drive of beams, and when the fast atoms resulting from CX strike the device wall they can cause damage to sensitive plasma-facing components, impurity sputtering and wall erosion. Based on observations in the first physics campaign (MU01) of MAST Upgrade (MAST-U), significant CX losses of beam ions are suspected in the upgraded device as well. MAST-U is equipped with two neutral beams, one injected in the geometric midplane of the vacuum vessel ("on-axis"), the other around 65 cm above the midplane ("off-axis"). The generation of fusion neutrons by the off-axis beam was lower than that of the on-axis beam, with a difference somewhat larger than that expected from considerations of geometry, plasma density and temperature, and magnetohydrodynamic (MHD) activity [2]. Fast-ion driven mode activity was much lower when only the off-axis beam was on compared to when only the on-axis beam was on. These observations suggest that, in nearly MHD-quiescent plasmas, ions from the off-axis beam, which are more susceptible to being ionized on loss orbits and to CX, have worse confinement than ions from the on-axis beam, with a difference larger than expected from direct orbit losses alone. This motivated detailed analysis of the impact of CX on beam ions in MU01 by means of modelling. In this article, we investigate the extent to which the ASCOT orbit-following code [3, 4] can be used in conjunction with other modelling tools to model the impact of CX with edge neutrals and other background neutrals on beam ions in MAST-U. We also discuss what this investigation reveals about the impact of CX on the beam-ion distribution.

The fast-ion CX model of ASCOT, which models the redistribution of fast ions due to neutralizing CX reactions with background neutrals, has previously been analytically verified and benchmarked against the fast-ion module NUBEAM [5] of the transport code TRANSP [6, 7], and its capabilities have been demonstrated in a predictive MAST-U scenario [8]. In this article, predictions by the ASCOT CX model are compared to experimental results from the MU01 campaign. ASCOT is used in conjunction with the ASCOT Fusion Source Integrator (AFSI) [9] to predict neutron emission from the plasma, and predictions are compared to measurements carried out with a fission chamber on MAST-U [10]. In addition, ASCOT is used in conjunction with the synthetic diagnostics code FIDASIM [11] to simulate passive Fast Ion Deuterium-Alpha (FIDA) spectra, and the simulated spectra are compared to measurements carried out with the MAST-U FIDA system [12, 13, 14]. Comparison to TRANSP modelling is made as well.

The background neutral density in MAST-U, a key input for fast-ion CX simulations, is reconstructed from radial deuterium-alpha (D-alpha) measurements carried out using the HOMER camera [15] and from plasma electron density and temperature measurements performed using a Thomson scattering system [16]. The

model developed for this reconstruction of neutral density was dubbed BATS1D. Due to uncertainty in the spatial calibration of HOMER, its major-radius abscissa is shifted such that the D-alpha emissivity peak aligns with the separatrix, whose location is estimated using Thomson scattering data. Due to uncertainty in the brightness calibration of HOMER, the calculated D-alpha emissivity is scaled based on comparison to D-alpha emissivity estimated by the kinetic transport code KN1D [17, 18] in experiments from the second physics campaign (MU02) of MAST-U. The outer-midplane neutral density reconstructed using BATS1D is compared to modelling carried out using the SOLPS-ITER code [19].

In section 2, the CX models of the previous and the latest versions of ASCOT are benchmarked, both of which were used in analysis. In section 3, a method is presented for reconstructing the outer-midplane neutral density and comparison is made to SOLPS-ITER modelling for MAST-U. In section 4, results are presented on the comparison of predicted and measured neutron emission from an MU01 plasma. In section 5, results are presented on the comparison of simulated and measured passive FIDA signals. In section 6, results are summarized and conclusions are drawn and discussed.

2. Benchmarking charge exchange in ASCOT4 and ASCOT5

The 5th version of the ASCOT orbit-following code (ASCOT5), which is a complete rewrite in the C programming language for the purpose of taking advantage of the parallelization capabilities of modern supercomputers, has previously been benchmarked against the 4th version (ASCOT4 [3]) without the inclusion of CX reactions [4]. The fast-ion CX model of ASCOT4 has previously been analytically verified, benchmarked against the fast-ion module NUBEAM [5] of the transport code TRANSP [6, 7], and demonstrated as a modelling tool [8]. The same model was adapted to and implemented in ASCOT5. In this section, predictions of the ASCOT5 CX model are benchmarked against those of the ASCOT4 CX model.

This benchmark was performed for MU01 shot number 44623: a double-null scenario with 750 kA of plasma current, a conventional divertor configuration and two beams. The chosen time was at 0.39 s when both the on- and off-axis beams were on. First, a TRANSP simulation was run for this case. The TRANSP run was prepared using the OMFIT program [20] based on experimental data from the MAST-U database. The equilibrium was reconstructed using processed data from the MAST-U database, which in turn was originally calculated using the EFIT++ code [21, 22] using magnetic measurements. An EFIT++ reconstruction constrained by measurements carried out using the Motional Stark Effect (MSE) diagnostic was not available. TRANSP recalculates the equilibrium, resulting in minor differences from the original EFIT++ reconstruction, relative differences estimated at 1%. The electron density and temperature profiles were reconstructed based on Thomson scattering data [16] using a fitting tool in OMFIT. The main plasma and beam species was deuterium. A single impurity species was assumed: fully ionized carbon, the plasma having an

effective charge state (Z_{eff}) of 1.5. Both ion temperatures were assumed to be equal to the electron temperature, since the on-axis beam was off and there was no ion temperature measurement for most of this shot. The neutral background was assumed to consist purely of deuterium atoms. The atomic density and temperature inside the plasma were calculated using a neutral model in TRANSP, the analytic neutral transport model FRANTIC [23, 24]. In the scrape-off-layer (SOL), TRANSP assumes an atomic temperature of 0 eV and a constant atomic density when calculating the CX reaction probability. The atomic density was assumed to be $5 \cdot 10^{17} \text{ m}^{-3}$, which is the default value used for MAST-U. A 2D contour designed for use with EFIT++ calculations was used as the first wall.

Inputs for the ASCOT runs were copied over from TRANSP, with a few approximations. TRANSP includes fast ions in its quasi-neutrality condition, while the ASCOT plasma was quasi-neutral with the bulk particles alone, resulting in slightly higher bulk ion densities in ASCOT. Plasma temperatures and densities in ASCOT were assumed to be exponentially decaying in the SOL, such that between the separatrix and $\rho_p = 1.1$ values decrease by a factor of 100. The normalized poloidal flux is defined as $\rho_p = \sqrt{(\psi_p - \psi_{p,\text{ax}})/(\psi_{p,\text{sep}} - \psi_{p,\text{ax}})}$, where ψ_p is the poloidal flux, and $\psi_{p,\text{ax}}$ and $\psi_{p,\text{sep}}$ are its values at the magnetic axis and inside the separatrix, respectively. For these benchmark simulations, the TRANSP neutral density was imported into ASCOT and interpolated using a piecewise cubic Hermite interpolating polynomial [25, 26]. Since ASCOT4 does not support a separate neutral temperature, it was approximated as equal to the ion temperature in both ASCOT4 and ASCOT5. The fast-ion CX process is insensitive to the neutral temperature [8]. In ASCOT, the divertors of the 2D wall were not included because they were not covered by the domain of the equilibrium data from TRANSP, which is a technical requirement in ASCOT. The divertor openings were replaced with virtual, horizontal wall segments, as shown in Fig. 8. Since the detailed deposition of fast ions in the divertor is not investigated, this modification is inconsequential for the modelling reported in this article.

A population of 7155 beam-ion markers, extracted from the TRANSP run, were simulated with the CX model turned on until thermalization in both ASCOT4 and ASCOT5, where "thermalization" was defined as reaching a kinetic energy equal to 1.5 times the local ion temperature or an energy of 100 eV. The full gyro-orbits of the markers were followed. A time step $1 \cdot 10^{-9} \text{ s}$ was used, which is 0.3% of the gyro-period at the outer-midplane plasma edge. To account for statistical error, the random seeds for markers were varied and the simulation for each code version was rerun seven times for a total of eight simulations per code version for the same simulation case.

With eight sets of simulation results per code version, average results were calculated and margins of standard error estimated. ASCOT4 predicts that $0.93 \text{ MW} \pm 0.01 \text{ MW}$ of 2.5 MW, or $37\% \pm 0.4\%$ of beam power deposited in the plasma is lost to the wall, $8\% \pm 0.2\%$ hit the wall as ions and $28\% \pm 0.5\%$ as neutrals. ASCOT5 predicts that $0.91 \text{ MW} \pm 0.01 \text{ MW}$ of 2.5 MW, or $36\% \pm 0.4\%$ of deposited power is lost, $8\% \pm 0.1\%$ hit the wall as ions and $27\% \pm 0.4\%$ as neutrals. In ASCOT4, the

average of the median times for beam particles to hit the wall as a neutral is 0.28 ± 0.008 ms after being ionized in the plasma. In ASCOT5, this time is 0.26 ± 0.01 ms. The ASCOT5 results agree with the ASCOT4 results within margins of standard error.

All 16 simulations were also repeated with the CX model turned off to allow direct comparison of the agreement between ASCOT4 and ASCOT5 with and without the inclusion of CX processes. Figure 1 shows the beam-ion density as a function of ρ_p . Each profile is an average of the predictions from the eight corresponding simulations. The differences between the ASCOT4 and ASCOT5 predictions, with and without CX reactions, are shown on a separate axis. Good agreement is observed in both cases, but the difference is noticeably higher when CX reactions are included. Omitting only the very core and analyzing outside $\rho_p = 0.1$, when CX is turned off, the code versions agree to within 2%. When CX is turned on, the code versions agree to within 4%. The modestly weaker agreement observed when CX is included may be explained by a difference in how atomic reaction data is handled in the two code versions. Unlike ASCOT4, the current version of ASCOT5 does not support non-uniform grid spacing in the imported reaction data. Hence, reaction data had to be approximated onto uniformly spaced grids for ASCOT5 for these simulations. In any case, the agreement in this comparison is good enough to consider the code versions equivalent on the level of precision featured in the analysis reported in this article.

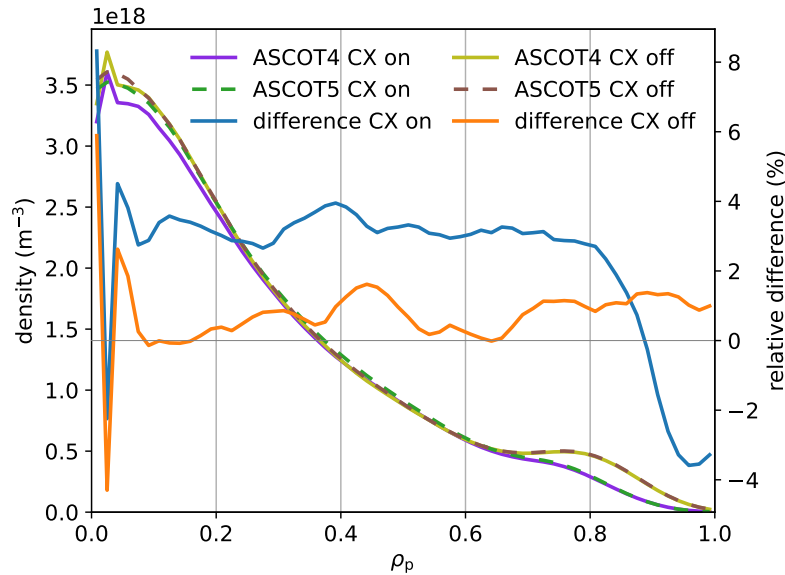


Figure 1. Simulated density of beam ions in MAST-U shot 44623 at 0.39 s as a function of the normalized poloidal flux $\rho_p = \sqrt{(\psi_p - \psi_{p,ax})/(\psi_{p,sep} - \psi_{p,ax})}$, where ψ_p is the poloidal flux, and $\psi_{p,ax}$ and $\psi_{p,sep}$ are its values at the magnetic axis and inside the separatrix, respectively. Each profile is an average of the predictions from eight simulations of the same case, differing only in the random seeds of the markers. Relative differences in the profiles from the two ASCOT versions are shown on a separate axis.

