

Simulating Energetic Particle Losses in JET Plasmas with a Reverse Integration Biasing Scheme

P. J. Bonofiglio¹, M. Podestà¹, R. B. White¹, V. Kiptily²,
V. Goloborodko³, N. N. Gorelenkov¹, M. Vallar⁴, F. E.
Cecil⁵, C. Giroud², S. Brezinsek⁶, and JET Contributors^{*}

¹Princeton Plasma Physics Laboratory, Princeton, NJ 08543, USA

²Culham Centre for Fusion Energy of UKAEA, Culham Science Centre,
Abingdon, UK

³Kyiv Institute for Nuclear Research, Kyiv 03680, Ukraine

⁴Ecole Polytechnique Fédérale de Lausanne, Swiss Plasma Center, CH-1015
Lausanne, Switzerland

⁵Colorado School of Mines, Golden, CO 80401, USA

⁶Forschungszentrum Jülich, Institut für Energie-und Klimaforschung -
Plasmaphysik, 52425 Jülich, Germany

^{*}See the author list of 'Overview of JET results for optimising ITER operation'
by J. Mailloux et al. to be published in Nuclear Fusion Special issue: Overview
and Summary Papers from the 28th Fusion Energy Conference (Nice, France,
10-15 May 2021)

E-mail: pbonofig@pppl.gov

4 October 2021

Abstract. An integrated energetic particle transport model has been constructed in JET plasmas constrained by experimental fast ion loss measurements. The model incorporates a synthetic fast ion loss detector identical to JET's thin-foil Faraday cup fast ion loss detector array. The loss model combines analyses from the TRANSP and ORBIT-kick codes with enhanced features for producing the synthetic diagnostic. Extensions to the ORBIT code framework allow a full-orbit representation within the vacuum region that can map particles directly to an installed detector geometry. Since synthetic fast ion loss detectors are plagued by weak loss statistics, a novel reverse integration biasing scheme has been implemented to boost computational efficiency. The model is validated against experimental loss measurements induced by long-lived kink modes and is found to be in good agreement. This confirms the development of a fully integrated transport/loss model which can be quantitatively verified against experiment allowing for future validation and predictive studies. The model is particularly useful for more complicated plasma scenarios that involve multiple fast ion species such as JET's 2021 DT-campaign.

Submitted to: *Nucl. Fusion*

1. Introduction

The study of energetic particles and their transport properties is well established. With JET's 2021 DT-campaign beginning[1], the construction of ITER nearing completion[2], and the design of compact devices like SPARC[3] becoming more prevalent, reactor relevant experiments have gained particular importance. While the good confinement of DT-alpha particles is necessary for achieving a sustained burning plasma, external heating sources, such as neutral beam injection (NBI) and ion cyclotron resonant heating (ICRH), establish further suprathermal ion populations that complicate our physical understanding.[4]

Energetic ions have been shown to interact strongly with various magnetohydrodynamic (MHD) instabilities. Energetic particles can destabilize MHD modes within a plasma that in turn drive fast ion transport.[5] In many cases, the outward transport induced by the wave-particle resonant interactions results in lost particles. Lost energetic particles are counter to the fusion mission and have become an active area of research.[6, 7] The direct transport mechanisms, however, are often complex in nature and can be composed of diffusive, convective, symmetry breaking, or phase-space stochasticization phenomena.[8] Consequently, numerical models have been employed to help study these interactions with the aim to minimize potential losses.[9, 10, 11, 12]

While experiment can verify the transport/loss model, the model can supply additional information lacking from measurement. By synergistically combining experiment and modeling, more information can be garnered from the physical scenario than with just measurement or simulation alone. In essence, experimental measurements often lack detail such as a complete picture of the fast ion distribution or wave-particle resonant interaction. Numerical modeling can include such features but requires verification and validation.[13]

Synthetic loss detectors have been produced to compare numerically modeled losses to experimental measurements.[14, 15, 16, 17, 18] It has been readily observed that fast ion loss detector (FILD) signals correlate with many loss mechanisms of interest.[19, 20, 21, 22, 23] Thus, it is possible to form a fully integrated energetic particle transport model which can then be quantitatively corroborated against experimental measurement. In theory, developing a synthetic FILD should be straightforward since one simply has to trace particles to the detector geometry. In practice, however, the relatively small number of lost ions captured by FILDs requires a large number (typically in excess of 10^6) of simulated particles to achieve reasonable signal levels. This has proven to be very

computationally inefficient and time consuming since the detection events follow weak Monte-Carlo counting statistics.[16]

Previous efforts[24, 25, 26] have utilized forward and backward time integration techniques to observe overlapping phase-space regions of interaction.[27] This report describes a more thorough methodology in which forward modeled particles are fully tracked to the detector geometry after biasing against the distribution of particles reverse integrated from the detector. This procedure combats the poor loss statistics common in other synthetic FILDs by weighting the initial particle sampling toward particles that are more likely to reach the detector while ignoring those deeply confined.

This paper presents an integrated transport model capable of generating synthetic energetic particle loss fluxes for comparison to one of JET's FILDs. A description of the FILD used to validate the model will conclude the introduction. Section 2 will begin with a general description of the transport/loss model which will incorporate both forward and backward orbit integration. An explanation of a novel biasing scheme will follow as well as the final methodology for achieving synthetic lost particle fluxes. An experimental discharge with strong kink mode induced losses is given as an example to prove the simulations efficacy where the real loss measurements are compared to those numerically produced. Lastly, this paper will conclude with a brief discussion of the results obtained, the merits of the loss model, and comments on improvements.

1.1. JET's Faraday Cup Fast Ion Loss Detector Array

JET contains an array of Faraday cup FILDs capable of providing energy and spatially resolved energetic particle losses.[28] (Note that pitch resolution is crudely limited to the aperture hole geometry.) The array consists of five structures, previously referred to as "pylons", each containing up to three radially displaced Faraday cups. The pylons span a large poloidal extent below the midplane and cover a slight radial range on the outboard side of the machine. The overall detector structure is shown in Figure 1.

Each Faraday cup is composed of conductive Ni foils separated by insulating mica layers. This establishes an energy resolution for an incident ion depending on its penetration depth. The insulating layers provide dead spaces and create separate energy channels within the foils where the captured lost ions are read as raw current. In short, the deeper the ion deposition, the higher the energy. Reference [22] provides a complete table of the resolveable energy range per foil per ion species.

The Faraday cup array has undergone recent




Figure 1. CAD schematic of JET's Faraday cup fast ion loss detector array with labeled pylon structures, midplane, toroidal direction (ϕ), and poloidal direction (θ). Each pylon can house up to 3 Faraday cups radially displaced by 2.5 cm. Each Faraday cup has a 8×17 grid of 3mm circular apertures (not visible). Reproduced from Bonofiglio 2020 *Rev. Sci. Instrum.* **91** 093502, with the permission of AIP Publishing.

improvements to its data acquisition system that have resulted in enhanced, and new, measurements as well as improved analysis techniques.[22] In particular, losses due to low frequency ($\omega < 100$ kHz) MHD activity are particularly well resolved. In that sense, the Faraday cup array is well suited to provide reliable, quantitative, data to validate the transport model. While the measurements provide decent energy resolution and cover a wide spatial breadth to support the model, the Faraday cups are incapable of distinguishing ion species since the signal is simply raw current. Such features, however, can be inferred from the loss model once validated.

JET also contains a scintillator probe fast ion loss detector capable of providing additional loss measurements.[29] While this report focuses on the synthetic losses measured by the Faraday cup array, the entire model methodology can be implemented for the scintillator probe. The only difference would be incorporating the scintillator probe geometry in place of the Faraday cups. Experimentally, the scintillator probe can provide supplemental information to further constrain and verify the Faraday cup losses and vice-versa.

