Plasma Phys. Control. Fusion 55 (2013) 124033 (6pp) doi:10.1088/0741-3335/55/12/124033

Advanced control of neoclassical tearing modes in DIII-D with real-time steering of the electron cyclotron current drive

A S Welander[1], E Kolemen[2], R J La Haye[1], N W Eidietis[1], D A Humphreys[1], J Lohr[1], S Noraky[1], B G Penaflor[1], R Prater[1] and F Turco[3]

[1] General Atomics, PO Box 85608, San Diego, CA 92186-5608, USA
[2] Princeton Plasma Physics Laboratory, PO Box 451, Princeton, NJ 08543-0451, USA
[3] Department of Applied Physics and Applied Mathematics, Columbia University, 116th St and Broadway, New York, NY 10027, USA

E-mail: welander@fusion.gat.com

Received 26 July 2013, in final form 15 October 2013
Published 28 November 2013
Online at stacks.iop.org/PPCF/55/124033

Abstract
The system for controlling neoclassical tearing modes (NTMs) in DIII-D now catches the NTM the moment it becomes unstable by turning on the stabilizing electron cyclotron current drive (ECCD) and promptly bringing it back to stable before it has grown to a large size. Between NTMs, the ECCD can be turned off to save power, which will improve the fusion gain, Q, when used in ITER. This technique, named 'catch and subdue' (C&S), has been made possible by several advancements over the years at DIII-D. Firstly, ECCD must be very accurately aligned to the NTM; this is achieved by algorithms that probe how the NTM responds to changes in the alignment. Secondly, the alignment must be maintained even when the NTM is gone so that the ECCD will immediately stabilize when turned on in response to a new NTM. This is made possible by real-time equilibrium reconstructions that include measurements of the motional Stark effect and by a refraction estimator. Thirdly, real-time steerable mirrors are now fast and accurate actuators for the alignment adjustments. Fourthly, early NTM detection is made possible by a real-time mode analysis that filters noise to minimize false positives. These various control elements will be described and followed by a discussion of the further development needed for NTM control on ITER.

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

1. Introduction

Tearing mode (TM) control is expected to be essential for ITER in order to maximize confinement and prevent potentially disruptive island growth. The ITER design has been equipped with an electron cyclotron current drive (ECCD) system shown to be nominally capable of the stabilization of saturated islands of both 3/2 and 2/1 TMs [1]. However, requirements for producing robust stabilization and minimizing average power use in the application of ECCD both remain to be established in operating experiments. In addition to this ITER operational motivation, present-day devices can benefit from effective control systems for TM suppression in the same ways as ITER. Effective control can also provide a powerful tool for physics programs to better understand TM physics. We describe the control capabilities and experiments in DIII-D investigating advanced control approaches for TMs, relevant to the ITER TM control system. In particular, the DIII-D system includes real-time steerable mirrors to align ECCD with the TM island, or with the rational surface in the absence of an island. DIII-D launcher positions are similar to those of the upper launchers in ITER. Section 2 contains a brief description of the expected prevalence of TMs in ITER and their implications. The TM control at DIII-D is described in section 3. Finally, section 4 outlines future directions to meet the needs for TM control in ITER. Conclusions are drawn in section 5.

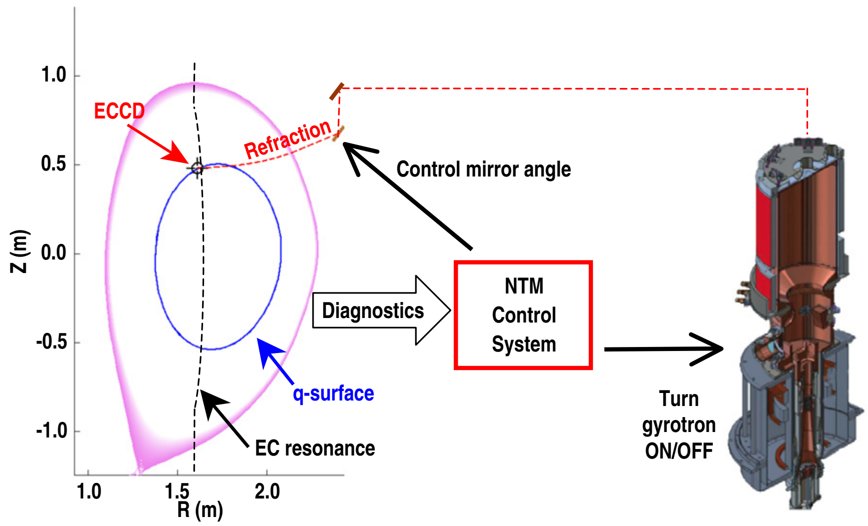

Figure 1. Diagnostics are sent into the NTM control system where they are analyzed to detect NTMs and find the positions of the ECCD and q-surface. The beams from the gyrotrons are reflected by focusing mirrors and controlled steerable mirrors into the plasma where they are refracted by the plasma and absorbed at the EC resonance line to produce ECCD. The mirrors are typically controlled even when the gyrotrons are off so that the ECCD will be immediately aligned when the gyrotrons are turned on.

