

JRT 2016 4th Quarter Report and Final Summary

Steve Sabbagh, Roger Raman, Nick Eidietis, Ted Strait, Robert Granetz

This report summarizes the work completed during the final quarter of FY2016 on the three major US tokamak facilities as part of this past year's joint research target. Following this summary section are detailed reports from the three facilities. We begin by re quoting the overarching JRT goals, including the particulars for the 4th quarter milestones:

FY16 DOE Research Target (JRT)

Summary: Conduct research to detect and minimize the consequences of disruptions in present and future tokamaks, including ITER. Coordinated research will deploy a disruption prediction/warning algorithm on existing tokamaks, assess approaches to avoid disruptions, and quantify plasma and radiation asymmetries resulting from disruption mitigation measures, including both pre-existing and resulting MHD activity, as well as the localized nature of the disruption mitigation system. The research will employ new disruption mitigation systems, control algorithms, and hardware to help avoid disruptions, along with measurements to detect disruption precursors and quantify the effects of disruptions.

4th Quarter Milestone: Use the disruption prediction algorithm to characterize the reliability of predicting a few types of common disruptions from at least two devices. Test on at least one facility to detect in real time an impending disruption and take corrective measures to safely terminate the plasma discharge. Report on capability to reduce disruption rate through active improvement of plasma stability. Test newly designed ITER-type massive gas injection valve to study the benefits of private flux region massive gas injection vs. mid-plane injection.

Complete the required experiments and analysis. Prepare a joint report summarizing the contributions toward the understanding of the plasma response to massive gas injection in discharges with pre-existing magnetohydrodynamic modes, runaway electron physics, private flux region injection, the reliability of the disruption prediction algorithm for detecting a few types of common disruptions, and avoidance methods for reducing the disruption rate.

With regard to the 4th quarter milestones and the JRT overall goals, we note that the MGI comparison between plasmas with and without pre-existing MHD modes (specifically n=1 locked modes) was successfully carried out on C-Mod and DIII-D, with conclusive results that will be presented in a combined paper at the upcoming IAEA conference in Kyoto.

Extensive work on reducing disruption likelihood has been carried out on DIII-D. This includes the identification of features that seem to enhance the likelihood of disruptions (q-profile near the q=2 surface; rotation locking of different helicity tearing modes), as well as development of

several techniques to determine, in real time, the proximity to stability limits (active MHD spectroscopy; driven excitation). DIII-D has also carried out experiments to actively expand the range of stable operation by several different methods, including control of error fields, ECCD stabilization of NTMs, and advanced controllers to feedback stabilize RWMs. A sophisticated, state-based event handling system has also been developed to respond to off-normal and fault events.

Newly designed ITER-type MGI valves were installed in three locations on NSTX-U, including the private flux region and the mid-plane. These valves were successfully tested in a 1-tesla ambient field, and also into the empty NSTX-U vacuum chamber.

Large databases of plasma parameters that may be useful for real-time prediction of upcoming disruptions have been established on C-Mod and DIII-D. Initial studies of a few of the parameters that show some predictive capability on DIII-D, such as ℓ/q_{95} and I_p error, do *not* show similar promise on C-Mod. This could mean that it may not be feasible to develop a robust real-time disruption warning algorithm that will work on multiple machines. On NSTX-U, there has been quite a bit more development of the DECAF code for disruption forecasting.

Finally, both C-Mod and DIII-D have interesting new data shedding light on the physics of relativistic runaway electrons. Experiments on both machines were carried during the flattop of MHD-quiescent, low density discharges, when runaways can be established in a reproducible, well-diagnosed manner. A new 2-D array of gamma-ray detectors on DIII-D measures the energy spectra of bremsstrahlung arising from interaction of the RE population with the background plasma, helping to resolve RE behavior as a function of energy. On C-Mod, synchrotron spectra have been obtained at 3 different magnetic fields (2.7, 5.4, and 7.8 tesla). For fixed RE energy and pitch angle, the synchrotron spectra should have very different shapes, but in fact, they are essentially identical.

Looking back at the entire year's worth of work, an overwhelming majority of the planned tasks have been successfully carried out and completed, resulting in a number of conference presentations and publications. Due to the unplanned interruption of NSTX-U plasma operation, the comparison of MGI performance between the private flux and mid-plane injection locations was not able to be carried out. Much effort has been focused on disruption prediction, with some inter-machine comparisons, but a real-time disruption predictor has not been tested in an actual plasma control system as part of this JRT. However, intense work on this is continuing, even though the JRT has officially concluded.

Detailed reports on the JRT work in the 4th quarter at all three facilities follow. (Note that the DIII-D report summarizes work from the first 3 quarters as well.)

DIII-D Contribution to FY 2016 Joint Research Target – Final Report
Disruption Avoidance and Mitigation

1.0 Experimental Results and Analysis – Disruption Avoidance

1.1 Understanding stability limits

Analysis of DIII-D data in 2016 revealed key features of internal plasma profiles associated with increased stability. These include specific features of the current density profile near the $q=2$ surface, and differential rotation between rational surfaces. The conditions under which a locked tearing mode is likely to evolve to a disruption were also identified.

Database analysis has shown that in DIII-D discharges that simulate the ITER baseline scenario, a minimum of the current density profile often develops in the outer part of the plasma near the $q=2$ surface. As shown in Fig. 1, a deeper minimum is associated with greater instability to tearing modes at the $q=2$ surface [2], suggesting that stability could be improved through modification of the current density profile. Work is in progress to understand this result through MHD stability modeling.

It has recently been observed that in discharges with low input torque, MHD modes at different rational surfaces tend to lock together (see Fig. 2). The locking reduces the rotational shear in the plasma, though not necessarily the overall rotation, and this leads to wall locking, growth of instabilities, and possible disruptions [3]. This is confirmed by additional database analysis showing that differential rotation between the $q=3/2$ and $q=2$ surfaces improves stability to tearing modes at the $q=2$ surface in “ITER baseline” plasmas. These observations suggest a possible warning signal for disruptions based on internal mode coupling, as well as a possible approach to disruption prevention by maintaining sufficient rotation shear.

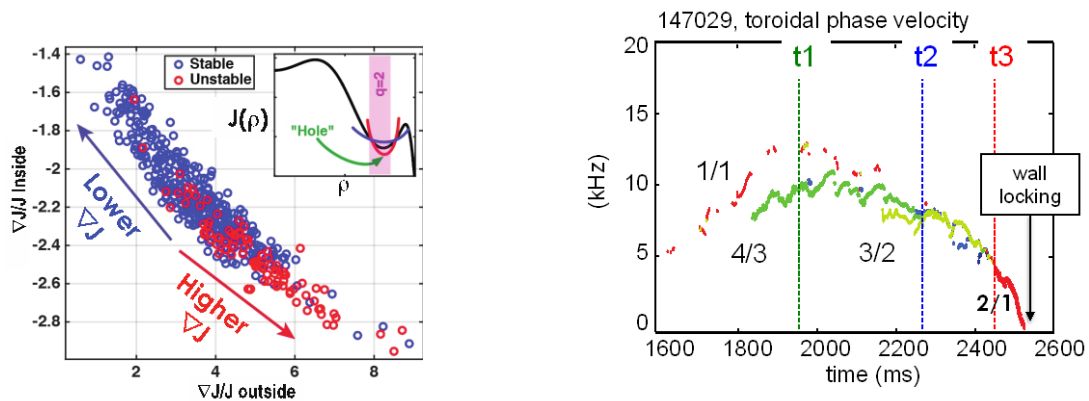


Fig. 1. The current density gradient to the inside (smaller radius) of the minimum is plotted vs. the gradient to the outside of the minimum, showing stable (blue) and unstable (red) cases. Current density profiles with a deeper minimum lie toward the lower right-hand corner of the plot.

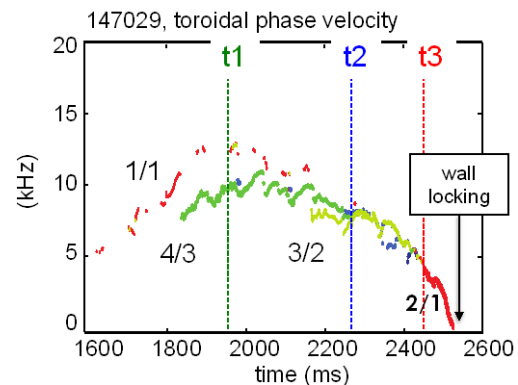


Fig. 2. Time evolution of the toroidal rotation rate (f/n) of several MHD modes. Phase locking is followed by growth of an $m/n=2/1$ tearing mode and wall locking.

Analysis of a large database is in progress to ascertain the conditions under which a locked tearing mode is likely to result in a disruption [4]. As seen in Fig. 3, the parameter l_i/q_{95} (ratio of internal inductance to safety factor at the 95% flux surface) provides a reasonably good separation of the dataset between initially-rotating locked modes that end in disruption and those that do not. This parameter is related to the form of the current density profile. Another parameter taking into account the island's size and radial location is found to provide comparable separation of the data when evaluated 20 ms before the disruption. Such information could be used in real time to decide on the appropriate remedial action, including recovery, “soft landing”, or rapid shutdown.

1.2 Measuring proximity to stability limits

DIII-D experiments and analysis in 2016 investigated several approaches to direct sensing of plasma stability limits. These include the familiar technique of low-frequency active MHD spectroscopy for measurement of kink mode stability, and MHD spectroscopy at higher frequencies characteristic of tearing modes, as well as more novel techniques for detection of tearing mode stability.

Active MHD spectroscopy, utilizing the plasma response to magnetic perturbations at very low frequencies, may be a useful tool in probing the stability limits of ITER-relevant discharges. This technique has proven successful in measuring the properties and stability limits of ideal-MHD resistive

wall modes at high beta. A similar low-frequency plasma response has been found experimentally in moderate-beta plasmas, increasing in amplitude as the torque and rotation are reduced (Fig. 4). Modeling is in progress to understand the relationship between this ideal MHD-like kink response at 10's of Hz, and the resistive tearing instability at a few kHz that typically occurs in these plasmas. [2]




Fig. 3. Locked-mode database plot, with disrupting and non-disrupting cases plotted as the current profile parameter l_i/q_{95} vs. normalized radius of the $q=2$ surface ρ_{q2} . Data are evaluated 100 ms before the termination of the locked mode (either disappearance of the locked mode or a disruption).

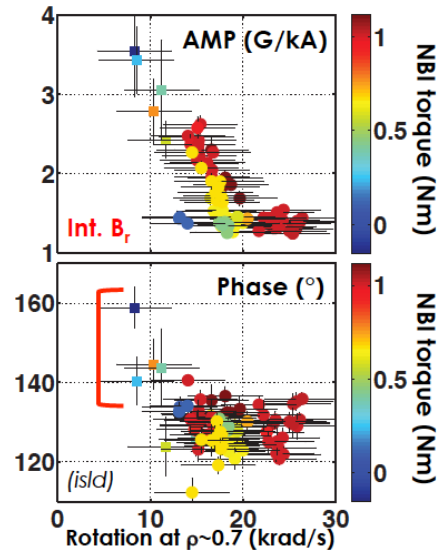


Fig. 4. Low-frequency (20 Hz) $n=1$ plasma response amplitude (top) and phase (bottom) vs. plasma rotation in ITER-like discharges. At low NBI torque and low plasma rotation, a rising amplitude and sudden phase jump at rotation frequencies below 10 krad/s (1600 Hz $\sim n=1$ tearing frequency) indicates a possible damped resonance.

A preliminary experiment [5] tested four methods of directly probing the plasma's stability to tearing modes. If the observed damping rate of a stable mode decreases with time, approaching zero, it would provide an early warning of an approaching instability. As summarized below, two methods aimed at resonant excitation at the frequency of the stable mode may be worth pursuing, while two methods aimed at momentarily destabilizing the tearing mode seem less promising. The $m=2/n=1$ tearing mode was selected for this test, and the experiment was carried out in a plasma with ITER-like shape (lower single null), safety factor ($q_{95} \sim 3.1$), beta ($\beta_N \sim 1.8$), and plasma rotation.

In the example shown in Fig. 5, resonant excitation of the 2/1 mode was provided by modulated electron cyclotron heating (ECH) at the $q=2$ surface. The gyrotrons were modulated at a fixed frequency while the neutral beam torque was varied in order to sweep the tearing mode frequency across the gyrotron frequency (see Fig. 5(a)). When the frequencies match, the stable mode should be driven resonantly to a finite amplitude. As seen in Fig. 5(b), a resonant peak in the plasma's magnetic response does seem to occur. The relative phase of the response changes by 180° across the peak, as expected in a resonant response. The time resolution of the plasma rotation measurement is not sufficient to translate the 25 ms peak width into a frequency width and damping rate, but the result provides motivation to repeat the experiment while sweeping the gyrotron modulation frequency instead.

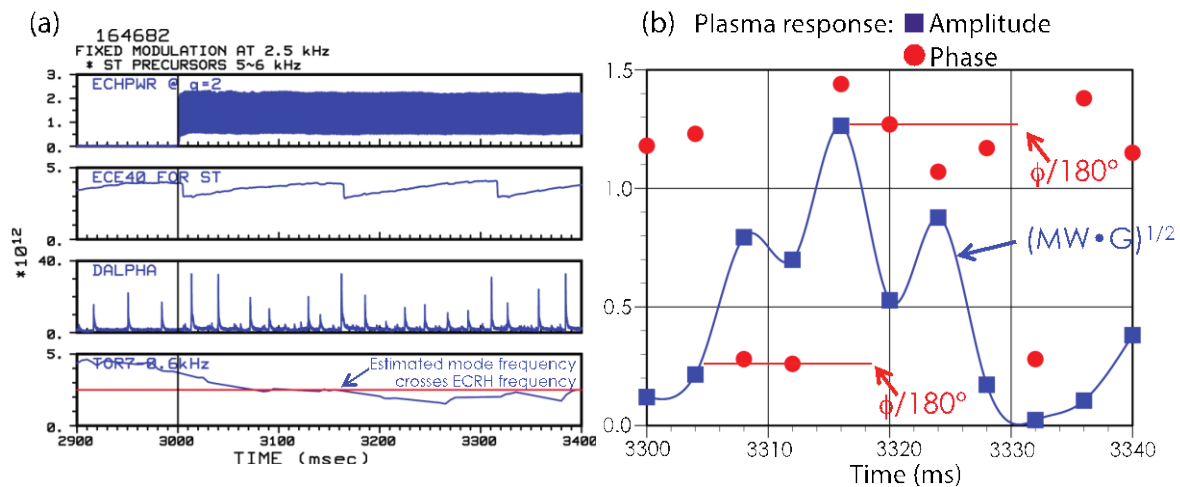


Fig. 5. Time evolution of discharge #164682. (a) 500 ms interval showing (top to bottom) ECH power at $q=2$ modulated between 0.5 and 2.2 MW, central electron temperature, divertor D-alpha emission, plasma rotation at the $q=2$ surface (blue curve) with 2.5 kHz ECH modulation frequency (horizontal red line). (b) Time-expanded plot of the $n=1$ magnetic response of the plasma cross-correlated with ECH power, showing amplitude in $MW^{1/2} \cdot G^{1/2}$ (blue squares) and relative phase in degrees/180 (red circles).

A similar experiment used the I-coils to apply a rotating $n=1$ magnetic perturbation, with frequency swept through the range 1-3 kHz. This technique has been well established at

lower frequencies (5-50 Hz) to measure the damping rate of the $n=1$ resistive wall mode ideal kink. In the present experiment, the coil current was limited by the inductive reactance of the coils and leads at kHz-range frequencies. In addition, at these frequencies which well exceed the inverse wall time, eddy currents induced in the wall greatly reduce the $n=1$ applied field at the $q=2$ surface (factor of 6 at 2 kHz). However, $n=1$ magnetic resonances can be measured (Figure 6) which increase in amplitude going towards eventual instability. The frequency steps used (0.2 kHz) were too coarse for good resolution, so this experiment will be repeated with continuous frequency sweeps.

