Fast-ion transport by Alfvén eigenmodes above a critical gradient threshold

W. W. Heidbrink,[1,a),b)] C. S. Collins,[2] M. Podestà,[3] G. J. Kramer,[3] D. C. Pace,[2] C. C. Petty,[2] L. Stagner,[1] M. A. Van Zeeland,[2] R. B. White,[3] and Y. B. Zhu[1]
[1]*University of California Irvine, Irvine, California 92697, USA*
[2]*General Atomics, San Diego, California 92186, USA*
[3]*Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543, USA*

(Received 1 December 2016; accepted 1 February 2017; published online 9 March 2017)

Experiments on the DIII-D tokamak have identified how multiple simultaneous Alfvén eigenmodes (AEs) lead to overlapping wave-particle resonances and stochastic fast-ion transport in fusion grade plasmas [C. S. Collins *et al.*, Phys. Rev. Lett. **116**, 095001 (2016)]. The behavior results in a sudden increase in fast-ion transport at a threshold that is well above the linear stability threshold for Alfvén instability. A novel beam modulation technique [W. W. Heidbrink *et al.*, Nucl. Fusion **56**, 112011 (2016)], in conjunction with an array of fast-ion diagnostics, probes the transport by measuring the fast-ion flux in different phase-space volumes. Well above the threshold, simulations that utilize the measured mode amplitudes and structures predict a hollow fast-ion profile that resembles the profile measured by fast-ion D_α spectroscopy; the modelling also successfully reproduces the temporal response of neutral-particle signals to beam modulation. The use of different modulated sources probes the details of phase-space transport by populating different regions in phase space and by altering the amplitude of the AEs. Both effects modulate the phase-space flows. *Published by AIP Publishing.* [http://dx.doi.org/10.1063/1.4977535]

I. INTRODUCTION

Adequate confinement of fast ions is essential for a practical magnetic fusion reactor. Fast-ion transport by instabilities, especially instabilities driven by the fast-ion population itself, is of particular concern, as the resonant interactions that destabilize the waves necessarily alter the fast-ion orbits. The resultant spatial transport affects the heating profile and can cause concentrated losses that damage vessel walls. A dangerous class of instabilities is the various types of Alfvén eigenmodes (AEs) that are driven primarily by the fast-ion spatial gradient.[1] AEs driven unstable by fast ions have been observed in most magnetic fusion experiments[1] and are predicted to be unstable in ITER.[2] If fast ions drive AEs unstable in ITER, it is likely that multiple AEs will be unstable, as the linear stability threshold for modes with different toroidal mode numbers has similar values.

Regimes with multiple unstable AEs are experimentally accessible in the DIII-D tokamak. Neutral-beam ions constitute the fast-ion population. Qualitatively, fast-ion transport in these regimes reproduces familiar features of thermal electron transport. These features are illustrated schematically in Fig. 1. One feature, called "profile resiliency" or "profile consistency" is the tendency of the electron temperature profile to adopt roughly the same shape, regardless of the applied heating profile.[3] Initially, as the heating power increases, the temperature increases proportionally but, as a critical gradient is approached, further increases in heating power have no effect. A second feature is the rapid increase of the amplitude of the limiting instabilities above a threshold (Fig. 1(b)). Once the driving gradients exceed the linear stability threshold, further increases in power only drive the modes harder. The effect of the modes on transport is evident in two ways. If one performs a power-balance analysis, the inferred diffusivity steadily rises with increasing power (Fig. 1(c)). The degraded overall confinement is caused by the enhanced fluctuations. Alternatively, one can measure the incremental flux associated with modulation of the heating power about a particular operating point. Below a critical threshold (when the modes are still stable), the modulated flux is small but, above threshold, the flux increases rapidly (Fig. 1(d)). This feature is called "stiff" transport. If one infers a transport coefficient from the modulated flux, it can exceed the power-balance coefficient by an order of magnitude.[4]

Section II of this paper summarizes recent DIII-D experiments that exhibit the four features associated with stiff transport. However, there is an important distinction between the fast-ion behavior and thermal transport. In the conventional picture of stiff transport, the threshold is associated with the linear stability threshold of an instability such as the electron temperature gradient mode (e.g., Ref. 5). For the fast ions, recent work[6] shows that the threshold for appreciable transport occurs at a threshold that is well beyond the linear stability threshold. The threshold occurs when orbits become stochastic in the portion of phase space interrogated by the fast-ion diagnostic. Because different fast-ion diagnostics are sensitive in different parts of phase space, the observed threshold occurs at different power levels for different diagnostics. Details of the comparison between the measured threshold and stochasticity theory are the focus of a forthcoming paper.[7] Because different modes interact with fast ions in different parts of phase space, the measured threshold depends sensitively on both the diagnostic and on the fast-ion population that drives the modes.[7]

Note: Paper BI2 1, Bull. Am. Phys. Soc. **61**, 20 (2016).
[a)]Invited speaker.
[b)]Bill.Heidbrink@uci.edu

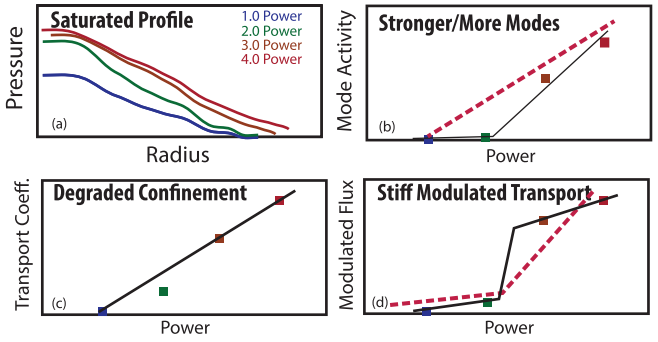

FIG. 1. Schematic illustration of the critical-gradient behavior. (a) As the auxiliary heating power is increased on successive shots, the profiles saturate with a shape that becomes insensitive to further changes in power or deposition profile. (b) The amplitude of instabilities increases with increasing power. The solid line with points indicates the behavior when a linear stability threshold is crossed. The dashed line illustrates the behavior that is actually observed for Alfvén eigenmodes. (c) Transport coefficients such as the electron thermal diffusivity inferred from power balance steadily increase with increasing power. (d) The modulated flux measured at different average heating powers jumps suddenly above a critical threshold. The solid line indicates a typical critical-gradient behavior, while the dashed line sketches the observed behavior for fast-ion transport by AEs.

Detailed probing of transport in phase space utilizes the combination of a modulated source and an array of fast-ion diagnostics. Similarities and differences between thermal-plasma and fast-ion modulation experiments are discussed in a recent paper.[8] The key difference is that, for thermal modulation experiments, the distribution function is a local Maxwellian while, for fast-ion experiments, the distribution function is a complicated function of phase-space variables. A related difference is that, while thermal diagnostics measure a moment of the distribution function such as the temperature or density, fast-ion diagnostics have a complicated phase-space sensitivity known as a "weight function."[9] Consequently, flows in both velocity space and configuration space affect the measured signals. Indeed, for a diagnostic with narrow sensitivity in velocity space, such as a neutral particle analyzer (NPA), changes in velocity can cause larger changes in signal than changes in position. The measured quantity is the divergence of fast-ion flux from the volume interrogated by the fast-ion diagnostic, $\nabla \cdot \tilde{\Gamma}$. Here, the modulated flux $\tilde{\Gamma}$ includes flows in both velocity and configuration space.

The focus of the present paper is fast-ion transport above the stochastic threshold. Transport in this regime has been studied previously. Sigmar *et al.*[10] first discussed stochastic transport by AEs decades ago. White *et al.*[11,12] performed an extensive analysis of DIII-D discharge #122 117, a discharge with strong AE activity that flattened the fast-ion profile.[13] The observed modes were matched to AEs calculated by a linear ideal magnetohydrodynamic (MHD) code and the amplitudes were adjusted to match the experimental values (as in Ref. 14), then the particle orbits were followed in the Hamiltonian guiding center code ORBIT.[15] Beam deposition and Coulomb collisions were included in the analysis. The analysis showed that, since the mode amplitudes are small ($\delta B/B \lesssim 10^{-3}$), the island widths associated with individual resonances are quite modest, $\Delta\Psi/\Psi_n \sim 2\%$.[11] (Here, Ψ is the poloidal flux, and Ψ_n is the difference between the flux at the magnetic axis and the last closed flux surface (LCFS).) However, because there are hundreds of resonances, and the modes have different toroidal mode numbers n, the small islands overlap and can cause orbit stochasticity. As a result, the calculated fast-ion density profile is flattened in a manner similar to the experiment.[12] Subsequently, White developed a much more efficient way to calculate orbit stochasticity.[16] The new method successfully reproduced the old results for discharge #122 117.[17]

Todo *et al.*[18,19] have analyzed a different DIII-D discharge that is above the stochastic threshold, discharge #142 111. The simulations employ a kinetic treatment for the fast ions and a resistive magnetohydrodynamic (MHD) model with artificially enhanced dissipation[20] for the background plasma. To efficiently model the evolution of the fast-ion distribution function, the simulations alternate between periods with and without AEs. Excellent agreement with observed AE amplitudes and mode structures is obtained, and the calculated fast-ion profile is close to the measured profile.[18] Surfaces of section plots demonstrate that the resonance overlap of multiple eigenmodes is responsible for the sudden increase in fast-ion transport above threshold.[19] Intermittent avalanches take place with contributions from the multiple eigenmodes.

A comparison of a reduced critical gradient model with DIII-D data by Waltz *et al.* found agreement with the fast-ion pressure but not the fast-ion D-alpha (FIDA) data, causing the authors to conclude that a proper treatment of different regions of phase space is essential to reproduce the data.[21]

In this paper, the use of a modulated source provides new information about the fast-ion transport by AEs above the stochastic threshold. Section II summarizes the experimental conditions from the perspective of the critical-gradient paradigm. Section III documents a set of three similar discharges that are well above the transport threshold. In Section IV, analysis by a recently developed "kick" model[22] is compared quantitatively with measured signals. Section V discusses the data in terms of flows in phase space. Conclusions appear in Section VI.