2. The need for TM control in ITER

The TM is a magnetic island caused by a perturbation to the plasma current of the same helicity as the field line on a flux surface where $q = m/n$ is a rational number, m being the poloidal mode number and n the toroidal. The neoclassical tearing mode (NTM) is a mode that is classically stable but can be destabilized by a helical perturbation of the bootstrap current. Saturated NTM islands with $m/n = 3/2$ degrade the energy confinement by typically 10%–30%, while modes with $m/n = 2/1$ lead to severe energy loss and frequently to disruption [2]. Evidence from experiments and numerical simulations show that, without control, ITER will have unstable 2/1 islands. The results suggest that the growth of a 2/1 island will produce a loss of H-mode, then lock to the wall and lead to a disruption. Hence, a system for suppression of TMs is essential for ITER.

TMs can be stabilized by driving current at the resonant surface. This increases linear stability and replaces the missing bootstrap current. A figure of merit for control of TMs is how much current density can be driven at the resonant surface. ECCD by gyrotrons is a suitable choice since the current deposition is narrow and can therefore be focused on the region around the resonant surface where it is stabilizing [3]. Since the ECCD is narrow, it must be aligned very accurately with the resonant surface.

3. Control of TMs in DIII-D

3.1. Hardware

The experiments described here had five gyrotrons capable of delivering 3 MW to the plasma. The mirrors that reflect the beams into the plasma can be steered in real time to adjust the polar angles of the beams. This moves the ECCD mainly in the vertical direction, since the major radius of the EC resonance depends only weakly on vertical position. A schematic is shown in figure 1.

3.2. Real-time analysis

The NTM control system uses a real-time version of the MHD equilibrium reconstruction code EFIT [4] to track the location of NTM-resonant surfaces (often referred to as q-surfaces). The equilibrium reconstructions are based on external measurements made by flux loops, low frequency Mirnov probes and a Rogowski coil; they are also based on internal measurements of the motional Stark effect (MSE). The MSE diagnostic gives the field line pitch at points inside the plasma and helps constrain the equilibrium reconstructions so that the resulting q-profile is sufficiently accurate for tracking changes to the q-surface. A high frequency toroidally distributed Mirnov probe array is used to infer the sizes of magnetic islands.

The gyrotron beam refraction estimator is based on shot reproducibility and uses the correlation between interferometer density and refraction in earlier discharges. A linearization is made around a nominal, most typical, density profile. The accuracy is typically sufficient ($<$3 mm error averaged over 50 ms) for plasmas with small density variations [5]. In 2013, this simple estimator will be replaced by the real-time ray tracing code TORBEAM, which will use detailed density profiles from Thomson scattering to compute ECCD positions [6].

Further details on NTM control system hardware can be found in [7].

3.3. Alignment techniques

Since good alignment of ECCD to the rational q-surface is crucial for NTM suppression, an early focus of the NTM control program at DIII-D has been to develop means of achieving the required high accuracy [8]. It is useful to distinguish between the analysis that only requires information about the 2D equilibrium and the analysis that uses measurements related to the NTM under control (when it is present).

The algorithm that uses 2D equilibrium information is known as 'active tracking' (AT). With this algorithm, the q-surface is tracked by the real-time EFIT analysis and the refraction is tracked by the refraction estimator. These calculated values for the q-surface and ECCD positions are typically not sufficiently accurate by themselves to be used for NTM stabilization; the primary purpose of the AT algorithm is therefore only to maintain correct alignment once it has been established by a fine-tuning algorithm.

The working hypothesis is that the alignment error that would result if AT was used alone consists of a high frequency random error and a systematic error, which remains the same or possibly drifts slowly. The high frequency noise is removed by a suitable filter. The systematic error (which can be caused by systematic errors in inferred positions of both the q-surface and the ECCD) is, as mentioned, corrected by a fine-tuning algorithm. In an automatic system, fine-tuning will only be used whenever the mode grows because of poor alignment. Several algorithms for fine-tuning alignment have been developed. These all require that an NTM be present in the plasma.