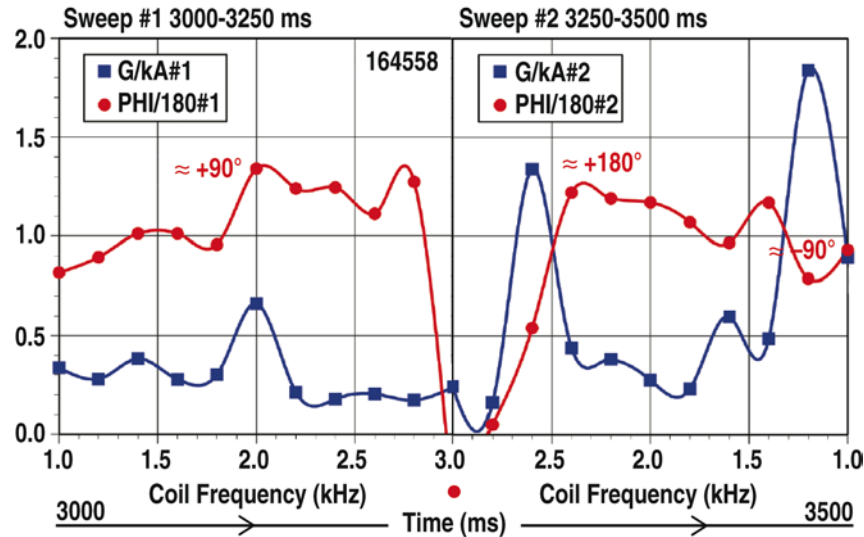


Fig. 6. Example of multiple $n=1$ I-coil current frequency sweeps across presumed ~ 2 kHz natural stable $m/n=2/1$ tearing frequency in a low-torque DIII-D ITER baseline scenario discharge. Sweep #1 is 1 to 3 kHz and has one resonance at ~ 2.0 kHz. Sweep #2 is from 3 to 1 kHz and has two (or three?) resonances at ~ 2.6 , $1.6(?)$ and 1.2 kHz. The $m/n=2/1$ mode goes unstable (not shown) at $t=3485$ msec with initial growing mode frequency of 1.8 kHz, just before the end of sweep #2 and the start of sweep #3.

Another approach tested was to temporarily destabilize the tearing mode at low amplitude, and then remove the destabilizing influence and directly measure the decay rate of the tearing mode. The destabilization was provided by pulses of counter-current electron cyclotron current drive (ECCD) at the $q=2$ surface, or a pulse of increased neutral beam heating to raise beta. In the first case no destabilization was observed, while in the second case the tearing mode was readily destabilized, but then grew to large amplitude and locked to the wall before beta could be reduced. This prevented a measurement of the damping rate.

1.3 Expanding the range of stable operation

Several experiments this year have advanced active means of improving plasma stability, including control of error fields with toroidal mode numbers greater than unity, stabilization of neoclassical tearing modes (NTMs) with electron cyclotron current drive (ECCD), and feedback stabilization of resistive wall modes (RWMs) with a new model-based controller.

Recent DIII-D experiments have explored the role of error fields with toroidal mode number $n=2$, and the importance of correcting such error fields [6, 7]. It has long been known that $n=1$ error fields can cause braking of plasma rotation and drive reconnection (i.e. a locked $n=1$ tearing mode or magnetic island), in some cases leading to a disruption. Recent experiments in low-density ohmic plasmas have shown that an externally applied $n=2$ field can drive an $n=2$ tearing mode in the same way. Growth of an $n=1$ mode often follows in DIII-D (see Fig. 7). The $n=2$ amplitude threshold for penetration and island formation scales almost linearly with density, similar to the $n=1$ threshold and with a comparable magnitude (see Fig. 8). Work is in progress to further quantify the threshold and compare it to measurements in other devices.

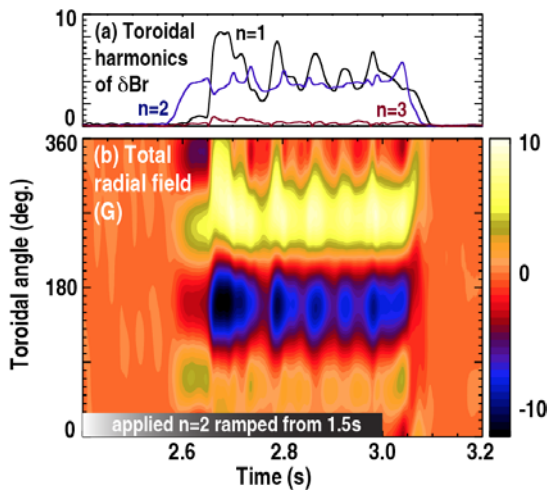


Fig. 7. (a) Time evolution of $n=1$, $n=2$ and $n=3$ non-rotating mode amplitudes; (b) contour plot of δB vs. time and toroidal angle. Shot 164950, q95-3, $T < 1$ Nm, with increasing $n=2$ applied field.

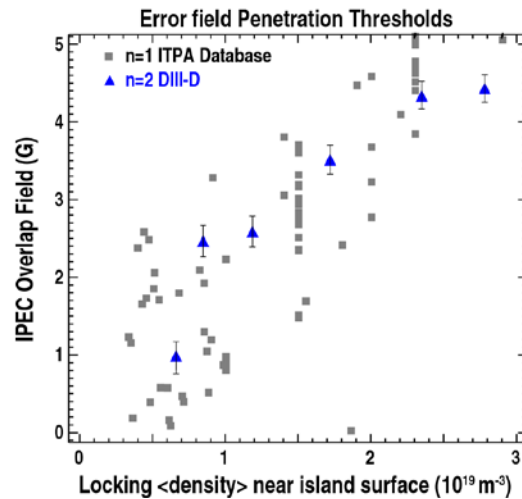


Fig. 8. The density dependence of the $n=2$ error field threshold (blue, triangles) in DIII-D along with the $n=1$ threshold (grey squares) from a multi-machine ITPA database. The thresholds are parameterized using the IPEC “overlap” field, which measures the plasma-mediated coupling of the external field to the resonant field that drives magnetic reconnection.

The growth of the $n=2$ tearing mode is almost always followed by growth of a larger $n=1$ tearing mode, as in the example of Fig. 7. The combination of the two island chains can cause strong reduction of confinement, and in some cases leads to a disruption. Thus, although the $n=2$ mode in isolation may not present a large risk of disruption, its synergy with the $n=1$ mode suggests that $n=2$ error field control may be desirable, particularly in low torque discharges. Work is in progress to identify the source of the $n=2$ error field in DIII-D and to develop algorithms for control.

Control of $n=2$ error fields may be more challenging than $n=1$ error fields. DIII-D experiments have shown that although the ITER Test Blanket Mockup Module generates a broad spectrum in toroidal mode number n , compensation of the $n=1$ component alone can recover most of the loss in plasma performance [8]. However, other recent experiments have shown that error fields with toroidal mode number $n=2$ can generate a multi-mode plasma response, making it more difficult to ameliorate all effects of the error

field simultaneously [7]. Future work will continue to investigate issues of multi-mode error field compensation with $n > 1$.

DIII-D research has also continued to develop methods for $n=1$ error field control. A new real-time algorithm for optimizing error field control based on maximizing the measured plasma rotation has been successfully tested [9]. This method is well suited to an integrated control approach in disruption sensitive devices, since it does not require triggering locked tearing modes or plasma current disruptions.

The use of electron cyclotron current drive (ECCD) for quickly suppressing $m/n=2/1$ neoclassical tearing modes (“catch & subdue”) before they grow to large enough amplitude to lock to the resistive wall was studied in the low-torque (1.3 Nm) low $q_{95}=3.1$ $\beta_N=1.8$ ITER baseline scenario in DIII-D. With gyrotrons in standby, the growing mode is detected by real-time magnetic fluctuation analysis in the plasma control system (PCS). The gyrotrons are turned on with mirrors aiming at the $q=2$ surface, using real-time equilibrium reconstruction to obtain the $q=2$ location and real-time ray-tracing calculations of the ECCD location for mirror alignment. However, suppression before locking was not achieved. Mode detection at the low torque did not allow enough time to do an alignment optimization (“target lock”). A subsequent experiment developed an improved real-time magnetic fluctuation analysis to better distinguish the $2/1$ mode from $1/1$ sawteeth precursors for example (Fig. 9). The new algorithm will be used for control in future experiments.




Fig. 9. Comparison of rt-NEWSPEC (a) and ODDEVENSPEC (b) which additionally uses 180 degree poloidally separated probes to distinguish odd m from even n . Without ECCD the $2/1$ mode (in blue in b.) grows to large amplitude and locks (at 1.3 Nm in this IBS DIII-D discharge) in less than an energy confinement time.

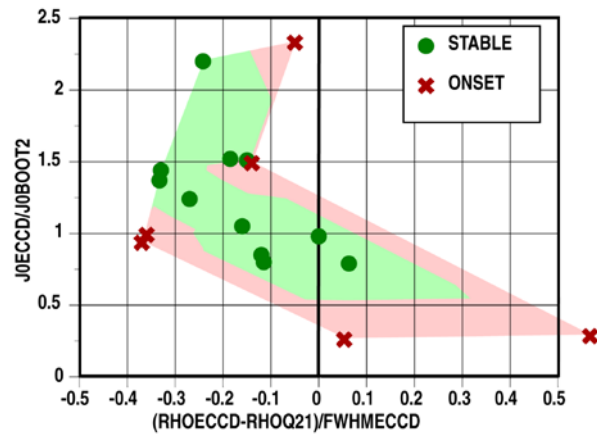


Fig. 10. At mode onset (red x) or if no mode at end of flattops (green circle) in a subset of low-torque IBS discharges in DIII-D during the “catch & subdue” run day. The vertical axis is the relative current drive strength at $q=2$ (ECCD from TORAY-GA and bootstrap from ONETWO). The horizontal axis is the relative misalignment with ECCD position and full width half maximum from TORAY-GA and $q=2$ position from MSEEFIT. Green contour is stable and red contour is unstable.

Data from the “catch and subdue” low-torque experiment was further analyzed for the possibility of preemptive stabilization. Control discharges with gyrotrons turned on and mirrors actively tracking *before* a mode onset as well as gyrotrons in low power standby *before or without* a mode were identified. Data points (in general one per discharge) are taken at the mode onset or at the end of constant discharge conditions if no mode occurs. The results are shown in Fig. 10. While not a systematic scan, this does indicate that a relatively low level of EC current drive, i.e. greater than about 0.5 of the local bootstrap current density, can preempt the 2/1 mode destabilization. The skew in the contours (one expects left-right symmetry about a misalignment of zero) suggests a systematic error in mirror alignment that is not taken into account in the real-time aiming algorithm. The contours are symmetrized (not shown) by allowing a linear offset of only +0.03 in the normalized radius of the $q=2$ surface with no ECCD, and of 0 for an ECCD level equal to 0.8 of the local bootstrap current density. As the equilibrium reconstructions do not resolve narrow radial features of localized ECCD, this is plausible although speculative. Further experiments will investigate the hypothesized alignment offset and associated requirements for 2/1 mode stabilization.

Accurate control of the toroidal position of an island that has magnetically locked to the wall allows stabilization by synchronized pulses of ECCD at the island’s O-point [10]. A new algorithm has been developed to control the position and slow rotation of a locked magnetic island based on its toroidal phase, independent of the island size, using an applied resonant magnetic perturbation [11]. As shown in Fig. 11, the new controller allowed the island position to track the reference waveform with dc steps and also with a programmed 20 Hz rotation rate. In a further test using the new phase controller, 20 Hz pulses of modulated ECCD were applied in phase with the island rotation. As shown in Fig. 12, the island amplitude decreased during each pulse, but increased again when the ECCD was reduced to very low power. Further tests will be needed to determine the requirements of power and pulse rate for full stabilization.




Fig. 11. Time evolution of the measured toroidal phase of an $n=1$ locked mode (black) and the reference phase (red) programmed in the phase controller. Shot 166560, $q_{95} \sim 4.5$, $\beta_N \sim 3.2$ before mode onset.

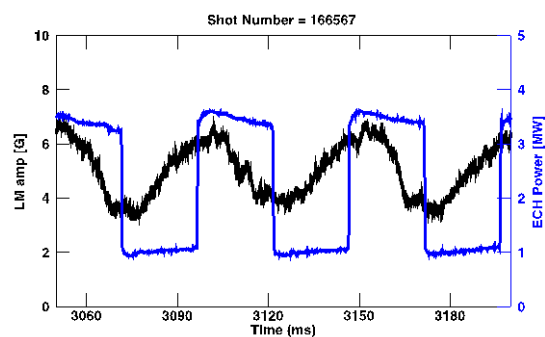


Fig. 12. Time evolution of the measured amplitude of an $n=1$ locked mode (black) and ECCD power (blue). Shot 166567, similar to 166560 in Fig. 11.

The possibility of stabilizing a tearing mode island by electromagnetic torque and forced rotation, without ECCD or other sources of heating and current drive, remains a topic of ongoing research [12]. The applied magnetic perturbation and the tearing mode locked to it can couple to other m/n rational surfaces, with the result that several rational surfaces may be locked together and rotating with the very low frequency of the applied perturbation. However, other factors (e.g. the redistribution of momentum injected by neutral beams, “intrinsic rotation” and radial distribution of momentum loss such as due to NTV) can lead to toroidal flow shear at the rational surfaces (see Fig. 13). According to flow shear stabilization models summarized by La Haye et al. [13], the observed flow shear is one order larger than the critical value for stabilization, thus, it is possible for this situation to lead to avoiding a major disruption. However, further research is needed in order to understand, predict, and control these highly nonlinear processes.

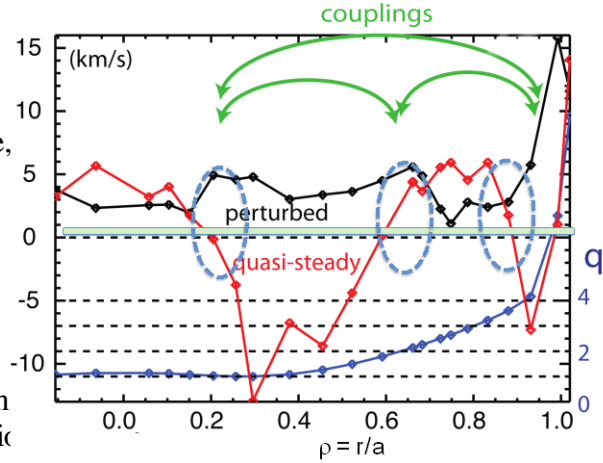


Fig. 13. Profiles of the time-averaged toroidal rotation (red) and modulation of the rotation by a slowly rotating locked mode and applied magnetic perturbation (black). The q profile is also shown (blue). Green arrows indicate coupling between $q=1, 2,$ and 4 surfaces.

A new high-speed controller for resistive wall mode stabilization [14] has been successfully tested on DIII-D, using the external control coils (“C-coils”). Magnetic feedback stabilization is likely to be needed in future devices that operate above the no-wall stability limit for ideal MHD kink modes. However, future devices are likely to be restricted to external control coils, where the speed of control is limited by the electromagnetic time constant of the vacuum vessel.

The control algorithm is executed on a graphics processing unit (GPU) for rapid computation by leveraging its massively parallel structure. In the recent test, a state-space control algorithm incorporating a model of the plasma instability, coils, power supplies, and vacuum vessel wall effectively compensated for the control delay of the vessel wall. In the experiment, the plasma beta was ramped up until an instability occurred. As shown in Fig. 14, the combination of the external C-coils with the GPU controller achieved the same performance as the faster internal coils (I-coils) with the more conventional proportional-integral-derivative (PID) control.