2. Description of the Transport and Loss Model

The transport/loss model incorporates experimentally constrained magnetic equilibria, realistic energetic particle distribution functions, magnetic pertur-

bative effects, and the FILD detector geometry to produce synthetic lost ion flux.

The TRANSP code[30] forms the base on which the overall model is constructed. TRANSP provides a time-dependent simulation of the equilibrium and fast ion distribution within a large physics framework (evolving profiles, atomic physics effects, transport and heating, etc.). The equilibria and heating effects are constrained to experimental profiles and parameters providing the most realistic scenario possible.

The equilibrium and fast ion distributions are passed into the ORBIT code[31] which performs the main particle pushing for calculating the fast ion transport. The model incorporates the ORBIT-kick variation[10, 32] which encodes the affects of the resonant wave-particle interactions via the fast ion's constants of motion (energy - E , canonical toroidal momentum - P_ζ , magnetic moment - μ). The magnetic moment is assumed to be conserved for low frequency activity ($\omega \ll \omega_{ci}$ where ω_{ci} is the ion cyclotron frequency) which is true for most modes of interest (kinks, fishbones, tearing modes, Alfvén eigenmodes, etc.), so that the wave-particle interaction occurs through changes, or “kicks,” in the particle energy, ΔE , and canonical momentum, ΔP_ζ .

The ORBIT-kick model contains a wealth of information unobtainable in experiment which makes it ideal for validation. ORBIT characterizes the resonant interactions based on the fast ion phase-space orbit topology. The topological transitions resulting from the supplied magnetic perturbation can be directly examined allowing a holistic view of the fast ion transport. The ORBIT outputs provide extemporaneous information on the wave-particle interaction and energetic particle species dependence which permits a better defined interpretation of the losses than could be with the Faraday cup FILD alone.

Since the model aims to construct a synthetic fast ion loss detector, one must first consider that ORBIT is a Hamiltonian guiding-center code that has, traditionally, only functioned on the domain within the last closed flux surface. As such, recent advancements to ORBIT have been made. Namely, calculations are now possible within the vacuum region and the gyrophase and Larmor radius are tracked as functions of time. In doing so, finite-Larmor radius effects can be included in a simple fashion via equations 1a and 1b where (R_{FLR}, Z_{FLR}) is the full orbit corrected position, (R_{GC}, Z_{GC}) is the guiding-center position, ρ_L is the Larmor radius, and ϕ_{gyro} is the gyrophase:

$$R_{FLR} = R_{GC} + \rho_L \cos(\phi_{gyro}) \quad (1a)$$

$$Z_{FLR} = Z_{GC} + \rho_L \sin(\phi_{gyro}) \quad (1b)$$

Possessing a full-orbit representation of the particle position and extending the domain to the vacuum region, a synthetic FILD can be constructed by adding

the detector geometry. The Faraday cup housing panels, shielding, and aperture geometries were all added to ORBIT within their measured engineering constraints. With these new added features, the full motion of energetic particles can be calculated in ORBIT and tracked to a replica FILD for synthetic detection.

2.1. Supplemental Model Information

In addition to a magnetic equilibrium and particle pusher, the model requires an initial representation of the fast ion population and magnetic perturbation. A realistic description of the energetic particle distribution function is required to initialize our particles in ORBIT, and the displacement vector for any MHD mode(s) of interest is(are) necessary to perturb well-confined orbits to losses.

The NUBEAM submodule[33] of the TRANSP code provides the energetic particle distribution for neutral beam injected ions as well as fusion products. The resolveable energy ranges of the Faraday cup foils are of order MeV, so beam-born particles (120 keV typical injection energy) are only resolveable if they have undergone further ion cyclotron resonance heating. The TORIC submodule[34] can be used in conjunction with NUBEAM to include the RF-heating effects. Selecting a time of interest, the full NBI RF-heated and fusion product distributions in (R, Z, Λ, E) -space, where $\Lambda = v_{\parallel}/v$ denotes the particle pitch, can be pulled from a TRANSP run for use in the model.

The RF-heated tail population of the beam born ions, however, is often weak in counts and gives rise to a sparse distribution. TRANSP runs typically include anywhere from 32,000-256,000 neutral beam Monte-Carlo markers. Since the deepest foil in the Faraday cup stacks can measure up to 2.0 MeV deuterons, which is well above the maximum beam injection energy of 120 keV, the number of ions in the RF-tail exponentially decreases rapidly with foil depth. As such, it is necessary to add additional statistics to fill in the distribution. Performing more TRANSP runs or simply increasing the number of neutral beam particles is too computationally costly and time consuming. Instead, stand-alone executions of the NUBEAM submodule can be executed from the TRANSP plasma state files at the time of interest to produce the fast ion distributions. These runs can be done in parallel and are faster than the full TRANSP computations. Figure 2 presents an example of added RF-tail statistics from additional NUBEAM calculations relative to the original distribution produced from the TRANSP run. After ensembling the NUBEAM runs together, the previously absent RF-heated tail deuteron population is now present and better resolved up to the maximum measureable foil energy (2 MeV for deuterons).




Figure 2. The original deuteron fast ion distribution from a TRANSP simulation with NBI+ICRH heating, (a), and the statistically enhanced distribution after ensembling 30 stand-alone NUBEAM runs, (b). The ensembled distribution shows a more elongated and filled in RF-heated tail population.

The magnetic perturbation can originate from almost any source. The radial structure is most commonly found from some MHD eigenmode solver, gyrokinetic code, or analytic theory. The radial structure is then passed to ORBIT to examine the resulting fast ion response. While the radial structure is well-defined from theory or numerical computation, the mode amplitude remains largely unknown. On rare occasion, experimental measurements may constrain the mode amplitude, but this is typically only near the edge far from the actual mode location. Comparisons between global parameters such as the measured and TRANSP calculated neutron rates have been a common procedure[32] to estimate the mode amplitude, but uncertainties around 20% in JET's measured neutron rate and TRANSP output make this difficult. Instead, a procedure utilizing the model output and Faraday cup measurements is discussed toward the end of this section as an alternative.

2.2. Reverse Integration Biasing Scheme

Full particle orbits are integrated backwards in time originating from the Faraday cup apertures. This is accomplished with an ad-hoc code that computes the full orbit motion and is featured in references [25, 26]. The orbit integration is carried out in the presence of only the equilibrium magnetic field, so the resulting motion is absent of any perturbative effects. Figure 3 presents an example orbit for a 1.5 MeV deuteron integrated backwards from the detector to the inner plasma region. The resulting orbit trajectory is