The first fine-tuning method that was developed is known as 'search and suppress' [9]. With this method, the ECCD is turned on to suppress the mode. If it turns out that the mode is not suppressed at a sufficient rate then an adjustment of the alignment is made (typically ~1 cm) after a dwell time (typically 50 ms). If the initial step was in the wrong direction, the algorithm will discover this after a few steps and go back in the other direction. In our implementation, the step sizes and dwell times are selected manually before the shot; however, in a fully automated controller, these parameters would be a function of a growth rate analysis based on the modified Rutherford equation (MRE).

The second fine-tuning algorithm is known as 'target lock' (TL) [5]. This turns on the ECCD and sweeps the alignment back and forth at a fairly high speed to see where the ECCD has the maximum effect on the mode amplitude and then returns to that point, thus locking on the target. In our implementation, the length and duration of the sweep are selected manually, but in a fully automated controller, this would be a function of growth rate analysis based on the MRE. Figure 2 is one of several tests of the TL algorithm (discharge 149713). At 3.2 s into the discharge, the controller decides that an alignment correction is necessary since the ECCD has been turned on without suppressing the mode. The controller begins by moving the mirrors into the starting position for a sweep. The sweep begins at 3.25 s and ends at 3.65 s. The controller then returns to the position that was passed 50 ms before the minimum in the mode amplitude (the dip at 3.6 s). At 3.7 s, one of the gyrotrons was lost but with the improved alignment, the mode is suppressed despite the lower ECCD.

Prior to these tests, the correct alignment of the ECCD was established by multiple sweeps of ECCD back and forth over a saturated island. The ECCD was then deliberately initially misaligned to check if the TL algorithm would find the right

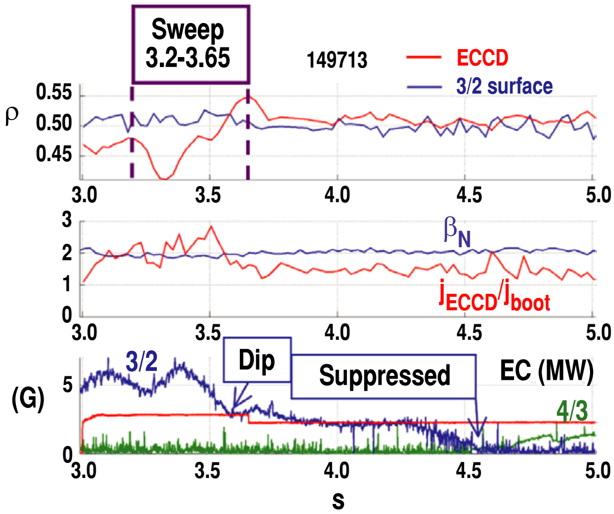

Figure 2. Alignment correction with the TL algorithm: (*a*) rho is the normalized minor radius with both ECCD location (red) and $q = 3/2$ (blue) shown; (*b*) normalized beta (blue) was controlled to a value of 2 and the ratio $j_{\mathrm{ECCD}}/j_{\mathrm{boot}}$ (red) was around 2 until 3.7 s and then dropped to 1.5 with the loss of a gyrotron; and (*c*) EC-injected power (red), 3/2 mode amplitude (blue) and 4/3 amplitude (green).

correction. The result was an improved alignment in every case.