Figure 2 compares the predicted beam-ion distributions, integrated over the whole spatial domain, as functions of energy and pitch when CX is accounted for. Each distribution is an average of the predictions from the eight corresponding simulations. Good agreement is observed. Noticeable deviations are found only in the bottom right corner where the energy is high and the pitch low, which corresponds to the start of a simulation, i.e. just after a particle has been injected into the plasma. On the far-right, where the largest deviations are, statistics are low.




Figure 2. Distribution functions predicted by ASCOT4 (solid) and ASCOT5 (dashed) for beam ions in MAST-U shot 44623 at 0.39 s, integrated over the whole spatial domain. Each distribution is an average of the predictions from eight simulations of the same case, differing only in the random seeds of the markers. Pitch = v_{\parallel}/v , where v_{\parallel} is the velocity component parallel to the magnetic field line and v is the total speed.

Figure 3 shows the spectra for the final kinetic energies of markers that hit the wall when CX was accounted for. Each spectrum is an average of the predictions from the eight corresponding simulations. Estimated margins of standard error are shown with shaded colour bands. The predictions by ASCOT4 and ASCOT5 agree within margins of standard error. The six injection energies featured in the two-beam experiment are indicated in the figure. Peaks in the loss spectra correspond to the injection energies. A significant fraction of the losses are direct orbit losses, but most are from CX, as implied by the charge-specific power losses listed above. Peaks in the loss spectra just below the injection energies are expected, since maximal gyroradii maximize the likelihood of orbits intersecting the wall and expose the beam ions to maximal neutral density. Furthermore, beam-ion populations on orbits intersecting the wall and on orbits exposed to high neutral density are quickly depleted, explaining the sharpness of the peaks.

The two code versions ASCOT4 and ASCOT5 are both used in the analysis reported




Figure 3. Simulated spectra for the final kinetic energies of beam-particle markers that hit the wall in MAST-U shot 44623 at 0.39 s. Each spectrum is an average of the predictions from eight simulations of the same case, differing only in the random seeds of the markers. The shaded colour bands indicate estimated margins of standard error, with colours matching the corresponding average spectrum. Neutral-beam injection energies are indicated using gray dashed lines.

in this article, with emphasis on the CX models of the code versions. The benchmark reported in this section shows good agreement between the code versions in a range of tests representative of analysis of fast ions in the presence of CX reactions in the MAST-U spherical tokamak. The level of agreement is deemed sufficient to render the codes equivalent in the analysis that follows.

3. Reconstructing neutral density in MAST-U

3.1. D-alpha and inverted KN1D

The kinetic transport code KN1D [17, 18] can be used to calculate 1D atomic and molecular density profiles in an ionizing plasma. However, the code requires as input the neutral pressure at the device wall. Since no fast ion gauges (FIG) were operational during MU01 in autumn 2021, measurements of the neutral gas pressure were not available for this analysis. Another method for calculating neutral density was devised using deuterium-Balmer-alpha (D-alpha) measurements and an inverted KN1D algorithm.

Once KN1D has predicted the neutral densities, it estimates the Balmer-alpha emission from atoms excited by electron impact. The rate of photon emission in units

of $\text{m}^{-3}\text{s}^{-1}$ is calculated as [18]

$$R_{\text{ph}} = A_{\text{H}\alpha} \cdot \left(r_{03} + \frac{r_{13} \cdot n_0}{n_{\text{H,Saha},1}} \right) \cdot n_{\text{H,Saha},3}, \quad (1)$$

where $A_{\text{H}\alpha}$ is the spontaneous emission coefficient for protium-Balmer-alpha, which differs from that of D-alpha by only 0.03%, r_{03} and r_{13} are the recombination and ionization coefficients, respectively, from energy level 3, n_0 is the ground-state atomic density, and $n_{\text{H,Saha},1}$ and $n_{\text{H,Saha},3}$ are the Saha equilibrium population densities (m^{-3}) for atomic hydrogen at the energy levels 1 (ground-state) and 3, respectively. The photon emissivity, in accordance with the Bohr model, is converted into emissivity in units of Wm^{-3} as

$$\epsilon = E_1 \cdot \left(\frac{1}{4} - \frac{1}{9} \right) \cdot R_{\text{ph}} \cdot e, \quad (2)$$

where E_1 is the ground-state ionization energy of hydrogen in eV and e is the elementary charge in C. D-alpha emission on the midplane was measured in MU01 using the linear CCD camera HOMER [15], which features horizontal sightlines with tangency points ranging from the SOL through the plasma and through the centre column to the other side of it, allowing for the calculation of a radially resolved D-alpha emissivity profile. This provides the emissivity ϵ in equation (2). Equations (1) and (2) were rearranged such that ground-state atomic density could be calculated from D-alpha emissivity. First the emissivity is converted into a photon emission rate as

$$R_{\text{ph}} = \frac{\epsilon}{eE_1} \left(\frac{1}{4} - \frac{1}{9} \right)^{-1}. \quad (3)$$

Then the ground-state atomic density is calculated as

$$n_0 = \left(\frac{R_{\text{ph}}}{A_{\text{H}\alpha} n_{\text{H,Saha},3}} - r_{03} \right) \cdot \frac{n_{\text{H,Saha},1}}{r_{13}}. \quad (4)$$

The atomic density on the outer midplane in MAST-U can be estimated using eqs. (3) and (4) by inserting Thomson scattering data for the electron density and temperature for the midplane, and D-alpha measurements in the midplane obtained using the HOMER camera. This model, which uses Balmer-Alpha and Thomson Scattering data in inverted equations from the KN1D code, was dubbed as BATS1D. A similar method has been used to estimate the outer-midplane neutral density in NSTX-U [27]. In eq. (4), the electron density and temperature are used to evaluate the parameters r_{03} , r_{13} , $n_{\text{H,Saha},1}$ and $n_{\text{H,Saha},3}$.

Atoms excited by molecular dissociative processes contribute to the D-alpha emission, especially in the SOL [27]. In the above described model, specifically the inverted KN1D equations, the atomic density is estimated based on the assumption that the full emissivity measured by HOMER results from atoms excited by electron impact. Therefore, the calculated atomic density is an upper estimate for the true atomic density. In the NSTX-U study, there was evidence to suggest that a similar

upper estimate tracked closely the true atomic density in the plasma, where molecular density is expected to be low, and that the upper estimate tracked closely the sum of atomic and molecular density in the SOL, where molecular density is expected to be significant [27]. It is reasonable then to assume that a BATS1D-calculated upper estimate for the atomic density is, in fact, an estimate for the total neutral density. Furthermore, at temperatures typical in the plasma edge and SOL, between 1 eV and 100 eV, in an energy range relevant for fast-ion CX in MU01, between 10 keV and 70 keV, the reaction probability is similar for CX with deuterium atoms and molecules. Under these conditions, the reaction rate coefficients for CX between a deuteron and a Maxwellian deuterium atom, and between a deuteron and a Maxwellian deuterium molecule, when the molecule starts in the vibrational ground state and ends up in any vibrational state, agree to within 50%, such that the probability for CX with a molecule is 0–50% higher [8, 28]. Therefore, when simulating beam-ion CX in MAST-U, it would be a good approximation to consider a neutral background consisting of atoms and molecules as a neutral background consisting purely of atoms, assuming an atomic density equaling the sum of the true atomic and molecular densities. The BATS1D model, which includes D-alpha light from all sources, produces an approximation of such a total neutral density that includes both atoms and molecules. It should be noted that the plasma edge and SOL contain molecular ions as well. However, according to KN1D modelling performed for MU02 experiments for this article and according to literature, their density is typically at most one tenth of the neutral density, so they can be neglected [27].

3.2. Working around calibration uncertainties

There was uncertainty in both the spatial and brightness calibrations of the HOMER camera in MAST-U. Therefore, the BATS1D model introduced in section 3.1 could not be directly used. An effort was made to circumvent the uncertainties with ad hoc solutions to allow progress in reconstructing the neutral density.

There was uncertainty in the accuracy of the major radius coordinate R for the HOMER data, i.e. the tangency radii corresponding to the sightlines of the camera channels. A spatial calibration to check if the real view matches the virtual view in the HOMER CAD (computer-aided design) model had not yet been conducted. However, four alternative mappings of the HOMER channels to tangency radii had been calculated. They all appear misaligned, as they would place the D-alpha peak either deep inside the plasma or far outside the separatrix. In a simple virtual reconstruction of the tokamak, made in Matlab independently of the HOMER CAD model, these candidate abscissae were compared to a design schematic that had been used for the installation of the HOMER camera in the vessel wall. The schematic shows two coordinates for the camera, determining its position and viewing direction in the device. For each of the four candidate abscissae, sightlines from the camera lens to the left-most and right-most tangency points, when viewed from above, were calculated in Matlab. Then

the bisector was calculated as an estimate of the central sightline and compared to the central sightline based on the camera schematic. The set of tangency radii whose estimated central sightline best matched the schematic was chosen. The chosen abscissa still placed the D-alpha peak too far outside the plasma. Therefore, the abscissa was shifted such that the D-alpha peak aligned with the separatrix. The exact details of how this shift was calculated are given below.

Typically, an absolute brightness calibration is applied to the raw HOMER data to translate counts registered by the sensor to power emitted by the plasma. Then the inverse Abel transform is applied to the line-integrated brightness signal to acquire the emissivity as a function of major radius. Finally, the signal is multiplied by 4π sr to account for emission in all directions from a plasma volume element. If BATS1D could be directly applied to the obtained emissivity, this would be the extent of the data processing for the brightness measurement. However, there was uncertainty in the brightness calibration. Therefore, extensive comparison of HOMER data and KN1D modelling was performed for MU02 experiments, for which there were neutral pressure measurements, to find a scaling factor to be applied in addition to the above described typical data processing to raise the D-alpha emissivity obtained from HOMER to a level consistent with KN1D modelling.

Unlike in MU01, in MU02 there was an operational FIG to measure neutral gas pressure at the vessel wall. This allowed the use of KN1D to simulate D-alpha emissivity profiles for MU02 experiments, based on the Thomson scattering data for the electron density and temperature and the neutral pressure measured by the FIG. The KN1D prediction could then be compared to HOMER measurements of the same experiments. The configuration of the HOMER camera was the same in MU02 as it had been in MU01.