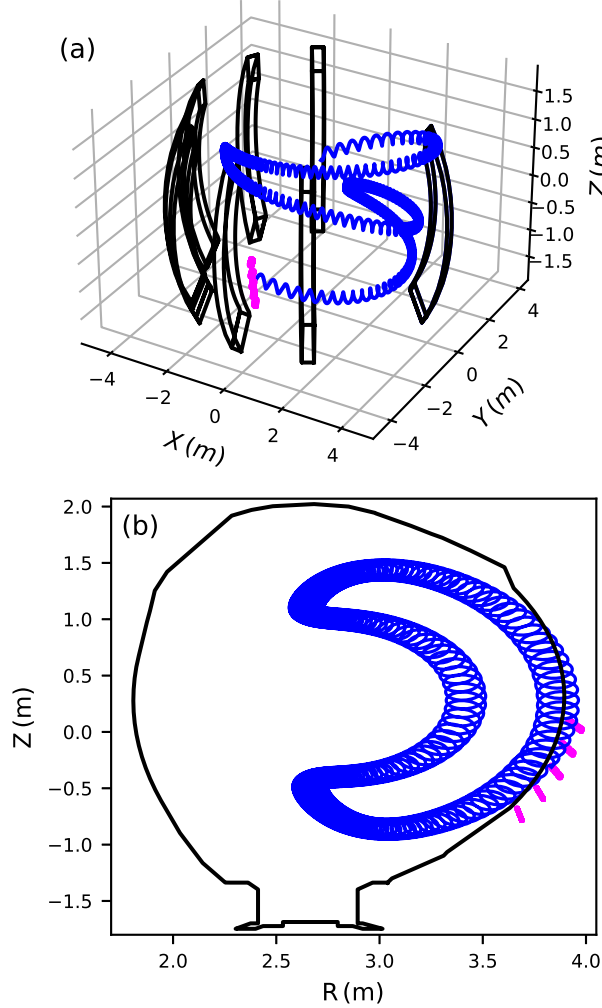


Figure 3. The full orbit, (a.), and poloidal projection, (b.) of a 1.5 MeV deuteron orbit integrated backward in time starting from one of the Faraday cup apertures. The black vertical structures are the poloidal limiters, magenta represents the Faraday cup aperture locations, and blue is the particle motion. Note that the Faraday cups are within the limiter's shadow.

calculated until a set end time (~ 0.1 ms) or until the particle is relost. For the example of Figure 3, the ion is relost to a collision with one of the poloidal limiters (black finger-like structures).

By starting the fast ion at the detector and integrating backwards in time, we are ensuring that the particle will be exactly lost and recorded by the Faraday cups when moving forward in time. This feature may be exploited by observing that if a particle were to originate anywhere along this path, then there is an increased likelihood that it will be a measured loss. Therefore, by starting fast ions across all of the Faraday cup apertures with varied pitches and energies, a complete distribution in (R, Z, Λ, E) can be ensembled to represent particles strictly lost to the detector. Figure 4 presents such a distribution for 150-2,000 keV deuterons integrated backwards from all of the Faraday cup apertures. The intensity is of arbitrary




Figure 4. The reverse integrated intensity distribution for deuterons in the RZ-plane, (a), and (energy,pitch)-space, (b). The magenta points represent the Faraday cup apertures. Deuteron energies ranged from 150-2,000 keV with pitches valid for the aperture geometry.

units but highlights regions where strong losses should be expected. Unsurprisingly, this is near the plasma edge localized near the detectors. Furthermore, Figure 4 shows that measurable losses are more likely for particles near the trapped-lost boundary (see Figure 3 as an example) with a slight preference toward counter-passing trajectories. While Figure 4 is computed for

deuterons, the same trends hold for the other ion species as well.

The reverse integrated loss distribution is treated in a binary fashion, i.e. either existence or non-existence of an orbit, and biased against the initial particle deposition used in ORBIT. Simply put, particles that lie within the exact loss distribution are considered for synthetic loss detection while all others are ignored. Instead of randomly sampling from the energetic particle distribution function, the reverse integrated distribution will weight the sampling towards particles which are more likely to be lost to the Faraday cups. This is done to mitigate wasted computational effort on fast ions that are deeply confined or lost elsewhere in the machine. If one were to simply choose a particle at random from the fast ion distribution, then the probability for recording a detection event is extremely unlikely. This problem arises from the simple fact that FILDs sample a very narrow parameter space relative to the entire energetic particle distribution. Consequently, previous efforts on synthetic loss detectors have required many particles to achieve reasonable loss statistics and long computational run times.[15, 16, 18]

Since the NUBEAM computed fast ion distribution only consists of confined orbits, the sampling needs to extend across the confined-lost topological boundary so that the biasing can occur. This is accomplished by “smearing,” or displacing, the initial sample points. That is, our particle sampling in (R, Z, Λ, E) will be altered to $(R \pm \delta R, Z \pm \delta Z, \Lambda \pm \delta \Lambda, E \pm \delta E)$ where the δ ’s correspond to the respective displacement in each dimension. By altering the particle initialization in such a manner, two problems are solved: the orbits are displaced from well-confined phase-space areas to loss regions, and the discrepancy between ORBIT’s guiding-center calculations and the full orbit lost distribution is broken. This transformation breaks the integraton and averaging constraints imposed by NUBEAM, but the relative ion flux calculated is conserved. Each sampled particle maintains the original NUBEAM particle flux and carries it until synthetic detection (discussed further in the next subsection).

The displacements, i.e. the δ ’s, represent a transition from well-confined to potentially lost orbits, so their respective magnitudes will be on order of the wave-particle interactions, i.e. the ORBIT-“kicks.” An initial ORBIT-kick simulation with a uniform sampling of particles in phase-space can be conducted to calculate the $(\Delta E, \Delta P_\zeta)$ kicks. The ΔE kick directly converts to the δE displacement while the ΔP_ζ can be translated to the positional shifts δR and δZ . Since there is no direct translation to the pitch displacement, $\delta \lambda$, a conservative value of 0.2 is taken. The initial ORBIT sampling for “smeared” and biased

particle markers is shown in Figure 5 overlaid on the NUBEAM ensembled distribution. While not as




Figure 5. The deuteron distributions in (R, Z) -space, (a), and $(\text{energy}, \text{pitch})$ -space, (b), where the black markers denote the initial sampling to be used in ORBIT with applied “smearing” and biasing towards the reverse integrated loss distribution (Figure 4). The plotted number of sample points has been reduced for visibility.

evident in $(\text{energy}, \text{pitch})$ -space, the biasing effect is more readily observed in the (R, Z) -plane. The markers are more skewed toward the outboard side and Faraday cup positions.

When applying the displacements, the Larmor radius and position of the original particle sampled are, effectively, randomized according to a deterministic probability (the “kicks”). This shift in the initial guiding-center and full orbit positions helps break any conflict arising between the two positional viewpoints among the codes. This permits a reasonable connection among the ad-hoc reverse integrated distribution, the NUBEAM distribution, and ORBIT. Discussed in the previous section, the particle pushing calculations will then be conducted within a guiding-center framework while an effective full-orbit position will be accounted from the particle’s instantaneous gyrophase and Larmor radius. In summary, after sampling

particles from the NUBEAM distribution, “smearing” their distributional positioning, and biasing against the exact loss distribution, the energetic particle trajectories can be monitored in ORBIT to the Faraday cup geometry for synthetic detection.

2.3. Synthetic Flux Calculation

ORBIT calculates the particle trajectories until the ion guiding-center positions approach within a Larmor radius of any edge boundary (in this case the outer wall, limiter, or any Faraday cup detector structure). Discussed at the beginning of Section 2, the guiding-center motion may cross into the vacuum region if needed. The full-orbit position is established from the Larmor radius and gyrophase, and the final coordinates are checked against the Faraday cup geometry. Examples of loss positions found from the effective full orbit position in ORBIT are shown in Figure 6 perturbed from a supplied kink mode (discussed in section 3). Lost ions that land within the Faraday cup apertures (red points in Figure 6) are considered for synthetic detection while all others are discarded. It is important to note that while the biasing scheme preferentially selects particles more likely to be synthetically detected, they are not guaranteed to strictly track to the Faraday cup apertures. The interaction with the magnetic perturbation as well as the initial toroidal position, which is randomized from $[0, 2\pi]$, may alter the particle motion from the apertures. Figure 6 successfully conveys the biasing scheme due to the high number of losses around the detector (red and orange) such that the geometry portrayed in Figure 1 is clearly visible.

Further discrimination occurs in pitch and energy. The aperture dimensions (3mm depth, 3mm radius) define an acceptance angle for incident particles. The acceptance cone can be rewritten with respect to the magnetic field, so that only a window of particle pitches can pass through the aperture into the foil stack. Owing to the overall structure design, the accessible pitches are near the perpendicular as can be seen in Figure 4. Discussed in the introduction, the alternating stack of conductive foils and insulators produces a distinct energy resolution per ion species from which additional lost ions are filtered out.