The experiments reported here are the first to use mirror steering as an actuator for alignment control. For this reason, a relatively simple version of the TL algorithm was employed to avoid the risk of sending commands that could not be executed. Plans are under way to refine both the interpretation of the data gathered during a sweep and the logic for the sweep itself. In order to obtain a complete data set for suppression rate versus alignment correction, the TL sweep should extend from one side of the optimum alignment to the other. When the sweep begins, the initial direction will be toward the optimum (if any information about that direction is available from other diagnostics). During the sweep, the MRE is used to calculate what the evolution of the NTM would be based on the available information about the parameters that affect the growth rate of the NTM. The actual NTM evolution will differ from the calculation, because of inaccuracies in these parameters and in the MRE model itself. The most significant uncertainties are in the value of Δ' and in the optimum alignment (which is what will be determined). For this reason, several MRE calculations are done in parallel with a range of different corrections to these two parameters. The different calculations are then compared to the actual NTM evolution; the corrections that make the calculations best match the measured data are used to infer a corrected Δ' and the optimum alignment. Although there are other uncertainties, they can, to a fair approximation, be lumped into the corrections of Δ' and the optimum alignment. As the sweep progresses and data is gathered at more alignments, the certainty regarding the optimum alignment improves, i.e. the range of possible values (or error bar) decreases. If the initial sweep was in the wrong direction, the alignment will eventually be outside the range for where the optimum could possibly be and a decision will be taken to switch and go back in the other direction. When the sweep goes through the optimum, the speed must be high enough to get to the other side before the mode is fully suppressed and slow enough to have a visible effect on the mode. A suitable speed is selected based on the predicted growth rate (which becomes more accurate as more data is collected). Once the alignment is moving away from the range of possible values for the optimum on the other side, the direction is again switched. At this point, the controller may decide to simply apply the inferred correction or continue gathering data with successively smaller sweeps around the optimum to home in on it with better accuracy until the mode is suppressed. There is a trade-off to be made between speed and accuracy with this method. At one extreme, the sweep could end as soon as the optimum alignment is known with sufficient certainty (without sweeping through to the other side). The NTM will then be suppressed with adequately good alignment. At another extreme, the alignment could be swept back and forth across the optimum in a way that never allowed the ECCD to suppress the NTM but that gathered a large statistical data set on where the optimum alignment is. The experience from more than a decade of NTM experiments on DIII-D is that once an alignment correction has been found, it remains applicable for as long as the MSE calibration remains intact, which is on the order 1 year. It is therefore a good investment to spend some extra time on the sweep and achieve good accuracy. A sweep of about 1 s can deliver an alignment correction with uncertainty below 5 mm.

There are also two methods under investigation at DIII-D that utilize measurements of electron cyclotron emission (ECE) to fine-tune alignment. ECE measurements can be used to detect the location of an NTM if it is rotating, since it then gives rise to temperature fluctuations that have a phase shift between the inside and the outside of the island [10]. One elegant method is to view the ECE from the same major radius and height and along the same polar and azimuthal angles as the gyrotron beams [11, 12]. In this case, temperature fluctuations originating from inside the deposition will be seen on channels with a higher frequency than the gyrotron frequency and the outside will be viewed at lower frequencies. The goal then becomes to simply move the diagnostic mirror in tandem with the gyrotron mirrors until the phase shift on the diagnostic occurs at the gyrotron frequency. The second ECE method involves pulsing the gyrotron power (typical pulse frequency is 100 Hz) to generate localized heat pulses in the plasma that can be detected by the ECE diagnostic. The mirrors are then steered until the heat pulses appear on the same ECE channel as the phase shift in temperature fluctuations caused by the rotating NTM (a small shift is added to account for the difference between the centers of the heat and current depositions). In the experiments reported here, data was collected to evaluate this second ECE method [13]. This method can potentially be used to inform the TL algorithm about what initial direction to take for the sweep and to serve as a complement to the analysis of how the growth rate depends on the alignment correction made by the TL algorithm.

3.4. Catching NTMs

Prior to 2012, NTM control strategies on DIII-D involved driving ECCD continuously on the resonant surface to preemptively suppress the NTM. In some cases, fine-tuning of the alignment was made initially by sweeping the ECCD back and forth across a saturated NTM. This can be thought of as a manual version of TL; it gave a correction term to AT that then worked well for all similar discharges. In 2012, studies began of a strategy envisioned for ITER [14] and termed 'catch and subdue' (C&S) at DIII-D. In this case, the gyrotrons remain off until an NTM is detected. When the NTM is detected, the gyrotrons turn on immediately to suppress the NTM. If the detection and alignment (the 'catch') is early enough, this technique can promptly suppress the NTM (the 'subdue'). The scenario can lead to increased fusion gain, Q, since the EC system can be off when the NTM is not present.