II. CRITICAL GRADIENT BEHAVIOR

According to a recent review of self-organized criticality,[23] the ingredients for sandpile-like transport are a slowly-driven, open system with a fast relaxation mechanism mediated by a local threshold. Alpha transport in future reactors is a candidate for this type of behavior,[24] as is AE-induced fast-ion transport in DIII-D plasmas. Fusion reactions or neutral-beam heating gradually replenish the system, the system is open to fast-ion losses at the edge, local profile relaxation takes place quickly when wave-particle resonances overlap,[24] and the fast-ion gradient drives the modes. The goal of this section is to show that DIII-D plasmas manifest the four qualitative features of critical gradient behavior that are illustrated in Fig. 1. Since much of this material is already published, only a brief summary is given.

The experiments are performed in the DIII-D tokamak. DIII-D is heated by deuterium neutral beams with typical

injection energies of 70–81 keV. The beams inject at different angles (Fig. 2): Co- and counter-current, on- and off-axis, and tangential and perpendicular directions. The plasmas with multiple unstable AEs generally have either reversed-shear or very weak-shear q profiles. Unstable toroidal Alfvén eigenmodes (TAE) and reversed-shear Alfvén eigenmodes (RSAE) are common.[14] Conditions with multiple unstable AEs occur both in L-mode plasmas during the current ramp[6,14,25] and in H-mode steady-state scenario plasmas with elevated q profiles.[26,27] Fast-ion diagnostics include the volume-averaged neutron rate,[28] which is dominated by beam-plasma reactions in the L-mode plasmas, FIDA diagnostics with oblique sightlines,[29] and solid-state neutral-particle analyzers (NPA) operated in current mode[30] (Fig. 2).

The first feature of critical-gradient behavior is profile resiliency (Fig. 1(a)). Already in 2007,[13] FIDA and fast-ion pressure profile measurements showed that the fast-ion profile is very flat in the presence of many unstable AEs, much flatter than "classically" expected. (Here, the fast-ion distribution function that would occur in the absence of wave-induced transport is termed "classical.") Then, in 2013,[25] an experiment that utilized different combinations of on-axis and off-axis injection reported that, once the AE activity was strongly driven, the FIDA profile was invariant to the heating location (Fig. 3(a)). Most recently,[6] the FIDA profiles were measured during a systematic power scan. Initially, as the power increased, the fast-ion profile grew larger but, for higher power levels where the AE activity was strong, the FIDA profile was "clamped" (Fig. 3(b)).

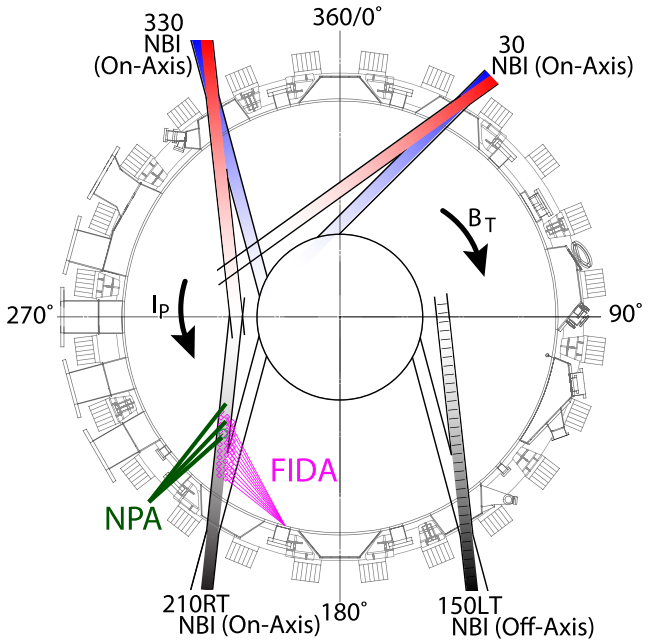

FIG. 2. Plan view of the DIII-D tokamak, showing the direction of injection for the eight neutral beam sources and the sightlines for the FIDA and NPA data shown in this paper. The FIDA (NPA) sightlines view the active beam from a port above (below) the midplane. For the recent critical-gradient experiments, the 210RT beam is on continuously for active FIDA and NPA measurements. On all but two discharges, the off-axis 150LT beam is the modulated source. The four sources at 30° and 330° are used to vary the average heating power shot-to-shot.

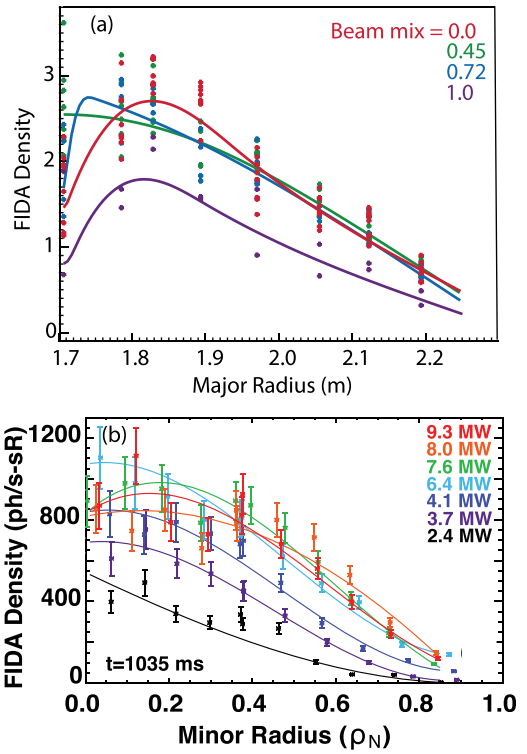

FIG. 3. Resiliency of FIDA profiles. (a) In experiments with different combinations of on-axis and off-axis neutral beam injection, the profiles are invariant to the deposition profile except for pure off-axis injection (called "beam mix = 1.0.") The AE activity is below the fast-ion transport threshold for pure off-axis injection, but above threshold for the other cases. Reproduced with permission from Heidbrink *et al.*, Nucl. Fusion **53**, 093006 (2013). Copyright 2013 International Atomic Energy Agency. (b) In experiments using the configuration shown in Fig. 2, the FIDA profile initially increases when the AE activity is below the transport threshold but saturates for powers above the threshold. Reproduced from Phys. Rev. Lett. **116**, 095001 (2016). Copyright 2016 AIP Publishing. In both figures, the ordinate is the FIDA brightness divided by the injected neutral density.

The second feature of critical-gradient behavior is steadily increasing mode amplitude above a threshold (Fig. 1(b)). Figure 4 shows the overall strength of AE activity, the average number of unstable AE modes, and the average amplitude per mode in the same power scan as Fig. 3(b). The "AE amplitude" plotted in the figure is the sum of the amplitudes of coherent modes between the geodesic acoustic mode (GAM) and TAE frequencies detected by the electron cyclotron emission (ECE) radiometer diagnostic,[31] using the algorithm described in Ref. 25. The observed behavior differs from the typical critical gradient behavior sketched in Fig. 1(b): Even at the lowest heating power, some unstable AE activity is observed. The observed amplitude increases approximately linearly with the heating power.

The third feature is increasing transport with increasing power (Fig. 1(c)). The transport coefficient for this comparison is associated with the *total* heating power. In the AE experiment of Fig. 3(b), the neutron rate deviates from the classical prediction when the AEs are strongly driven. The NUBEAM module[32] of the TRANSP code[33] can apply spatial diffusion to the fast-ion population. This *ad hoc* diffusion coefficient D_B does *not* model the actual physics of the fast-ion interaction with AEs but is useful as a rough estimate of the severity of transport. Recently, the TRANSP code has

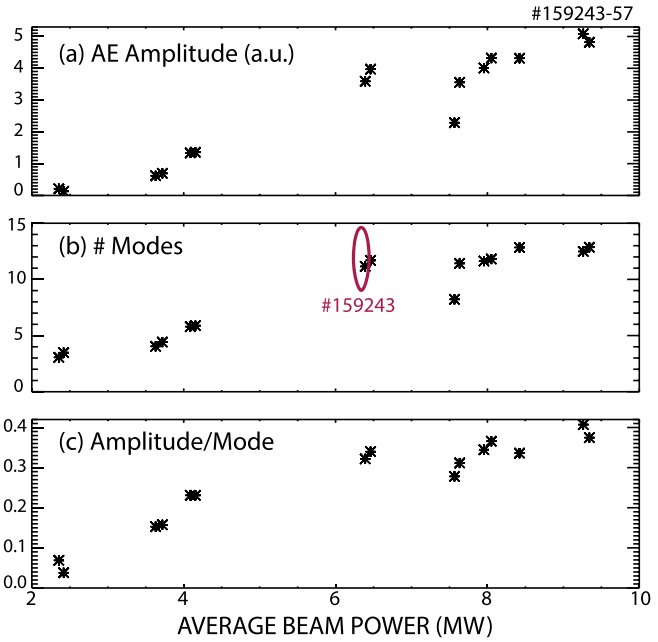

FIG. 4. (a) Total amplitude of coherent AEs $\sum \delta T_e/T_e$, (b) time-average number of unstable modes, and (c) average amplitude per mode for the power scan that shown in Fig. 3(b). The AEs are measured by a 40-channel ECE radiometer.[31] The data are averaged over the time interval 624–897 ms. The oval encircles the discharge analyzed in detail in Secs. III–V.

incorporated the ability to monitor the ratio of measured-to-predicted neutron rate and feedback on the value of D_B to maintain agreement between the measurement and prediction. Application of this capability to the power scan of Fig. 3(b) yields Fig. 5. The fast-ion diffusion coefficient inferred from the neutron deficit increases approximately linearly with the AE amplitude.