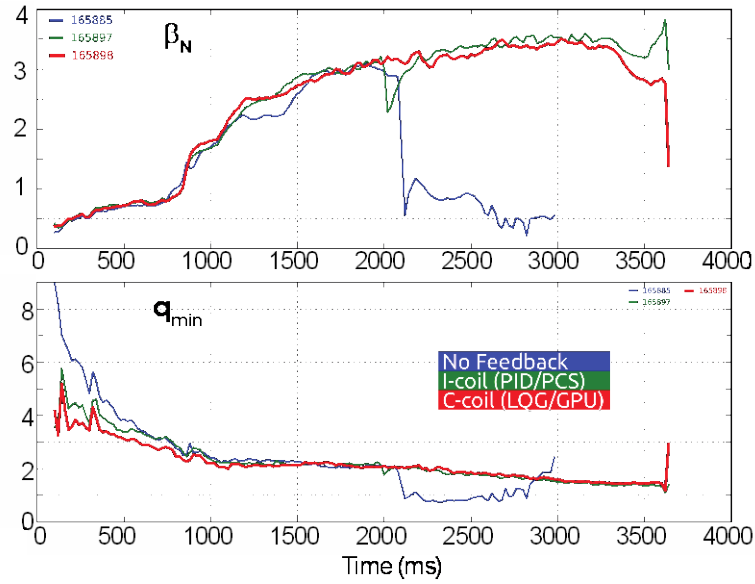


Fig. 14. Time evolution of (a) normalized beta and (b) minimum q , comparing cases with no feedback control (blue, shot 165885), I-coils with conventional PID control (green, shot 165897), and C-coils with GPU state-space control (red, shot 165898).

1.4 Recovering from an instability

DIII-D experiments have also explored remedial actions that could avoid the need for disruption mitigation after the plasma becomes unstable, including a controlled termination of the discharge or recovery of high performance operation. Once a tearing mode has locked, it is still possible to control or stabilize it. As discussed above, locked $n=1$ tearing modes can be positioned or rotated at low frequency using applied $n=1$ magnetic perturbations in order to postpone disruptions or to aid ECCD stabilization of the locked mode [10]. In the example shown in Fig. 15, feedback control of the I-coils is enabled after the tearing mode locks. The feedback gain is raised in several steps, and at sufficient gain the island begins to rotate at about 30 Hz. The rotation frequency depends on the gain, but is limited by the time constant of the resistive wall [12]. In general, forced rotation of the island simply limits its growth, avoiding an immediate disruption, unless other measures such as ECCD are used to stabilize and remove the island (see Sec 3.3).

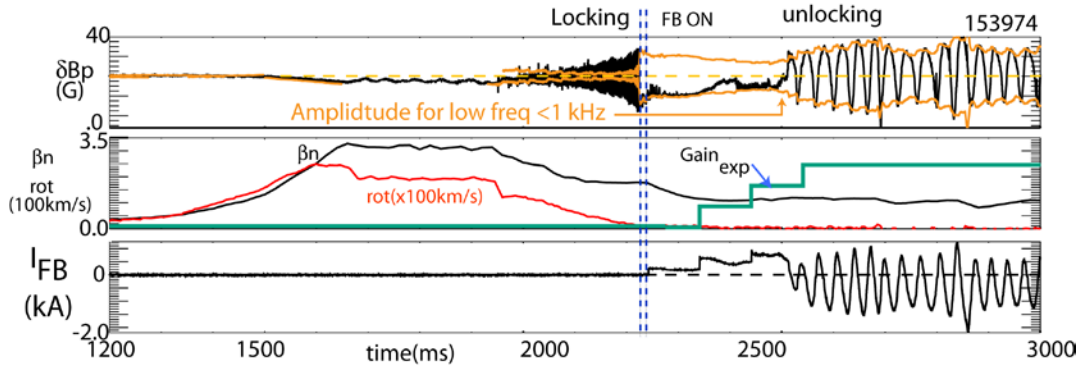


Fig. 15. Time evolution of discharge 153974, showing (first panel) magnetic perturbation of the mode, (second panel) normalized beta, plasma rotation, and feedback gain, (third panel) feedback-controlled I-coil current.

A recent experiment shows that it is possible to remove the island and recover good operation without the use of ECCD. This preliminary experiment was carried out with $q_{95} \sim 4.3$ and $\beta_N \sim 3$, both somewhat higher than for the ITER baseline scenario. In the first discharge shown in Fig. 16, a rotating tearing mode begins to grow at about $t=3.0$ s, and locks at about $t=3.4$ s. At this point the control system drops the neutral beam power to zero, while initiating feedback control of the I-coil to force the mode to rotate as it decays. When neutral beam heating resumes a short time later, the tearing mode has vanished. The plasma is then able to recover H-mode operation and return to the previous value of beta. The second discharge shown in Fig. 16 is similar, except that the beam power is reduced to 2 MW instead of zero. This is enough to sustain the island at a low amplitude, and it promptly grows and locks again when the beam power is raised.




Fig. 16. Time evolution of two discharges with and without successful recovery from a locked mode. Shown (top to bottom) are plasma current, neutral beam power, normalized beta, $n=1$ mode amplitude, and I-coil current. Red ellipses highlight the difference in minimum neutral beam power.

This comparison points to a possible recovery scenario in which the plasma is “reset” by a return to ohmic operation similar to the beginning of the discharge, before raising beta again.

1.5 Integrated control

An event handling system has been developed as part of the DIII-D plasma control system (PCS), with integrated logic for detecting instabilities (off-normal events) and system failures (faults), making decisions, and responding with appropriate actions. This off-normal and fault response (ONFR) system provides the features critical to a robust event handling system, including:

1. Sequential response to cascading faults
2. Event recovery
3. Simultaneous handling of multiple events
4. Actuator prioritization

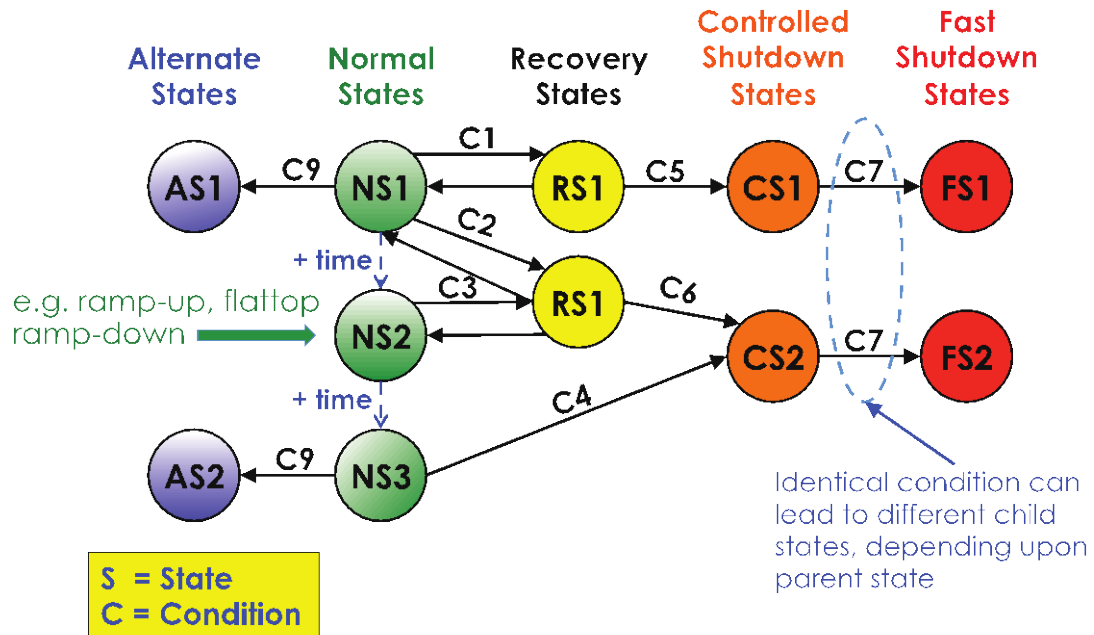


Figure 17: Conceptual diagram of finite state machine. The system exists within a given state (circles) at all times, transition between states only when one or more conditions (“C” notation above transition arrows) are valid. This diagram only depicts exclusive states

The ONFR system is implemented as a finite state machine, conceptually pictured in Fig. 17. The condition of the tokamak is at all times unambiguously described by one or more states, which determine the control actions taken. Transitions between states occur when a valid event condition is detected. Control actions only change when the state changes, not in direct response to an event. Identical events at different times during the discharge may lead to completely different responses (or lack thereof) depending upon the ONFR state at the time of the event and the available transitions from that state.

Topologically, ONFR states can be exclusive (only one active at a time), parallel (multiple states can be active simultaneously), or nested (substates within a parent state), allowing for extremely complex supervisory logic. In order to maintain the ONFR in a compact, human understandable, easily modifiable format, it has been implemented using the Stateflow® finite state extension to the Matlab® Simulink® graphical programming environment. The intuitive graphical interface enables rapid design and debugging of the finite state logic before it is automatically exported as a C library for use in the real-time PCS using the Matlab Embedded Coder®. An initial test of the ONFR system on DIII-D that demonstrates staged responses, event recovery, actuator prioritization, and simultaneous event handling leading to the safe sustainment and ramp-down of an unstable discharge is shown in Fig. 18.

At time (1), a large 2/1 NTM arises, triggering an NTM event as the $n=1$ magnitude rises above the trip threshold and shifting the ONFR into a “NTM” state parallel to the “Baseline” state. This activates ECCD suppression of the NTM, shrinking the NTM.

At time (2), the density threshold is artificially dropped to simulate a high density event, moving the ONFR into a High Density state in parallel with the Baseline and NTM states. The High Density state requests the gyrotrons to turn off for their own protection, in conflict with the NTM state. However, the High Density state (being for machine protection) has higher actuator priority than the NTM state, so its request prevails and the gyrotrons turn off. This leads to uncontrolled growth of the original NTM, and subsequent locking.

The locked mode event at time (3) causes a transition to the Locked Mode state in parallel with the Baseline and High Density states, requesting both rotating $n=1$ fields from the I-coils for mode entrainment and synchronous ECCD deposition on the O-point. The High Density state initially blocks the gyrotron request from the Locked Mode state, but between times (3) and (4) the density recovers, causing a transition out of the High Density state and ceding gyrotron control to the lower priority Locked Mode state. This leads to a brief recovery from the locked mode condition. At time (4), but the locked mode re-appears, causing another transition into the Locked Mode state to control the locked mode until the successful end of the discharge.

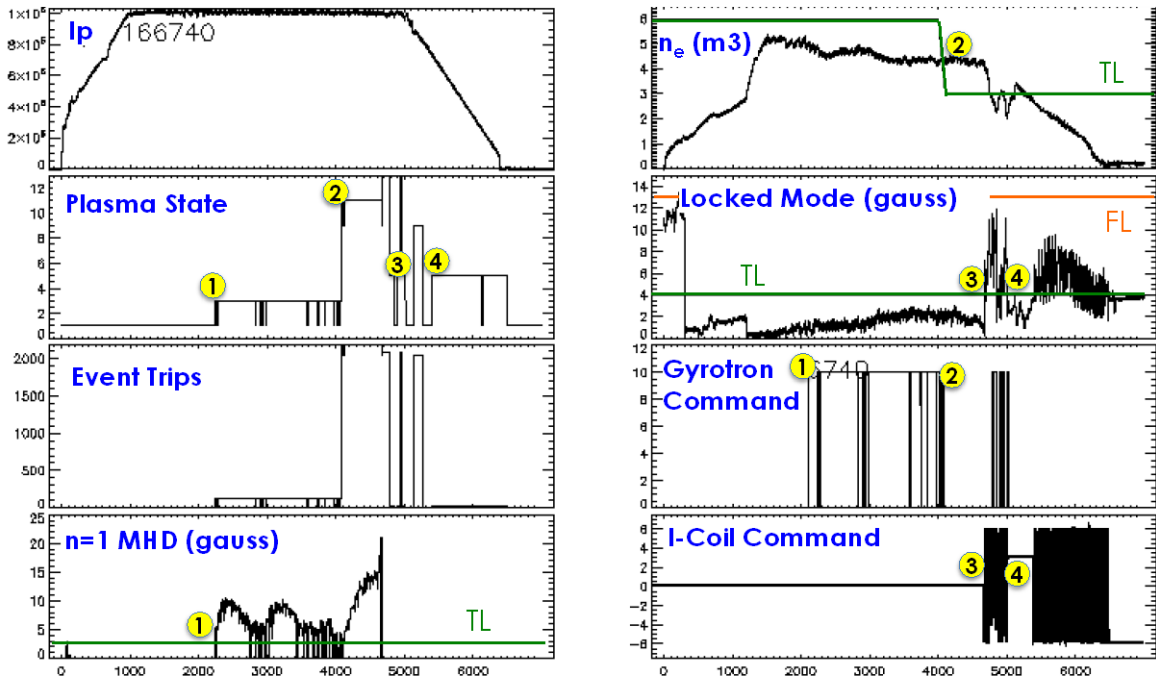


Figure 18: ONFR handling of NTM→locked mode discharge. Time points referred to in text are denoted by yellow circles. Green lines denote threshold levels above which an event is tripped. The Plasma State and Event Trips signals are bit masks identifying the active ONFR states and events, respectively.

The above case demonstrates the complex behavior that can be extracted from even a fairly simple finite state ONFR architecture. Further development will focus upon adding hysteresis to the event recovery to avoid dithering in and out of states and implementing a system for actuator sharing in addition to actuator prioritization.

2.0 Experimental Results and Analysis – Disruption Mitigation

2.1 Thermal quench mitigation in plasmas with pre-existing MHD activity

Experiments on DIII-D studied the effect of pre-existing MHD upon the effectiveness of both shattered pellet injection (SPI) and massive gas injection (MGI). An ITER baseline scenario plasma served as the target, as that scenario reliably produces 2/1 NTM that slow into a locked mode prior to disrupting when rotation is slowed using balanced neutral beam torque. Newly implemented plasma control system (PCS) control over the impurity injection trigger enabled asynchronous injection based upon tearing mode amplitude or locked mode amplitude, with variable delay, providing a continuous scan of injection timing relative to the start of MHD (see Fig. 3 of [15]).

Overall, experimental results indicate that the presence of pre-existing MHD does not impede the effectiveness of the impurity injection. Fig. 19 displays the assimilation of SPI, as measured by the peak line-integrated density following injection. No significant difference in impurity assimilation is observed between SPI during the healthy H-mode or L-mode period versus after the MHD begins, either as a rotating or locked mode. Similarly, protection of the divertor from thermal quench heat flux, which is the end goal of thermal quench mitigation, is unaffected by the existence of a mode. As shown in Fig. 20b, the divertor heat flux after injection remains fairly constant regardless of the state of the plasma at the time of SPI, whereas the unmitigated disruptions result in

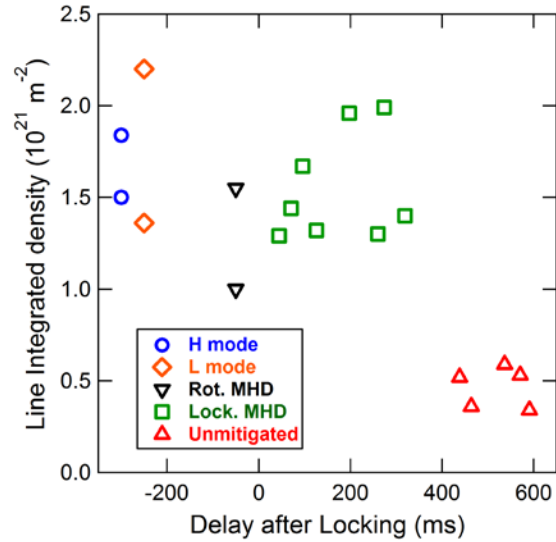


Fig. 19 Peak line integrated density following SPI injection, as a function of SPI delay from onset of locking (except unmitigated cases). Symbol coloring indicates the MHD state of the plasma, as noted in the legend. [From Ref. 15]

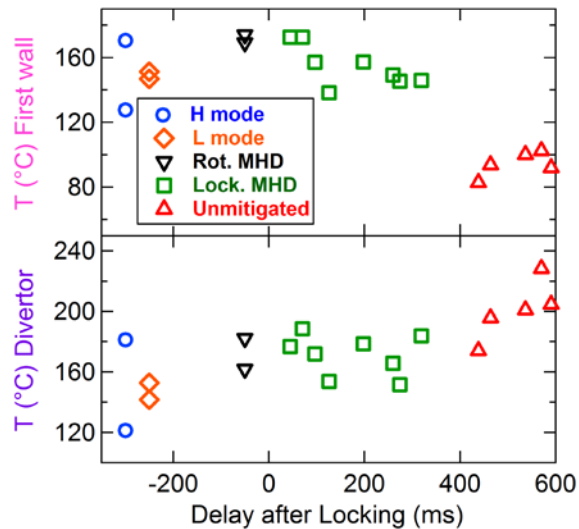


Fig. 20. (a) Peak first wall temperature following SPI, as a function of SPI delay from onset of locking. (b) Peak divertor temperature. [From Ref. 15]

noticeably higher divertor temperature. Conversely, the first wall temperature is consistently higher for all mitigated cases to the unmitigated cases, indicating that the SPI is equally effective at radiating plasma thermal energy to the wall (avoiding conduction to the divertor) regardless of the presence or absence of MHD activity in the plasma (Fig. 20a).