A number of high-confinement mode (H-mode), low-confinement mode (L-mode) and ohmic MU02 shots and time points were chosen for a study to estimate a scaling factor for the HOMER signal for use in modelling MU01 experiments. The aim was to choose MU02 cases similar to the MU01 cases that are the main focus of this article, in an attempt to maximize comparability and, therefore, the physical accuracy of the resulting scaling factor. Several MU02 cases were chosen for each MU01 case, but single, most similar cases were also selected for each MU01 case for the purpose of having case-specific best-estimate scaling factors for the analysis reported in later sections. Similarity was judged based on plasma scenario, confinement mode, time point during the shot, both line-integrated and full-profile electron density, both core and full-profile electron temperature, fuelling scheme, beam configuration, and line-integrated midplane D-alpha emissivity. Table 1 lists the chosen shots and times as well as the relevant results of the study. The H-mode cases were chosen to be similar to the MU01 H-mode case of shot 44623 at times 0.31–0.33 s and 0.39–0.41 s whose neutron emission is investigated in section 4. Across the investigated time ranges for 44623, the time point 0.32 s, when the plasma is in a relatively steady state, was chosen as the time point for the neutral density reconstruction. Shot 46957 at 0.32 s was selected as the MU02 case most similar

Table 1. A list of the MU02 cases, in terms of shot number and time point, that were used to determine the distance shift needed to correct the spatial calibration of HOMER and the scaling factor needed to complement the brightness calibration of HOMER. The plasma type is indicated as a confinement mode or heating scheme. The fourth column records the distance shift needed to align the HOMER peak with the separatrix location. The fifth column records the ratio between the KN1D-simulated and HOMER-measured D-alpha peak amplitudes.

type	shot	time (s)	shift (m)	ratio
H-mode	46957	0.32	-0.113	5.0
		0.40	-0.112	4.9
	47036	0.32	-0.106	5.6
		0.40	-0.095	5.5
L-mode	46618	0.20	-0.135	3.4
		0.22	-0.118	4.1
	46624	0.35	-0.115	4.1
		0.37	-0.114	5.0
	46828	0.25	-0.122	3.1
	46875	0.35	-0.114	2.2
ohmic	47113	0.35	-0.112	4.1
		0.38	-0.107	3.8
		0.43	-0.108	4.8
		0.54	-0.114	3.7
	47121	0.20	-0.110	3.1

to 44623 at 0.32 s. The L-mode cases were chosen to be similar to the MU01 L-mode case of shot 45091 at time 0.35 s whose passive FIDA is investigated in section 5. Shot 46618 at 0.22 s was selected as the most similar MU02 case. The ohmic cases were chosen to be similar to the MU01 ohmic case of shot 45469 at time 0.40 s for which neutral density estimates by BATS1D and SOLPS-ITER are compared in section 3.3. Shot 47113 at 0.43 s was selected as the most similar MU02 case.

Before comparing D-alpha profiles from HOMER and KN1D to circumvent the uncertainty in the brightness calibration for HOMER, a distance shift for the HOMER data was estimated to circumvent the uncertainty in the spatial calibration for HOMER. The HOMER camera should not move between experiments, so the error in the spatial calibration should be consistent. Therefore, a single distance shift, to be used for all modelling cases in this article, was estimated as the average of the distance shifts needed to move the Abel-inverted HOMER peak to the separatrix in each of the individual MU02 cases. The HOMER peak location was defined as the midpoint of the full width at 95% of the maximum emissivity. To estimate the major-radius coordinate of the separatrix at the outer-midplane, $R_{\text{sep,OMP}}$, the upstream electron temperature at the separatrix, $T_{\text{e,sep}}$, was estimated. To calculate $T_{\text{e,sep}}$, the expected parallel electron heat flux density is integrated along the separatrix derived by EFIT++ (see for example

eq. (4) in [29]). The $R_{\text{sep,OMP}}$ was then assumed to be where the linearly interpolated Thomson scattering electron temperature profile crosses the estimated $T_{\text{e,sep}}$. When applying this model, the power crossing the separatrix was assumed equal to the ohmic heating power, which causes slight underestimation of $T_{\text{e,sep}}$ in the cases with one or two beams. The error introduced by this was estimated for shot 44623 at 0.32 s, where TRANSP predicts 0.7 MW of input ohmic power and 1.5 MW of deposited beam power. In this case, approximating the power crossing the separatrix as equal to the sum of ohmic and beam heating power instead of only ohmic heating power gives a result three times as large. The expected value of $T_{\text{e,sep}}$ then becomes $3^{2/7} \approx 1.4$ times as high. This translates into a 3 mm smaller $R_{\text{sep,OMP}}$. Moreover, the radial decay length for the heat flux in the SOL was assumed to be 10 mm, a common default value based on studies performed for MAST data [30, 31]. The sensitivity of $R_{\text{sep,OMP}}$ on this parameter was tested for shot 46957 at 0.32 s. When the decay length was varied between 5 mm and 15 mm, $T_{\text{e,sep}}$ varied between 33 eV and 24 eV, which translates into a variation in $R_{\text{sep,OMP}}$ of 4 mm. The spatial coordinates of the Thomson scattering data at the separatrix are precise to 1 cm. The distance shifts needed to move the Abel-inverted HOMER peak to the separatrix in the investigated MU02 cases are listed in Tab. 1. The average distance shift is -11.3 cm along the major radius and was applied to HOMER data used in modelling in this article. Since the above estimated uncertainty in $R_{\text{sep,OMP}}$ is small compared to the spatial uncertainty in the Thomson scattering data and compared to the variations in the distance shift and D-alpha peak ratio in Tab. 1, error margins are not calculated based on the uncertainty in $R_{\text{sep,OMP}}$.

For each MU02 case, the D-alpha emissivity profiles measured by HOMER on the outer midplane were compared to profiles predicted by KN1D. The inputs for the KN1D simulations were the scalar neutral-gas pressure measured by the FIG at the vessel wall, time-averaged 10 ms in each direction, as well as the electron density and temperature profiles measured using Thomson scattering. The main result of each comparison is the ratio between the peak values of the KN1D and HOMER D-alpha profiles at the separatrix. The ratios are listed in Tab. 1. The minimum and maximum ratios were 2.2 and 5.6, respectively. These will be used as limiting-estimate scaling factors for HOMER D-alpha measurements in the analysis in later sections, the expectation being that the true D-alpha value is somewhere between these limits. In addition, best-estimate scaling factors were selected for the different MU01 cases based on their most similar MU02 counterparts, which were mentioned above. The best-estimate scaling factor for shot 44623 at times 0.31–0.33 s and 0.39–0.41 s is 5.0, for shot 45091 at time 0.35 s it is 4.1, and for shot 45469 at time 0.40 s it is 4.8.

The calculated ratios between D-alpha peaks do not reflect possible mismatch in the exact peak locations, but it was confirmed that HOMER and KN1D are in reasonable agreement about the peak locations along the major radius. Using the same definition as above for a D-alpha peak location, in 12 of the 15 cases, the two peaks were within 1 cm both of each other and of the separatrix estimate. In two cases, the peaks were within 1 cm of each other but within 2 cm of the separatrix, outside it in one of the cases

and inside it in the other. In one case, the peaks are within 1 cm of the separatrix, on either side of it, with almost 2 cm between the peaks. These observations imply that the distance shift calculated for and applied to the HOMER data, -11.3 cm, is a successful shifting correction to the error in the spatial calibration. However, a possible error in the spacing between the HOMER R points is not corrected by the shift. Indeed, when comparing HOMER and KN1D D-alpha profiles, it was observed that the HOMER peak tended to be wider, which could be explained by an overestimation of the spacing between the HOMER sightlines.

The KN1D prediction of D-alpha emissivity only considers atoms. As discussed above, molecules contribute to D-alpha as well, and the HOMER measurement is expected to include this. However, based on the KN1D modelling and SOLPS-ITER modelling, discussed more in section 3.3, atomic density is typically about an order of magnitude higher than molecular density at the separatrix. In the SOL, the molecular density typically becomes higher than the atomic density within a few centimeters of the separatrix and settles at a level about an order of magnitude higher than the atomic density. Since the scaling factors calculated for the HOMER emissivity signal are based on the D-alpha peaks at the separatrix, error from the omission of D-alpha from molecules in KN1D is expected to be low. What error there is will lead to an underestimation of the total neutral density and its impact on fast ions, as far as this part of the reconstruction is concerned. Since the whole HOMER D-alpha emissivity profile is scaled based on the peak ratio at the separatrix, a molecular contribution to the HOMER profile in the shape of a thicker tail in the SOL will be included. However, the signal drops rapidly in the SOL with the drop of the electron density. The reconstruction of the neutral density in the SOL is further discussed in section 3.4. A molecular contribution is another possible explanation for why the HOMER peaks tend to be wider in the SOL than the KN1D peaks. The CX losses due to molecules interacting with the fast ions is further underestimated since the molecular density that is accounted for is approximated as additional atomic density. At the time of analysis, the models for the neutral background in ASCOT4 and ASCOT5 were limited to atoms, for which the CX reaction probability is modestly lower than for molecules, as was explained in section 3.1.

When comparing HOMER and KN1D D-alpha profiles, two additional noteworthy observations were made. Looking carefully at the beam timings and comparing the HOMER D-alpha from the different MU02 cases, there is evidence that the interaction between beam particles, especially those from the off-axis beam, and background neutrals generates enough additional D-alpha to significantly inflate the D-alpha peak measured by HOMER, both in height and in width. The transmission function for the wavelength filter of HOMER is approximately a Gaussian where the central wavelength is 655 nm and the full width at half maximum is 10 nm, which means that Doppler-shifted fast-ion D-alpha is not excluded. Hence, a beam-ion contribution is indeed expected in the HOMER measurement. KN1D does not simulate any fast-ion contribution. Therefore, this is a third possible explanation for why the HOMER D-

alpha peak tends to be wider than the one simulated by KN1D. This could also explain the higher peak ratios obtained for H-mode compared to L-mode, as shown in Tab. 1, since the fast ions make a larger relative contribution in L-mode where the bulk density is lower at the edge. However, it should be noted that the ohmic cases, where there is no beam-ion contribution, also yield lower ratios than the H-mode cases. An impact by beam particles was observed also in the line-integrated midplane D-alpha measurements. The other noteworthy observation is that measurement noise in the SOL electron density and temperature from Thomson scattering, which was allowed to propagate through to the KN1D inputs, can cause KN1D to predict significant D-alpha in the SOL, which, in turn, results in a lower peak at the separatrix. This could cause KN1D to underestimate the D-alpha peak at the separatrix.

There are a number of details about the Thomson scattering and HOMER data used in this work that were not mentioned above but should be noted. In all the analysis done related to reconstructing the neutral background, HOMER brightness data were time-averaged 10 ms in each direction, and Thomson scattering electron density and temperature data were time-averaged 20 ms in each direction. Thomson scattering data points were considered spurious and filtered out if they had a value more than twice the value of the previous point when moving radially outward along the outer midplane. Beyond this, noise in the Thomson scattering measurements, which is substantial especially in the SOL, was allowed to propagate through to the neutral density calculations. The used Thomson scattering data were interpolated at the HOMER R abscissa, which has significantly higher spatial resolution, for the final insertion into the inverted KN1D equations. The HOMER camera is located 18.3 cm above the midplane. This is significant since the separatrices of equilibria in MU01 and MU02 plasmas tend to have a bulge at the outer midplane. The value of the separatrix R coordinate on the HOMER plane is typically 2–3 cm smaller than on the Thomson scattering plane, which is at 1.5 cm above the midplane, i.e. essentially on the midplane. This vertical misalignment of the HOMER and Thomson scattering sightlines was accounted for by mapping the R coordinates from the HOMER plane via the flux coordinate ρ_p , which was based on an EFIT++ equilibrium reconstruction, to the Thomson scattering plane. This mapping was performed in the opposite direction for the separatrix location, which was calculated based on Thomson scattering data as explained above, when it was compared to the HOMER peak location to calculate the distance by which to shift HOMER data.