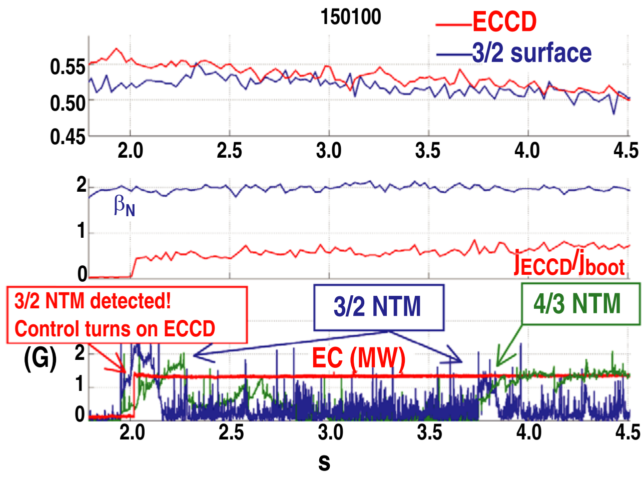

Figure 3. Prompt suppression of the 3/2 NTM using C&S. Normalized beta was controlled to a value of 2. The ratio $j_{\mathrm{ECCD}}/j_{\mathrm{boot}}$ was only 0.5 during the first NTM suppression at 2.2 s (partly due to high density) and then drifted up to 0.65 at 3.7 s when a second (preemptive) NTM suppression occurred.

Figure 3 shows a case of prompt 3/2 suppression with C&S. The 3/2 NTM is detected at about 2 s. The gyrotrons are turned on and the mode is suppressed within 0.2 s. In this shot, the ECCD stayed on after suppression since the function to turn off had not yet been implemented. A second NTM was triggered at 3.75 s and suppressed promptly, since the ECCD was already on.

Early detection is crucial for prompt suppression with C&S. Figure 4 compares the evolution of the NTM for five similar discharges. In the two cases shown in red, the ECCD was turned on when the mode had exceeded a critical amplitude, whereas for the three cases in blue, the ECCD was turned on with the mode still below this amplitude. Suppression is significantly delayed when the catch is late.

The minimum power needed for complete NTM suppression is also found to vary depending on when the NTM is intercepted. For these plasmas, 1.5 MW of injected EC power was needed to obtain a peak ECCD density (j_{ECCD}) equal to the bootstrap current density at the $q = 3/2$ surface (j_{boot}). For fully saturated 3/2 NTMs, such as the case shown in figure 2, more than 2.25 MW of ECCD was needed (corresponding to $j_{\mathrm{ECCD}}/j_{\mathrm{boot}} > 1.5$), whereas only 0.5 MW sufficed under similar conditions if the ECCD was on continuously for preemptive suppression. The minimum

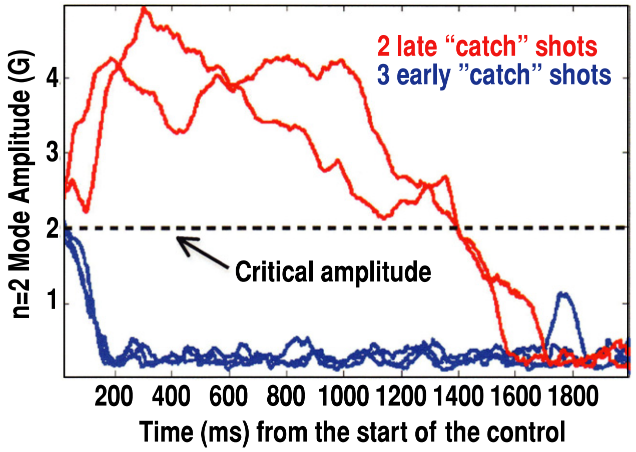

Figure 4. Showing the 3/2 mode amplitude for five cases. The ECCD is turned on at $t = 0$ in all cases. For the cases shown in red, the mode has then already exceeded the critical amplitude for prompt suppression which then takes significantly longer.

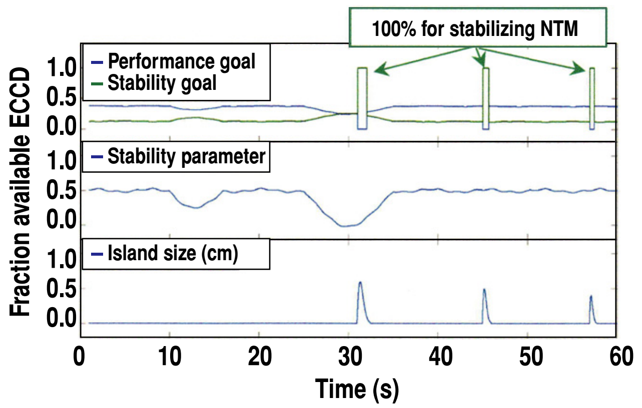

Figure 5. Schematic demonstration of ECCD use in an envisioned ITER scenario where 50% of the available power is assumed shared between profile control for performance and stability control. When NTMs strike, 100% of the available power goes to suppressing the NTM. The NTM stability parameter is yet to be clearly defined, but is an area of active research.

power needed for complete suppression with C&S has not yet been tested, but complete suppression was seen with powers below 1.5 MW, as the case in figure 3 shows.