Similar results are exhibited for neon MGI. The total radiated energy during the thermal quench stays fairly level for all mitigated cases (Fig. 21a), even as the total plasma thermal energy content drops precipitously in the presence of a locked mode (Fig. 21b). This indicates that a large fraction of the plasma thermal energy is being radiated to the wall in all mitigated cases, regardless of mode activity. Unmitigated disruptions exhibit a much lower level of radiated power, demonstrating that even mitigation late in a locked mode is far preferable than no mitigation at all.

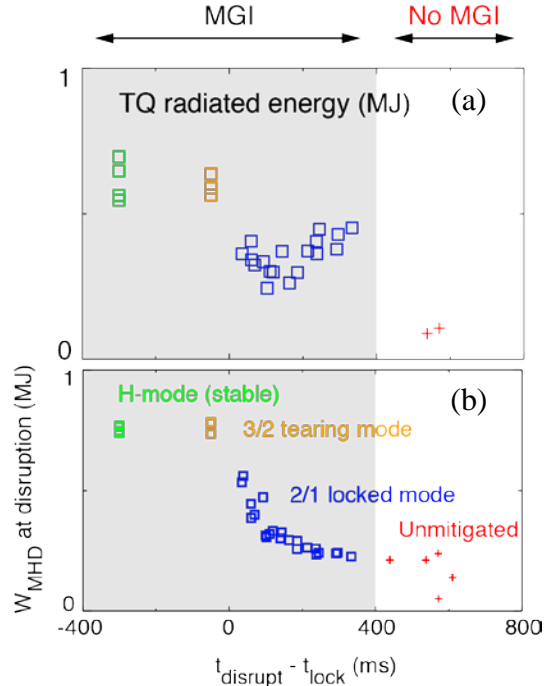


Fig. 21. (a) Total energy radiated during the thermal quench. (b) Plasma thermal energy at time of MGI (except unmitigated cases). Coloring indicates the plasma MHD state at time of MGI, as noted by labels. [From Ref. 16]

2.2 Mitigation of runaway electrons

2.2.1. Dissipation of runaway electrons using SPI

An initial experiment [17] was undertaken to compare the effectiveness of shattered pellet injection (SPI) vs massive gas injection (MGI) for dissipation of a post-disruption runaway electron (RE) plateau. Due to its compact injector size and rapid impurity transport relative to massive gas injection (MGI), SPI is the primary disruption mitigation system technology option for ITER. However, DIII-D is presently the only device with SPI capability. Until this time, SPI had never been tested for the critical DMS task of dissipating runaway electron beams.

Figure 22 displays a comparison of 1200 Torr-L neon MGI (blue) and SPI (red) fired into similar RE plateaus, created by injection of a small (~ 16 Torr-L) argon pellet. After injection of the impurities around 1400ms, both the MGI and SPI initially exhibit an almost identical rapid dissipation of RE current, as indicated by the parallel dotted lines in Fig. 22(a). However, whereas the dissipation by neon MGI continues linearly to 0A, the dissipation in the SPI case flattens after 30-50ms, resulting in a lower plateau current at the saturated loop voltage of ~ 10 V (Fig. 22(b)). This post-SPI plateau exhibits a high density ($\sim 1E20 \text{ m}^{-3}$) as expected (Fig. 22(c)). However, it also exhibits a complete lack of residual Argon (Ar) emission from the original argon pellet that is evident prior to the SPI emission (Fig. 22(d)).

The fact that the initial dissipation of the RE plateau by SPI and MGI is equivalent, as well as the roughly equivalent initial jumps in density, indicate that both mitigation technologies are equally effective for RE dissipation. However, the subsequent reduction of the long time-scale SPI dissipation indicates more complicated processes are at work. The sustained high density after the SPI injection indicates that the SPI material is resident in the RE beam. However, at the same time the complete absence of Ar light indicates removal of the highly dissipative residual Ar left over from the initial Ar pellet.

One significant difference between the neon MGI and the equivalent neon SPI is that the SPI includes a small (estimated at $\sim 10\%$) amount of deuterium (D_2). That D_2 is required as a lubricant on the outside of the neon pellet so that the SPI can fire in a simple pipe-gun configuration without the pellet shattering from the propellant gas. It is postulated that the minority D_2 in the neon SPI displaces the higher Z impurities, resulting in a much lower resistance RE beam than the pure neon MGI.

Purging of high-Z impurities by lower Z impurities has been also observed during MGI dissipation of RE beams [18]. As shown in Fig. 23, equivalent amounts of pure Ar and 90%/10% Ar/ D_2 MGI fired into a runaway beam result in two very different evolutions. Similar to SPI, both the pure and mixed MGI cases initially dissipate the RE beam at the same rate, but then the dissipation rate in the mixed species case slows significantly. For reasons not yet understood, the low-Z impurities, even when a small minority, appear to displace the higher-Z species. If a pure low-Z impurity injection is used (i.e. pure D_2 or helium), this “cleanout” effect is even more pronounced. This effect does not appear to be a function of injector technology (SPI or MGI). Rather, it appears to be a basic transport process that is not yet well understood.

These initial results may have ramifications for the ITER SPI design. If the small amount of required D_2 lubricant significantly impedes the efficacy of SPI for RE dissipation, then the simple pipe-gun design

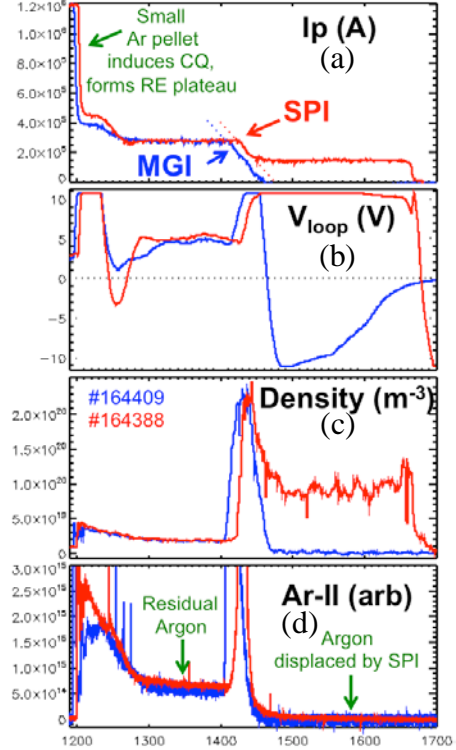


Fig. 22: Comparison of RE plateau dissipation by MGI (blue) and SPI (red).

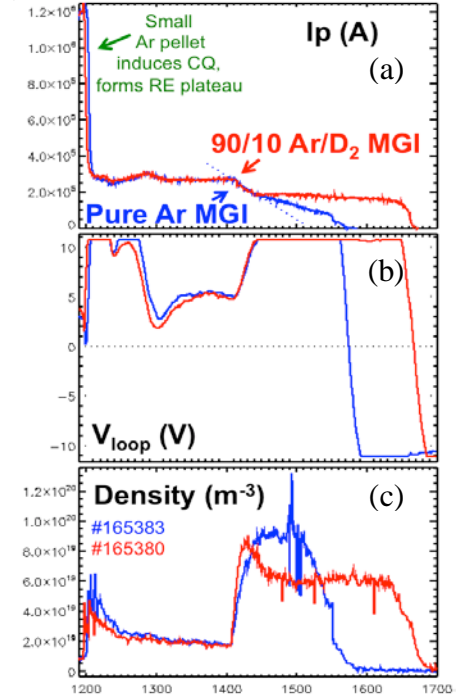


Fig. 23: Comparison of RE plateau dissipation by pure Argon MGI (blue) and 90%/10% Ar/ D_2 MGI (red).

may not be adequate for ITER. Instead, a mechanical punch system that does not require D2 lubricant, similar to that used for the Ar pellet launcher on DIII-D, may be required even for neon injection. Further experiment and analysis is required to confirm these recent results.

2.2.2. Affect of impurities upon runaway electron energy distribution function

In order to assure rapid dissipation of a runaway electron beam in ITER, it is critical to understand the physics governing that dissipation. It has been shown in [19] that RE plateau current dissipates significantly faster than expected from standard avalanching theory [20] in the presence of massive impurity injection. Similarly, it was demonstrated in [21] that quiescent runaway (QRE) population, formed during the plateau of low-density ohmic discharges, also dissipates faster than expected from avalanche theory when the RE population was measured using an external hard X-ray (HXR) detector.

Recent results using the new DIII-D Gamma Ray Imager (GRI) [22] to measure the QRE energy spectrum via pulse height analysis of its bremsstrahlung emission reveal that the RE dissipation rates are actually a function of RE energy. The dissipation/growth of the RE population in the presence of impurities is not uniform across energy spectrum, and the divergence from conventional avalanche theory is also not uniform. As shown in Fig. 24, the growth rate curves approaches that expected from avalanche theory (crossing zero growth rate at $E/E_{crit} = 1$) as the energy of the sampled portion of the energy spectrum increases.

The external HXR data from, shown in black, indicates that the detector is primarily sensitive to the low energy portion of the RE spectrum, which exhibits the greatest divergence from avalanche theory.

These energy-resolved results allow for detailed, direct comparison with various models for RE evolution in order to verify and improve those models.

2.2.3. RE magnetic to kinetic energy conversion during the final loss

In ITER, the RE beam magnetic energy will be several times larger than the kinetic energy. It has been demonstrated [23, 24] that the RE magnetic energy can be rapidly converted into RE kinetic energy during the final termination of the RE beam, drastically

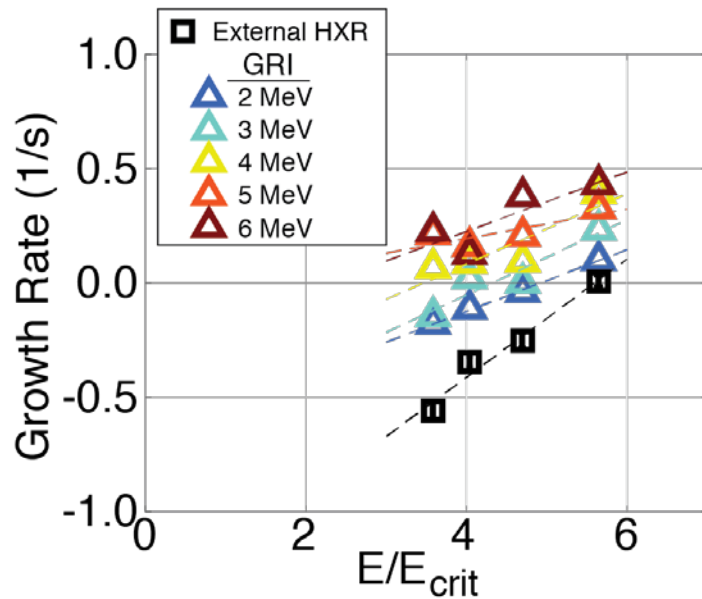


Figure 24. Quiescent runaway growth rate as a function of applied electric field normalized to the critical electric field for avalanching. Colors represent different portions of the runaway energy distribution. Dashed lines are linear fit to each energy group.

increasing the destructive capability of the beam several time over. Some models indicate this process could potentially get worse in the presence of high-Z impurities used to rapidly dissipate RE beams.

Recent studies on DIII-D systematically examined the affect of injected impurities upon this energy conversion process. Initially centered RE plateaus under Ip control were injected with varying species and quantities of impurities and then kicked into the lower divertor where they disrupted and underwent their final loss. The energy deposited to the wall during the RE final loss was then

estimated using infrared measurements of wall heating and HXR measurements. As shown in Fig. 25(b), the deposited energy does not increase monotonically with injected Z but rather reaches a maximum and then starts to decrease. This non-monotonic behavior has two contributors. First, at high Z the kinetic energy of the RE beam is significantly reduced at the time of the final loss (see Fig 25(a)), reducing the total energy available for deposition to the wall. Secondly, shorter final loss durations at high Z also reduce the efficiency of the magnetic to kinetic transfer. The low-Z data is still not well understood, but the presence of a deposited energy maximum at medium Z indicates that high Z (e.g. Argon, Xenon) dissipation of the RE beam can be undertaken without fear of worsening the energy transfer process.

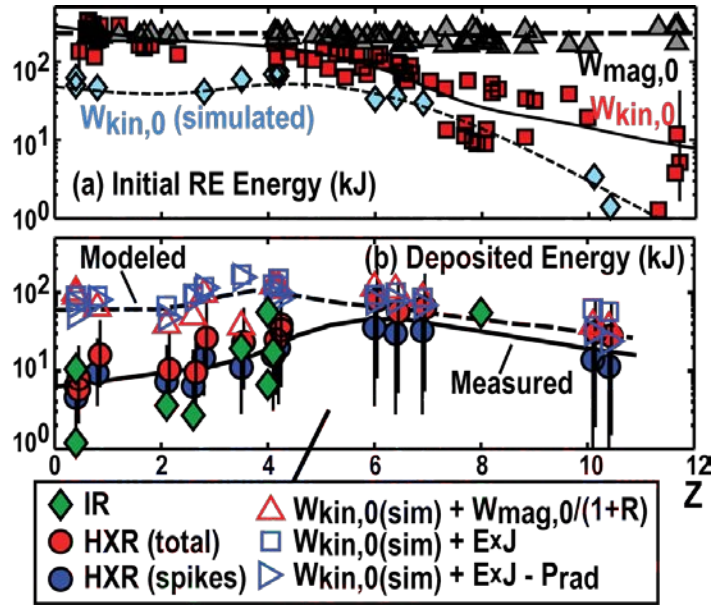


Figure 25. (a) Magnetic and kinetic energy content of RE beam at final loss as function of Z. (b) Energy deposited to wall during final loss, as function of Z.