3.3. Comparing to SOLPS-ITER

To test the BATS1D model, including the extra corrections described in section 3.2, it was compared to SOLPS-ITER modelling performed for ohmic L-mode shots from MU01 [19]. The case modelled in SOLPS-ITER was a double-null L-mode shot with a toroidal magnetic field strength of 0.5 T on the magnetic axis, plasma current of 650 kA and ohmic heating of 0.6 MW. Fuelling was done from the high-field side, and the

outer-midplane separatrix density was $1.0 \cdot 10^{19} \text{ m}^{-3}$, corresponding to a fairly high-density L-mode from MU01. The SOLPS-ITER simulations are 2D in the poloidal plane, with the grid based on the full equilibrium reconstruction by EFIT++. The SOL plasma is modelled as a fluid while the neutrals are modelled kinetically. Various neutral reactions are considered, including recycled molecules and their breakup. The model is considerably more complete than BATS1D and, therefore, serves as a good benchmark. However, due to the lack of pressure gauges in MU01, the SOLPS-ITER modelling is not experimentally constrained and the absolute values of the neutral densities predicted cannot be assumed physically accurate. In that sense, the BATS1D modelling for MU01 is more experimentally relevant, since the D-alpha profile is scaled based on KN1D modelling constrained by pressure measurements from MU02. Indeed, as is further discussed below, it is more useful here to compare BATS1D to SOLPS-ITER in terms of the shape of the neutral density profile rather than its magnitude.

One of the MU01 shots on which the SOLPS-ITER modelling is based, and the case chosen to be modelled using BATS1D for this comparison, is shot 45469 in the time range 0.25–0.75 s. Instead of simulating a single time point, SOLPS-ITER runs to steady state. For the BATS1D modelling, the time 0.40 s was chosen because of the relatively steady state of the plasma. Figure 4 shows the data that BATS1D uses to estimate the neutral density: the outer-midplane Thomson scattering electron density and temperature data as well as the HOMER signal that has been absolutely calibrated, Abel-inverted, shifted, mapped to the Thomson scattering plane and scaled to complement the uncertain absolute calibration. From a set of SOLPS-ITER simulations where the assumed fuelling rate was varied, that case was chosen whose electron density at the outer-midplane separatrix best matches the Thomson scattering measurement for shot 45469 at 0.40 s.

The BATS1D and SOLPS-ITER neutral density estimates near the separatrix are shown in Fig. 5. SOLPS-ITER estimates both the atomic and molecular deuterium densities. Both the separate profiles and the total neutral density are shown. As explained in section 3.2, BATS1D, scaled based on KN1D modelling for MU02, includes molecules a limited distance into the SOL if the molecular density at the separatrix is small, which SOLPS-ITER suggests it is in this case. Therefore, the BATS1D result should be interpreted as a total neutral density and compared to the total neutral density estimate by SOLPS-ITER. However, the Thomson scattering noise plays a large role in the shape and maximum value of the BATS1D profile outside the separatrix. The BATS1D estimate, even its upper limit, is significantly lower than the SOLPS-ITER neutral density, by a factor of two or more at the separatrix. In the range 4–5 cm outside the separatrix, the difference approaches an order of magnitude, but in the range 5–6 cm there is a sharp increase in the BATS1D profile, which returns the difference to about a factor of two. Further out, the BATS1D profile starts dropping with the electron density. As explained, the SOLPS-ITER modelling is not experimentally constrained and could be overestimating the neutral density. A more useful comparison to make is that of the profile slopes. When plotted on a logarithmic scale, the slope of the BATS1D neutral density profiles agrees with that of SOLPS-ITER at the separatrix,

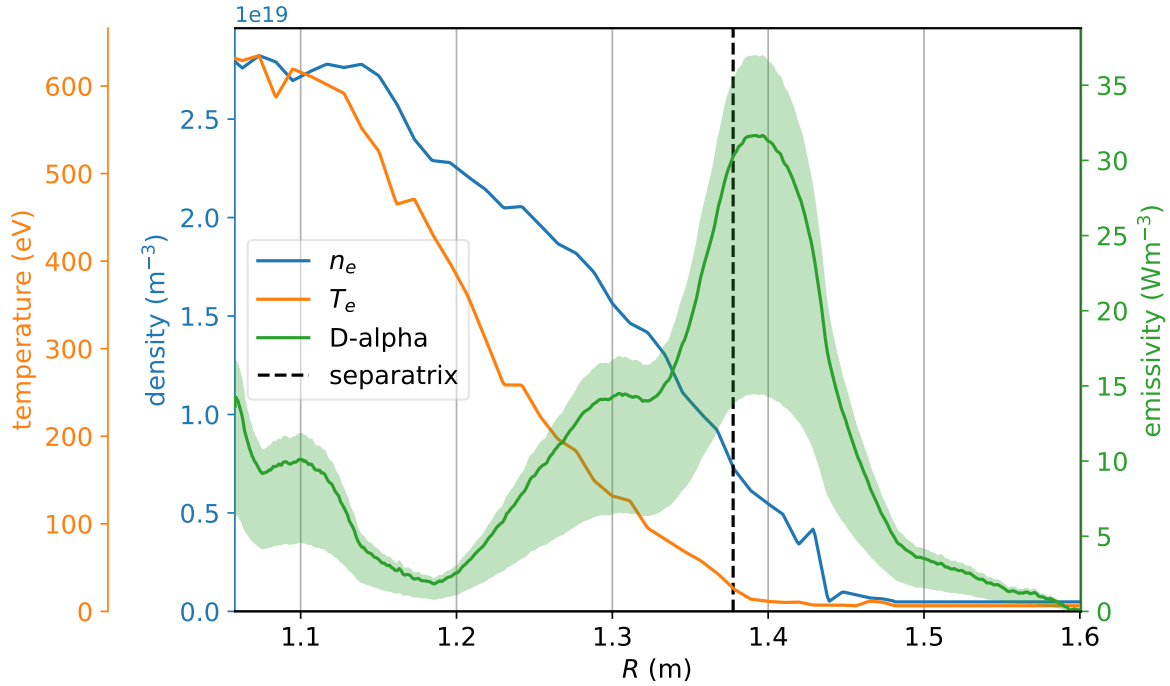


Figure 4. Outer-midplane HOMER D-alpha and Thomson scattering electron density n_e and temperature T_e data used for estimation of the neutral density in shot 45469 at 0.40 s, plotted as functions of major radius. For the D-alpha profile, the result when scaling with the best-estimate scaling factor for this case, 4.8, is shown with the solid green line. The results using the minimum and maximum scaling factors, 2.2 and 5.6, respectively, are indicated using a shaded green band.

and a few centimeters into the plasma. This implies that the BATS1D model is able to correctly capture the steep gradient of the neutral density at the outer-midplane separatrix in a MAST-U plasma. Deeper inside, the BATS1D slope levels out while that of SOLPS-ITER remains steep. This implies that the BATS1D prediction becomes unreliable inside the plasma. This may be explained by how the edge is very bright, but little D-alpha from deeper inside the plasma reaches the HOMER lens. Via an imperfect Abel-inversion of the line-integrated brightness, the neutral density may then be overestimated deeper inside the plasma. This observation is taken into account in section 3.4, when constructing the full neutral density profile to be used in ASCOT and other modelling.

3.4. Constructing neutral density profiles

BATS1D and the TRANSP internal neutral model FRANTIC were used to reconstruct the outer-midplane neutral deuterium density in two MU01 shots: in shot 44623 at times 0.31, 0.32, 0.33, 0.39, 0.40 and 0.41 s, and in shot 45091 at time 0.35 s. The neutral density estimates were used in ASCOT4 modelling of shot 44623 to compare predicted and measured neutron rates, which is reported in section 4, as well as in ASCOT5 and




Figure 5. Estimates by BATS1D and SOLPS-ITER of the outer-midplane neutral deuterium density near the separatrix in shot 45469 at 0.40 s as a function of major radius minus the major-radius coordinate of the separatrix. For SOLPS-ITER, atomic (n_D) and molecular (n_{D2}) densities are also shown separately. For BATS1D, the result when scaling with the best-estimate scaling factor for this case, 4.8, is shown with the solid blue line. The results using the minimum and maximum scaling factors, 2.2 and 5.6, respectively, are indicated using a shaded blue band.

FIDASIM modelling of shot 45091 to compare predicted and measured passive FIDA signals, which is reported in section 5. The final neutral density profiles for shot 44623 at the example time 0.40 s and for shot 45091 at time 0.35 s are shown in Figs. 6 and 7, respectively, and the reconstruction process is explained in detail below.

Due to an interruption in the camera operation, there are no HOMER data for shot 44623. To circumvent this problem, the similar case of shot 44578 at time 0.32 s was used as a proxy. Similarity was judged in the same way as in section 3.2. Thomson scattering electron density and temperature as well as HOMER D-alpha emissivity data from shot 44578 at time 0.32 s were used in BATS1D to obtain an approximate reconstruction of the neutral density in 44623 at time 0.32 s, which is used for the modelling of all the six time points of interest. There are HOMER data for shot 45091, so BATS1D could be applied directly for the case of shot 45091 at time 0.35 s. The part of the neutral density profile that was calculated by BATS1D is shown in blue in Figs. 6 and 7. Especially in shot 45091, the Thomson scattering noise, which was allowed to propagate through the calculation, gives the profile an erratic and unphysical shape. However, beam-ion gyroradii are typically large compared to the spatial scale of the irregularities in the profile, meaning that the impact of the noise is averaged out. Each reconstruction was calculated with all three alternative D-alpha profiles resulting from using the three different scaling factors estimated in section 3.2: the best-estimate factor




Figure 6. Outer-midplane neutral deuterium density reconstruction for shot 44623 at time 0.40 s. BATS1D part based on proxy shot 44578 at 0.32 s. Extrapolation of maximum BATS1D value in yellow. Best-estimate profile shown using solid lines, limiting-estimate profiles using shaded colour bands. TRANSP-part version that was not scaled to BATS1D part is shown using dashed orange line.

and the limiting-estimate factors. This results in three distinct neutral density profiles for each case: a best-estimate profile, shown using solid lines in Figs. 6 and 7, and limiting-estimate profiles to function as error margins, shown using shaded colour bands.