The final step in the loss model is to calculate the relative synthetic energetic particle flux. This is relatively straightforward after performing the initial sampling and biasing. Marker weights equal to the energetic particle density at the initial position are assigned to every particle. The energetic particle density is taken directly from the NUBEAM distribution with units $[\#/cm^3/eV/(d\omega/4\pi)]$ where cm^3 refers to the given volume element of the distribution based on the 2D grid used in NUBEAM,

eV the energy, and $d\omega$ the solid-angle in velocity space. The marker weight can be easily translated to a relative raw number count, $[\#]$, by renormalizing against the initial volume element, energy bin spacing, and velocity space solid-angle of the NUBEAM distribution. Thus, by accumulating the total, renormalized, marker weights to a given Faraday cup, the relative fast ion flux in NUBEAM can be recast to a relative lost ion flux in ORBIT.

Synthetic lost ion flux for every energetic particle species relevant to the experiment needs to be individually calculated and summed together since the experiment measures total current. The procedure is exactly the same for each species with the only differences being the response to the MHD perturbation (ORBIT-kicks), the reverse integrated exact loss distribution, and the foil energy deposition. Computationally, each ORBIT simulation can run simultaneously, so each ion species can execute concurrently and any additional time is often minimal at the operational level. More formally, the synthetic flux can be defined as a modification of the relative NUBEAM energetic particle flux by some weighting function.[35] The synthetic lost ion flux, per species, is mathematically represented in equation 2 for a given Faraday cup foil energy resolution (E_{dep}) and ion species (i):

$$\Gamma_{ORB} = \sum_i \int w_i(R, Z, \Lambda, E) \Gamma_{NUB,i} dR dZ d\Lambda dE \quad (2)$$

where Γ_{ORB} is the total calculated synthetic lost ion flux to the Faraday cups, $\Gamma_{NUB,i}$ denotes the relative energetic particle flux taken from the NUBEAM distribution for species i , and $w_i(R, Z, \Lambda, E)$ represents the corresponding weighting function. The weighting function shifts the relative NUBEAM flux to a more realistic flux measured by the Faraday cups. This encompasses the exact loss distribution biasing, aperture cone, pitch-angle discrimination, and energy deposition:

$$w_i(R, Z, \Lambda, E) = f_{bias,i}(R, Z, \Lambda, E) \cdot f_{cone}(\Lambda) \cdot \delta(E - E_{dep,i}) \quad (3)$$

where δ denotes the Kronecker-delta function, $f_{cone}(\Lambda)$ represents an arbitrary function that converts the particle pitch relative to the aperture normal to compare against the detector acceptance cone, and $f_{bias,i}(R, Z, \Lambda, E)$ accounts for the biasing effect against the reverse integrated distribution. Recall that the reverse integrated exact loss distribution is treated in a binary fashion so that:

$$f_{bias,i}(R, Z, \Lambda, E) = \begin{cases} 1, & \text{if } f_{rev,i}(R, Z, \Lambda, E) \neq 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

and constrains the relative lost ion flux to the intersection of the NUBEAM and reverse distributions.

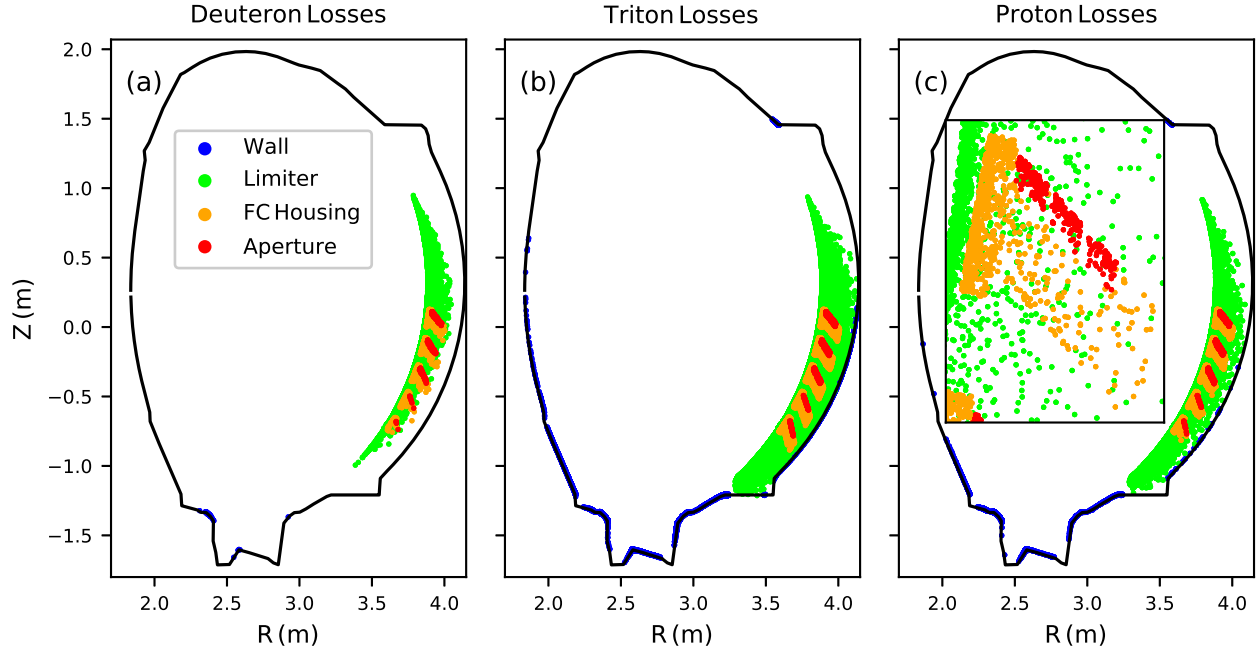


Figure 6. Fast deuteron, (a), triton, (b), and proton, (c), loss positions after including finite-Larmor effects due to a supplied kink mode. The relevant loss locations are the wall (blue), limiter (green), Faraday cup housing/shielding (orange), and the Faraday cup apertures (red). See Figure 1 for comparison. The inset portion in subplot (c) focuses on the uppermost pylon and highlights the three distinct radial cups.

Since particles initialized in proximity to the loss region have an increased likelihood for being lost, the biasing artificially inflates the relative lost flux. Weaker modes, which translate to smaller ORBIT-kicks, smear/displace the biased particles closer to the reverse integrated loss distribution while stronger modes place particles farther away. This leads to the counter-intuitive result that weaker modes can produce higher synthetic losses. To counteract this effect, the relative synthetic flux is scaled according to the proportion of phase-space covered by the ORBIT-kick displacement:

$$\Gamma_{ORB} \rightarrow \Gamma_{ORB} \frac{\int f_{NUB} dV_{kick}}{\int f_{NUB} dV_{NUB}} \quad (5)$$

where Γ_{ORB} denotes the synthetic loss flux calculated in equation 2, f_{NUB} is the NUBEAM fast ion distribution, $\int dV_{kick}$ represents the volume integral over the phase-space covered by the wave-resonant ORBIT-kick (i.e. the smearing region), and $\int dV_{NUB}$ is the total fast ion distribution. Scaling the flux by the ratio of the kick-space relative to the distribution space will bolster higher amplitude modes and weaken smaller amplitude modes counteracting the artificial biasing inflation. In the limit of no kicks ($dV_{kick} = 0$), all losses are classical and should be ignored. As the kick sizes approach the entire phase-space ($dV_{kick} \rightarrow dV_{NUB}$), uniform sampling is recovered and the scaling approaches unity. One must make sure

to not overestimate the mode strength to unphysical levels though.