The fourth feature is stiff transport above a threshold (Fig. 1(d)). In this case, the transport measurement is associated with the *modulated* heating power. Figure 6 reproduces

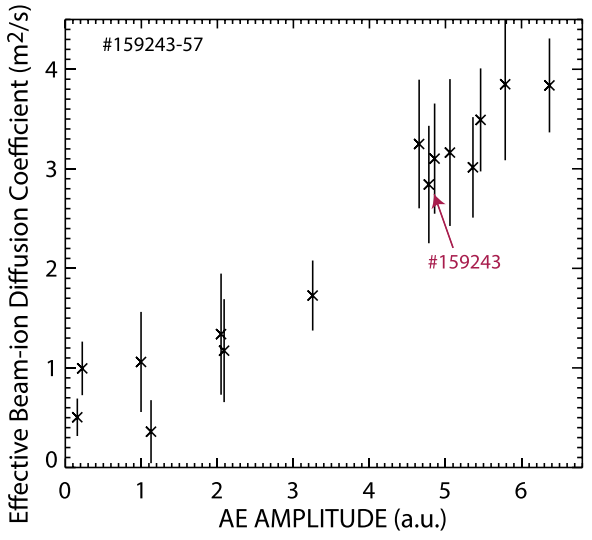

FIG. 5. Ad-hoc beam-ion diffusion coefficient D_B vs. average AE amplitude $\sum \delta T_e/T_e$ for the same scan as Fig. 3(b). D_B is found by matching the measured and calculated neutron rate as a function of time. The error bars represent the standard deviation of D_B between 516 and 897 ms. The discharge analyzed in detail in Secs. III–V is indicated.

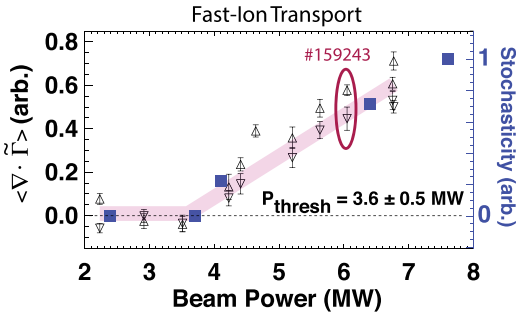

FIG. 6. Divergence of the modulated flux $\nabla \cdot \tilde{\Gamma}$ inferred from the outer NPA channel for the first half (triangles pointing up) and second half (triangles pointing down) of the modulation period. Error bars are the standard deviation of the time average over the half period. The solid symbols represent the number of broken KAM surfaces in the phase-space volume interrogated by the NPA diagnostic. The onset of transport correlates with the theoretical onset of stochasticity. Reproduced from Phys. Rev. Lett. **116**, 095001 (2016) . Copyright 2016 AIP Publishing. The oval encircles the discharge analyzed in detail in Secs. III–V.

a figure that was published in Ref. 6. The modulated flux measured by an NPA diagnostic jumps suddenly above a threshold of about 4 MW. This threshold coincides with the onset of orbit stochasticity in the phase-space volume interrogated by the NPA diagnostic. Other examples of a transport threshold will appear in a forthcoming publication.[7] Evidence that the transport threshold exceeds the linear stability threshold is also found in steady-state scenario plasmas. Figure 4 of Ref. 26 shows a pair of successive discharges, one with modest AE activity and the other with strong AE activity. The fast-ion measurements in the discharge with modest AE activity is consistent within experimental errors with classically predicted FIDA, neutron, and fast-ion stored energy signals, but the discharge with strong AE activity exhibits strong deviations on all fast-ion signals. Evidently, one discharge resides below the stochasticity threshold in most of the phase space while the discharge with strong AE activity is above the stochasticity threshold.

A fifth feature often associated with the critical-gradient behavior is bursts of non-Gaussian transport events above threshold. Bursts of this sort are detected by a fast-ion loss detector.[6] Detailed discussions of these data will appear in a future publication.

III. THREE SIMILAR "ABOVE-THRESHOLD" DISCHARGES WITH DIFFERENT MODULATED SOURCES

The remainder of the paper focuses on discharge #159 243 and two similar discharges. This section documents those discharges. Discharge #159 243 has an average injected beam power of 6.4 MW, ~10 unstable AEs (Fig. 4(b)), effective fast-ion diffusion of ~2.8 m^2/s (Fig. 5), and is well above the observed transport threshold for all fast-ion diagnostics, including the NPA (Fig. 6). Figure 7 shows profiles of electron density n_e and temperature T_e, ion temperature T_i, toroidal rotation Ω_{rot}, Z_{eff} inferred from the carbon density, and q profile for this discharge. Two very similar discharges that differ only by which neutral-beam source is

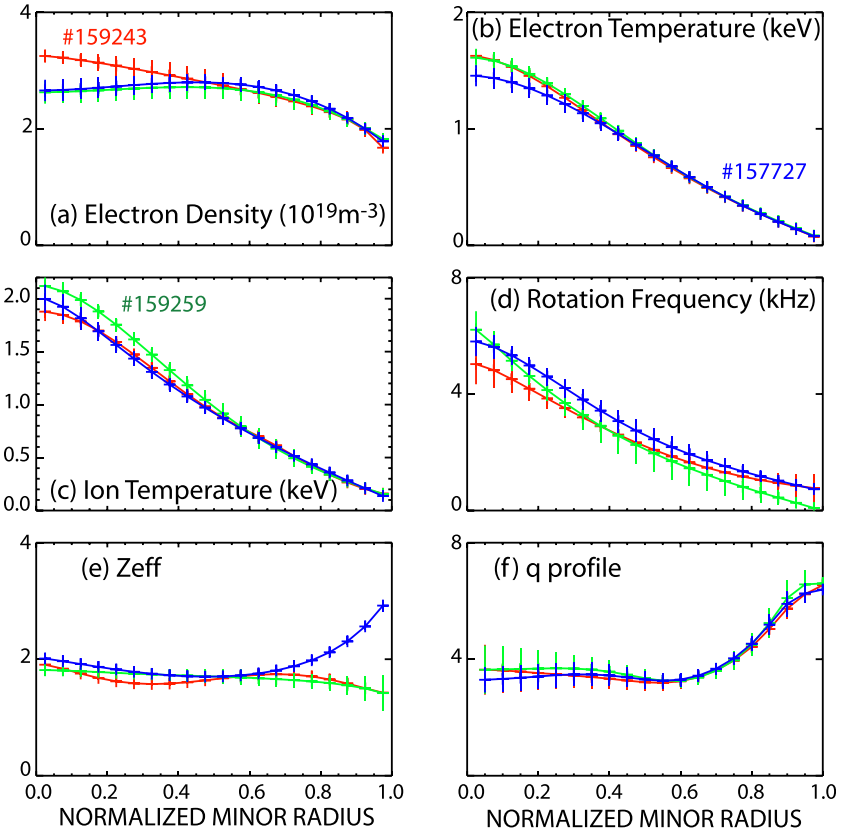

FIG. 7. Profiles of (a) electron density, (b) electron temperature, (c) ion temperature, (d) toroidal rotation, (e) Z_{eff}, and (f) safety factor vs. the square root of the normalized toroidal flux for three discharges with different modulated sources. The profiles are averaged between 516 and 897 ms; the error bars represent the rms variation over this time interval.

modulated are also shown. The differences between the three discharges are comparable to the fitting uncertainties.

The plasma has an oval cross-section and is limited on the inner wall (Fig. 8). In this paper, fast-ion signals are conditionally averaged over an integer number of periods of the modulated beam between 624 and 897 ms. During this time interval, the minimum value of the safety factor q_{min} drops from 3.5 to 2.7, and the normalized radius of q_{min} shrinks from $\rho_{qmin} = 0.48$ to 0.42. (Here, ρ is the square root of the toroidal flux normalized to the value at the last closed flux surface.)

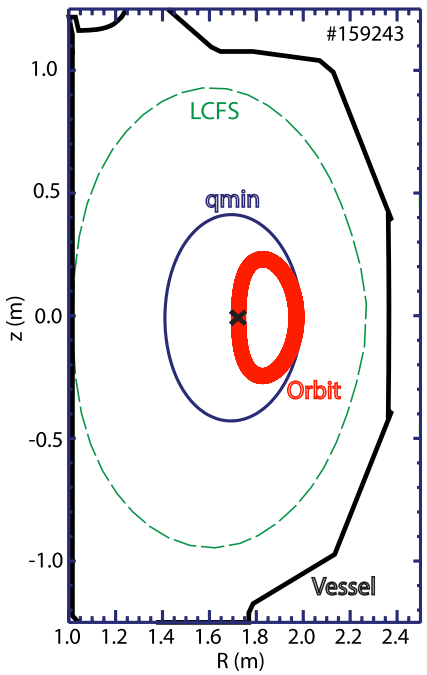

FIG. 8. Elevation of the DIII-D tokamak showing the vessel wall (thick solid line), the last-closed flux surface (dashed line), the minimum q surface (solid line), and the magnetic axis (X) for the equilibrium at 805 ms. The projection of a 75-keV fast-ion orbit is also shown.

Figure 9 compares the AE activity in the same three discharges. Because the average modulated power is a small fraction of the total (10%, 15%, and 14% for the three cases), the three discharges all have similar AE activity. Table I provides details. The modulated beam in discharge #159 243 is a tangential off-axis source; this is the modulated source indicated in Fig. 2. The modulated source in discharge #159 259 is an on-axis tangential source. The modulated source in discharge #157 727 is an on-axis perpendicular source. Quantitatively, for the three cases, the total AE activity and average number of unstable modes differ by 6% and 7%, respectively (Table I).

Section V discusses the measurements of phase-space flows in these three discharges. The presented results are reproducible. To acquire additional FIDA data, the power scan of Figs. 5 and 6 included repeat shots at every power level. The reproducibility of AE amplitude and fast-ion signals on the repeat shots is high throughout the scan. Similarly, although discharge #157 727 is from a different day than the other two discharges in Table I, a companion discharge (#157 726) with an off-axis modulated source is very similar to the three presented discharges. The conditionally-averaged neutron and NPA signals in this discharge closely resemble the signals in the presented off-axis case (#159 243).