Based on the comparison to SOLPS-ITER in section 3.3, the BATS1D neutral density profile is deemed invalid deeper than a few centimeters inside the plasma. For the plasma region deeper than the BATS1D domain of validity, the corresponding TRANSP-estimated atomic density profile was used in a manner conditional to the BATS1D profile. These TRANSP parts of the profiles are shown in orange in Figs. 6 and 7. For shot 44623, the full atomic density from TRANSP was used, which means that atoms were included from all three of the simulated sources: atoms recycled from the wall, beam-halo atoms and atoms born in recombination reactions in the plasma. Since the analysis of shot 45091 is based on measurements by the passive FIDA diagnostic, whose sightlines are on the opposite side of the plasma from the neutral beam sightlines by design, the beam-halo atoms were omitted from the TRANSP atomic density profile for this shot. Furthermore, for a more comprehensive investigation, two of the remaining alternative versions of the TRANSP atomic density were used: one with only wall atoms, and one with wall atoms and recombination atoms. In Fig. 7, the best-estimate profile for the version with only wall atoms is shown using a dashed orange line, while the solid orange line and shaded orange colour band show the version with wall and recombination atoms according to the plotting scheme explained above. For the data shown in Fig. 5, the domain of validity of the BATS1D neutral density profile in the

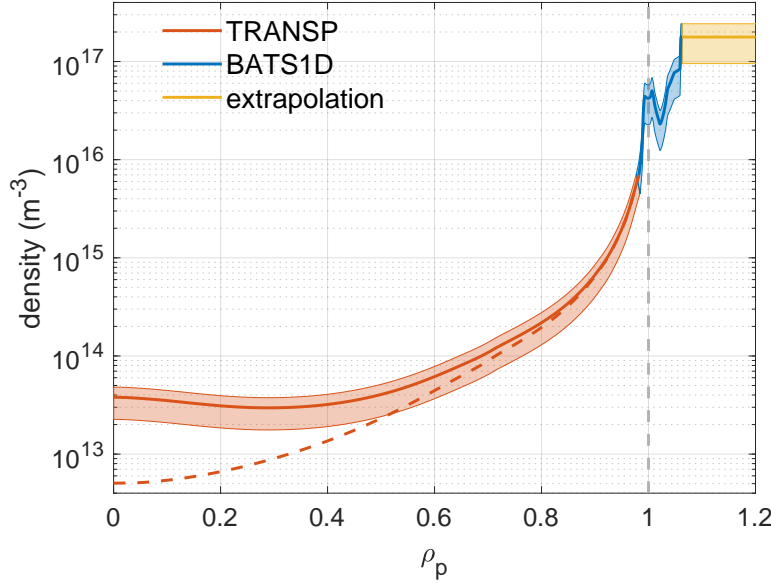


Figure 7. Outer-midplane neutral deuterium density reconstruction for shot 45091 at 0.35 s. Extrapolation of maximum BATS1D value in yellow. Best-estimate profile shown using solid lines, limiting-estimate profiles using shaded colour bands. TRANSP-part version with wall and recombination atoms shown using solid orange line and orange colour band. Best-estimate TRANSP-part version with only wall atoms is shown using dashed orange line.

plasma edge inside the separatrix is assumed to end where the slope diverges from that of the SOLPS-ITER profile, when the profiles are plotted on a logarithmic scale. For shots 44623, or its proxy shot 44578, and 45091, SOLPS-ITER modelling has not been performed. Instead, the TRANSP neutral model is used as the point of comparison for BATS1D. The BATS1D profile is cut off and replaced by the TRANSP profile at the point where the slopes of the BATS1D and TRANSP profiles match best. This point of matching slopes is searched from the region between $\rho_p = 0.95$ and $\rho_p = 0.99$, which is where the inner limit of the validity of the BATS1D model is expected to be based on the comparison to SOLPS-ITER. While the BATS1D estimate for the time point 0.32 s was used for the BATS1D parts of the neutral density profiles for shot 44623 at all six time points, the TRANSP parts of the profiles were taken from the correct time points in the TRANSP simulation for shot 44623. At all six times of the TRANSP simulation for shot 44623, and with all three scaling factors for the BATS1D reconstruction for shot 44578 at 0.32 s, the profile slopes best match at $\rho_p = 0.97$. For shot 45091 at time 0.35 s, for all six combinations of the two versions of the TRANSP profile and the three scaling factors for the BATS1D reconstruction, the profile slopes match best at $\rho_p = 0.98$. After choosing the junction, the TRANSP profiles were scaled with a constant factor such that they equal the corresponding BATS1D profile at the junction.

The described use of the TRANSP neutral density deeper inside the plasma has not been validated thus far. However, since most CX losses of beam ions result from

neutralizations in the SOL and the very edge of the plasma, the part of the neutral density profile that is based on TRANSP is expected to be of little consequence for the CX losses. In other words, the choice of the neutral density profile deeper inside the plasma is not expected to have a significant impact on the beam-ion confinement, as long as the neutral density is not greatly overestimated. To test this, versions of the best-estimate neutral density profiles for shot 44623 were constructed where the TRANSP part was not scaled to the BATS1D part. In the example of shot 44623 at 0.40 s in Fig. 6, this is shown using a dashed orange line. The impact of using these profiles instead of those where the TRANSP part is scaled are reported in section 4. In the analysis of shot 45091, neutrals impact not only the redistribution of beam ions but also the generation of FIDA light. That makes the analysis more sensitive to the reconstruction of the neutral density deeper inside the plasma, but also provides an opportunity to validate the reconstruction. This is further discussed in section 5.

In the SOL, the validity of the BATS1D model ends where the electron density drops to near-zero values. For the neutral density reconstruction, the profile calculated by BATS1D was used up to its maximum value. Further out, this maximum value was used as a constant extrapolate. This extrapolated part of the profile is shown in yellow in Figs. 6 and 7. Both for shot 44578 at 0.32 s and 45091 at 0.35 s, the maximum is reached within 5 cm of the separatrix at the outer midplane, where 5 cm outside the separatrix corresponds to about $\rho_p = 1.1$. According to the KN1D modelling done for MU02 cases in section 3.2, the neutral density profile flattens out somewhere in the range 5–10 cm outside the separatrix, which approximately corresponds to the range $\rho_p = 1.1$ –1.2. Based on these observations, the extrapolated parts of the reconstructed neutral density profiles are underestimations. This SOL region is relevant, since beam ions on the outer midplane in MAST-U have gyroradii of up to 10 cm.

The neutral density distributions constructed as described above and used in the modelling reported in sections 4 and 5 were assumed to be poloidally uniform. Each reconstructed outer-midplane density was used for the whole plasma and SOL as a 1D density function along ρ_p . The assumption of poloidal uniformity is probably not correct, as strong recycling from the divertors and neutral gas from the fuelling valves are likely to yield a poloidally non-uniform neutral density distribution. However, because of a pronounced Shafranov shift, strong orbit drifts and strongly major-radius-dependent gyroradii, fast ions in MAST-U have orbits that are weighted towards the low-field side of the plasma [8]. This is illustrated in Fig 8, which shows the beam-ion distribution in the (R, z) plane in shot 44623 at 0.39 s when both the on- and off-axis beams are on, as predicted by ASCOT4 when CX reactions are not accounted for. The beam-ion density is essentially zero in the very top and very bottom of the plasma as well as near the edge on the high-field side. This is expected to limit the interactions of the beam ions with the higher neutral densities towards the divertors and high-field-side fuelling valves. The beam-ion population overlaps with the relatively high neutral density levels of the plasma edge and SOL on the low-field side, in particular on the outer midplane, which is where the neutral density reconstruction was calculated. Finally, the

passive FIDA diagnostic, used for the analysis of shot 45091, views the outer midplane. These observations imply that error introduced by assuming a poloidally uniform neutral density is significantly mitigated. Therefore, the approximation is deemed good for the analysis performed in this article.




Figure 8. Beam-ion density predicted by ASCOT4 in shot 44623 at 0.39 s when CX reactions are not accounted for. Equilibrium contours for $\rho_p = 0.2, 0.4, 0.6, 0.8$ and 1.2 in gray and for $\rho_p = 1.0$ in red, and the 2D wall representation in black are indicated.

4. Comparing predicted and measured neutron rates

During the MU01 campaign, a fission chamber was operated outside the tokamak in toroidal sector 11 to measure the total neutron emission rate from the plasma [10]. Neutron rate predictions made by combined modelling using ASCOT4 and AFSI were compared to fission chamber measurements from shot 44623 to validate the ASCOT4 CX model and to investigate the impact of CX on beam-ion confinement. As a benchmark exercise, ASCOT4-AFSI and TRANSP predictions were also compared. The 4th version of ASCOT was chosen for this study since the corresponding AFSI version was ready for use and has been validated and widely used in previous work [9, 32, 33, 34].

Shot 44623 is interesting since there is a long, MHD-quiescent period 0.10–0.34 s when only the off-axis beam is on, followed by a period 0.34–0.42 s when both the on- and off-axis beams are on. Fast ions from the off-axis beam are deposited closer to the plasma edge, meaning they orbit closer to the device wall and are more susceptible to

CX with background neutrals. Hence, we expect direct orbit and CX losses to impact the neutron contribution from the off-axis beam more than that from the on-axis beam. Even without considering losses, fast ions from the on-axis beam are expected to be more effective at producing neutrons, since they are on average deposited closer to the plasma core where they have longer slowing-down times due to the higher electron temperature. In addition, the neutron rate from reactions between two beam ions is higher for ions from the on-axis beam because of the smaller volume of the core. The relative distances from various plasma regions to the fission chamber, the centre of which is on the midplane level, can also play a role, especially in highly elongated plasmas. Similar to the temperature profile, the density often peaks in the core, which would further increase neutron production in the core relative to other regions. However, in the case of shot 44623, the density profile is hollow in the analyzed time range, reaching its highest value between $\rho_p = 0.7$ and 0.8 .

There is uncertainty in the absolute calibration of the fission chamber. ASCOT4-AFSI and TRANSP both predict neutron rates that are higher by a factor of 2–4 than those measured by the fission chamber for shots of MU01. Based on a study comparing neutron rates measured by the fission chamber and rates predicted by TRANSP in MU01, experts at MAST-U derived an additional scaling factor of 2.6 for the fission chamber, which was applied to measurements from MU02. Since TRANSP commonly underestimates fast-ion losses, the experts expect this scaling factor to cause an overestimation of the measured neutron emission. This scaling factor was applied to fission chamber measurements from MU01 for the analysis reported in this article.

The ASCOT4-AFSI-estimated neutron rate, with and without CX reactions included in the ASCOT4 simulation, was compared to the rate measured by the fission chamber. As a code benchmark exercise, the ASCOT4-AFSI-estimated neutron rate was also compared to the rate estimated using TRANSP, with the TRANSP neutral density being used in ASCOT4. ASCOT4 simulations, which are typically performed for single time points when the background plasma parameters and magnetic field are assumed to be in a steady state, were performed for six time points: 0.31, 0.32, 0.33, 0.39, 0.40 and 0.41 s. The ASCOT4 simulations were prepared based on the same TRANSP simulation that was used for the benchmark in section 2. In fact, the ASCOT4 case used in section 2 was one of the cases modelled here, the case at time 0.39 s. In the beam-ion populations extracted from TRANSP for the six time points, the number of markers varied between 7155 and 14405. The exception to using TRANSP inputs in ASCOT4 in the neutron analysis reported in the present section was the neutral background input. The neutral density was estimated using BATS1D when comparing ASCOT4-AFSI neutron prediction to measurement. The neutral temperature was approximated as equal to the ion temperature, since ASCOT4 does not support a separate neutral temperature. Neutral and ion temperatures are typically estimated to be similar in TRANSP, and the fast-ion CX process is insensitive to the neutral temperature [8]. According to the Thomson scattering measurements, the electron density profile was hollow for all time points analyzed. Since the on-axis beam was turned on at 0.34

s and ASCOT4 estimates average beam-ion slowing-down times of 10 and 20 ms at times 0.33 and 0.39 s, respectively, it was assumed that the plasma had reached steady-state conditions at 0.39 s, 50 ms after the on-axis beam was turned on. Therefore, the application of ASCOT4 modelling was deemed valid.

Figure 9 compares the ASCOT4-AFSI and TRANSP neutron emission predictions in the case where the TRANSP neutral density is used also for the ASCOT4 CX model. The ASCOT4-AFSI predictions agree with TRANSP to within 7%.




Figure 9. Total neutron rate from the plasma in shot 44623 as predicted by ASCOT4-AFSI and TRANSP using the TRANSP neutral density and including CX.

Before comparing modelling results to neutron measurements, the sensitivity of the ASCOT4 simulation of beam-ion CX on the neutral density deeper inside the plasma was tested. This was done by use of alternative neutral density profiles, whose construction was explained in section 3.4. For each of the six time points, ASCOT4 simulations with CX included were compared for two variants of those neutral density profiles where the HOMER-measured D-alpha emissivity was scaled with the best-estimate factor of 5.0: one variant where the density deeper inside the plasma, the part that comes from TRANSP, is scaled to the edge density profile, the part calculated by BATS1D, and one variant where it is not. Examples of these variants are shown in Fig. 6 using a solid and a dashed orange line, respectively. Only minor differences were observed in the simulation results. Taking as an example the simulation of the time point 0.39 s, the predicted fast-ion density profiles as functions of ρ_p agree to within 3% outside $\rho_p = 0.3$. In the core, the fast-ion density is at most 8% lower for the case where the TRANSP neutral density was scaled. When the neutral density deeper in the plasma was scaled, the loss of beam power is 3% higher. Throughout the six time points simulated, the predicted neutron rates agree to within 1% when using the two variants of the neutral

density profile. In the following analysis, those neutral density profiles were used where the TRANSP part was scaled to the BATS1D part.