JET's Faraday cups have not been absolutely calibrated, so only the relative lost ion flux can be calculated. Had the detector been absolutely calibrated, such as the lost scintillator probe[36], then additional factors can be added to equations 2 and 3 that would convert the relative flux to raw current.[35] As such, since these calibration factors are unknown, a 1:1 comparison between the synthetic loss signals and experimental measurements cannot be made. Instead, the relative signals between model and experiment will be examined in Section 3.

2.4. Model Summary

While the synthetic loss model appears complex, the translation and calculation of the relative lost energetic particle flux is relatively straightforward. TRANSP/NUBEAM provides an initial fast ion distribution which we explicitly weight toward the loss detector. This is accomplished by biasing against orbit trajectories integrated backwards from the Faraday cup detector geometry. The particles receive marker weights proportional to the energetic particle density calculated in NUBEAM. ORBIT calculates the resulting motion affected by a supplied perturbation representative of some MHD instability. Finite-Larmor radius corrections match lost particles to the Faraday

cup apertures where only certain pitches and energies are tallied as acceptable losses. The process is repeated across all ion species of interest and summed together for a net Faraday cup foil signal.

Step-by-step, the process for obtaining the modeled fast ion losses to a Faraday cup foil in ORBIT is outlined as follows:

- (i) Precompute the ORBIT-kicks and reverse integrated exact loss distribution for the energetic particle species of interest
- (ii) Initialize particles in ORBIT by sampling from the NUBEAM computed distribution (with displacements on order of the ORBIT-kicks) and bias against the reverse integrated distribution
- (iii) Take the fast ion density as marker weights and track the resulting particle motion due to the supplied perturbation(s)
- (iv) Halt the particle motion within a Larmor radius of the wall boundary and establish a final, finite-Larmor, corrected position from the stored gyrophase
- (v) Tabulate all particles which map to the detector apertures, contain the requisit pitch for aperture acceptance, and possess a valid energy for foil deposition
- (vi) Sum over the renormalized marker weights to give the flux of lost ions and scale against the relative kick size
- (vii) Repeat steps (i)-(vi) for all fast ion species of interest and ensemble the results

While Section 2 describes the full methodology and procedure for modeling energetic particle transport and losses, the resulting synthetic lost ion flux needs to be compared against experiment. This allows a fine tuning of the model parameters (mode amplitude, mode structure, multimodal interactions, equilibrium, etc.). By validating the model against measurement, the model is held to a higher degree of certainty and confirms details unobtainable from experiment. The next section is dedicated to providing a full example of the transport/loss model associated with long-lived kink modes and verified against experimental loss measurements.

3. Validating Model Results Against Measurement

3.1. Reference Discharge and Experimental Loss Measurements

A reference discharge with high external heating power and strong low frequency MHD activity was selected to validate the model performance. JET pulse 97493, shown in Figure 7, features 30 MW of deuterium

fueled NBI and 5 MW ICRH injected into a deuterium plasma as part of JET's high performance, Ne seeded, ITER-like discharge.[37] The plasma contains a relative




Figure 7. Basic plasma parameters for JET pulse 97493: toroidal field on axis, (a), plasma current, (b), Thomson scattering core electron temperature, (c), Thomson scattering core electron density, (d), neutron rate, (e), and external heating powers, (f).

stable period in all parameters from about 12 – 16 sec which provides a suitable period for modeling. Additionally, the robust neutron rate demonstrates a high rate of fusion activity and sufficient population of energetic particle species for analysis. The various drops in neutron rate correlate with sawtooth crashes which help constrain the magnetic equilibrium, provide a fast ion loss mechanism, and are indicative of other MHD activity. This report will not focus on the sawtooth-induced transport and is left for future work.[38]

Strong, long-lived, low frequency modes are present as precursors to the sawteeth and produce measureable fast ion losses. Figure 8 presents spectrograms produced from an edge magnetic Mirnov coil, one of the photomultiplier tubes (PMTs) from the scintillator probe FILD, and a foil measurement from one of the Faraday cups. The magnetics show strong modes with various toroidal mode numbers present throughout the




Figure 8. Spectrograms for shot 97493 produced from an edge magnetic Minrov coil, (a), a photomultiplier tube from the scintillator probe FILD, (b), and a Faraday cup foil FILD measurement, (c). The toroidal mode numbers, n , found from magnetic phase analysis are labeled in subplot (a) for reference.

time domain of interest. The occasional breaks in the spectrogram signals denote the sawtooth crashes. Both the scintillator probe and Faraday cups display coherent losses with the observed MHD modes. In particular, the $n = 1$ mode appears to result in the strongest losses. Differences between the scintillator probe and Faraday cup measurements can be attributed to the difference in positioning (scintillator probe is located approximately 180° toroidally and poloidally equivalent to pylon 3 of the Faraday cup array) and energy sensitivity.

The low frequency modes were confirmed to be kink modes via experimental phase analysis in electron cyclotron emission (ECE) and soft x-ray (SXR) signals. Figure 9 presents the phase of the ECE and SXR signals at the $n = 1$ mode frequency for shot 97493. Since SXR measurements are line-averaged quantities, the radial values shown for each SXR chord were taken as the tangent point relative to the magnetic flux surfaces. A clear phase inversion for both signals exists near the magnetic axis. The noisy jumps at large and small radii can be attributed to edge measurements outside the core plasma region and poor signal-to-noise as the mode amplitude decreases. Considering the axial π -phase jump, low mode number, and presence prior to sawteeth crashes, the modes may be reasonably identified as $n = 1$ internal kink modes with dominant




Figure 9. Phase of electron cyclotron emission (ECE) and soft x-ray (SXR) measurements from JET shot 97493 at the $n = 1$ mode frequency shown in 8 with labeled magnetic axis.

poloidal mode number $m = 1$. [39, 40, 41, 42]

The discharge presents reliable loss measurements to validate the transport model and compare against the synthetic fast ion flux. As such, the model will be constructed during the kink growth phase between the sawtooth crashes where strong-coherent losses are prevalent. The model methodology outlined at the end of Section 2 is followed: a TRANSP simulation constrained by measurement provides the magnetic equilibrium, NUBEAM generates statistically useful energetic particle distributions, and reverse integrated distributions are constructed for biasing. The last free parameters are the choice for the kink mode structure and amplitude.

3.2. Mode structure and Amplitude

Since the $n = 1$ kink mode appears to drive the largest measureable losses, the other toroidal mode numbers will be ignored for simplicity.[43] In practice, however, the model can contain multiple modes. A representative radial displacement is calculated for the discharge from the NOVA code[44] to be used within the ORBIT simulations. Figure 10 displays the resulting normalized radial displacement for the $n = 1$ kink mode including the first three poloidal harmonics. The mode structure is in agreement with analytic kink-like representations[45, 43] and results from the KINX code[46] which exhibit a dominant, core-localized, $m = 1$ component with subdominant $m = 2, 3$ components extending radially outward.

An initial analysis of the mode amplitude can be performed with ORBIT and ECE measurements. The electron response to the NOVA mode structure can be monitored in ORBIT with reference to the initial




Figure 10. A representative mode structure given as the normalized radial displacement for the $n = 1$ kink mode with $m = 1 - 3$ components calculated by NOVA for shot 97493.