The absolute uncertainty in the neutron calibration for the discharges of Table I is $\gtrsim$15%. The FIDA diagnostic is

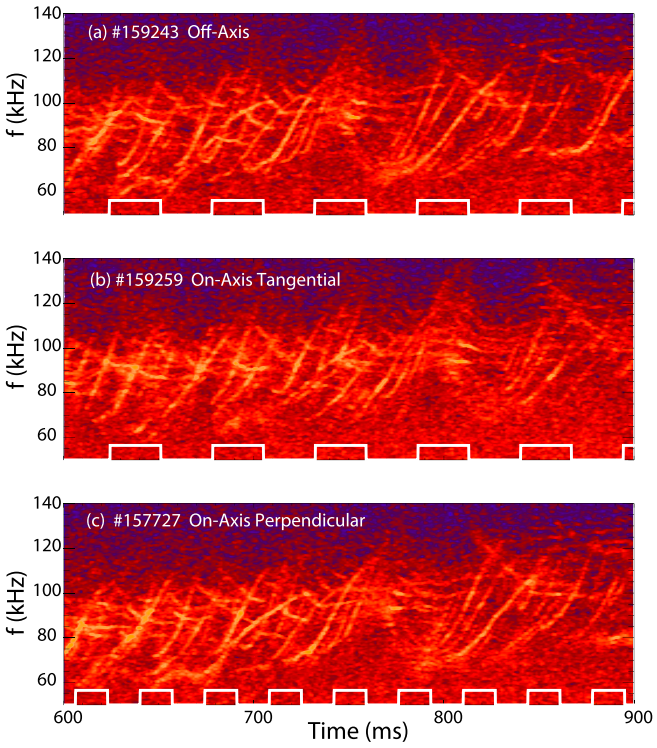

FIG. 9. Spectra from the cross-power of two CO_2 interferometer chords for the three discharges with different modulated sources. The color scale is logarithmic in the amplitude. A timing waveform for the modulated source is also shown. Coherent activity below 50 kHz and above 140 kHz is negligible.

absolutely calibrated with an in-vessel source before and after each campaign. The NPA diagnostic has not been absolutely calibrated. In this paper, a low-power (1.8 MW) calibration shot (#157 725) is used to calibrate the neutron scintillator, to provide an independent check on the FIDA photometric calibration, and to provide a relative cross-calibration of the three NPA channels. This discharge has no detectable MHD in the period just prior to the onset of sawteeth. The expected signals are computed by TRANSP[33] for the neutrons and by FIDASIM[34] for the FIDA and NPA diagnostics.

In an ideal perturbation experiment, the modulated source does not perturb the plasma. This experiment deviates from the ideal. After conditional averaging, the variation in plasma parameters is $\lesssim$2% (Ref. 8) but the AE amplitude is strongly affected by the modulated source (Fig. 10). As discussed below, these variations have a strong impact on the fast-ion signals. The variation is smaller for the on-axis perpendicular source than for the other two sources. (Note that the conditional averaging employed here is relative to the mean value of the signal during the cycle, so the modulated difference signal is both positive and negative. Throughout the paper, linear detrending is applied to the signals prior to averaging.)

TABLE I. Modulated beam parameters (direction, injection energy, power, and modulation period) and time-averaged AE activity between 624 and 897 ms when using different modulated sources.

Shot	Mod. beam	E_{inj} (keV)	P_B (MW)	Period (ms)	AE (A)	# Modes
159 243	Off-axis	71	1.3	54	3.6	10.8
159 259	On-axis tang	81	2.3	54	3.6	10.4
157 727	On-axis perp	69	1.8	34	3.4	11.2

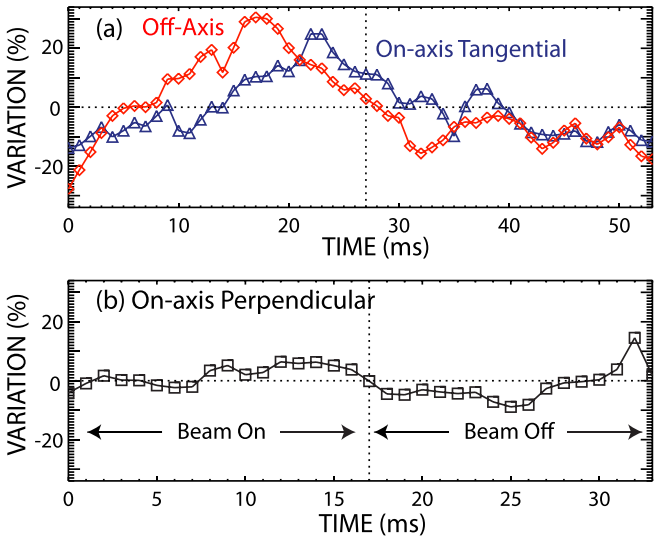

FIG. 10. Fractional variation in the AE amplitude for the discharges in Table I. The modulated sources are (a) off-axis (diamonds), on-axis tangential (triangles), and (b) on-axis perpendicular (square). The data are conditionally averaged over (a) the 54-ms beam period for 5 cycles between 624 and 894 ms and (b) the 34 ms beam period for 7 cycles between 640 and 878 ms.

IV. "KICK" ANALYSIS OF DISCHARGE #159 243

In this section, data from the above-threshold discharge with an off-axis modulated beam (#159 243) are compared with predictions of the "kick" model[22] that has been incorporated into the TRANSP NUBEAM code. Unlike the *ad hoc* diffusion model of Fig. 5, the kick model gives fast ions Monte Carlo displacements in phase space that mimic the actual orbital displacements by AEs. Since AEs have small amplitudes and frequencies that are far below the ion cyclotron frequency, the magnetic moment μ is conserved. Phase space is represented by three constants of motion of the equilibrium orbits: μ, energy E, and the toroidal canonical angular momentum P_ϕ. A consequence of Hamiltonian theory is that, in an interaction with a single mode with frequency ω and toroidal mode number n, the change in energy ΔE is related to the change in toroidal angular momentum ΔP_ϕ by

$$n\Delta E = \omega \Delta P_\phi. \quad (1)$$

This relation is the basis of the kick model. The modes are entered into the ORBIT code and the probability of correlated changes in E and P_ϕ are calculated throughout phase space. These probabilities are inserted into the TRANSP NUBEAM code to calculate the evolution of the distribution function.

The first step of the analysis is to use magnetic, MSE, and thermal pressure data to prepare "kinetic" EFIT equilibria[35] that are consistent with MHD spectroscopy.[36] The ideal MHD code NOVA[37] calculates AEs that are matched by frequency, toroidal mode number, and radial structure to electron cyclotron emission (ECE) profiles, as in Ref. 14. The

NOVA mode amplitude is scaled by the measured amplitude $\delta T_e/T_e$. The scaled NOVA modes are entered into ORBIT[15] to compute probability matrices. Three TAEs with $n = 4$, 5, and 6 are grouped together as a single mode. Probability matrices for the 8 RSAEs with $n = 2$–6 are computed individually. Tests indicate that the transport associated with $n \geq 7$ RSAEs is negligible, so these modes are neglected. For simplicity, the transport probabilities are inferred assuming a fixed q profile and constant frequencies and eigenfunctions for the modes. The implies that the wave-particle resonances are not correctly tracked as the modes evolve and likely underestimate the actual transport, although adjustment of the mode amplitudes (described below) partially compensates for the associated error. As an indication of the significance of this approximation, the magnitude of classical prompt loss calculated by TRANSP changes by 10% between 624 and 897 ms.

The next step is to evolve the mode amplitudes in time. To that end, the kick probabilities are multiplied by amplitude scaling factors that vary in time. Figure 11(c) shows the employed scaling factors for the 9 different modes (or group of modes). These factors are selected to match the measured neutron rate (Figure 11(a)), with the RSAE temporal evolution also taken into consideration. To match the neutrons, the utilized mode amplitudes average ~77% of the experimentally measured values. The utilized kick probabilities are based on ORBIT calculations that neglect the finite Larmor radius (FLR) effects. Tests indicate that, with FLR effects included, the mode amplitudes must be ~30% larger to produce the same kick probabilities. Thus, the utilized amplitudes are reasonably consistent with the measured mode amplitudes.

After conditional averaging, the modeled neutron rate is not an exact match to the measured rate but does follow the general trends (Fig. 11(b)). Figure 11(d) shows an independent check on the modeled amplitude variation. The conditionally-averaged mode scaling factors are close to the measured variation in AE amplitude in the negative portion of the modulation cycle but differ from the measurement in the positive phase. These discrepancies are due (at least in part) to difficulty in matching the measured neutron rate with the kick TRANSP prediction (Figs. 11(a) and 11(b)), even after several iterations.

The output of the kick TRANSP analysis is the fast-ion distribution function f. The forward-modeling code FIDASIM[34] uses f to predict FIDA and NPA signals. The calculated fast-ion density profile is slightly hollow (inset of Fig. 12). This distribution function yields a calculated FIDA profile in excellent agreement with experiment (Fig. 12). The central FIDA signal is only ~1/3 of the classical prediction. Physically, the hollow profile may be associated with the finite orbit size of the fast ions. As shown in Fig. 8, fast ions on orbits that produce appreciable FIDA signal can traverse the magnetic axis and ρ_{qmin}, so RSAE-induced transport can easily modify the central fast-ion density.

Another check on the calculated f is the FIDA spectral shape. Through the Doppler shift, the FIDA spectrum is sensitive to the component of the fast-ion velocity along the FIDA sightline. The spectral shape predicted using the kick model is in excellent agreement with the experimental spectra; Fig. 13 shows an example. The agreement is quite good for all channels except the outermost one—the channel that is most likely to be corrupted by errors in background subtraction. Although the magnitude of the prediction is too large, the spectral shape predicted by the classical calculation also agrees well with the experiment. This suggests that

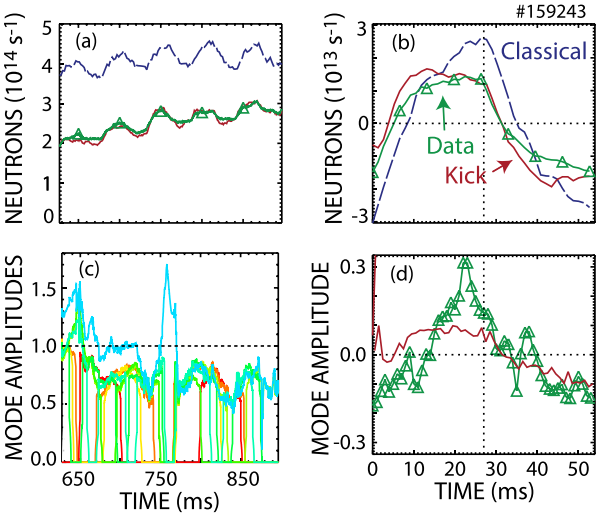

FIG. 11. Comparison of the measured neutron rate (△) with the classical prediction (dashed) and prediction of the kick model (solid) (a) as a function of time and (b) after conditional averaging over 5 cycles of the modulated beam. (c) Evolution of mode amplitudes used in the kick modeling. The dashed horizontal line indicates the approximate value of the measured amplitude at 780 ms. (d) Fractional variation of the measured AE amplitude (△) and of the sum of the modelled amplitudes (solid) after conditional averaging. The vertical dashed lines in panels (b) and (d) show when the modulated beam turns off.