References

- [1] C.M. Greenfield, et al., “Transient Events in Tokamak Plasmas,” Report of Fusion Energy Sciences Workshop (June 2015).
- [2] F. Turco, et al., “Confinement and Stability of the ITER Baseline Scenario in DIII-D,” submitted to 26th IAEA Fusion Energy Conf. (Kyoto, 2016).
- [3] B. Tobias, et al., “Rotation profile flattening and toroidal flow shear reversal due to the coupling of magnetic islands in tokamaks” *Phys. Plasmas* 23, 056107 (2016).
- [4] R. Sweeney, et al., “Statistical analysis of m/n=2/1 locked and quasi-stationary modes with rotating precursors in DIII-D,” GA-A28126 (2016), submitted to *Nucl. Fusion*.
- [5] R.J. La Haye, et al., “Exploration of the means for real-time probing of m/n=2/1 tearing mode stability evolution in the ITER baseline scenario in DIII-D,” 43rd EPS Conference on Plasma Physics (Leuven, Belgium, 2016), paper P5.026.
- [6] M.J. Lanctot, “Impact of toroidal and poloidal mode spectra on the control of non-axisymmetric fields in tokamaks,” to be submitted to *Phys. Plasmas* (2016).
- [7] C. Paz-Soldan, et al., “Decoupled recovery of energy and momentum with correction of n=2 error fields,” *Nucl. Fusion* 55, 083012 (2015).
- [8] M.J. Lanctot, et al., “A Path to Stable Low Torque Plasma Operation in ITER with Test Blanket Modules,” GA-A28009, 2015, submitted to *Nuclear Fusion*.
- [9] M.J. Lanctot, et al., “Error field optimization in DIII-D using extremum seeking control,” 2016 *Nucl. Fusion* 56 076003
- [10] F. A. Volpe, et al., “Avoiding Tokamak Disruptions by Applying Static Magnetic Fields That Align Locked Modes with Stabilizing Wave-Driven Currents,” *Phys. Rev. Lett.* 115, 175002 (2015).
- [11] W. Choi, et al., “Feed-back control of 2/1 locked mode phase: experiment on DIII-D and modeling for ITER,” 58th Annual Meeting of the APS Division of Plasma Physics (2016), <http://meetings.aps.org/Meeting/DPP16/Session/TO4.11>
- [12] M. Okabayashi, et al, “Avoidance of Tearing Mode Locking and Disruption with Electro-Magnetic Torque Introduced by Feedback-based Mode Rotation Control in DIII-D and RFX-mod,” 25th IAEA Fusion Energy Conference, Oct. 13-18, 2014, EX/P2-42.
- [13] R. La Haye, D.P. Brennan, R.J. Buttery and S.P. Gerhardt, "Islands in the stream: the effect of plasma flow on tearing stability", *Phys. Plasmas* **17**, 056110 (2010)
- [14] M.D. Clement, et al, “Resistive Wall Mode feedback on DIII-D using Linear Quadratic Gaussian control and a GPU powered control system,” 57th Annual Meeting of the APS Division of Plasma Physics (2015), <http://meetings.aps.org/link/BAPS.2015.DPP.JP12.90>

- [15] N. Commaux *et al.*, “Shattered Pellet Injection Mitigation Applied To MHD Unstable Plasmas On DIII-D,” presented at the APS-DPP conference, Nov. 16-20, 2015.
- [16] D. Shiraki *et al.*, “The impact of pre-existing MHD instabilities on disruption mitigation by shattered pellet and massive gas injection,” to be submitted to Nuclear Fusion.
- [17] N. Commaux, *et al.*, “Runaway Electron - SPI Dissipation in Current Quench,” DIII-D Miniproposal D3DMP 2015-01-05, Dec. 7, 2015.
- [18] E. Hollmann, *et al.*, “RE final loss dynamics,” DIII-D Miniproposal D3DMP 2016-25-01, Jan. 29, 2016.
- [19] E. Hollmann *et al.*, “Effect of applied toroidal electric field on the growth/decay of plateau-phase runaway electron currents in DIII-D” Nucl. Fusion **51**, 103026 (2011)
- [20] M.N. Rosebluth and S.V. Putvinski, “Theory for avalanche of runaway electrons in tokamaks” Nucl. Fusion **37**, 1355-1362 (1997)
- [21] C. Paz-Soldan *et al.*, “Growth and decay of runaway electrons above the critical electric field under quiescent conditions” Phys. Plasmas **21**, 022514 (2014)
- [22] D.C. Pace *et al.*, “Gamma ray imager on the DIII-D tokamak” Rev. Sci. Instrum. **87**, 043507 (2016)
- [23] A. Loarte *et al.*, “Magnetic energy flows during the current quench and termination of disruptions with runaway current plateau formation in JET and implications for ITER” Nucl. Fusion **51**, 073004 (2011)
- [24] E. Hollmann *et al.*, “Control and dissipation of runaway electron beams created during rapid shutdown experiments in DIII-D” Nucl. Fusion **53**, 083004 (2013)

Alcator C-Mod contribution to FY 2016 Joint Research Target – 4th quarter R. Granetz and A. Tinguely

4th quarter JRT activities for the Alcator C-Mod group primarily encompassed two main thrusts. One is a comparison of synchrotron spectra from relativistic runaway electrons at three different magnetic fields. We have also been doing more analysis of the synchrotron emission images. The other major thrust has been our continuing analysis of disruption forecasting using our established databases at C-Mod, DIII-D, and EAST.

Runaway electron synchrotron spectra and image studies

The theoretical formula^[1-3] for the spectrum of synchrotron light emitted by a relativistic electron in a tokamak configuration varies significantly with the amplitude of the toroidal magnetic field. Alcator C-Mod is capable of running fusion-relevant plasmas over a very large range of toroidal magnetic fields, from 2.7 to 8.0 tesla. Figure 1 shows the theoretically predicted synchrotron




Fig. 1 – Theoretically calculated spectra of synchrotron emission from a relativistic electron in a tokamak configuration for three different values of toroidal field.

spectra for toroidal fields of 2.7, 5.4, and 7.8 tesla from a single runaway electron at a fixed energy of 30 MeV and a fixed pitch (v_{\perp}/v_{\parallel}) of 0.1. As the B-field is raised, the overall radiated power per electron obviously increases dramatically for fixed energy and pitch. But for the C-Mod experiment, we are particularly interested in the relative shapes of the spectra. As the B-field is raised, the calculated synchrotron spectrum shifts towards shorter wavelengths. This is a little difficult to see in Fig. 1 due to the concomitant increase in power, and also because the peak of the emission for a 30 MeV electron is in the infrared, off the right side of the plot, beyond the wavelength range of our spectrometers. The shift towards the blue can be better seen in Figure 2, where we have taken the ratio of each field's spectrum to the 5.4 tesla spectrum, after first normalizing their intensities to their value at 600 nm respectively. This removes the effect of the increasing amplitude, and leaves just the effect of the change in relative spectral shape. It is clearly seen that as the B-field increases from 2.7 to 5.4 to 7.8 tesla, the relative spectra increase at shorter wavelengths (left side of Fig. 2) and decrease at longer wavelengths. *Again, this shift of the synchrotron spectra towards shorter wavelengths assumes that the runaway energy and pitch remains fixed.*

Experiments to generate runaway electrons during the flattop of very low density, MHD-quiet discharges at 2.7, 5.4, and 7.8 tesla were successfully carried out during the last quarter

of FY2016. (These were the final disruption-related experiments on Alcator C-Mod. The tokamak is now permanently shut down. These high-field runaway experiments are likely to be the last ones for decades to come.) A pair of dedicated spectrometers, which have been absolutely calibrated in-situ, were used to measure the synchrotron spectra. The parameters that determine runaway electron population, and therefore the absolute intensity of the synchrotron emission, are not well known, and the absolute intensity is not reproducible. That is why we are particularly interested in the spectral shapes, and not their relative

amplitudes. Figure 3 shows the measured brightness spectra for the three B-fields. The brightness spectra are obtained by convolving the raw data with the absolute calibrations. This accounts for the wavelength-dependent absorption in the quartz transmission fibers, as well as the detector array sensitivity versus wavelength. There are also a few deuterium Balmer series lines (i.e. D α at 656.1 nm) from the background plasma which have been removed or reduced manually. Unlike, the theoretical spectra, the measured spectra at the three different B-fields

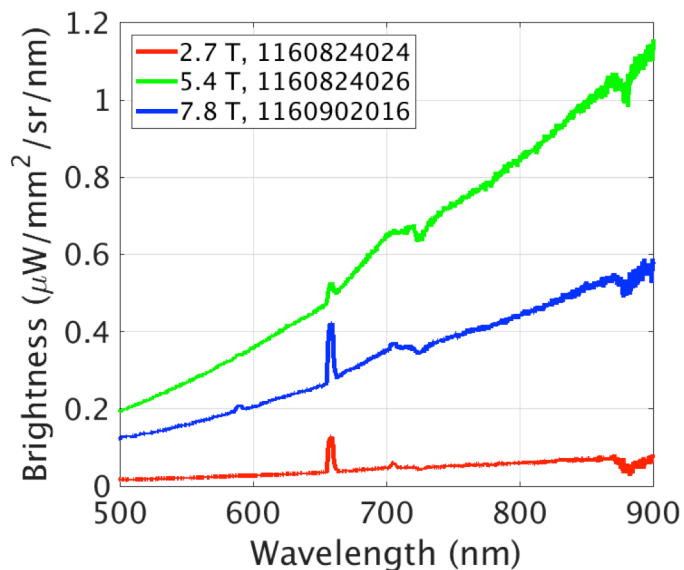


Fig. 3 – Synchrotron emission spectra measured on Alcator C-Mod at three different magnetic fields. The spectra are absolutely calibrated. The spectra are very similar in shape.

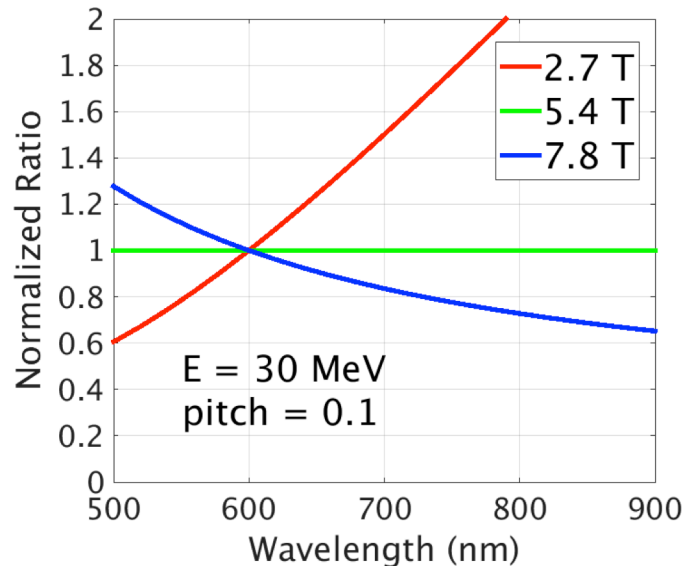


Fig. 2 – Ratios of the theoretically calculated spectra to the 5.4 tesla spectrum, normalized in amplitude at $\lambda = 600$ nm. The shift towards shorter wavelengths at higher B-field is clearly seen.

have very nearly identical shapes.

This can be seen more clearly in Figure 4, which shows the ratio of each spectrum to the 5.4 tesla spectrum, after normalizing their amplitudes at 600 nm. This is exactly the same procedure used with the theoretical spectra to produce Fig 2, and the vertical scale is the same in both plots. It is clearly seen that the measured synchrotron spectra are essentially identical in shape, even though the toroidal field varies by a factor of 3.

One possible explanation for these results is that the relativistic electron energy does not remain constant as the B-field is changed. According to




Fig. 4 – Ratios of the experimentally measured spectra to the 5.4 tesla spectrum, normalized in amplitude at $\lambda = 600$ nm. (Same processing as done for Fig 2.) The measured spectra at the three different B-fields have remarkably shapes.

field is raised, the synchrotron power radiated per electron increases theoretically. If the maximum relativistic energy of each electron is actually limited by its own synchrotron emission, then its energy will decrease at higher field. This intriguing hypothesis is being explored in detail as part of a doctoral thesis.

In addition to the spectral measurements, we have also been studying in more detail the images of synchrotron emission in C-Mod flattop runaway discharges that have been seen with our video cameras. A particularly interesting image is shown in Figure 5. For a highly relativistic electron that is traveling parallel to a pure toroidal field, the synchrotron emission cone sweeps out a pattern that is not difficult to describe. But a relativistic electron that is traveling along the helical field

theory, the synchrotron spectrum increases in amplitude and shifts toward shorter wavelengths as the electron energy increases. Conversely, the spectrum shifts towards longer wavelengths as the electron energy decreases. Therefore, if the relativistic electron energy decreases as the B-field is raised, the blue shift due to the higher field can be countered by the red shift due to lower electron energy, perhaps leaving the spectral shape nearly unchanged. But in order to be consistent, there needs to be a mechanism that causes this to happen. That mechanism could be the synchrotron power itself, *if it is the primary energy loss mechanism that sets the maximum energy of the runaways*. Theoretically, as the B-

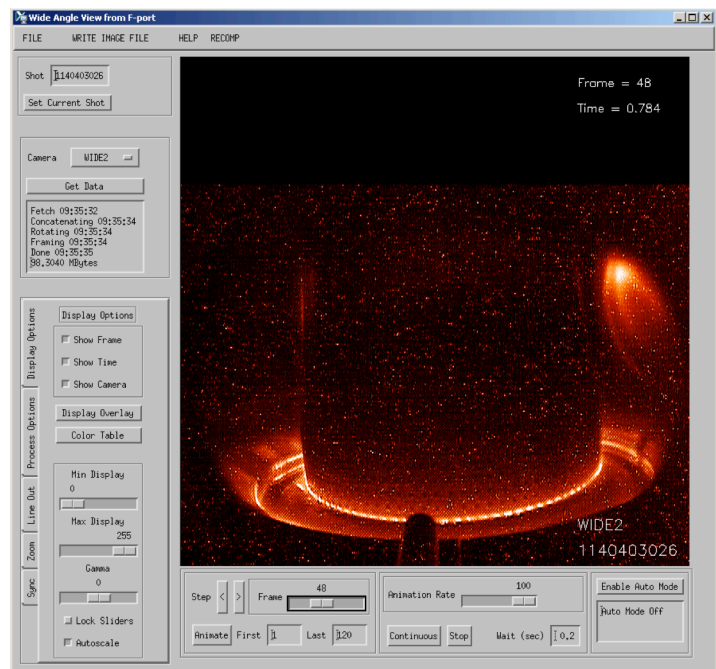


Fig. 5 – Camera image of synchrotron emission from runaway electrons in Alcator C-Mod. The emission is only on the right side of the torus, due to the extreme forward-peaked cone from the highly relativistic electrons. Note the complexity of the emission pattern.




Fig. 6 – Synthetic image of synchrotron emission, as calculated by M. Hoppe and O. Embréus (Chalmers Univ.) using the actual C-Mod geometry and plasma equilibrium fields. There are notable similarities with the measured image shown in Fig 5.

lines of a tokamak, and that has a fraction of its velocity in the direction perpendicular to the magnetic field (i.e. finite pitch angle), has a much more complicated emission pattern. This gets even more complicated if the imaging camera is well off the mid-plane, as it is in C-Mod. We have been working with the theory group at Chalmers, in particular Mathias Hoppe and Ola Embréus to understand the images. Mathias has written a program that calculates a synthetic image that an observer would see for a uniform spatial distribution of runaways in a tokamak, with a given energy and pitch. Figure 6 shows one of the synthetic images using parameters that mimic the C-Mod geometry, including the camera location and viewing direction. A mono-energetic population at 30 MeV, with a pitch of 0.1, and a

uniform distribution out to a minor radius of $r/a \sim 0.75$), except for a small void on the magnetic axis, was assumed. The actual magnetic fields from an equilibrium reconstruction are used in the calculation. It seems that just with these very simple parameters we can get quite close to the observations, and it may be possible to extract information about the energy, pitch, spatial distribution, etcetera from detailed comparisons of the images with the synthetic calculations.

Disruption warning database development

One of the key tasks of this year's JRT is to determine if a disruption warning algorithm can be developed that works on multiple tokamaks. We have set up SQL disruption warning databases for Alcator C-Mod, EAST, and DIII-D. These databases contain relevant plasma parameters, mostly taken from work by S. Gerhardt on NSTX^[4], sampled repetitively throughout both disruptive *and* non-disruptive discharges for the entire 2015 campaign on each machine. These SQL databases are primarily designed to allow queries about the behavior of certain plasma parameters as a disruption approaches, *and* to compare this to the behavior of these parameters on shots that do not disrupt. Previously, we had populated about half of the parameters for the C-Mod and EAST databases, and during this last quarter we have begun populating many of the parameters in our DIII-D database. In order to expedite this process, we have hired a post-doc, Dr. Cristina Rea, who started on Aug 1st, and is sited at GA. Dr. Rea has extensive experience in the application of advanced machine learning algorithms^[5-8], such as support vector machines (SVM), random forests, and deep neural networks, to large datasets in order to organize the data

into distinct classes for predictive purposes. The machine learning work is just getting started as this JRT is wrapping up, but this is a long-term goal that will continue unabated with strong support.