T_e profile. Any alteration of the electron position can then be correlated to a fluctuation in temperature. This method follows that used in reference [47] which estimates the sawtooth amplitude from the electron mixing. Figure 11 (a) displays the corresponding δT_e fluctuations found from ORBIT with the NOVA supplied perturbation with an amplitude of $\tilde{b}/B \sim 1e^{-3}$. The $m = 1$ and $m = 3$ contributions are clearly evident with a more subtle $m = 2$ component. The variances in harmonic amplitudes can be attributed to the local T_e gradient at each harmonic's location. The corresponding RMS amplitude is plotted in subplot (b) along with the experimental ECE temperature fluctuations. The error in the ORBIT simulations was found by altering the input T_e profiles within experimental error bars. The $m = 1$ and $m = 3$ harmonics are clearly observed in experiment. The experimental error bars are large near the plasma edge, but a noticeable bump corresponding with the $m = 3$ mode is present. Overall, both the ORBIT and ECE temperature fluctuations appear consistent with the NOVA calculated mode structure. The agreement between experiment and modeling indicates that the mode amplitude can be estimated as $\tilde{b}/B \sim 1e^{-3}$.

The mode amplitude can be independently constrained with the Faraday cup measurements themselves as well. The synthetic losses should vary with mode amplitude and provide a check on the amplitude itself. Stronger modes should exhibit larger radial transport, and higher observed losses. By performing a scan in the ORBIT supplied mode amplitude, the model output can be compared against measurement until a convergence is met. Such a comparison is shown in Figure 12 for the third radial




Figure 11. The ORBIT calculated δT_e fluctuations in (θ, ψ_p) , (a), and the RMS amplitude fluctuations from (a) (green) compared against the experimental δT_e fluctuations found from ECE measurements (blue), (b)

Faraday cup of pylon 1 due to the kink mode presented in Figure 10 in orders of magnitude of \tilde{b}/B . The ratio of foil signals is shown due to the absence of absolute calibration factors. Furthermore, Foil 1 within the Faraday cup stacks is sensitive to capacitive pickup[48] and anomalous current noise[49]. This makes the experimental front foil untrustworthy and is instead used to subtract off noise features present from the deeper foils within the respective stack.[22] The mode amplitudes were kept constant throughout the ORBIT simulations.

A convergence is observed between model and experiment for a mode amplitude around $\tilde{b}/B \sim 6e^{-3} - 6e^{-4}$. The trend of higher losses correlating with larger mode amplitude is not immediately clear. This trend is disguised by the examination of the ratio of foil signals as well as the examination of a single Faraday cup. It seems apparent, however, that a reasonable estimate for the mode amplitude is somewhere within the range of $\tilde{b}/B \sim 6e^{-3} - 6e^{-4}$. This is in agreement with the ORBIT analysis and ECE measurements presented in Figure 11. The amplitude of $\tilde{b}/B = 6e^{-4}$ was taken as our saturated kink mode amplitude for further

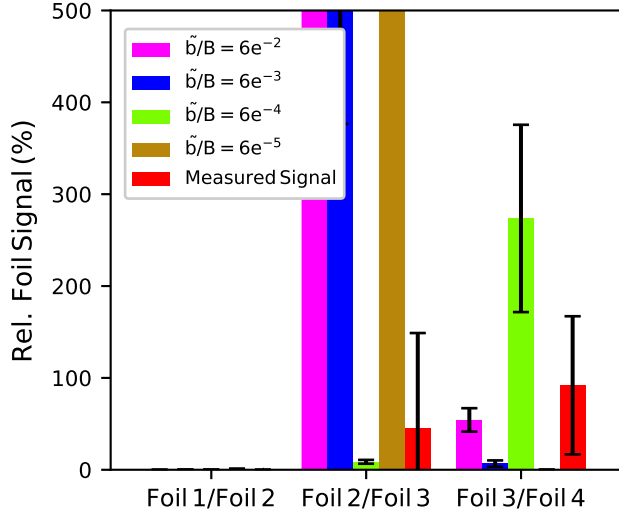


Figure 12. Comparison of the relative FILD signals produced from the synthetic loss model (blue bars) and experimental measurement (red bars) in the third Faraday cup of pylon 1 as a function of the kink mode amplitude. The mode amplitude is normalized to the equilibrium magnetic field on axis.

analysis.

The mode amplitude should be further constrained by examining multiple Faraday cups. Figure 12 could be produced for all of the Faraday cups within the diagnostic array providing an additional constraint from the losses' spatial sensitivity. That is, the losses should have a poloidal sensitivity detectable by the Faraday cup array. Shot 97493, however, occurred during diagnostic upgrades and only seven foil channels were fitted with the suitable data acquisition hardware.[22] Therefore, an experimental examination/comparison of the poloidal dependence on losses is not possible. The synthetic model, however, is well-suited to handle all Faraday cups, as evidenced by the losses in Figure 6, and will be briefly discussed in the next section.

3.3. Comparison of Synthetic and Real Losses

The kink mode perturbation was passed into ORBIT with particle markers from the NUBEAM fast ion distribution to compute the final synthetic flux shown in Figure 12. The case of $\tilde{b}/B = 6e^{-4}$ is examined in more detail here. Beam-injected RF-heated deuterons and DD-fusion born tritons and protons were considered for analysis. Each species was analyzed with 300,000 particles and integrated 2.0 ms in time relative to the mode period of 0.088 ms. Limiter losses were ignored to slightly boost count rates assuming toroidal symmetry but should be included when considering losses due to 3D effects such as resonant magnetic perturbations or toroidal field ripples. Losses that occurred prior to half a wave

period (~ 0.04 ms) were discarded to remove any orbits that were initialized directly on lost trajectories (i.e. prompt losses).

For added clarity, only the relevant portions with $\tilde{b}/B = 6e^{-4}$ mode amplitude from Figure 12 are shown below in Figure 13. The experimental signals were




Figure 13. Comparison of the relative FILD signals produced from the synthetic loss model (blue bars) and experimental measurement (red bars) in the third Faraday cup of pylon 1. The signals shown correlate with the $n = 1$ kink mode present in JET pulse 97493 with a normalized mode amplitude of $6e^{-4}$.

taken at the time of interest and filtered around the $n = 1$ mode frequency of 11.3 kHz to remove any contribution from the other modes. The frequency filtering also removes any contribution from non-resonant prompt losses which may occur from the NBI deposition. Lastly, the signals were averaged over a 0.5 ms time window to achieve the final values and error bars. The model error bars stem from the counting statistics for each foil and ion species.

Examining Figure 13, the synthetic signals are within reasonable agreement to the measured signals signifying a quantitative verification of the transport model with experiment. Furthermore, agreement has been achieved with only 300,000 particles (per species) which is relatively low for a synthetic FILD. This can be attributed to the successful implementation of the reverse biasing scheme which skews the particle sampling toward synthetic detection events. If the model is repeated without the biasing scheme and a uniform sampling across the NUBEAM distribution, then the synthetic count rates plummet to only a few particles. Since the aperture area is vanishingly small with respect to the wall surface area, either a large number of marker particles must be introduced, or the sampling must be weighted towards an increased likelihood of detection (i.e. the biasing scheme).