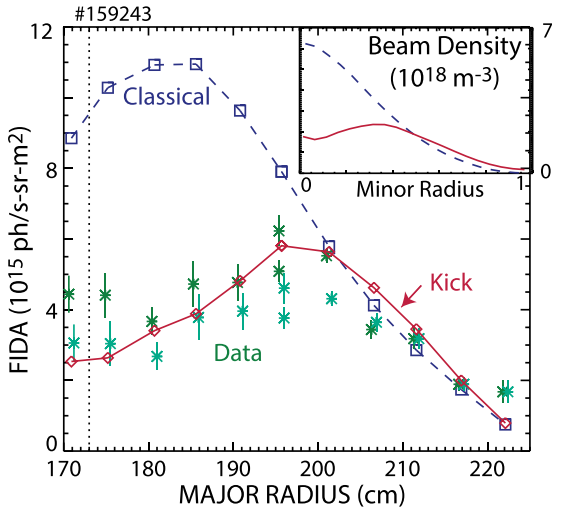

FIG. 12. Comparison of measured (*) FIDA brightness profile at 785 ms with simulated profiles. The dashed line with the square symbols is the classical prediction and the solid line with diamonds is the prediction of the kick model. The inset shows the total beam-ion density vs. normalized minor radius for the two calculations. The FIDA spectra are integrated between 650.7 and 653.25 nm. The error bars represent the uncertainty associated with systematic uncertainty in the background subtraction. The two sets of data utilize different intensity calibrations: the lower data points are from an absolutely calibrated light source during a vent and the upper data points are calibrated relative to a discharge with negligible MHD. The dotted vertical line indicates the location of the magnetic axis.

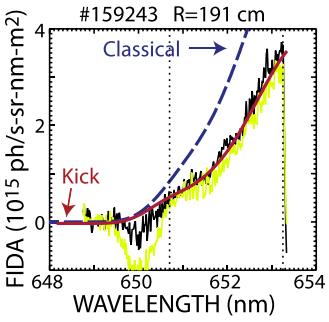

FIG. 13. FIDA spectrum for the channel at $R = 191$ cm. The jagged curves are the data, using two different methods to measure the background. The solid line is the FIDASIM prediction of the kick model; the dashed line is the classical prediction. The spatial profile is obtained by integrating the spectrum between the vertical dotted lines. The theoretically predicted spectral shape agrees well with the experiment. (The data are unreliable near 650 nm due to an oxygen impurity line.)

the AE activity does not cause large distortions in the shape of the velocity distribution function.

The predicted NPA signals are close to the experimental measurements. Figure 14 compares the conditionally averaged measurements for the three NPA sightlines with the kick predictions after processing by FIDASIM. Both the magnitude and time evolution are close to the experimental values—as close, in fact, as the agreement with the neutron signal that was used to fit the mode amplitudes. The NPA diagnostics are sensitive to the trapped-particle population, while the neutron diagnostic is sensitive to all pitch angles. The good agreement of the kick prediction with the NPA signals suggests that the interaction of AEs with trapped particles is accurately treated by the modeling.

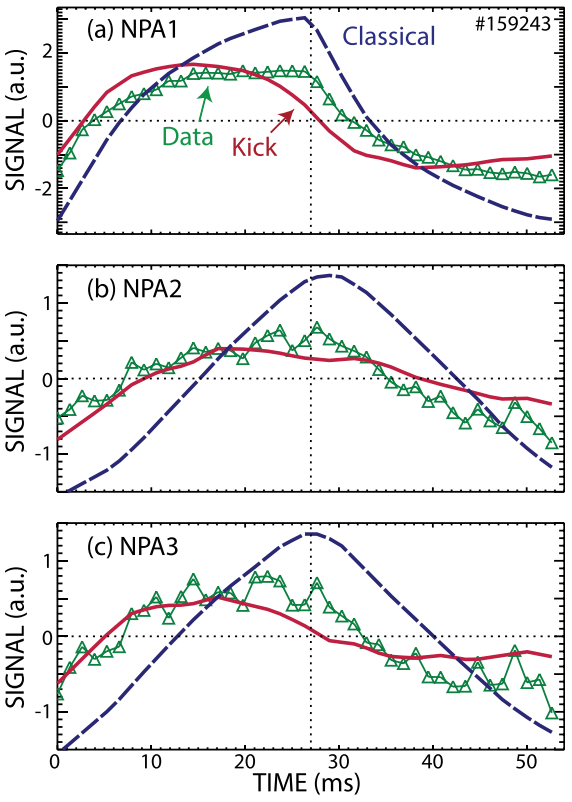

FIG. 14. Comparison of measured (△) signal with the classical prediction (dashed line) and with the prediction of the kick model (solid) for the NPA sightline that crosses the midplane at (a) 183 cm, (b) 165 cm, and (c) 150 cm after conditional averaging over 5 cycles between 624 and 897 ms. The vertical dashed lines show when the modulated beam turns off.

Calculated fast-ion losses are qualitatively consistent with the loss detector measurements. Figure 15(a) shows calculations of the power lost to the wall and to charge exchange outside the last-closed flux surface (LCFS). Figure 15(b) shows the signal measured by a fast-ion loss detector (FILD) photomultiplier. The temporal evolution of the calculation is very similar to the measurement, despite the fact that the TRANSP calculation effectively sums all of phase space outside the LCFS, while the FILD diagnostic measures a small portion of phase space. In the calculations, in order to isolate the effect of one type of mode, it is possible to "turn off" either the RSAEs or TAEs. The prediction with TAEs alone is close to the prediction that utilizes all of the measured modes (Fig. 15(a)), showing that TAEs play a more important role in losses than the RSAEs. Indeed, analysis of the FILD signal reveals coherent oscillations at TAE frequencies but not RSAE frequencies (Fig. 15(c)). Nevertheless, the calculated losses are largest when both RSAEs and TAEs are included (Fig. 15(a)). A likely explanation is that RSAEs redistribute confined ions into regions where TAEs cause actual losses.

Two distinct fast-ion populations determine the time evolution of the conditionally-averaged signals. In the absence of instabilities, the fast-ion signals depend only on the modulated source. Although steadily-injected ions contribute to the signal, since their signal is constant in time,

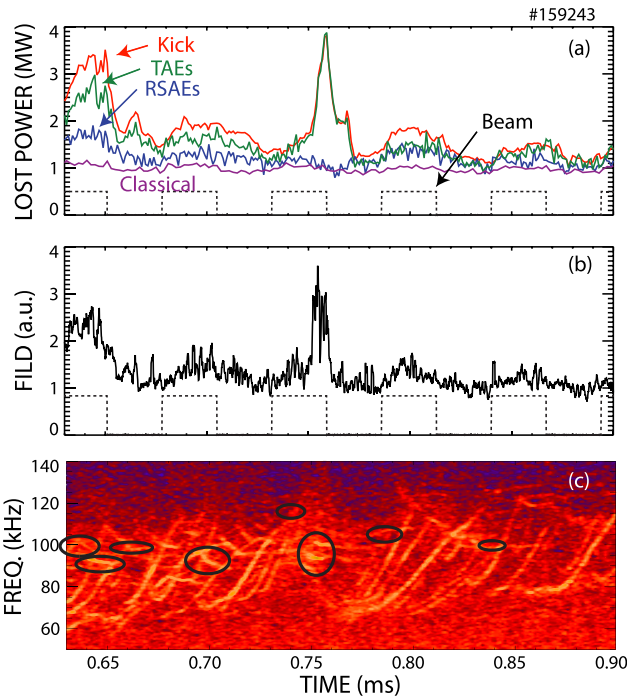

FIG. 15. (a) TRANSP calculations of beam power lost to the walls and to charge exchange. The four different calculations are the kick model with all of the modes, the kick model with only TAEs, the kick model with only RSAEs, and the classical calculation without any modes. The dotted line shows the timing of the modulated beam. (b) Fast-ion loss signal from a portion of the FILD scintillator. (c) Same spectrum as in Fig. 9(a). The circled TAEs produce coherent losses in the FILD signal.

they make no contribution to the conditionally-averaged difference signal. This is no longer the case in the presence of AEs. The AEs transport fast ions deposited by the modulated source, altering their contribution. But, because the modulated source alters the AE amplitude (Fig. 10), it also modulates the contribution of the steadily-injected fast ions to the conditionally-averaged signal. These effects are illustrated in Fig. 16. Classically, the neutron signal rises when the modulated beam is on and falls when the modulated beam is off. In a kick simulation with only the modulated source, the predicted signal resembles the classical signal but the amplitude is reduced due to AE-induced transport. In contrast, in a kick simulation with the same modes but only the steady sources, the predicted signal has the opposite slope. In the positive phase of the cycle, the modes become stronger, which causes additional transport of previously confined fast ions and reduction in the signal. In the negative phase of the cycle, the modes become weaker, the transport is reduced, and the signal rises. The full response is the sum of the opposing responses of these two populations. The competing tendencies produce a flat signal in the latter half of the positive and negative phases. This flat signal is observed experimentally but, as shown in Ref. 8, cannot be reproduced by diffusion of fast ions from the modulated source alone.