In the ASCOT4 simulations, of the beam power deposited in the plasma, i.e. power in beam particles ionized in the plasma, practically all loss occurs as direct orbit losses and, when the CX model is on, due to CX with background neutrals. When only the off-axis beam is on, ASCOT4 predicts that 8–9% of deposited beam power is lost when CX is off and 45–46% when CX is simulated using the best-estimate neutral profiles, 34–47% when the limiting-estimate cases are included. When both beams are on, ASCOT4 predicts that 10–11% of deposited beam power is lost when CX is off and 32% when CX is simulated using the best-estimate neutral profiles, 25–33% when the limiting-estimate cases are included. As expected, the relative CX losses are higher when only the off-axis beam is on, since ions from the off-axis beam are more susceptible to CX. All else being equal, it would also be expected that the relative amount of direct orbit losses is higher when only the off-axis beam is on. However, in the latter analyzed time range, when both beams are on, the low-field side separatrix is closer to the device wall, by 3 cm on the outer midplane. This explains higher direct orbit losses in the latter time range.

Figure 10 compares the neutron predictions by ASCOT4-AFSI, with and without CX included, to the neutron measurement by the fission chamber. When CX was included, the neutral density reconstructions based on BATS1D modelling were used. The predictions when CX was simulated using the best-estimate neutral profiles are shown using blue dots, and the limiting-estimate cases are indicated with blue error bars. When only the off-axis beam is on, ASCOT4-AFSI overestimates the measured neutron rate by 59–70% when CX is off and by 22–28% when CX is simulated using the best-estimate neutral profiles, by 19–44% when the limiting-estimate cases are included. When both beams are on, ASCOT4-AFSI overestimates the measured neutron rate by 49–64% when CX is off and by 36–48% when CX is simulated using the best-estimate neutral profiles, by 36–54% when the limiting-estimate cases are included. When calculating the above quoted differences at the simulated time points, the fission chamber data were time-averaged 5 ms in each direction.

The predicted neutron emission rates are consistently higher than the measured, as shown in Fig. 10. This could mean that the fission chamber calibration is still incorrect and that an additional scaling factor higher than the factor 2.6 used here is needed to raise the signal to the true level. It could also mean that the ASCOT4 modelling is underestimating beam-ion losses. Perhaps the estimated neutral density is too low, or perhaps significant beam-ion losses are caused by a mechanism not captured by the modelling. As was discussed in section 3, there are several reasons why the density of neutrals or their effect on fast ions may be underestimated. In the analyzed time range when only the off-axis beam is on, no MHD modes were observed. When the on-axis beam turns on and all neutron signals get a large absolute increase, the absolute difference between the predictions with and without CX only increases modestly. This is expected, since fast ions from the on-axis beam are less susceptible to CX. Similarly, the relative differences between the predictions and the measurement were expected




Figure 10. Total neutron rate from the plasma in shot 44623 as measured by the fission chamber, with an additional scaling factor of 2.6, and as predicted by ASCOT4-AFSI, with and without the inclusion of CX. CX was simulated using the neutral density reconstructions based on BATS1D calculations. Predictions when CX was simulated using the best-estimate neutral profiles are shown using blue dots, and the limiting-estimate cases are indicated with blue error bars.

to decrease substantially when the on-axis beam was turned on, since fast ions from the on-axis beam are less susceptible both to direct orbit and CX losses. Instead, the difference decreased only modestly for the CX-off case and increased substantially for the CX-on case. This suggests significant beam-ion losses from an unknown mechanism, which is induced by the on-axis beam. Indeed, while the time range with only the off-axis beam appeared MHD-quiescent, a weakly-excited mode at a frequency of about 5 kHz was observed in the time range with both beams. The mode would indeed not be captured by the ASCOT4 modelling, or the proceeding TRANSP modelling, and may be causing significant additional beam-ion losses that explain the lower than expected neutron measurements.

While aspects of shot 44623 remain unclear, when CX is accounted for, the simulated neutron rate agrees quantitatively better with the measurement. This is particularly the case before the on-axis beam is turned on, when the plasma is in a steady state with little to indicate other fast-ion loss mechanisms than direct orbit losses and CX losses. These results indicate that CX has a significant impact on beam-ion confinement in MAST-U, and that the ASCOT4 CX model is able to reproduce this effect.

5. Comparing predicted and measured passive FIDA

MAST-U is equipped with an 11-channel FIDA system with both an active view, in the same plane as and intersecting with the on-axis beam, and a passive view, which does not overlap with any beam [12, 13, 14]. When the bremsstrahlung background is accounted for and a relevant wavelength interval is investigated, FIDA light from CX between fast ions and edge neutrals can be observed if the overlap between the fast-ion population and edge neutrals is large enough [35, 36, 37]. In MAST-U, this overlap is larger than in, for example, ASDEX Upgrade, due to the larger fast-ion gyroradii in MAST-U. Hence, significant passive beam-ion FIDA light is expected in MAST-U. Moreover, since the redistribution and loss of beam ions due to CX with background neutrals, i.e. the topic of this article, predominantly affects the beam-ion distribution function near the edge, analysis of passive FIDA can be expected to capture these effects. In this study, passive FIDA intensity was predicted based on ASCOT5 modelling and compared to measurement to validate the ASCOT5 CX model and to investigate the impact of CX on the beam-ion distribution.

For this analysis, the MU01 shot 45091 was chosen: a double-null, L-mode plasma with 600 kA of plasma current and a conventional divertor configuration. The time point 0.35 s, when both beams are on, was chosen because of its relative MHD-quiescence. The case was first reconstructed using OMFIT and simulated using TRANSP. Same as in section 2, the TRANSP inputs were imported into ASCOT5. This TRANSP run was mostly configured similarly as the TRANSP run used in sections 2 and 4, but this time, since the on-axis beam was on for the whole shot, there was an ion temperature measurement from the Charge Exchange Recombination Spectroscopy (CXRS) system on MAST-U [38]. The measured ion temperature was reconstructed in OMFIT and used both in TRANSP and ASCOT5 simulations. The beam-ion population consisted of 11880 markers. The exception to using TRANSP inputs in ASCOT5 was the neutral density input. Even the TRANSP neutral temperature was imported, since, unlike ASCOT4, ASCOT5 does support a separate neutral temperature. However, the fast-ion CX process is insensitive to the neutral temperature [8]. The neutral density profiles reconstructed using BATS1D were used in ASCOT5 when comparing to measurement. ASCOT5 was also run using the TRANSP neutral density to allow code comparison. ASCOT5 was run with and without including CX. Using an ASCOT5-FIDASIM interface program from the ASCOT5 code library [4], the ASCOT5-predicted beam-ion distribution functions and other inputs were inserted into the FIDASIM code to predict passive FIDA spectra [11]. The 5th version of ASCOT was chosen for this study because of the pre-existing interface program for coupling ASCOT5 to FIDASIM. There exists an established methodology for using TRANSP and FIDASIM modelling to compare simulated fast-ion distributions to experiments based on measurements of FIDA intensity in MAST and MAST-U, and the absolute calibration of the FIDA system has been validated [12, 13, 14]. To verify the ASCOT5-FIDASIM modelling scheme, its passive FIDA predictions were compared to those of the established TRANSP-FIDASIM

scheme. The results are reported below.

As an example, Fig. 11 shows the ASCOT5-FIDASIM-simulated and measured passive FIDA spectra for channel 9 out of 11, with the bremsstrahlung background subtracted from the measured spectrum to reveal the underlying beam-ion FIDA signal. The major radius coordinate of the tangency point of the sightline for channel 9, i.e. its tangency radius, is 1.29 m, which corresponds to $\rho_p = 0.70$ in shot 45091 at time 0.35 s. The bremsstrahlung emissivity, which is not expected to vary significantly over the entire range of the plotted spectrum, i.e. down to 656 nm, was estimated by taking the signal average between 662.0 and 662.4 nm, which is beyond the possible wavelengths from beam-ion D-alpha in this case. The maximum beam-ion energy, i.e. the full energy of the off-axis beam, was 72.5 keV, which yields a maximum Doppler-shifted D-alpha wavelength of 661.9 nm. Indeed, the simulated and the bremsstrahlung-corrected measured spectra drop off below that wavelength. The peaking in the measurement at wavelengths below 659.0 nm is caused by the two passive carbon lines at 657.8 and 658.3 nm. The measurement error is based on the known photon statistics properties, and is part of the established methodology of FIDA measurement at MAST-U [12, 13, 14]. Error from the bremsstrahlung subtraction, accounting for the averaging using the formula for the standard error of the mean, was added to the final error margins. The simulated results shown in Fig. 11 are based on those best-estimate and limiting-estimate neutral profiles whose TRANSP parts include recombination atoms. At channel 9, between the carbon lines and the upper limit of the Doppler shift, the prediction when CX is simulated using the best-estimate neutral profile in ASCOT5 mostly agrees with the measurement. If the predictions based on the limiting-estimate neutral profiles are viewed as error margins, there is complete agreement between prediction and measurement. When CX is not accounted for in ASCOT5, the best-estimate neutral profile is used in FIDASIM, and wavelengths between 659.0 and 661.0 nm are considered, the prediction is about three times as high as the measurement. To avoid the two carbon lines and the high noise-to-signal ratio in the drop of the spectrum past 661.0 nm, the spectra from each channel are averaged over the relatively flat range between wavelengths 659.0 and 661.0 nm in the following analysis.