The losses can be mainly attributed to confined

trapped orbits which are transported outward by the $n = 1$ kink mode to lost orbit trajectories. Figure 14 displays the initial orbit types within the fast ion's constants of motion space[8] for all ions lost across the detector array. The boundaries are representative since

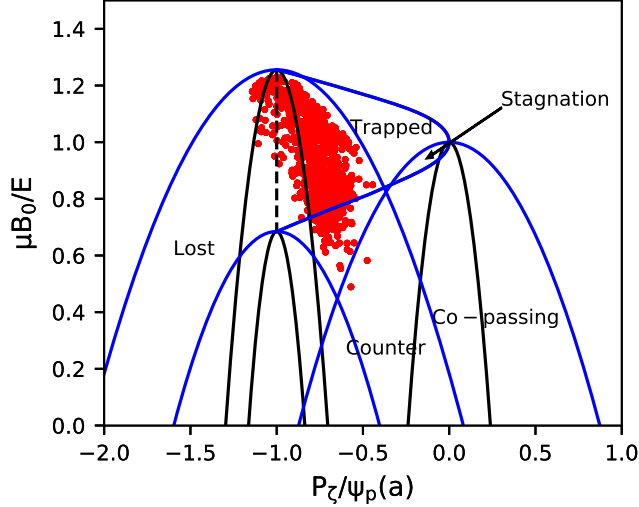


Figure 14. Fast ion orbit topology plot showing the initial orbit positions of the aperture detected particles across all pylons and species. The topological areas denoting co-passing, counter-passing, trapped, stagnation, and lost orbits are labeled for reference. The black boundaries represent the minimally allowed energy (150 keV proton) while the blue boundaries represent a maximally allowed energy (2000 keV deuterons)

the plot contains varying ion species and energies, but the overall deposition remains similar. Furthermore, since orbits are allowed to transition back and forth across the last closed flux surface, the traditional orbit boundaries in (E, P_ζ, μ) -space are no longer strictly defined by the magnetic geometry. The loss boundary in Figure 14 should be shifted to the left (radially outward) to account for full orbit effects and the plasma-vacuum transitions. The biasing effects from Figures 3 and 4 previously hinted that low pitch orbits were the most susceptible for detection, and Figure 14 confirms this. The synthetically recorded orbits predominantly originate from trapped, stagnation, and counter-passing orbits.

It is difficult to make any statement regarding the fast ion's impact on mode stability for the studied example. The ORBIT particles experience no net energy exchange; the ions that gain energy from the kink mode are approximately equal to those that damp the mode. It appears that the particles are not contributing energy to the mode while experiencing strong radial transport. Kink mode stability, however, is more strongly tied to the q -profile shape and plasma beta than any fast ion affect, so any analysis from ORBIT lacks the full stability calculations.

Purely energetic particle driven modes, such as Alfvén eigenmodes, would be a more interesting example in regards to mode stability and transport. The ORBIT-kick code, and thus the synthetic loss model, are fully capable of analyzing such modes in future work.

While an experimental comparison among the poloidal dependence cannot be made for this discharge, the trends can be observed from the model results. Figure 15 presents the poloidal distribution of losses to all loss surfaces among all ion species as a function of mode amplitude. A log scale is used because




Figure 15. The poloidal distribution of the relative synthetic loss flux for all particle species as a function of mode amplitude.

the losses associated with the $\tilde{b}/B = 6e^{-2}$ mode amplitude are orders of magnitude higher than the other cases. This is a consequence of the large mode amplitude allowing core-localized particles with high density to be lost and indicates an overestimation of the mode amplitude. Overall, the expected result that higher mode amplitudes correlate with increased losses is clearly present. Losses peak closer to the midplane (nearer to Pylon 1 and the Faraday cup of interest) and decrease as the poloidal angle approaches the divertor region. Figure 6 (a) displays the same trend among the deuterons which are the dominant contributor to loss particle flux. The $\tilde{b}/B = 6e^{-2}$ & $6e^{-3}$ smearing/displacement factors did not drastically differ which resulted in similar profiles in Figure 15. In practice, after all hardware upgrades have been completed, Figure 15 can be constructed using only the aperture collected losses and compared against experiment. This would provide a further spatial constraint on the mode amplitude in addition to other physical insights.

After experimental validation, the various model attributes and outputs can be analyzed for physical intuition not present in experiment. This includes the

species dependencies absent from the experimental raw current measurements. Figure 16 presents the distribution of each ion species signal across the four Faraday cup foils. The results are in good qualitative




Figure 16. The distribution of the ion signal as a function of Faraday cup foil. The DD-triton birth energy is 1.01 MeV and detectable within Foil 2, the DD-proton birth energy is 3.02 MeV which is greater than the resolveable energy range of the Faraday foil stacks, and the 120 keV NBI deuterons are measureable by Foil 1.

agreement with the expected energy populations of the respective species. As shown in Table 1, the DD-triton birth energy (1.01 MeV) is detectable by Foil 2, and the proton birth energy (3.02 MeV) is greater than the foil stack energy range. This means tritons

Faraday Foil	^1_1H	^2_1H	^3_1H
1	0.0-0.49	0.0-0.49	0.0-0.50
2	0.68-0.96	0.79-1.10	0.84-1.20
3	1.10-1.32	1.35-1.60	1.48-1.76
4	1.45-1.65	1.78-2.00	2.00-2.25

Table 1. Fast ion energy deposition in MeV for the fast ion species of interest as a function of Faraday cup foil.

should largely be expected in Foils 1 and 2 and protons should be mostly present in Foil 4 and follow a slowing distribution. The ICRH deuterons are so far along the RF-tail, that the relative difference between foil energy ranges is minimal and the signal is approximately constant across foil depth. Of note, 150 keV was taken as the minimal energy for all ion species which would exclude all beam ions. This was intentional since the front-foils are experimentally untrustworthy and emphasis was placed on the higher energy losses. It is theorized that this is the reason why the synthetic signal is relatively absent in Foil 1.

The synthetic raw counts are presented in Table 2 for each ion species and Faraday foil. The raw

markers are converted to synthetic flux via the procedure outlined in Section 2. While the deuteron

Ion Species	Foil 1	Foil 2	Foil 3	Foil 4
Deuterons	0	0	18	12
Tritons	48	432	0	0
Protons	0	0	36	66

Table 2. Kink mode induced synthetic loss counts for shot 97493 as a function of ion species and foil depth.

counts are small relative to the proton and triton numbers, the relative signal/flux is much higher since the beam-born density is much greater than the fusion product density. Overall, Figure 16 displays the above noted trends providing further confidence in the model efficacy and confirming features absent in the experimental Faraday cup signals.

4. Discussion and Conclusion

An energetic particle transport model capable of producing loss signals consistent with experimental measurements from JET’s Faraday cup FILD has been constructed. The model is built on the TRANSP/NUBEAM and ORBIT-kick code frameworks. New features such as the inclusion of an effective finite-Larmor radius position and the ability to operate beyond the last closed flux surface have been introduced to the ORBIT code. The energetic particle distributions are produced from the NUBEAM submodule of TRANSP and bolstered with additional statistics. An ad-hoc, full orbit, reverse integration code is used to construct an exact loss distribution that is biased against the NUBEAM distribution skewing the particle sampling toward loss detection. The energetic particle density is taken as marker weights for the particles which are tracked to a realistic Faraday cup FILD geometry in ORBIT as a result of a supplied perturbation. The marker weights are ensembled across all ion species of interest and foil energy ranges to produce the final synthetic lost ion flux.

The model has been successfully verified against a deuterium JET discharge with high D-NBI and ICRH power. The JET pulse exhibits robust kink mode induced losses which are modeled with a representative mode structure obtained from NOVA. The mode amplitude was constrained by finding a convergence in the model output and experimental measurements. The experimental loss measurements are compared to the model output and found to be in good quantitative and qualitative agreement. Furthermore, the analysis of loss signals in fast ion phase-space, poloidal angle, and ion species demonstrates the model robustness and capabilities not present in experiment.

A key feature of the synthetic transport/loss

model is the biasing of forward integrated particles against reverse integrated trajectories stemming from the detector apertures. This methodology discards strongly confined orbits unlikely to be lost and favors those that do. This allows for a lower number of simulated particles to be used for improved computational efficiency. The example shown in Section 3 utilized only 300,000 particles for each species which is approximately a factor of 3 in reduction to the typical 10^6 used in forward-loss models. The computational time varied around 160-768 computer hours per run and typically used 16 or 32 CPUs. The variance in compute time arises from the varied mode amplitudes and ion species. Considering all species (3) and mode amplitudes (4), the total computational time can be estimated around 7500 compute hours which is significantly less than a typical forward loss model when considering the multiple ion species and mode amplitude scan.[18] Additionally, submitting multiple runs as simultaneous batch jobs produces further parallelization (i.e. all species can be ran concurrently instead of consecutively) and enhances workflow.