In order to accurately capture the pre-disruption behavior of relevant plasma parameters, we believe these parameters must be sampled relatively frequently during a finite window of time before disruptions occur, *but* not as frequently throughout the entirety of discharges. This is important as a practical matter, since it significantly reduces the overall number of records required in the database to a manageable level. Even with this assumption, our DIII-D disruption warning database has more than a half-million records for its 2146 plasma discharges in the 2015 campaign, and our C-Mod database has 167,500 records for its 1821 plasmas in 2015. For C-Mod, we sample the parameters at 20 ms intervals throughout the full duration of all plasma discharges (corresponding to the default EFIT analysis times), and we augment these with samples at 1 ms intervals for the 20 ms period before each disruption. For DIII-D, the full-shot intervals are every 25 ms, supplemented by 2 ms sampling for the 100 ms period before each disruption. These sampling rates and pre-disruption windows were determined by a general survey of disruptions on each machine, but as we continue our analysis, we may decide in the future to alter these sampling specifications. Since many of the plasma parameters are derived from EFIT^[9], we elected to run our own dedicated EFITs for C-Mod, DIII-D, and EAST, with reduced filtering of magnetics input data, at the desired sampling times. This eliminates the need to interpolate EFIT quantities at high sampling rates from the less frequent default analysis rates.

At this time, our disruption warning databases on the different machines, although not completed, have enough key parameters populated to begin investigating the differences between the machines. These include parameters such as I_p error (the difference between the actual plasma current and the pre-programmed current) at each sampling time, loop voltage, radiated power fraction ($p_{\text{rad}}/p_{\text{input}}$), l_i , q_{95} , vertical position error (z error), etc. But we do not yet have other notable disruption-relevant parameters such as $n=1$ δB amplitude and/or locked mode amplitude. It should be emphasized that for all of the plasma parameters, it is paramount to avoid using signals that have been processed with excessive non-causal smoothing/filtering, since that can artificially introduce post-disruption behavior into pre-disruption data samples.

One of the clear differences between machines can be seen by looking at the behavior of I_p error. Figure 7 shows the data for I_p error plotted versus time

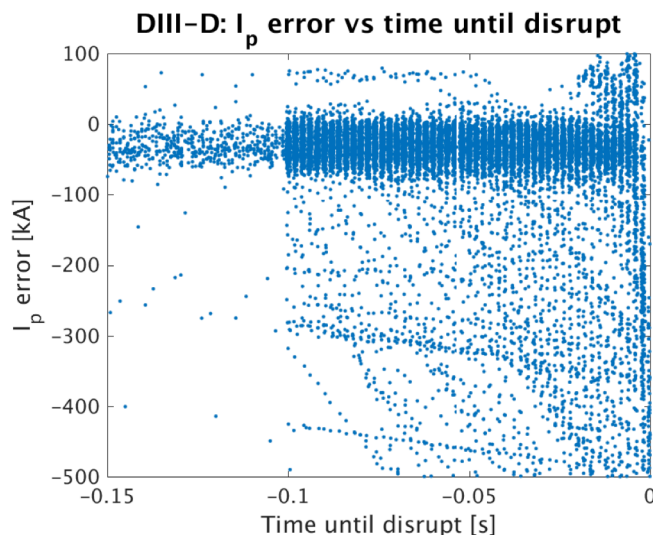


Fig. 7 – I_p error versus time before disrupt, for all shots that disrupted during flat-top on DIII-D in the 2015 campaign. A noticeable fraction of the data points have a growing I_p error well before the time of disruption.

before disrupt for all DIII-D discharges in the 2015 campaign that disrupted during the flattop. (The increase in sampling rate from every 25 ms to every 2 ms during the 100 ms pre-disruption period is clearly seen.) A noticeable fraction of the data points have a growing I_p error well before the time of disruption. The error is predominantly negative, showing that the actual I_p tends to droop below the pre-programmed I_p prior to disruptions. Figure 8 shows the histogram of the I_p error during the flattop for all non-disruptive discharges in the DIII-D 2015 campaign, and Figure 9 is the histogram for I_p error during the period between 30-50 ms before the time of disruption. The probabilities are normalized so that the integrals are unity. The narrow width in Fig 8 essentially shows the quality of the plasma control system's I_p feedback. If we specify an error of -60 kA as a threshold for declaring an impending disruption, then the histogram data tells us that 25.6% of all flattop disruptions on DIII-D in 2015 would have been predicted at least 30 ms before they disrupt, with a false positive rate of 3.0%.




Fig. 8 – Histogram of I_p error on DIII-D for all flattop times from all non-disruptive shots in 2015.




Fig. 9 – Histogram of I_p error on DIII-D for the period between 30 and 50 ms before the time of disruption, for all flattop disruptions in 2015.

In contrast, the I_p error parameter does not seem to be similarly successful at providing an indication of impending disruption in Alcator C-Mod. Figure 10 shows the I_p error data prior to disruption for all flattop disruptions in the C-Mod 2015 campaign. (The increase in sampling rate from every 20 ms to every 2 ms during the 20 ms pre-disruption period is clearly seen.) Not much change in the dataset is seen until 10 ms or less before the disruption. Figure 11 shows the histogram of the I_p error during the flattop for all non-disruptive discharges in C-Mod (showing the quality of I_p feedback control), and Figure 12 is the histogram for I_p error during the period between 30-50 ms before the time of disruption, corresponding to Fig 9 for DIII-D. For the same threshold level of -60 kA, the histogram data tell us that only 8.5% of all flattop disruptions in C-Mod in 2015 would have been predicted at least 30 ms before the disruption occurred, which is not nearly as good as the 25.6% success rate for DIII-D. Even if we take the I_p error data from 10-20 ms before the disruption time, the success rate only rises to 13.5%. A warning time of 10-20 ms may be barely enough time to trigger a disruption mitigation scenario, but not likely enough time to switch to an avoidance scenario.

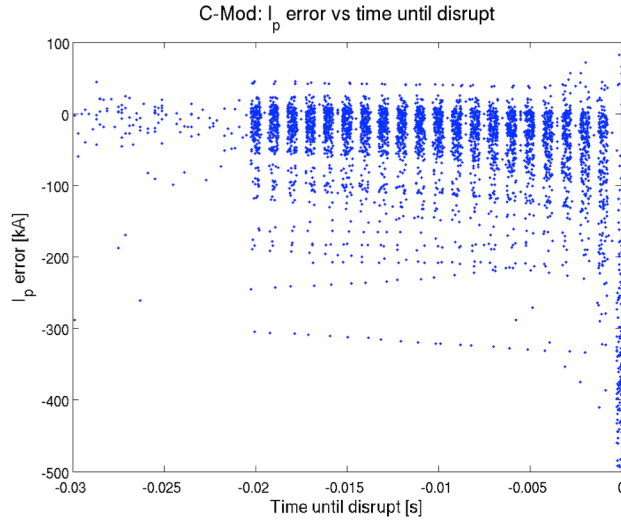


Fig. 10 – I_p error versus time before disrupt, for all shots that disrupted during flattop on C-Mod in the 2015 campaign. Not much change in the dataset is seen until 10 ms or less before the disruption.




Fig. 11 – Histogram of I_p error on C-Mod for all flattop times from all non-disruptive shots in 2015.




Fig. 12 – Histogram of I_p error on C-Mod for the period between 30 and 50 ms before the time of disruption, for all flattop disruptions in 2015.

The radiated power fraction is presumably another useful indicator of an approaching disruption. We have the $p_{\text{rad}}/p_{\text{input}}$ data in the C-Mod database, but it will be a few more weeks before that data is available in the DIII-D database. (The standard P_{rad} signals available on DIII-D have unacceptable non-causal smoothing, and we will have to re-calculate these signals from the raw data for all the plasmas in the 2015 campaign.) But this information is available in our EAST database, which will be used as a comparison here. Figure 13 shows the radiated power fraction in EAST versus time until disrupt, for all flattop disruptions in their 2015 campaign. The number of points with radiated fraction > 0.35 is increasing as the discharges approach disruption. Figure 14 shows a histogram of radiated power fraction during the flattop of all non-disruptive shots in




Fig. 13 – Radiated power fraction on EAST vs time until disrupt for all flattop disruptions in the 2015 campaign. A noticeable increase in points with value > 0.35 is seen as the disrupt time is approached.




Fig. 14 – Histogram of radiated power fraction for all flattop times from all non-disruptive shots in EAST.

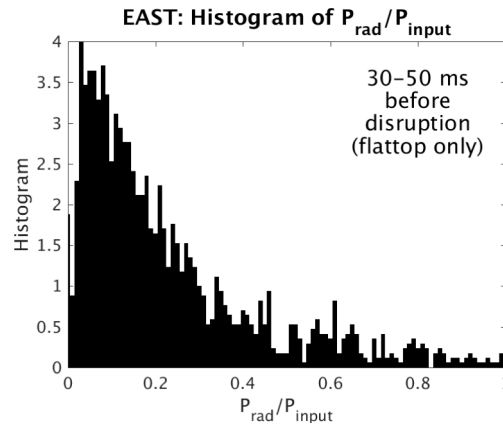


Fig. 15 – Histogram of radiated power fraction from 30-50 ms before the time of disruption on EAST.

EAST. Specifying a value of 0.35 as the threshold for triggering a disruption warning would result in a false positive rate of 1.0%. Figure 15 shows the histogram for the 30-50 ms period before the time of disruption, for all flattop disruptions in the EAST 2015 campaign. There is a very obvious enhancement in the fraction of points with high radiated fractions. The 0.35 threshold value would detect 24.9% of disruptions on EAST at least 30 ms before they occur. In contrast, Figure 16 shows the corresponding radiated fraction vs time until disrupt on C-Mod during the 20 ms period prior to the time of disruption. There does not seem to be an obvious change in behavior as the disruption time is approached. This is made clearer by comparing the histogram in Figure 17, for all non-disruptive shots, with the histogram in Figure 18, which is for the range of just 5-10 ms before the time of disrupt. It's difficult to discern a qualitative

difference, let alone a quantitative difference, implying that the radiated fraction parameter does not show much promise for predicting disruptions in C-Mod.




Fig. 16 – Radiated power fraction on C-Mod vs time until disrupt for all flattop disruptions in the 2015 campaign. No obvious change in behavior occurs as the disruption time is approached.




Fig. 17 – Histogram of radiated power fraction for all flattop times from all non-disruptive shots in C-Mod.




Fig. 18 – Histogram of radiated power fraction from 5-10 ms before the time of disruption on C-Mod.

The previous two comparisons give the impression that disruptions on Alcator C-Mod may be more difficult to predict with a useful warning time than on DIII-D or EAST. A further example of this is given in Figure 19, which shows the loop voltage as a function of time before disrupt, for all disruptions that occur during the discharge flattop in C-Mod in its 2015 campaign. The loop voltage is also thought to be a likely disruption warning indicator, but the dataset on C-Mod does not show any capability of providing a warning on a useful timescale. (The standard loop voltage signal on DIII-D, which is derived from EFIT, has an unacceptable amount of non-causal

filtering, but there are several alternative options that we are investigating.)




Fig. 19 – Loop voltage on C-Mod vs time until disrupt for all flat-top disruptions in the 2015 campaign. The loop voltage does not show much sign of an impending disruption until just 2-3 ms before disruptions occur.

To date, we have not yet identified any individual plasma parameters on C-Mod that provide useful disruption warning times. We note that we are continuing to populate additional parameters that could prove useful, such as locked mode signals. But our ‘control room’ experience on C-Mod is that most disruptions do not show any obvious precursors, and many appear to be triggered very quickly by small, high-Z injections from its molybdenum tiles. These small, high-Z injections are very efficient radiators compared to carbon and other low-Z impurities, and can therefore disrupt a plasma extremely rapidly. It may be that high-Z, high energy density tokamaks such as C-Mod, and presumably ITER, generally disrupt with much less warning time than low-Z, lower energy density tokamaks.

References

- [1] R. Jaspers, N.J. Lopes Cardozo, A.J.H. Donné, H.L.M. Widdershoven, K.H. Finken, *Rev. Sci. Instrum.* **72**, 466 (2001)
- [2] I.M Pankratov, *Plasma Phys. Rep.* **25**, 145 (1999)
- [3] J. Schwinger, *Phys. Rev.* **75**, 1912 (1949)
- [4] S.P. Gerhardt, D.S. Darrow, R.E. Bell, B.P. LeBlanc, J.E. Menard, D. Mueller, A.L. Roquemore, S.A. Sabbagh, and H. Yuh, *Nucl. Fusion* **53** (2013) 063021
- [5] G.A. Ratta', et al *NF* **50** (2010) 025005
- [6] B. Cannas, et al *NF* **44** (2004) 68-76

[7] B. Cannas, et al Fusion Eng Des **82** (2007) 1124-30

[8] C.G. Windsor, et al NF **45** (2005) 337-350

[9] L.L. Lao, *et al*, Nucl. Fusion, **25**, 1611 1985

JRT-16: Progress on Milestones for Disruption Prediction, Avoidance, and Mitigation (NSTX-U – S.A. Sabbagh and R. Raman – FY16 4th quarter report)

During the first three quarters, a disruption prediction algorithm was established, and a related computer code DECAF (Disruption Event Characterization and Forecasting code) [1] was created, and initial analysis of disruption event characterization was reported.

In the present quarter, attention was placed on creating and performing initial analyses using the first predictive models implemented in the code. These physics models are (i) a kinetic resistive wall mode (RWM) predictive stability model, and (ii) a tearing mode characterization model. Initial results from both of these models are summarized below. Technical improvements of the code are also summarized briefly below.

Progress on the installation of the newly designed ITER-type massive gas injection valve to study benefits of private flux region massive gas injection vs. mid-plane injection will be described in detail after this section.

Technical Improvements to the DECAF code

An important designed capability of the DECAF code is the ability to read data from several tokamak devices. Coding has been implemented to allow reading data from the DIII-D tokamak, which supplements the initial ability of the code to read data from the NSTX database, and newly-created NSTX-U plasmas. The first analysis results of the newly developed kinetic resistive wall mode (RWM) predictive stability model, and the tearing mode characterization model are shown below for NSTX, and NSTX-U plasmas. Analysis of DIII-D plasmas will occur as a next step, as the tearing mode characterization model was thought to be a pre-requisite for analysis of DIII-D. In addition, the DECAF code will now also be able to read data from the KSTAR tokamak with a simple change to the workbook input files. This specifically means the simple change of signal names in the EXCEL workbook that DECAF uses to define the signal tag names. In addition, the code now has the ability to read and process plasma profiles, which was a necessary capability for the kinetic RWM predictive model analysis.

Initial DECAF Analysis of a Set of Resistive Wall Mode Disruptions in NSTX

To put the following results on the initial kinetic RWM predictive stability model in context, and as the present report is the final report for the Joint Research Target this year, the present section provides a brief review of initial DECAF results that were performed on a 44 shot database of NSTX plasma that disrupted due to unstable RWMs. This database was subsequently used to test the new predictive model shown in the next section. An unstable, growing RWM can be detected by examining when an exponential growth in the $n = 1, 2, \text{ or } 3$ Fourier decomposition of signals from low frequency poloidal magnetic sensors located between the plasma and the vacuum vessel exceeds a pre-set threshold. In the next, the more sophisticated predictive kinetic model will be described.

The range of flat-top normalized beta, β_N , for the 44 discharges analyzed here was roughly 4-6.5, although it should be noted that β_N has been shown to not be a good predictor of disruptivity by itself [2]. Many of these discharges also had $n = 3$ magnetic braking applied, which slowed the toroidal plasma rotation [3]. Tearing modes were stable during these discharges.