Figure 17 shows the calculated effect of the AEs on the distribution function. The reduction in fast-ion density is largest in the core (Figs. 17(a) and 17(b)). The appearance of a "peninsula" of ions with a pitch of ~ -0.3 shows that some ions change orbit types. At mid-radius (Figs. 17(c) and 17(d)), the changes in f are relatively modest, with the most important difference being an increase in co-passing ions with a pitch of ~ 0.6. In the outer part of the plasma (Figs. 17(e) and 17(f)), the transport causes an increase in fast-ion density, particularly of low-energy co-passing ions. The calculations indicate that, on average, fast ions lose energy to the waves, as expected for fast-ion driven instabilities.

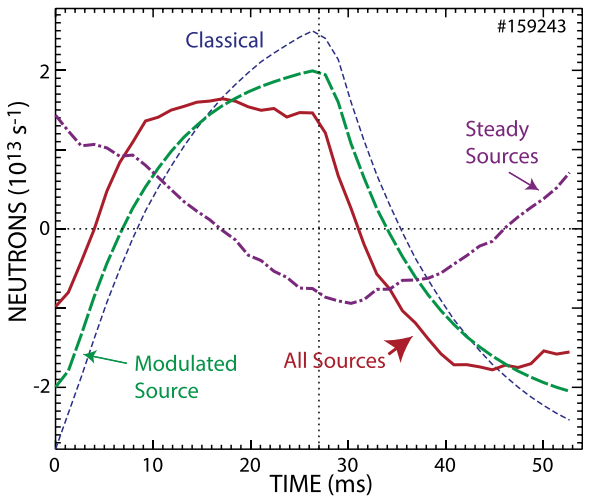

FIG. 16. Calculated conditionally-averaged response of the neutron signal for four different simulations: classical (short dash line), kick model with actual beam sources (solid line), kick model with only the modulated off-axis source (long dash line), and kick model with only the steadily injected sources (dashed-dotted line). All three kick simulations use the mode evolution shown in Fig. 11(c).

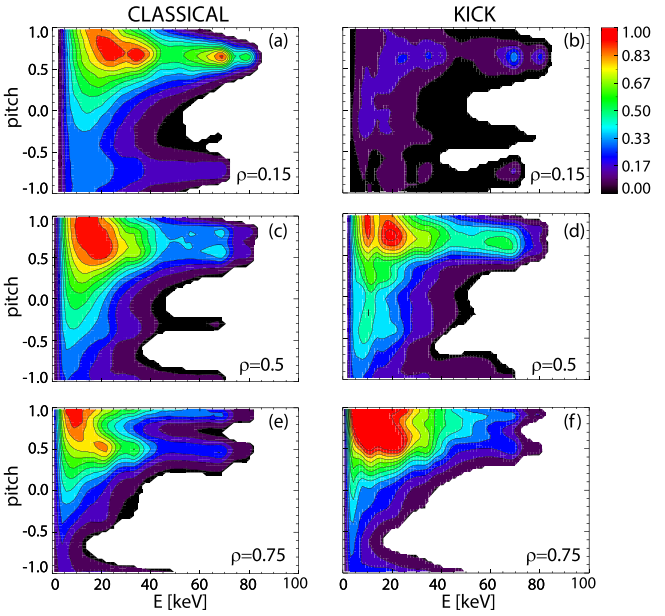

FIG. 17. Distribution functions calculated by TRANSP classically (left column) and by the kick model (right column) at three different normalized minor radii ρ. Each row uses the same normalization. The normalization factors are 2.22×10^8 for the top row, 7.98×10^7 for the middle row, and 2.79×10^7 for the bottom row. The pitch variable is $v_\parallel/v$, with a positive value indicating circulation in the direction of the plasma current.

V. PHASE-SPACE FLOWS

Section IV showed that calculations, which use empirically-fitted modes and orbital displacements, compare favorably with measured fast-ion signals. This section presents initial exploration of an alternative approach. In principle, stiff transport above threshold can be characterized by concepts developed in nonlinear dynamics. As orbits become stochastic, transport increases in some regions of phase space but not in others. A regime of island-induced diffusion with non-Gaussian statistics and Lévy flights occurs.[38] Because the orbits and wave-particle resonances are phase-space dependent, these flows differ in different parts of phase space.

As explained in Ref. 8, the conditionally averaged signals contain information about flows into or out of the portion of phase space interrogated by the fast-ion diagnostic. The experimentally-measured quantity is the divergence of the flux, $\nabla \cdot \tilde{\Gamma}$. At low beam power where the AE activity is weak, the classical and measured signals are virtually identical, so, within experimental uncertainty, the inferred phase-space flows are zero. At a higher power, AE-induced fast-ion transport causes differences between the classical predictions and the measurements. When the slope of the measured signal is greater than expected, ions are flowing into the volume interrogated by the fast-ion diagnostic and $\nabla \cdot \tilde{\Gamma} < 0$; when the slope of the measured signal is less than expected, ions are leaving the volume and $\nabla \cdot \tilde{\Gamma} > 0$. Note that, since the conditional averaging is about the *mean* value, the source during the positive and negative phases of the cycle are mirror images. Consequently, during the negative phase of the cycle, signs flip: $\nabla \cdot \tilde{\Gamma} > 0$ implies a flow into the volume and $\nabla \cdot \tilde{\Gamma} < 0$ implies a flow out of the volume.

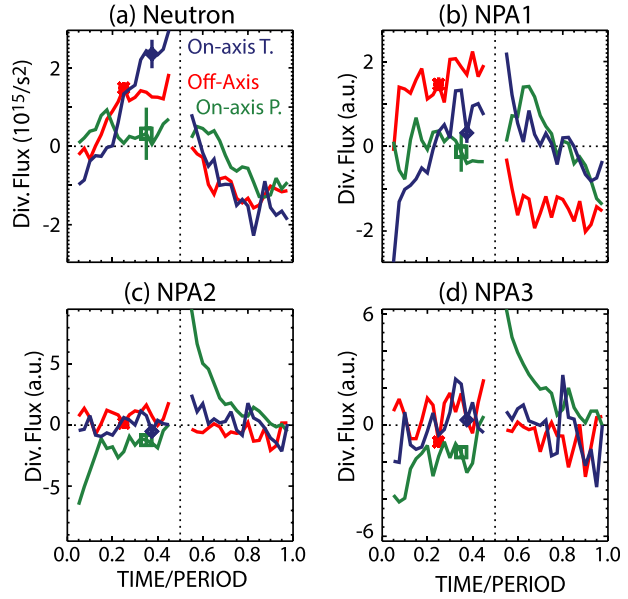

FIG. 18. Divergence of the flux $\nabla \cdot \tilde{\Gamma}$ inferred from the difference between the measured signal and the classical prediction for four different diagnostics in three different discharges. The four diagnostics are (a) neutrons, (b) the outermost NPA chord, (c) the middle NPA chord, and (d) the innermost NPA chord. The three discharges are the nearly identical discharges of Table I with different modulated beams: off-axis (*), on-axis tangential (diamond), and on-axis perpendicular (square). The error bar on each curve is an estimate of the systematic uncertainty in $\nabla \cdot \tilde{\Gamma}$ associated with fitting the classical prediction to the model equations;[8] random error can be assessed from the deviation of $\nabla \cdot \tilde{\Gamma}$ from smooth curves.

Figure 18 shows measurements of $\nabla \cdot \tilde{\Gamma}$ for the three discharges of Table I for four different diagnostics. The most striking aspect of this figure is the tremendous variety of responses. With few exceptions, every case is different. For a given diagnostic, the response is quite sensitive to the modulated source; for a given modulated source, the response is quite sensitive to the fast-ion diagnostic. The obvious qualitative conclusion is that the flows differ markedly in different parts of phase space.

In some cases, the inward flow is very large when the phase flips polarity (for example, the NPA2 and NPA3 signals with on-axis tangential injection). Inward flows occur when the measured signal changes more rapidly than classically expected. Since the AE-induced transport of newly deposited ions cannot instantaneously generate a significant flow, deposition by the modulated source cannot produce this effect. The effect is produced by modulation of the steady beam population through modification of the AE amplitude, as in Fig. 16.

To understand these differences between diagnostics and sources, consider the phase space. The measured signal s is the product of the diagnostic weight function W and the distribution function f, integrated over the phase space Ω, $s = \int W(\Omega) f(\Omega)\, d\Omega$. Modulated signals are produced by AE-induced changes in the distribution function, $\tilde{s} = \int W(\Omega)\, \tilde{f}(\Omega)\, d\Omega$. As mentioned above, AE-induced transport is expected to produce correlated changes in energy and toroidal canonical angular momentum without changing the magnetic moment. Accordingly, in Fig. 19, we plot the diagnostic weight functions for the four diagnostics of Fig. 18 in (P_ϕ, E) space for various values of μB_0. (B_0 is the magnetic field at the magnetic axis.) The neutron weight function is primarily sensitive to energy with weaker dependencies on P_ϕ and μ. In contrast, the NPA weight functions are quite sensitive to both μ and P_ϕ; like the neutrons, they are larger at a higher energy.