ASCOT5-FIDASIM agrees with TRANSP-FIDASIM on the passive FIDA signal to within 30%. Figure 12 shows the passive FIDA signals predicted by ASCOT5-FIDASIM and TRANSP-FIDASIM as functions of the tangency radii of the FIDA channels. For each model and each channel, the passive FIDA spectrum has been averaged between wavelengths 659.0 and 661.0 nm. In this ASCOT5 simulation, the TRANSP neutral density was used for comparability between codes. In the FIDASIM simulations for this code comparison, beam-halo and recombination atoms from the TRANSP neutral background were omitted, and regions outside the separatrix were omitted completely using artificial masking, i.e. the signal contribution from outside the separatrix was set to zero, as has been common practice in TRANSP-FIDASIM modelling of active FIDA in MAST and MAST-U [12, 13]. While active FIDA is predominantly produced by CX from the excited state $n = 3$ (n is the principal quantum number) in a fast beam neutral




Figure 11. ASCOT5-FIDASIM-simulated and measured passive FIDA spectra (logarithmic scale) for channel 9 (out of 11 channels), with tangency radius 1.29 m ($\rho_p = 0.70$), in shot 45091 at 0.35 s. Bremsstrahlung has been estimated as the signal average between 662.0 and 662.4 nm and subtracted from the measured spectrum. Predictions using the best-estimate neutral profile whose TRANSP part includes recombination atoms are shown using solid lines. The corresponding limiting-estimate cases are indicated with shaded colour bands.

or a beam-halo neutral to the state $n = 3$ in a neutralized fast ion, passive FIDA is predominantly produced by CX directly from the ground state ($n = 1$) in a background neutral to the state $n = 3$ in a neutralized fast ion. Therefore, a substantial contribution to the passive FIDA signal can come from the SOL even when the electron density and temperature are negligible. In the comparison to measurement, which is reported below, passive FIDA from the SOL is accounted for. In the code comparison, the ASCOT5-FIDASIM prediction for the passive FIDA signal agrees with the TRANSP-FIDASIM prediction to within 30%, except at the outermost point, where the ASCOT5-FIDASIM prediction is 50% lower. The channel numbering, 1–11, runs from lowest to highest major radius. In shot 45091 at 0.35 s, channels 1–4 have tangency points on the high-field side, while channels 5–11 have tangency points on the low-field side, as shown in Fig. 14. Based on the data shown in Figs. 12 and 13, channel 1 appears spurious. On the low-field side, the ASCOT5-FIDASIM and TRANSP-FIDASIM predictions shown in Fig. 12 start to diverge at channel 6. The ASCOT5-estimated fast-ion density as a function of ρ_p agrees with that of TRANSP to within 10% across ρ_p values corresponding to the tangency points of channels 6–10. However, at the tangency point of channel 11, the ASCOT5 value is only 30% of the TRANSP value. Both fast-ion density functions have very low values this close to the edge when CX is accounted for. The values at the tangency point of channel 11, which corresponds to $\rho_p = 0.95$, are 0.3% and 0.9%

of the values at $\rho_p = 0.50$ for ASCOT5 and TRANSP, respectively. This means that the simulation statistics are poor in this part of the distribution function. Since the passive FIDA signal is dominated by FIDA light from the edge for any channel [35], the significant relative difference in the fast-ion density predictions on the very edge likely explains most of the differences in the synthetic passive FIDA signals not only at channel 11 but also at channels 6–10. While the bright edge affects all channels, a lower tangency radius still implies a more radial sightline, which implies more signal contribution from deeper inside the plasma [12]. This may explain why the passive FIDA predictions differ noticeably but not strongly at channels 6–10. The demonstrated level of agreement with the established TRANSP-FIDASIM model and the identification of an explanation for the modest differences that were observed suggest that the ASCOT5-FIDASIM modelling works and can give physically accurate predictions of the passive FIDA signal in MAST-U.




Figure 12. Passive FIDA signals in shot 45091 at 0.35 s predicted by ASCOT5-FIDASIM and TRANSP-FIDASIM as functions of the tangency radii of the FIDA channels (corresponding ρ_p coordinates range from $\rho_p = 0.47$ on the high-field side to $\rho_p = 0.95$ on the low-field side). For each model and each channel, the passive FIDA spectrum has been averaged between wavelengths 659.0 and 661.0 nm.

Figure 13 shows the ASCOT5-FIDASIM-predicted, with and without CX, and bremsstrahlung-corrected measured passive FIDA signals as functions of tangency radius such that the spectra have been averaged between wavelengths 659.0 and 661.0 nm. The measurement error from the spectra was propagated through the calculation, accounting for the averaging using the formula for the standard error of the mean. The ASCOT5-FIDASIM modelling, both with and without CX included in ASCOT5, was performed using the six different neutral profiles constructed in section 3.4. In these FIDASIM simulations, artificial masking was applied only to regions outside $\rho_p = 1.2$, which

should be out of reach of any beam ions and, therefore, irrelevant. Predictions with CX on are shown in blue and those with CX off in orange. The predictions using the best-estimate neutral profile whose TRANSP part includes recombination atoms are shown using solid lines. The corresponding limiting-estimate cases are indicated with shaded colour bands to reflect that they are interpreted as error margins. The predictions using the best-estimate neutral profile whose TRANSP part excludes recombination atoms are shown using dashed lines. The corresponding limiting-estimate cases are indicated with dotted lines. The cases with the minimum-estimate neutral densities consistently yield the lowest passive FIDA signals and the cases with the maximum-estimate neutral densities yield the highest passive FIDA signals. When CX reactions are included in the ASCOT5 simulation, the estimated spectra are quantitatively more consistent with the measurements, in particular at higher tangency radii. As mentioned above, based on the data shown in Figs. 12 and 13, channel 1 appears spurious. At the outermost channel, channel 11, all predictions shown in Fig. 13 diverge from the measurement due to peaking in the predicted signals. It is unclear why the predictions overestimate peaking at the edge, but possible explanations were identified. As was discussed in section 3, there are several reasons why the density of neutrals or their effect on fast ions may be underestimated in the SOL, which may explain an overestimation of the beam-ion density and, therefore, FIDA emissivity on the very edge of the plasma and in the SOL. Moreover, the predicted signal is most sensitive at the outermost channel to possible, edge-localized error in the neutral density reconstruction. An overestimation of the neutral density in a narrow radial region around the tangency radius of channel 11 would likely have a larger impact on the production of FIDA light than on the confinement of beam ions whose gyroradii typically span several such radial regions. This may explain why the signal peaking is overestimated both with and without accounting for CX losses. Since the focus here is on the impact of accounting for CX in the ASCOT5 modelling, and the excessive edge peaking is similar with and without CX, channel 11 will be disregarded in the following analysis. If we consider the best-estimate cases with and without recombination atoms as two alternative predictions whose error margins are determined by the corresponding limiting-estimate cases, then one or both of the predictions agree with the measurements at all channels 2–10 within error margins when CX is accounted for. This result suggests that the neutral density reconstructions are good approximations of the true neutral background. Specifically, it lends credibility to the decision in section 3.4 to scale the TRANSP part of the neutral density profile to the BATS1D part since, unlike in the neutron emission analysis in section 4, here the accuracy of the neutral density deeper inside the plasma is important because of its effect on the production of passive FIDA light. For shot 45091 at 0.35 s, the TRANSP-part neutral density was higher before scaling, even in the case of the maximum estimate from BATS1D. The result of the comparison of simulated and measured passive FIDA further shows that, given an accurate neutral background density, the ASCOT5 CX model is able to capture the impact of CX on the beam-ion distribution.

Proceeding to more detailed analysis of the results shown in Fig. 13, predictions




Figure 13. ASCOT5-FIDASIM-simulated and bremsstrahlung-corrected measured passive FIDA signals (logarithmic scale) in shot 45091 at 0.35 s as functions of the tangency radii of the FIDA channels. For each channel, the simulated and measured passive FIDA spectra have been averaged between wavelengths 659.0 and 661.0 nm. Predictions using the best-estimate neutral profile whose TRANSP part includes recombination atoms are shown using solid lines. The corresponding limiting-estimate cases are indicated with shaded colour bands. Predictions using the best-estimate neutral profile excluding recombination atoms are shown using dashed lines. The corresponding limiting-estimate cases are indicated with dotted lines.

where ASCOT5 accounts for CX agree with measurement at channels 2–4 and 6–10 when recombination atoms are included. There is good agreement at channel 5 when recombination atoms are excluded. In addition, there is considerably better agreement at channels 6 and 7 when recombination atoms are excluded. These results suggest that there is a significant population of background neutrals from recombination, or some other source, deeper inside the plasma, but that the TRANSP neutral model is only partially successful in reproducing it, or that the TRANSP neutrals have been incorrectly incorporated in the neutral reconstruction in this article. Alternatively, considering the 3D shape of the plasma that the FIDA sightlines pass through, perhaps the FIDASIM modelling is overestimating a signal contribution from the edge on the far side of the plasma by underestimating the opacity of the plasma. This may explain why the predicted signals show a downward trend from channel 7 towards lower radii to channel 2, whereas the measured signal is quite constant across these channels. The tangency points of channels 2–4 and 5–7 lie quite symmetrically on either side of the magnetic axis, as shown in Fig. 14. Overestimating a signal contribution from the far side of the plasma may also explain why the predictions peak excessively at the edge channels.

Compared to modelling where ASCOT5 accounts for CX, modelling where ASCOT5

does not account for CX significantly overestimates the measured passive FIDA signal, in particular at higher tangency radii. When the best-estimate neutral profile with recombination atoms is used in FIDASIM but CX is not accounted for in ASCOT5, the predicted passive FIDA signal at channels 2–10 is 1.4–3.8 times as high as the measured signal. At channels 5–10, i.e. the channels with tangency points on the low-field side, the ratio is 2.1–3.8. When CX is accounted for in ASCOT5, the ratio of prediction to measurement is 0.9–1.8, both at all channels 2–10 and at only channels 5–10. The better performance of ASCOT5 simulations where CX is accounted for is explained by the difference in the simulated beam-ion density near the edge. Figure 14 shows the ratio of the ASCOT5-estimated beam-ion density on the outer midplane when CX is off to the one when CX is on. Around the outer-midplane separatrix, the beam-ion density is estimated to be 2–4 times as high when CX reactions are neglected, which is consistent with the difference predicted in the passive FIDA signal. The figure shows the tangency points of the eleven FIDA channels. The set of tangency points runs from the high-field side, through the axis, and along the outer midplane to the low-field-side edge. One should keep in mind, however, that since passive FIDA light is edge-localized, even the signal in the inner channels comes predominantly from where their respective sightlines cross the edge [35]. Since this is the passive FIDA view in MAST-U, there is no contamination from active FIDA from interactions between fast ions and a beam line.




Figure 14. Ratio between the beam-ion density around the outer midplane in shot 45091 at 0.35 s predicted by ASCOT5 when CX is off and the one predicted when CX is simulated using the best-estimate neutral profile whose TRANSP part includes recombination atoms. The equilibrium is illustrated using contours for $\rho_p = 0.2, 0.4, 0.6, 0.8$ and 1.2 in gray and for $\rho_p = 1.0$ in red. The tangency points of the FIDA channels are indicated by black crosses.

6. Summary

Based on observations from the first experimental campaign in MAST-U (MU01), CX was suspected to cause significant beam-ion losses, making it an interesting environment for the study of fast-ion CX. A new model, dubbed BATS1D, was developed

to reconstruct the outer-midplane neutral density based on D-alpha measurements carried out using the HOMER camera and Thomson scattering measurements of electron density and temperature. Due to uncertainties in the calibration of the D-alpha measurements, they were scaled based on an extensive comparison of D-alpha measurements and KN1D modelling for experiments in the second MAST-U campaign (MU02). The ability of the ASCOT orbit-following code to simulate beam-ion CX in MAST-U was validated through comparison to MU01 experiments, and the impact of CX on beam-ion confinement was investigated. The fast-ion CX models of ASCOT4 and ASCOT5, both of which were used for analysis in this article, were also benchmarked and found to agree.

The neutral density reconstruction was found to agree with SOLPS-ITER modelling for MU01 experiments in terms of the slope of the neutral density profile at the separatrix when plotting on a logarithmic scale, suggesting that the simple BATS1D neutral density model can be used to make physically relevant predictions for the outer-midplane in MAST-U. Since beam ions in MAST-U orbit significantly further out in the plasma on the low-field side, especially around the outer midplane, the error introduced by assuming a poloidally uniform neutral density distribution is expected to be significantly mitigated. It was observed that when fast-ion CX is accounted for in modelling, predicted neutron emission from the plasma is more consistent with fission chamber measurements than when CX is neglected in modelling. However, quantitative agreement of the predicted neutron rates with measurements is still somewhat poor, possibly due to remaining errors in the fission chamber calibration. When fast-ion CX was accounted for in simulations with the neutral density reconstructed using BATS1D, the measured passive FIDA signal was reproduced. When CX losses were neglected, the predicted signal was 2–4 times as high as the measured signal at channels with tangency points on the low-field side, which is explained by the impact of CX losses on the beam-ion density at the plasma edge. The results suggest that the BATS1D reconstruction was a good approximation of the outer-midplane neutral density, and show that, given an accurate neutral background density, the ASCOT fast-ion CX model is able to predict the impact of CX with edge neutrals on the beam-ion distribution. The results further indicate that CX with edge neutrals has a significant impact on beam-ion confinement in MAST-U.