The present version of the DECAF code, with eight event tests, was run on the 44 selected discharges to gain insight into common event chains that result during RWM disruptions in NSTX. The RWM event was detected in all of the discharges, as were plasma current not meeting request (IPR) and the disruption itself (DIS). Additionally, loss of wall proximity control (WPC) and low edge safety factor (LOQ) warnings also resulted in each discharge. The pressure peaking warning (PRP) occurred on a majority of the discharges analyzed (34 of 44), but typically occurred with or after the RWM, not before. Loss of vertical stability control (VSC) was present in most of the discharges as well (31 out of 44). Low density (LON) warnings occurred less often in this database (11 out of 44).

With the RWM $B_p^{n=1}$ lower sensor amplitude threshold of 30G ($\delta B/B_0 \sim 0.67\%$) used here the RWM warning was typically found near the disruption limit. In 59% of the cases, the RWM event occurred within $20 \tau_w$ of the time of disruption (DIS) (τ_w is the time scale of penetration of magnetic flux through the conducting structure, taken here to be 5 ms). Additionally, many of the earlier RWM warnings could not be considered false positives; they cause significant thermal collapses or “minor disruptions”, with subsequent recovery (as illustrated in Fig. 1).

One way of seeing which events are commonly associated is to examine a histogram of some of the timing of the events before the time of disruption (DIS), shown in Fig. 2. Here only the




FIG. 1: β_N and $n=1$ signal on lower poloidal magnetic sensors from NSTX discharge 135131, showing that RWM warnings sometimes indicate minor disruptions that cause decreases in β_N , with subsequent recovery.

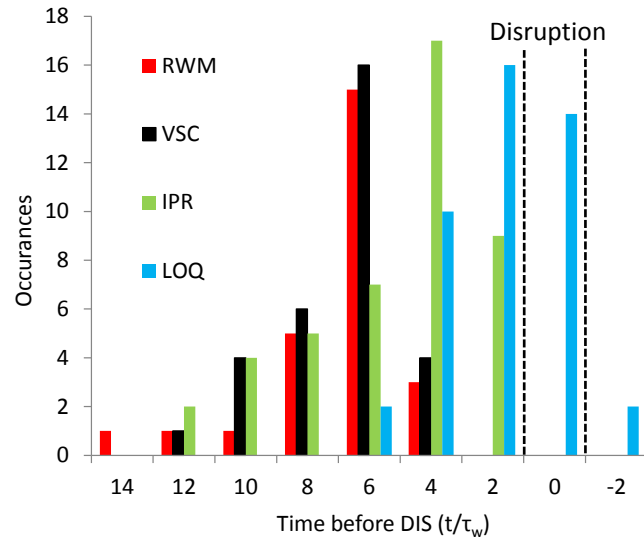


FIG. 2: Histogram of the timing of various disruption chain events in the 44 discharge NSTX database before the time of disruption, within 14 wall times.

events within $14 \tau_w$ of the disruption are shown. It is clear that LOQ and IPR events occur close to the time of disruption, and these are often preceded by VSC, and RWM events which peak around 30 ms prior to the disruption.

Examining the common chain of events can more closely provide insight into how to cue avoidance systems to return to normal plasma operations. The 44 RWMs were followed immediately by WPC and VSC (two events related to plasma motion) each 13 times, PRP 11 times, IPR 6 times, and LOQ once. The RWM event never proceeded to LON or DIS without another event happening first. Considering the two-event chains that happened directly after RWMs, we find that although there are 42 two-event combinations that could occur from the 7 currently tested for (in addition to RWM), six two-event chains accounted for 70% of the cases in this database (table 1).

Event chain	Percent
RWM \rightarrow VSC \rightarrow PRP	15.9%
RWM \rightarrow WPC \rightarrow PRP	13.6%
RWM \rightarrow IPR \rightarrow WPC	11.4%
RWM \rightarrow PRP \rightarrow IPR	11.4%
RWM \rightarrow WPC \rightarrow VSC	9.1%
RWM \rightarrow VSC \rightarrow WPC	9.1%

TABLE 1: The six most common two-event combinations that directly followed an RWM event in the 44 discharge NSTX database.

Reduced Kinetic Stability Model Implementation in DECAF

It was previously recognized that simplified model calculations based on physics insight from kinetic stability theory should be examined [2,4]. Now implemented in DECAF is a model based upon simplification of kinetic stability theory [5,6] with collisionality [7], implemented in the RWM dispersion relation: $\gamma\tau_w = \text{Re}[-(\delta W_\infty + \delta W_K)/(\delta W_b + \delta W_K)]$.

Models for ideal fluid stability δW terms have been previously developed for NSTX [4], and have now been implemented in DECAF and tested for NSTX-U discharges. For kinetic RWM stability, once the kinetic term δW_K is defined, the normalized growth rate $\gamma\tau_w$ can be calculated from the RWM dispersion relation. For the kinetic δW_K term, full calculations with codes such as MISK [8] cannot be performed in real time. Kinetic RWM stability theory has been developed to greater complexity in recent years, but here we wish to simplify the kinetic theory model to facilitate eventual real-time calculation.

For the present purposes, a functional form for δW_K was constructed that is easily and quickly calculable and that relies on a few important, measurable parameters. Any such model must capture the essential physics learned from the successful application of kinetic theory to experimental results in recent years. Namely, resonance between ExB frequency, ω_E , and precession drift frequency, ω_D , of trapped thermal ions at lower plasma rotation, and with bounce frequency, ω_b , at higher plasma rotation provides a stabilizing component to δW_K , but in between these the kinetic effects are weaker, allowing for instability [9]. Increased collisionality tends to damp the rotational resonance stabilization (see Fig. 3 of Ref. [7]) and shift it to slightly

lower rotation (see Fig. 6 of Ref. [9]). The imaginary terms of δW_K tend to peak at lower plasma rotation than the real parts (see Fig. 8 of Ref. [10]) so that plasmas move in kinetic stability space as rotation changes in looping paths (see Fig. 5 of Ref. [9]).

To that end, Gaussian functions were used to represent kinetic resonances (Fig. 3). The positions of the peaks in $\langle \omega_E \rangle$ are determined by typical experimental ranges of ω_D and ω_b and the height, width, and position all dependent on $\langle v \rangle$. In the following $\langle \omega_E \rangle$ and $\langle v \rangle$ represent average values for ExB frequency and collisionality as described in Ref. [2]. The bounce resonance contribution was allowed to continue to increase at high $\langle \omega_E \rangle$ to capture the many bounce harmonics and circulating particle contributions. Coefficients for the functions of this initial model were selected based on NSTX model validation studies. This model can now be tested against other tokamak data to determine if the coefficients need to be recomputed, or if the present values are sufficient to determine data from another device.

Recently, the DECAF code has been expanded to include the necessary measured and derived profiles for the reduced kinetic stability model. This is a significant update to DECAF, as previously it had only dealt with scalar quantities as a function of time. Profiles now included are electron temperature and density profiles from Thomson scattering measurements, ion temperature and density, and carbon ion rotation (ω_ϕ^C) profiles from charge exchange recombination spectroscopy. From these we derive the ion collision frequency profile and the carbon ion diamagnetic frequency, ω^{*C} . The ExB frequency is obtained via a carbon ion radial force balance from $\omega_E = \omega_\phi^C - \omega^{*C}$.

Once the form of δW_K versus these quantities is established in the model, it is left to implement the model by following the evolution of a plasma discharge in time through the space of these

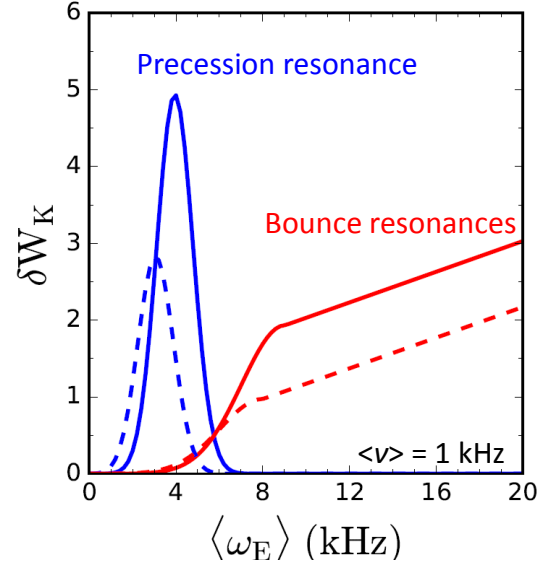


FIG. 3: Modeled real (solid) and imaginary (dashed) δW_K terms for precession (blue) and bounce (red) resonances for NSTX at $\langle v \rangle = 1$ kHz.




FIG. 4: Trajectory of NSTX discharge 139514 through $\text{Re}(\delta W_K)$ vs $\text{Im}(\delta W_K)$ space. The colored circles indicate the unstable region with increasing C_β .

quantities. The procedure is simply laid out as such: 1) Internal inductance, pressure peaking, and aspect ratio are used in the ideal beta limit model to calculate $\beta_{N,\text{no-wall}}$ and $\beta_{N,\text{with-wall}}$ [4]. 2) The ideal β limits and the measured β_N give $C_\beta \equiv (\beta_N - \beta_{N,\text{no-wall}}) / (\beta_{N,\text{with-wall}} - \beta_{N,\text{no-wall}})$. 3) Expressions for the fluid δW terms as functions of C_β that mimic DCON results give δW_b and δW_∞ . 4) The ideal δW terms give the fluid growth rate, $\gamma_f \tau_w$, and also set the unstable region in a $\text{Im}(\delta W_K)$ vs. $\text{Re}(\delta W_K)$ stability diagram. 5) Calculated $\langle \omega_E \rangle$ and $\langle v \rangle$ are used in the reduced kinetic model to calculate δW_K . 6) Finally, δW_K is used in the kinetic RWM dispersion relation to find $\gamma \tau_w$.

Through the changing levels of total (precession plus bounce terms) real and imaginary δW_K , one can plot the trajectory of the plasma in $\text{Re}(\delta W_K)$ vs. $\text{Im}(\delta W_K)$ space. This is shown in Fig. 4 along with the unstable regions for various levels of C_β (0.2-1.0). The circular lines indicating the unstable boundary correspond to $\gamma \tau_w = 0$. Inside these circles $\gamma \tau_w$ is positive, and therefore the RWM is unstable. As the plasma moves in time in the δW_K space, at the same time the size of the unstable region is changing as well. Within the plasma trajectory shown, colored circular markers indicate the times that the plasma crosses the corresponding C_β level. So, for example, in this case at the time in the discharge when $C_\beta = 0.2$ (cyan) the plasma is just outside the unstable region while by the time of $C_\beta = 0.4$ (green) and $C_\beta = 0.6$ (magenta) δW_K has decreased due to the changing $\langle \omega_E \rangle$ and $\langle v \rangle$ and additionally the unstable region has increased in size due to the fluid terms at the larger C_β . The combined effect is that the plasma is now inside the unstable region.

Alternatively, one can show a stability diagram in the $\langle \omega_E \rangle$, $\langle v \rangle$ space at a given level of C_β by plotting contours of $\gamma \tau_w$ (similar to Fig. 6 in Ref. [9]). Here we show the trajectory of the same plasma in this space as time increases, $\langle \omega_E \rangle$ increases, and $\langle v \rangle$ decreases (Fig. 5). Similarly to Fig. 4, in this diagram the unstable region changes




FIG. 5: Trajectory of NSTX discharge 139514 through $\langle \omega_E \rangle$ vs. $\langle v \rangle$ space. The colored contours represent the unstable region for various levels of C_β .




FIG. 6: Calculated ideal (blue) and kinetic (red) normalized growth rates for NSTX shot 139514.

with time as C_β changes.

Finally it is natural to simply plot the kinetic RWM growth rate vs. time. In Fig. 6 we do this for the same discharge for both the fluid and kinetic growth rate, where it is easy to see the transition into the unstable range at a time of around 0.75s.

Application of the kinetic RWM Model to the NSTX Database

It is useful now to apply the reduced model to a database of NSTX RWM discharges, which also include stable plasmas. For a large number of discharges we will presently show their trajectories




FIG. 7: Stability diagram (left) and forecast growth rate (right) for unstable (colored) and stable (black) NSTX discharges.

on a stability map as in Fig. 5. Also it is natural to simply plot the forecasted RWM growth rate as a function of time. Here we plot $\gamma\tau_w$ vs. time before DIS, the time of disruption (as determined by tests within DECAF). For discharges without an RWM induced disruption, the time DIS effectively indicates the natural end time of the discharge by other means.

These plots are shown in Fig. 7 for 20 discharges with unstable RWMs in NSTX (color) and 8 without (black). Unstable RWMs were determined to have occurred in these discharges by both independent assessment of relevant signals as well as a threshold test on a poloidal magnetic signal within DECAF. The colors indicate the warning time before disruption when the model indicates the RWM should be unstable ($\gamma\tau_w$ crosses zero). Red is for a single case of < 0.1 s warning, orange five cases with 0.1-0.2s, green eleven cases with 0.2-0.3s and blue three cases with 0.3-0.32s warning.

One can see a clear difference in the evolution in $\langle\omega_E\rangle$ vs. $\langle v\rangle$ space between the stable and unstable discharges. While all the discharges drop in collisionality with time during the shot, due to increasing temperature, in the unstable cases a turn towards higher $\langle\omega_E\rangle$ leads into the unstable region. This is avoided in all the stable cases shown here (in fact, some drop towards zero $\langle\omega_E\rangle$ leading one case to just barely touch $\gamma\tau_w = 0$).

In addition to the cases shown in Fig. 7, many others were analyzed. In fifteen additional RWM unstable cases, the model also showed $\gamma\tau_w$ crossing zero into the unstable region, but in these cases this occurred well before the disruption and in fact were all correlated with *minor* disruptions that occurred earlier in those shots. Here a minor disruption is defined as a 10% drop in both β_N and stored energy within 0.1s, that subsequently recovers. In each of the fifteen cases considered, $\gamma\tau_w$ crossed zero within 0.1s of a minor disruption. There were, however, other minor disruptions in the database that did not correlate with the reduced kinetic model warning; whether these are due to other causes will be further explored.

In any case, there were 35 discharges in the database where the RWM became unstable leading to a disruption in which the reduced kinetic model predicted instability within 0.32s of the disruption or 0.1s of an earlier minor disruption. Additionally in three experimentally RWM unstable cases, the model gave a warning 0.4s in advance without any related minor disruption, which is considered a false positive because it is so early. Finally, this initial model sometimes misses unstable RWMs. There was one case in which $\gamma\tau_w$ barely didn't cross zero, three cases with very low C_β disruptions that the model missed, and three cases where $\langle\omega_E\rangle$ was in what the model considered to be a stable range, yet an unstable RWM occurred. Altogether the model failed to predict an unstable RWM at all in 7 out of 45 experimentally unstable cases, or 15.6%. The success rate of this first model is surprisingly high given its relative simplicity. Further research will aim to improve the success rate. A summary of the statistics are shown in Fig. 8.

In addition to the eight successful predictions of stability for the experimentally *stable* discharges shown above, five more stable discharges were tested. In three of these cases the discharge evolution in $\langle\omega_E\rangle$ vs. $\langle v\rangle$ space was very similar to the *unstable* cases shown in Fig. 7, but nevertheless the discharge remained




Fig. 8. DECAF predicted kinetic RWM statistics for an NSTX database. A total of 84% of unstable shots are predicted unstable, surprisingly high for this initial model. Seven percent of the cases were determined to be false positives.

stable. It is possible that in these cases some other stabilizing effect not captured by the reduced model was present, but this remains to be determined. In two other experimentally stable cases, the RWM warning was triggered by the reduced model because $\langle\omega_E\rangle$ went to zero (hints of this behavior also appear in some of the shots in Fig. 7). The unstable region at $\langle\omega_E\rangle \sim 0$ is present in the model due to theory expectation, but has not yet usefully captured an unstable RWM in our NSTX analysis. This region could be eliminated in the model since we are interested in improving the model's usefulness whether or not it agrees perfectly with theory, but this requires further investigation. If those cases were eliminated then 10 out of 13, or 77%, of stable high β , long-pulse NSTX discharges analyzed were predicted stable in the reduced model.