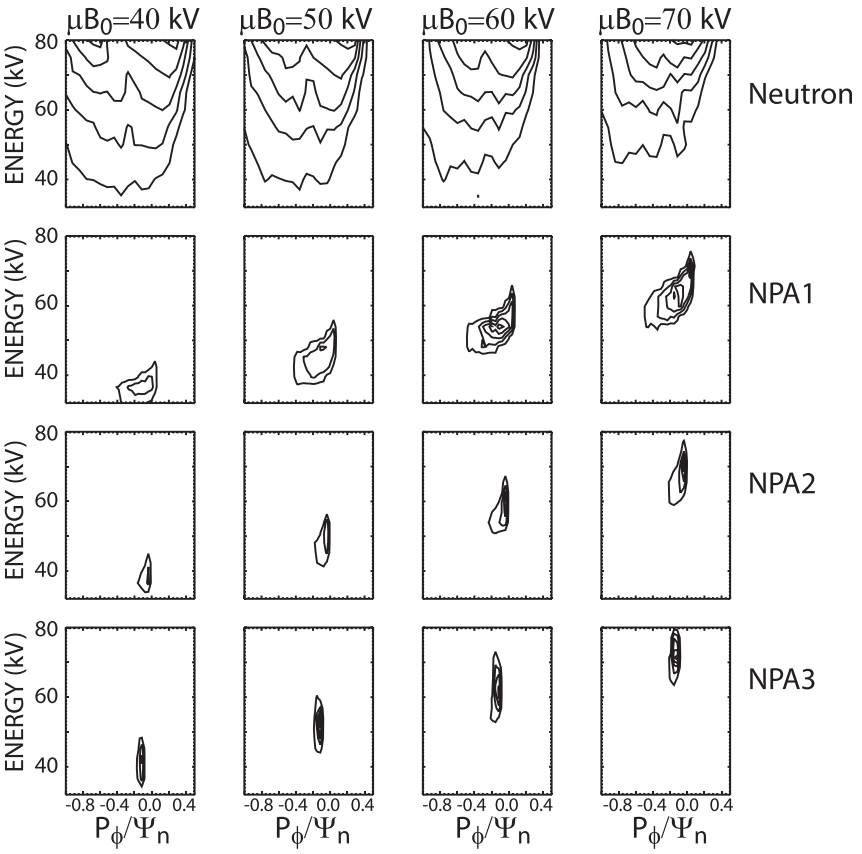

FIG. 19. Diagnostic weight functions for discharge #159 243 at 719 ms. The four columns are for $\mu B_0 = 40 \pm 5$, 50 ± 5, 60 ± 5, and 70 ± 5 keV; the four rows are for the neutron, outermost NPA chord, middle NPA chord, and innermost NPA chord. For each diagnostic, the 5 linear contours are plotted using the same scale for the entire row. The abscissa is the normalized toroidal canonical angular momentum P_ϕ/Ψ_n, and the ordinate is the energy.

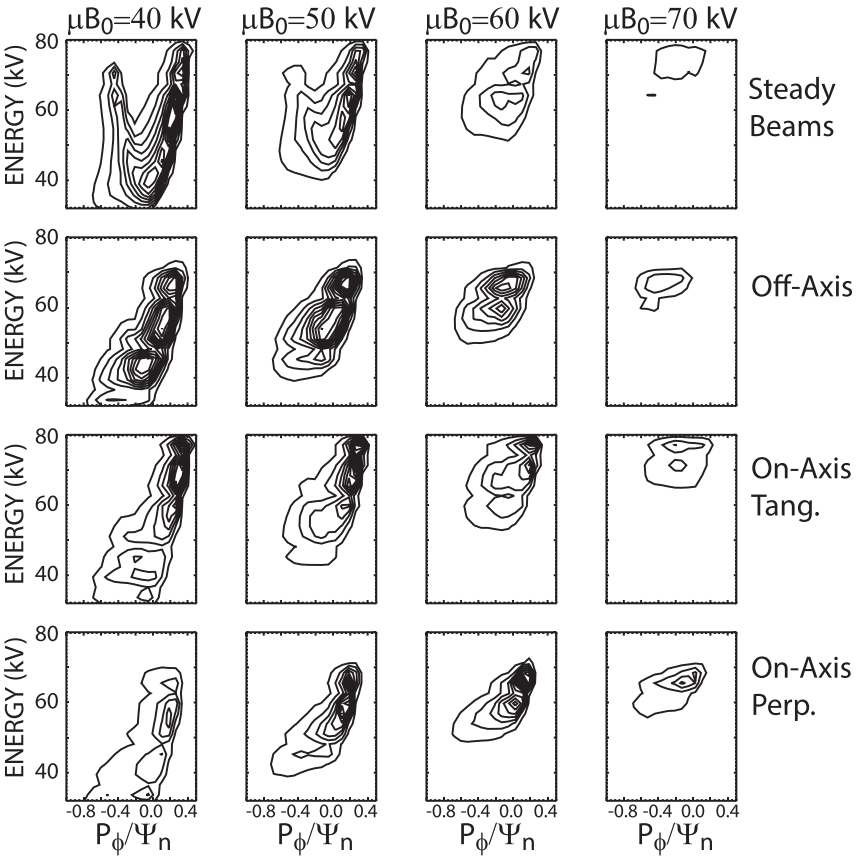

FIG. 20. Classical distribution functions. The four columns are for $\mu B_0 = 40 \pm 5$, 50 ± 5, 60 ± 5, and 70 ± 5 keV. The four rows are the average distribution function produced by the steadily injected beams, the modulated beam on discharge #159 243, the modulated beam on discharge #159 259, and the modulated beam on discharge #157 727. For each source, the 10 linear contours are plotted using the same scale for the entire row. The modulated-beam distribution functions are from the latter half of the positive cycle. The abscissa is the normalized toroidal canonical angular momentum P_ϕ/Ψ_n, and the ordinate is the energy.

Figure 20 shows the relevant unperturbed distribution functions plotted using the same phase space coordinates. The steady beams occupy larger regions in phase space than the modulated distribution functions; they also extend to lower energy. The modulated sources differ from one another but the differences appear modest in this representation.

To determine where the appreciable transport is expected, the fitted modes used in the kick analysis are used to compute broken Kolmogorov-Arnold-Moser (KAM) surfaces.[16] This analysis uses the experimental amplitudes at 780 ms. The result is shown in Fig. 21 for one value of μ; the results for other values of μ are similar. Although there are isolated regions with intact KAM surfaces, broken orbits fill most of the relevant phase space. Appreciable phase-space flows are expected.

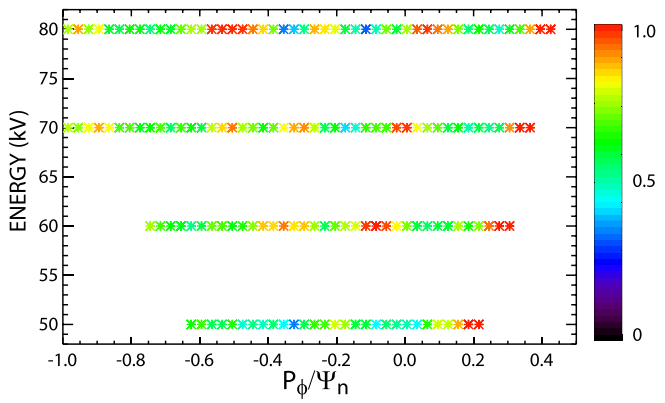

FIG. 21. Calculation of intact and broken KAM surfaces in various regions of phase space for $\mu B_0 = 40$ keV. Each point represents the fractional number of preserved surfaces in that region. The abscissa is the normalized toroidal canonical angular momentum P_ϕ/Ψ_n, and the ordinate is the energy.

Figure 22 presents an initial attempt to organize the data of Fig. 18. Each panel shows a phase-space plot of the diagnostic weight function, the modulated source, and the steady fast-ion population. Consistent with Eq. (1), the kick transport probabilities lie along lines with diagonal slope. Each panel contains an arrow showing a phase-space flow that is qualitatively consistent with the measured $\nabla \cdot \tilde{\Gamma}$. The overall pattern defies a simple explanation. In some cases, the flow is large when the source is outside the diagnostic weight function but, in others, the flow is largest when the source and weight function coincide. An additional complication is that Fig. 22 only shows a single value of μ but the observed signals integrate over μ. Understanding these complicated phase-space patterns is an outstanding challenge for future work.

VI. CONCLUSION

In the presence of multiple small-amplitude AEs, DIII-D plasmas manifest all of the features of critical gradient behavior. Above a threshold, orbits become stochastic. This qualitatively explains why the fast-ion profile is resilient, the equivalent spatial diffusion is large, and the modulated transport is stiff.

Using modes derived from experimental measurements and amplitudes that are adjusted to match the measured neutron rate, the TRANSP kick model satisfactorily reproduces many experimental features, including