The ability demonstrated here to account for CX in fast-ion modelling also provides an opportunity for more detailed analysis of the impact of fast-ion-driven mode activity on fast ions. If the impact of CX on the fast-ion distribution is accounted for in FIDASIM modelling, remaining differences between the synthetic and measured FIDA signals provide a more accurate measure of the impact of fast-ion-driven modes on fast-ion redistribution and loss. Going forward, as uncertainties are reduced and the neutral background reconstruction is improved, the ability to accurately simulate fast-ion CX will be useful in a wide range of increasingly detailed modelling activities in the endeavour to better understand fast ions in spherical tokamaks and other magnetic-confinement fusion devices. Such activities may be expected to improve confidence in

predictions of alpha-particle physics in fusion power plants, particularly with regard to redistribution and loss of these particles due to instabilities and 3D field perturbations.

Acknowledgements

We thank William Heidbrink and Christopher Beckley for helpful discussion. We acknowledge the contributions by Andrew Kirk and Andrew Thornton to the calibration of the fission chamber. This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No. 101052200 — EUROfusion) and also with support from the EPSRC [grant No. EP/W006839/1] and the US DoE Grant No. DE-SC0019007. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

References

- [1] K.G. McClements, K. Tani, R.J. Akers, Y.Q. Liu, K. Shinohara, H. Tsutsui, and S. Tsuji-Iio. The effects of resonant magnetic perturbations and charge-exchange reactions on fast ion confinement and neutron emission in the Mega Amp Spherical Tokamak. *Plasma Physics and Controlled Fusion*, 60(9):095005, July 2018. doi: <https://doi.org/10.1088/1361-6587/aad25210.1088/1361-6587/aad252>.
- [2] M Ceconello, I J Dolby, A Sperduti, J Rivero-Rodriguez, G Ericsson, I Fitzgerald, S Y Allan, J Voller, B Honey, B A Nizar, S D Elmore, and the MAST-U team. First observations of confined fast ions in MAST Upgrade with an upgraded neutron camera. *Plasma Physics and Controlled Fusion*, 65(3):035013, feb 2023.
- [3] E. Hirvijoki, O. Asunta, T. Koskela, T. Kurki-Suonio, J. Miettunen, S. Sipilä, A. Snicker, and S. Äkäslompolo. ASCOT: Solving the kinetic equation of minority particle species in tokamak plasmas. *Computer Physics Communications*, 185(4):1310–1321, April 2014. doi: <https://doi.org/10.1016/j.cpc.2014.01.01410.1016/j.cpc.2014.01.014>.
- [4] Jari Varje, Konsta Särkimäki, Joonas Kontula, Patrik Ollus, Taina Kurki-Suonio, Antti Snicker, Eero Hirvijoki, and Simppa Äkäslompolo. High-performance orbit-following code ASCOT5 for Monte Carlo simulations in fusion plasmas. *arXiv preprint arXiv:1908.02482*, 2019.
- [5] A. Pankin, D. McCune, R. Andre, G. Bateman, and A. Kritz. The tokamak Monte Carlo fast ion module NUBEAM in the National Transport Code Collaboration library. *Computer Physics Communications*, 159(3):157–184, 2004.
- [6] R.J. Goldston, D.C. McCune, H.H. Towner, S.L. Davis, R.J. Hawryluk, and G.L. Schmidt. New techniques for calculating heat and particle source rates due to neutral beam injection in axisymmetric tokamaks. *Journal of Computational Physics*, 43(1):61–78, September 1981. doi: [https://doi.org/10.1016/0021-9991\(81\)90111-X10.1016/0021-9991\(81\)90111-X](https://doi.org/10.1016/0021-9991(81)90111-X10.1016/0021-9991(81)90111-X).
- [7] PPPL TRANSP. <https://transp.pppl.gov/index.html>. Accessed: April 17, 2023.
- [8] Patrik Ollus, R Akers, B Colling, H El-Haroun, D Keeling, Taina Kurki-Suonio, R Sharma, Antti Snicker, Jari Varje, et al. Simulating the impact of charge exchange on beam ions in MAST-U. *Plasma Physics and Controlled Fusion*, 64(3):035014, 2022.
- [9] Paula Sirén, Jari Varje, Simppa Äkäslompolo, Otto Asunta, Carine Giroud, Taina Kurki-Suonio, Henri Weisen, et al. Versatile fusion source integrator AFSI for fast ion and

- neutron studies in fusion devices. *Nuclear Fusion*, 58(1):016023, November 2017. doi: <https://doi.org/10.1088/1741-4326/aa92e9>.
- [10] C Vincent, S Allan, G Naylor, R Stephen, S Bray, A Thornton, and A Kirk. Fission chamber data acquisition system for neutron flux measurements on the Mega-Amp Spherical Tokamak Upgrade. *Review of Scientific Instruments*, 93(9):093509, 2022.
- [11] WW Heidbrink, D Liu, Y Luo, E Ruskov, and Benedikt Geiger. A code that simulates fast-ion $D\alpha$ and neutral particle measurements. *Communications in Computational Physics*, 10(3):716–741, 2011.
- [12] AR Jackson, AS Jacobsen, KG McClements, CA Michael, and Marco Cecconello. Diagnosing fast ion redistribution due to sawtooth instabilities using fast ion deuterium- α spectroscopy in the mega amp spherical tokamak. *Nuclear Fusion*, 60(12):126035, 2020.
- [13] A.R Jackson. *Fast-Ion Deuterium-Alpha Measurements on MAST and MAST-U*. PhD thesis, University of Durham, 2022.
- [14] C A Michael, N Conway, B Crowley, O Jones, W W Heidbrink, S Pinches, E Braeken, R Akers, C Challis, M Turnyanskiy, A Patel, D Muir, R Gaffka, and S Bailey. Dual view FIDA measurements on MAST. *Plasma Physics and Controlled Fusion*, 55(9):095007, jul 2013.
- [15] J Storrs, J Dowling, G Counsell, and G McArdle. Real-time optical plasma edge detection and position control on MAST. *Fusion engineering and design*, 81(15-17):1841–1845, 2006.
- [16] R Scannell, MJ Walsh, MR Dunstan, J Figueiredo, G Naylor, T O’Gorman, S Shibaev, KJ Gibson, and H Wilson. A 130 point Nd: YAG Thomson scattering diagnostic on MAST. *Review of Scientific Instruments*, 81(10):10D520, 2010.
- [17] Brian LaBombard. KN1D: A 1-D space, 2-D velocity, kinetic transport algorithm for atomic and molecular hydrogen in an ionizing plasma. 2001. (Available at: <https://dspace.mit.edu/handle/1721.1/93211>).
- [18] LC Johnson and EJ Hinnov. Ionization, recombination, and population of excited levels in hydrogen plasmas. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 13:333, 1973.
- [19] D. Moulton, L. Xiang, J. Harrison, and the MAST Upgrade team. Interpretative modelling of the target ion flux rollover in Conventional and Super-X divertor configurations on MAST Upgrade. In *48th EPS Conference on Plasma Physics*. European Physical Society, 2022. <http://ocs.ciemat.es/EPS2022PAP/pdf/O3.108.pdf>.
- [20] Orso Meneghini and Lang Lao. Integrated modeling of tokamak experiments with OMFIT. *Plasma and Fusion Research*, 8:2403009–2403009, 2013.
- [21] LL Lao, H St John, RD Stambaugh, AG Kellman, and W Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611, 1985.
- [22] LC Appel, GTA Huysmans, LL Lao, PJ McCarthy, DG Muir, ER Solano, J Storrs, D Taylor, W Zwingmann, et al. A unified approach to equilibrium reconstruction. In *Proceedings-33rd EPS conference on Controlled Fusion and Plasma Physics*, pp. P–2.160, 2006.
- [23] S Tamor. ANTIC: A code for calculation of neutral transport in cylindrical plasmas. *Journal of Computational Physics*, 40(1):104–119, 1981.
- [24] TRANSP Help. <https://w3.pppl.gov/pshare/help/transp.htm>. Accessed: April 17, 2023.
- [25] Frederick N Fritsch and Ralph E Carlson. Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246, 1980.
- [26] David Kahaner, Cleve Moler, and Stephen Nash. *Numerical methods and software*. Prentice-Hall, Inc., 1989.
- [27] F Scotti, DP Stotler, RE Bell, BP LeBlanc, SA Sabbagh, VA Soukhanovskii, MV Umansky, and SJ Zweben. Outer midplane neutral density measurements and H-mode fueling studies in NSTX-U. *Nuclear Fusion*, 61(3):036002, 2021.
- [28] Ratko K Janev, Detlev Reiter, and Ulrich Samm. *Collision processes in low-temperature hydrogen plasmas*, volume 4105. Forschungszentrum, Zentralbibliothek Jülich, 2003.
- [29] C Cowley, B Lipschultz, D Moulton, and B Dudson. Optimizing detachment control using the magnetic configuration of divertors. *Nuclear Fusion*, 62(8):086046, 2022.

- [30] A J Thornton, A Kirk, and the MAST Team. Scaling of the scrape-off layer width during inter-ELM H modes on MAST as measured by infrared thermography. *Plasma Physics and Controlled Fusion*, 56(5):055008, apr 2014.
- [31] J.R. Harrison, G.M. Fishpool, and A. Kirk. L-mode and inter-ELM divertor particle and heat flux width scaling on MAST. *Journal of Nuclear Materials*, 438:S375–S378, 2013. Proceedings of the 20th International Conference on Plasma-Surface Interactions in Controlled Fusion Devices.
- [32] S. Äkäslompolo, G. Bonheure, G. Tardini, T. Kurki-Suonio, and The ASDEX Upgrade team. Adjoint Monte Carlo simulation of fusion product activation probe experiment in ASDEX Upgrade tokamak. *Journal of Instrumentation*, 10(10):P10012, oct 2015.
- [33] P. Sirén, J. Varje, H. Weisen, and L. Giacomelli. Role of JETPEAK database in validation of synthetic neutron camera diagnostics and ASCOT- AFSI fast particle and fusion product calculation chain in JET. *Journal of Instrumentation*, 14(11):C11013, nov 2019.
- [34] J Kontula, J P Koschinsky, S Äkäslompolo, T Kurki-Suonio, and the W7-X team. ASCOT simulations of 14 MeV neutron rates in W7-X: effect of magnetic configuration. *Plasma Physics and Controlled Fusion*, 63(3):035022, jan 2021.
- [35] Benedikt Geiger, Manuel Garcia-Munoz, WW Heidbrink, RM McDermott, Giovanni Tardini, R Dux, R Fischer, V Igochine, et al. Fast-ion D-alpha measurements at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 53(6):065010, 2011.
- [36] Nathan G. Bolte, William W. Heidbrink, David Pace, Michael Van Zeeland, and Xi Chen. Measurement and simulation of passive fast-ion D-alpha emission from the DIII-D tokamak. *Nuclear Fusion*, 56(11):112023, sep 2016.
- [37] G Z Hao, W W Heidbrink, D Liu, M Podesta, L Stagner, R E Bell, A Bortolon, and F Scotti. Measurement of the passive fast-ion D-alpha emission on the NSTX-U tokamak. *Plasma Physics and Controlled Fusion*, 60(2):025026, jan 2018.
- [38] NJ Conway, PG Carolan, John McCone, MJ Walsh, and Marco Wisse. High-throughput charge exchange recombination spectroscopy system on MAST. *Review of scientific instruments*, 77(10):10F131, 2006.