Initial Automated Tearing Mode Characterization Model in DECAF

Rotating MHD modes, often saturated tearing instabilities, were quite often suppressed in NSTX at high plasma elongation, or through lithium wall conditioning. However, such activity is ubiquitous in tokamaks in general and can lead to, or be part of a disruption event chain. Therefore, analysis to determine the appearance of these instabilities and to automatically characterize them is critical to implement in DECAF before analysis of a wider tokamak dataset is possible. Attention was placed to create automated code to characterize these modes in DECAF. Examples of automated analysis from the new coding are given below. Here, two plasmas (one from NSTX, and one from NSTX-U – shown in Fig. 9) are considered to illustrate the code capability.

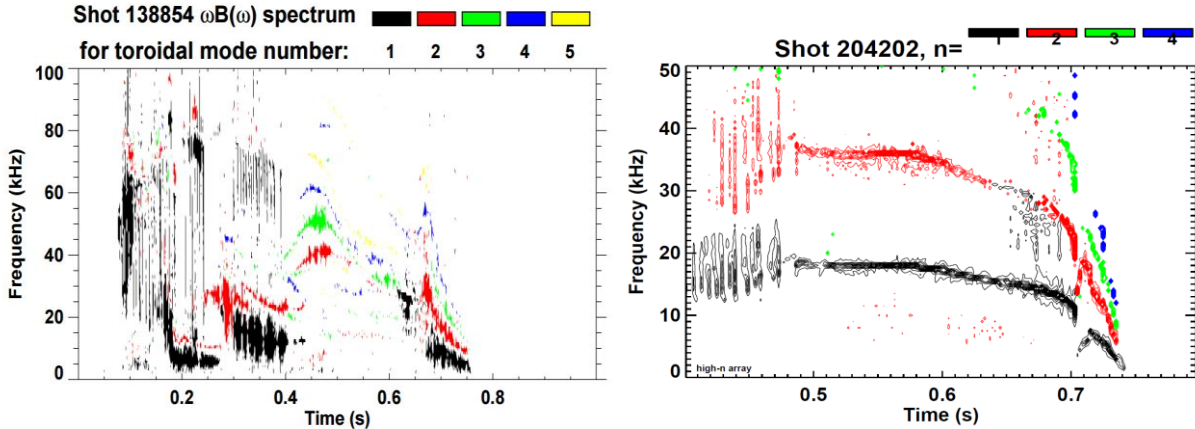


Fig. 9. Magnetic spectrograms illustrating rotating MHD modes in NSTX (left) and NSTX-U (right)

The present algorithm written for DECAF was created to automatically determine and define events related to rotating MHD that lead to disruption. The initial model targets rotating MHD that sets up an initial rotation frequency when born, and/or sets up a quasi-steady state frequency (at which time the plasma is typically non-disruptive) and subsequently slows (drops in frequency) and locks. Some period after the locking, the plasma very often disrupts.

The present version of the code to determine these events utilizes FFTs on magnetic probes to determine the existence of a mode. Presently the code can distinguish between odd and even modes, but will be expanded as a next step to fully distinguish the toroidal mode number of the

rotating MHD. The mode frequency is then tracked in time and smoothed by fitting a cubic polynomial to the time-advancing data points. This smooth functional form is then easily differentiated analytically, and both the mode frequency f , and df/dt are then used to determine the initial, or quasi-state frequency, the bifurcation time of the mode before it locks, and the locking time. These characteristics are shown in Fig. 10 for the plasmas shown in Fig. 9 near the time of mode locking in each case. The code presently has high success for determining these characteristics, and will next be tested on a wider data set, and on data from other tokamaks.




Fig. 10. Rotating MHD characteristics automatically determined by new algorithm in DECAF created to identify and predict events that typically lead to disruptions due to these modes.

Along with the definition of the events themselves, the code also identifies a status for the mode itself – if it is present at a given time, or not, and if the mode is locked. The code can also tell if the mode is beyond the bifurcation point in the locking evolution. The DECAF mode status for odd-n MHD is shown for NSTX-U shot 20402 (Figs. 9 and 10) in Fig. 11.

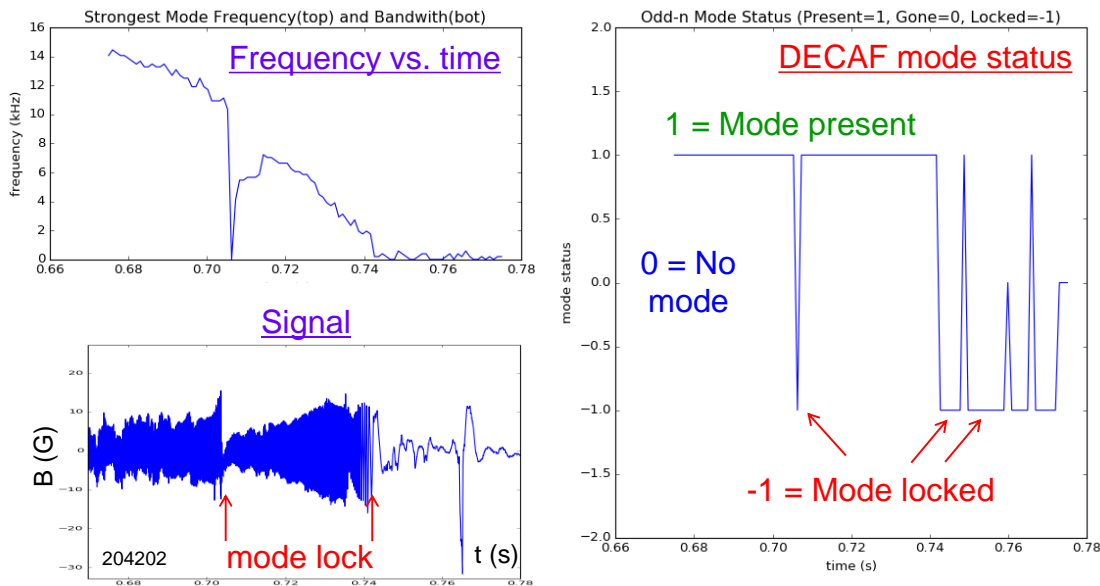


Fig. 11. DECAF mode status for $n = 1$ rotating MHD mode in NSTX-U plasma that locks.

As well as determining the proximity in time to a mode lock, and the time of the locking event itself, it is also important for the code to know that a rotating MHD mode exists, but does not lock. This determination is also a valid outcome. Such an outcome is illustrated in Fig. 12 which shows the result of the time evolution of the calculated mode status for the even- n activity as determined by the magnetic pickup loops in NSTX-U. In this automated analysis, the mode is found to be present, but is not found to be locked.




Fig. 12. DECAF mode status for $n = 1$ rotating MHD mode in NSTX-U plasma that does not lock.

-
- [1] S.A. Sabbagh, J.W. Berkery, J.M. Bialek, et al. “Global MHD Mode Stabilization for Tokamak Disruption Avoidance: Linear/Non-linear Considerations”, 20th Workshop on MHD Stability Control, Princeton, NJ, November 2015.
 - [2] J.W. Berkery, et al., Physics of Plasmas **21** (2014) 056112.
 - [3] W. Zhu, et al., Physical Review Letters **96** (2006) 225002.
 - [4] J.W. Berkery, et al., Nuclear Fusion **55** (2015) 123007.
 - [5] Y.Q. Liu, et al., Physics of Plasmas **15** (2008) 112503.
 - [6] J.W. Berkery, et al., Physics of Plasmas **18** (2011) 072501.
 - [7] J.W. Berkery, et al., Physical Review Letters **106** (2011) 075004.
 - [8] B. Hu, et al., Physics of Plasmas **12** (2005) 057301.
 - [9] J.W. Berkery, et al., Physical Review Letters **104** (2010) 035003.
 - [10] J.W. Berkery, et al., Physics of Plasmas **17** (2010) 082504.

MGI on NSTX-U Contributions JRT-16

Predicting and controlling disruptions is an important and urgent issue for ITER. In support of this activity, NSTX-U has employed three Massive Gas Injection (MGI) valves that are similar to the double flyer plate design being developed for ITER [1]. NSTX-U experiments will offer new insight to the MGI database by studying gas assimilation efficiencies for MGI gas injection from different poloidal locations, with emphasis on injection into the private flux region. Results from the operation of the valve, including tests conducted in 1 T external magnetic fields, are described in this report. The pressure rise in the test chamber is measured directly using a fast time response baratron gauge. At a plenum pressure of just 1.38 MPa (~200 psig), the valve injects 27 Pa.m³ (~200 Torr.L) of nitrogen with a pressure rise time of 3 ms. In support of next-term experiment on NSTX-U, which is to conduct a comparison of the mid-plane to private flux region injection, the valves located in the divertor region and at the mid-plane location have been fully commissioned on NSTX-U, and are ready to support plasma operations.




FIG. 1. MGI valve installation locations on NSTX-U. At present three valves are installed (shown by locations 1, 2 and 3). The fourth valve at location 4 is planned for a future installation. These locations are: (1a) private flux region injection, (1b) lower scrape-off-layer (SOL) and lower divertor injection, (2) conventional mid-plane injection, and (3) upper divertor injection. Two different plasma shapes are shown, that allow the lower divertor gas to be injected either into the private flux region or into the SOL located in the high-field side region.

The valve has similarities in design to several valves built for a Compact Toroid (CT) injector [2], but draws on design features used in the TEXTOR valve [3, 4], and is motivated by the work of Lehnen [5]. The valve operation is as follows. A pancake disk solenoid is rapidly charged, inducing eddy currents in a conducting plate (flyer plate); the resulting $\mathbf{J} \times \mathbf{B}$

Introduction

Massive Gas Injection (MGI) is the most developed disruption mitigation system to-date and will be implemented as a secondary disruption mitigation system on ITER to protect internal ITER components during unplanned tokamak disruptions. NSTX-U research will offer new insight by studying gas assimilation efficiencies for MGI injection from different poloidal locations using identical gas injection systems. At present three valves have been installed on NSTX-U corresponding to locations 1, 2 and 3 in Fig. 1. An unique aspect of NSTX-U MGI experiments is that the poloidal injection comparisons will be made using near-identical systems: identical valves, and nearly-identical piping configuration between the valve and the vacuum vessel will be used.

The ITER-type NSTX-U MGI valve

In support of NSTX-U MGI experiments, an electromagnetic MGI valve was designed, built, and tested. Fig. 2 is an internal view of the NSTX-U MGI valve. The valve operating principle is similar to that being considered for the for the ITER MGI

interaction pushes the plate away from the solenoid. The plate is connected to a plunger, such that as it moved, the valve is opened.




FIG. 2. Internal view of the NSTX-U double flyer plate MGI valve. Gas from the primary plenum is injected into the plasma discharge.



FIG. 3. Bench-test setup used for MGI valve testing in an external magnetic field. The valve is located between two external field coils.

coils are used to generate magnetic fields up to 1 T in the region of the MGI valve coils. As shown in Fig. 4, as the field increases to 1 T, there is a 10% reduction in the amount of injected gas. Fig. 4 also shows that the fast baratron begins to measure a pressure increase in the test chamber about 3 ms after the current pulse through the MGI valve coils. The baratron is located about 1 m away from the MGI valve. The sound speed in nitrogen is about 354

The initial valve concept used a single pancake coil, and a single conducting disk (flyer plate) [6]. Motivated by the work of Baylor, [1] for an ITER test valve, we incorporated a second pancake coil and a second conducting disk for valve operation [7]. The current in the second pancake coil is opposite in direction to that in the first coil. The effect of these oppositely driven currents is to nearly cancel the $J \times B$ torque the valve would experience in a magnetic field. All valves deployed on NSTX-U are of this dual pancake coil design.

An important observation was that, compared to the single flyer plate design, adding a second pancake coil did not substantially increase the size of the capacitor bank power supply or the operating voltage for injecting similar amounts of gas, and demonstrated similar gas pressure rise times in the test chamber.

The valve has been operated in two different configurations. In the first configuration, the coils are connected in series so that the same current passes through each solenoid. In the second configuration, both coils are connected in parallel to the power supply. This configuration reduces the total system inductance, reduces the coil current pulse width, but increases the peak power supply current by about a factor of two. Both these operating conditions inject similar amounts of gas with similar gas pressure rise times in the test chamber ($27 \text{ Pa}\cdot\text{m}^3$ of nitrogen with a 3 ms gas pressure rise time). This is because the longer current pulse duration for the series configuration results in the magnetic forces acting on the piston for a longer period.

The valve operation is not affected by fields $< 0.8 \text{ T}$. Fig. 3 shows the experimental setup used for tests during the presence of an external magnetic field. Two external field

m/s, which is consistent with the 3 ms response time for detecting the pressure increase in the chamber.

The valves have operated successfully on NSTX-U and are ready to support operations. Fig. 5 shows the amount of injected neon as function of operating voltage for valves located at location 1 (lower divertor) and 2 (mid-plane) on NSTX-U. The valves on NSTX-U inject over 400 Torr.L of neon at an operating voltage of less than 800 V, and a plenum fill pressure of 200 Psig. The power supply for operating these valves is rated for 1.5 kV. The power supplies connected to a dummy load have been tested at up to 1 kV on NSTX-U and up to 1.2 kV in off-line tests, when connected to an MGI valve.

Summary

An electromagnetic valve to support NSTX-U MGI experiments has been built and tested. Although neon or argon would be the impurity gas used in NSTX-U disruption mitigation




FIG. 5. Vessel pressure increase as a function of valve operating voltage for the lower divertor and mid-plane MGI valves on NSTX-U.




FIG. 4: Experimental traces from the operation of the valve with and without the presence of an external magnetic field. Shown are the current pulse duration through the solenoid (the external field generating coils) and the current pulse through the gas valve pancake coils. The gas valve is discharged 10 ms after the external field coil discharge is initiated. Shown are two gas pressure traces for cases in which there was no external magnetic field. Shown also are three gas pressure traces during the presence of a 1 T external magnetic field in a configuration in which the field is parallel to the pancake coil surface.

experiments, the valve has been tested using nitrogen gas. The valve has been calibrated for injecting $27 \text{ Pa}\cdot\text{m}^3$ of nitrogen, similar to the levels planned for NSTX-U experiments. The gas pressure rise time is about 3 ms, consistent with the gas sound speed for nitrogen gas. The valve for NSTX-U uses a double pancake coil configuration. This has the benefit of nearly cancelling the $\mathbf{J} \times \mathbf{B}$ torque that acts on the valve when it is operated in an ambient magnetic field. The valve has been tested in ambient magnetic fields up to 1 T, and found to operate well at these field levels. Two valves are now fully operational on NSTX-U, are ready to support plasma operations, and have been used to

inject over 400 Torr.L of neon into the NSTX-U vessel.

References

- [1] L.R. Baylor, C.C. Barboer, J.R. Carmichael, et al., Disruption mitigation system developments and design for ITER, *Fusion Science and Technology* **68** (2015) 211, <http://dx.doi.org/10.13182/FST14-926>
- [2] J.C. Thomas, D.Q. Hwang, R.D. Horton, J.H. Rogers, and R. Raman, *Review of Scientific Instruments* **64** 1410 (1983)
- [3] S.A. Bozhnikov, K-H. Finken, M. Lehnen, and R.C. Wolf, *Review of Scientific Instruments* **78** 033503 (2007)
- [4] K.H. Finken, M. Lehnen and S.A. Bozhnikov, *Nuclear Fusion* **48** 115001 (2008)
- [5] U. Kruezi et al., *Fusion Engineering and Design* **96-97** (2015) 286-289
- [6] R. Raman, T.R. Jarboe, B.A. Nelson, et al., Design and operation of a fast electromagnetic inductive massive gas injection valve for NSTX-U, *Review of Scientific Instruments* **85**, 11E801 (2014)
- [7] R. Raman, G. Plunkett, W-S. Lay, Massive Gas Injection Valve Development for NSTX-U, *IEEE Transactions on Plasma Science* **44**, No. 9, 1547 (2016). DOI: 10.1109/TPS.2016.2565658