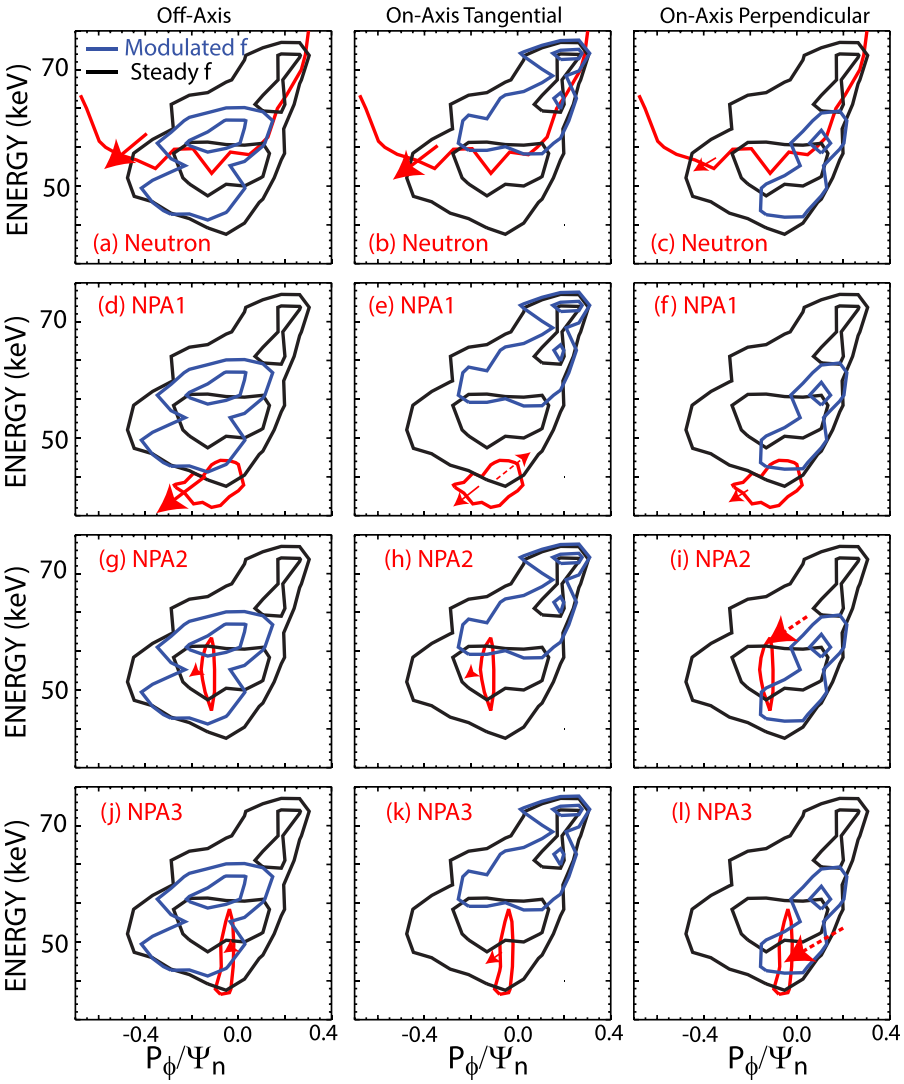

FIG. 22. Qualitative interpretation of the data shown in Fig. 18. Each figure shows a $(P_\phi/\Psi_n, E)$ phase-space plane with $\mu B_0 = 60 \pm 5$ kV. In each figure, one contour represents the weight function of a diagnostic, while two contours show the time-average distribution function and the modulated distribution function. The three cases are the three cases of Table I: (left column) modulated beam that is off-axis, (middle column) on-axis tangential, and (right column) on-axis perpendicular. The arrows placed beside each diagnostic weight function represent the observed values of $\nabla \cdot \tilde{\Gamma}$, with larger arrows representing stronger flows. Solid (dashed) arrows represent the direction of the flows near the end (beginning) of the positive beam pulse.

- the FIDA spatial profile (Fig. 12) and spectral shape (Fig. 13),
- the time evolution of the NPA signals (Fig. 14), and
- the correlation of the FILD signal with TAEs (Fig. 15).

The good agreement with the data suggests that, in this case, the fitted modes are reasonably well represented by ideal MHD eigenfunctions and that the correlated $(\delta E, \delta P_\phi)$ probabilities are an adequate treatment of the wave-particle interaction.

The kick analysis shows that, in addition to the fast ions from the modulated source, steadily injected fast ions respond to the beam modulation. The contribution of the steadily-injected ions to the signals is a nonlinear effect caused by correlated changes in mode amplitude. Both types of fast ions contribute to observed phase-space flows.

Analysis of phase-space flows is in its infancy. Conceivably, with an extensive diagnostic set and probe beams that populate different portions of phase space, reconstruction of phase-space flows in a manner similar to velocimetry is possible. Even the present measurements demonstrate that the transport varies markedly in different parts of phase space. These data provide stringent validation tests for predictive models of AEs. Indeed, based on these data, an ambitious verification and validation effort involving nine first-principles codes and three reduced models is already underway.

ACKNOWLEDGMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award No. DE-FC02-04ER54698. We thank the entire DIII-D team for their support. DIII-D data shown in this paper can be obtained in a digital format by following the links at https://fusion.gat.com/global/D3D_DMP.

[1]W. W. Heidbrink, Phys. Plasmas **15**, 055501 (2008).
[2]S. D. Pinches, I. T. Chapman, P. W. Lauber, H. J. C. Oliver, S. E. Sharapov, K. Shinohara, and K. Tani, Phys. Plasma **22**, 021807 (2015).
[3]B. Coppi, Comments Plasma Phys. Controlled Fusion **5**, 261 (1980).
[4]N. J. Lopes Cardozo, Plasma Phys. Controlled Fusion **37**, 799 (1995).
[5]J. C. DeBoo, C. C. Petty, A. E. White, K. H. Burrell, E. J. Doyle, J. C. Hillesheim, C. Holland, G. R. McKee, T. L. Rhodes, L. Schmitz, S. P. Smith, G. Wang, and L. Zeng, Phys. Plasma **19**, 082518 (2012).
[6]C. S. Collins, W. W. Heidbrink, M. E. Austin, G. J. Kramer, D. C. Pace, C. C. Petty, L. Stagner, M. A. Van Zeeland, R. B. White, Y. B. Zhu *et al.*, Phys. Rev. Lett. **116**, 095001 (2016).

[7]C. S. Collins, W. W. Heidbrink, M. E. Austin, G. J. Kramer, D. C. Pace, C. C. Petty, L. Stagner, M. A. Van Zeeland, R. B. White, Y. B. Zhu *et al*., "Critical gradient behavior of fast-ion transport by Alfven eigenmodes in phas space," Nucl. Fusion (submitted).
[8]W. W. Heidbrink, C. S. Collins, L. Stagner, Y. B. Zhu, C. C. Petty, and M. A. Van Zeeland, Nucl. Fusion **56**, 112011 (2016).
[9]W. W. Heidbrink, Y. Luo, K. H. Burrell, R. W. Harvey, R. I. Pinsker, and E. Ruskov, Plasma Phys. Controlled Fusion **49**, 1457 (2007).
[10]D. J. Sigmar, C. T. Hsu, R. White, and C. Z. Cheng, Phys. Fluids B **4**, 1506 (1992).
[11]R. B. White, N. Gorelenkov, W. W. Heidbrink, and M. A. VanZeeland, Plasma Phys. Controlled Fusion **52**, 045012 (2010).
[12]R. B. White, N. Gorelenkov, W. W. Heidbrink, and M. A. VanZeeland, Phys. Plasma **17**, 056107 (2010).
[13]W. W. Heidbrink, N. N. Gorelenkov, Y. Luo, M. A. Van Zeeland, R. B. White, M. E. Austin, K. H. Burrell, G. J. Kramer, M. A. Makowski, G. R. McKee, R. Nazikian, and DIII-D Team, Phys. Rev. Lett. **99**, 245002 (2007).
[14]M. A. Van Zeeland, G. J. Kramer, M. E. Austin, R. L. Boivin, W. W. Heidbrink, M. A. Makowski, G. R. McKee, R. Nazikian, W. M. Solomon, and G. Wang, Phys. Rev. Lett. **97**, 135001 (2006).
[15]R. B. White and M. S. Chance, Phys. Fluids **27**, 2455 (1984).
[16]R. B. White, Commun. Nonlinear Sci. Numer. Simul. **17**, 2200 (2012).
[17]R. B. White, Plasma Phys. Controlled Fusion **53**, 085018 (2011).
[18]Y. Todo, M. A. Van Zeeland, A. Bierwage, and W. W. Heidbrink, Nucl. Fusion **54**, 104012 (2014).
[19]Y. Todo, M. A. Van Zeeland, A. Bierwage, W. W. Heidbrink, and M. E. Austin, Nucl. Fusion **55**, 073020 (2015).
[20]A. Bierwage, Y. Todo, N. Aiba, and K. Shinohara, Nucl. Fusion **56**, 106009 (2016).
[21]R. E. Waltz, E. M. Bass, W. W. Heidbrink, and M. A. Van Zeeland, Nucl. Fusion **55**, 123012 (2015).
[22]M. Podesta, M. Gorelenkova, and R. B. White, Plasma Phys. Controlled Fusion **56**, 055003 (2014).
[23]R. Sanchez and D. E. Newman, Plasma Phys. Controlled Fusion **57**, 123002 (2015).
[24]H. L. Berk, B. N. Breizman, J. Fitzpatrick, and H. V. Wong, Nucl. Fusion **35**, 1661 (1995).
[25]W. W. Heidbrink, M. A. Van Zeeland, M. E. Austin, E. M. Bass, K. Ghantous, N. N. Gorelenkov, B. A. Grierson, D. A. Spong, and B. J. Tobias, Nucl. Fusion **53**, 093006 (2013).
[26]W. W. Heidbrink, J. R. Ferron, C. T. Holcomb, M. A. Van Zeeland, X. Chen, C. M. Collins, A. Garofalo, X. Gong, B. A. Grierson, M. Podestà, L. Stagner, and Y. Zhu, Plasma Phys. Controlled Fusion **56**, 095030 (2014).
[27]C. T. Holcomb, W. W. Heidbrink, J. R. Ferron, M. A. Van Zeeland, A. M. Garofalo, W. M. Solomon, X. Gong, D. Mueller, B. Grierson, E. M. Bass, C. Collins, J. M. Park, K. Kim, T. C. Luce, F. Turco, D. C. Pace, Q. Ren, and M. Podesta, Phys. Plasma **22**, 055904 (2015).
[28]W. W. Heidbrink, P. L. Taylor, and J. A. Phillips, Rev. Sci. Instrum. **68**, 536 (1997).
[29]C. M. Muscatello, W. W. Heidbrink, D. Taussig, and K. H. Burrell, Rev. Sci. Instrum. **81**, 10D316 (2010).
[30]Y. B. Zhu, A. Bortolon, W. W. Heidbrink, S. L. Celle, and A. L. Roquemore, Rev. Sci. Instrum. **83**, 10D304 (2012).
[31]M. E. Austin and J. Lohr, Rev. Sci. Instrum. **74**, 1457 (2003).
[32]A. Pankin, D. McCune, R. Andre, G. Bateman, and A. Kritz, Comput. Phys. Commun. **159**, 157 (2004).
[33]See http://w3.pppl.gov/transp/ for information about TRANSP.
[34]W. W. Heidbrink, D. Liu, Y. Luo, E. Ruskov, and B. Geiger, Commun. Comput. Phys. **10**, 716 (2011).
[35]L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer, Nucl. Fusion **25**, 1611 (1985).
[36]A. Fasoli, D. Testa, S. Sharapov, H. L. Berk, B. Breizman, A. Gondhalekar, R. F. Heeter, M. Mantsinen, and Contributors to the EFDA-JET Workprogramme, Plasma Phys. Controlled Fusion **44**, B159 (2002).
[37]C. Z. Cheng, Phys. Rep. **211**, 1 (1992).
[38]R. B. White, *The Theory of Toroidally Confined Plasmas*, 3rd ed. (Imperial College Press, London, 2014), Chap. 9.