Improvements can be made to further improve the loss statistics since Figure 6 portrays a relatively large number of particle trajectories tracking around the detector and Figure 13 has modest error bars. Focusing on a single Faraday cup could enhance both computational time and count statistics. Currently, the model examines losses across all Faraday cups/apertures. If one is interested in only a single cup (such as presented in this paper), then the biasing can be localized to the desired cup. However, the wide spatial resolution JET's Faraday cup array encourages the use and analysis of all cups. A more sophisticated toroidal deposition than the current random placement may be needed though. Lastly, reinjecting lost particles may also slightly boost statistics.

A purely reverse integrated modeling scheme could achieve substantially higher lost counts since all particles would be guaranteed losses. This would require further adaptations to the ORBIT code. It is also worth noting that recent work is being performed to model losses with a reverse Monte-Carlo scheme in which the random variations in an ion's motion can be described via a recursion relation to construct a trajectory probability map in phase-space.[50, 51, 52] The numerical computation of the recursion relation should be much faster than the tracing of many particle orbits.

Overall, the validation of an energetic particle transport model against experimental measurements presents vast opportunities. A more detailed analysis of the fast ion transport can be conducted with greater faith in the numerical results while features lacking in experiment can be revealed. The computational

efficiency and speed improves upon previous work and permits more rapid parameter/loss scans. This is particularly useful in multi-species scenarios where the overall complexity increases such as JET's DT-campaign in the fall of 2021. The distinction between alpha and beam-born losses will prove critical for transport interpretation and analysis. Lastly, the model provides a means for future scenario development to study expected fast ion transport and losses in planning experimental campaigns and device construction.

Acknowledgements

This manuscript is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, and has been authored by Princeton University under Contract Number DE-AC02-09CH11466 with the U.S. Department of Energy. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work was supported in part by the Swiss National Science Foundation.

Additionally, the authors would like to thank the work of all JET task force leaders, machine operators, and diagnosticians for the shot involved in this work. Marina Gorelenkova's and Nicola Bertelli's assistance with the functionality and use of stand-alone NUBEAM is also greatly appreciated.

References

- [1] Joffrin E *et al.* 2019 *Nucl. Fusion* **59** 112021
- [2] Aymar R *et al.* 2002 *Plasma Phys. Control Fusion* **44** 519
- [3] Creely A J *et al.* 2020 *J. Plasma Phys.* **86** 865860502
- [4] Heidbrink W W and Sadler G J 1994 *Nucl. Fusion* **34** 535
- [5] Heidbrink W W and White R B 2020 *Phys. Plasmas* **27** 030901
- [6] Zweben S J *et al.* 1995 *Nucl. Fusion* **35** 893
- [7] Zweben S J *et al.* 1995 *Nucl. Fusion* **35** 1445
- [8] Heidbrink W W 2008 *Phys. Plasmas* **15** 055501
- [9] Fredrickson E D *et al.* 2009 *Phys. Plasma* **16** 122505
- [10] Podesta M *et al.* 2014 *Plasma Phys. Control Fusion* **56** 055003
- [11] Collins C S *et al.* 2017 *Nucl. Fusion* **57** 086005
- [12] Yavorskij V *et al.* 2015 *J. Fusion Energ.* **34** 774
- [13] Pace D C *et al.* 2014 *Rev. Sci. Instrum.* **85** 11D841
- [14] Pinches S D *et al.* 2006 *Nucl. Fusion* **46** S904
- [15] Garcia-Munoz M *et al.* 2016 *Rev. Sci. Instrum.* **87** 11D829
- [16] Akaslopolo S *et al.* 2016 *Fusion Sci. Technol.* **69** 620
- [17] Oyola P *et al.* 2021 *Rev. Sci. Instrum.* **92** 043558
- [18] Sipila S *et al.* 2021 *Nucl. Fusion* **61** 086026
- [19] Chen X *et al.* 2012 *Rev. Sci. Instrum.* **83** 10D707
- [20] Kim J *et al.* 2012 *Rev. Sci. Instrum.* **83** 10D305
- [21] Darrow D S *et al.* 2013 *Nucl. Fusion* **53** 013009
- [22] Bonfigli P J *et al.* 2020 *Rev. Sci. Instrum.* **91** 093502
- [23] Gonzalez-Martin J *et al.* 2021 *Rev. Sci. Instrum.* **92** 053538

- [24] Perez von Thun C *et al.* 2011 *Nucl. Fusion* **51** 053003
- [25] Kiptily V *et al.* 2018 *Nucl. Fusion* **58** 014003
- [26] Fitzgerald M *et al.* 2019 *Nucl. Fusion* **59** 016004
- [27] Pace D C *et al.* 2010 *Rev. Sci. Instrum.* **81** 10D305
- [28] Darrow D S *et al.* 2004 *Rev. Sci. Instrum.* **75** 3566
- [29] Baeumel S *et al.* 2004 *Rev. Sci. Instrum.* **75** 3563
- [30] doi:10.11578/dc.20180627.4
- [31] White R B *et al.* 1984 *Phys. Fluids* **27** 2455
- [32] Podesta M *et al.* 2017 *Plasma Phys. Control Fusion* **59** 095008
- [33] Pankin A *et al.* 2004 *Comput. Phys. Commun.* **159** 157
- [34] Brambilla M *et al.* 1999 *Plasma Phys. Control Fusion* **41** 1
- [35] Galdon-Quiroga J *et al.* 2018 *Plasma Phys. Control Fusion* **60** 105005
- [36] Rivero-Rodriguez J F *et al.* 2021 *Rev. Sci. Instrum.* **92** 043553
- [37] Giroud C *et al.* 2021 *High performance ITER-baseline discharges in deuterium with nitrogen and neon-seeding in the JET-ILW* (28th IAEA-FEC Nice, France)
- [38] Teplukhina A A *et al.* 2021 *Nucl. Fusion* **61** Accepted
- [39] Papp G *et al.* 2011 *Plasma Phys. Control Fusion* **53** 065007
- [40] von Goeler S *et al.* 1974 *Phys. Rev. Lett.* **33** 1201
- [41] Buratti P *et al.* 2006 *Plasma Phys. Control Fusion* **48** 1005
- [42] Buratti P *et al.* 2012 *Nucl. Fusion* **52** 023006
- [43] Kim D *et al.* 2018 *Nucl. Fusion* **58** 082029
- [44] Cheng C Z and Chance M S 1986 *Phys. Fluids* **29** 3695
- [45] Farengo R *et al.* 2013 *Nucl. Fusion* **53** 043012
- [46] Degtyarev L *et al.* 1997 *Comput. Phys. Commun.* **567** 10
- [47] Podesta M *et al.* Submitted 2021 *Plasma Phys. Control Fusion*, “Development of a reduced model for energetic particle transport by sawteeth in tokamaks”
- [48] Darrow D S *et al.* 2010 *Rev. Sci. Instrum.* **81** 10D330
- [49] Cecil F E *et al.* 2010 *Rev. Sci. Instrum.* **81** 10D326
- [50] Hirvijoki E *et al.* 2013 *Phys. Plasmas* **20** 092505
- [51] Zhang G *et al.* 2017 *Phys. Plasmas* **24** 092511
- [52] Hirvijoki E 2019 *Eliminating poor statistics in Monte-Carlo simulations of fast-ion losses to plasma-facing components and detectors* (arxiv:1905.04952)