Experiments with control of the 2/1 NTM have also been carried out. In this case, more power was needed for suppression than for the 3/2. Saturated 2/1 NTMs could not be fully suppressed, even with all of the available 3 MW of ECCD. With preemptive suppression, 1.5 MW of ECCD was sufficient. Three cases of complete suppression with C&S were successful at a power of 2.5 MW; however, the catches were all done above the critical amplitude, since the noise level was too high to set the detection threshold below this amplitude.

The real-time NTM detection system on DIII-D is based on Fourier analysis of Mirnov probes. At present, only a toroidal decomposition is made and an NTM is considered detected when the amplitude of its toroidal mode number exceeds a preset threshold for a preset time interval. In 2013, a new bandpass filter was installed to reject noise and make it possible to detect NTMs earlier at a lower threshold without responding to false positives. In addition, the ECCD can now be turned off after each new NTM is suppressed. This means that occasional false positives are acceptable since the only consequence is that the ECCD is on for a brief time. A future upgrade is also being planned to a real-time analysis that will resolve both the toroidal and poloidal spectra of the Mirnov oscillations and further decrease the occurrence of false positives.

4. Implications and vision for ITER solution

The ITER TM control system must be part of a larger control function to provide robustly stable control (under normal operating conditions, including expected disturbances) and respond appropriately under fault conditions (known as 'exceptions' in ITER parlance). Robustly stable operation begins with real-time control to produce stable profiles (where possible and desirable), including real-time monitoring to predict the evolution of a potentially unstable profile. The ITER approach to this prediction is to run a faster-than-real-time simulation (FRTS) in the ITER PCS that is capable of predicting the current profile evolution and thereby the location of rational q-surfaces, as well as onset risk for NTMs. The simulation will also calculate the locations of ECCD. The simulation must be updated in real-time to ensure that the evolution of the simulated plasma, as measured by a comprehensive list of diagnostics, matches the actual evolution. Determining proximity to stability boundaries may make use of active techniques, such as MHD spectroscopy, in addition to real-time stability calculations based on projected profile evolution. Island response data from other active perturbations, such as TL action, may provide further constraints on the FRTS calculation.

The control actions will depend on the present and predicted state of the plasma. During normal NTM-free operation, the goal is to minimize the onset risk for NTMs. This may include methods that intentionally excite NTMs periodically while maintaining active ECCD stabilization [15, 16]. In the event that the NTM becomes truly unstable and a seed island is formed, the control system will try to catch the growing mode for prompt suppression. If the island nonetheless grows, the control system may decide to turn on more ECCD or correct a misalignment or go to exception handling. Methods for exception handling are also under active research at DIII-D but are outside the scope of this paper.

Figure 5 shows a schematic example of how EC power might be used during ITER operation. In the example, the EC power is shared between the dual objectives to optimize plasma performance (profile control) and to keep it stable to NTMs. It is assumed that the correct alignment of ECCD and an (as-yet undetermined) NTM stability parameter are known through comprehensive FRTS analysis. A low stability parameter means that NTMs are seeded frequently and are hard to suppress. When the stability is good, the EC power is focused on increasing plasma performance. When the stability parameter begins to falter, more EC power is shifted to improving it. Whenever an NTM strikes, extra EC power is immediately turned on to suppress the NTM.

5. Conclusion

Techniques developed on DIII-D for the control of neoclassical TMs using ECCD have demonstrated key features of the control strategies envisioned for ITER. It has been shown that the position of the resonant q-surface can be tracked with adequate precision by equilibrium reconstructions that include internal field measurements, such as the motional Stark effect (MSE). The refraction can be tracked with adequate precision by ray trace analysis. The ECCD is kept aligned to the q-surface using real-time mirror steering. Systematic errors in the alignment can be automatically corrected when necessary. NTMs can be promptly suppressed by turning on the ECCD before they exceed a critical mode amplitude. This allows the catch and subdue (C&S) strategy to potentially save significantly on the average EC power with a resulting increase of the fusion gain, Q, compared to using continuously applied ECCD.

Acknowledgment

This work was supported by the US Department of Energy under DE-FC02-04ER54698, DE-AC02-09CH11466 and DE-FG02-04ER54761.

References

[1] La Haye R J 2006 *Phys. Plasmas* **13** 055501
[2] Buttery R J *et al* 2000 *Plasma Phys. Control. Fusion* **42** B61–73
[3] Prater R 2004 *Phys. Plasmas* **11** 2349
[4] Ferron J R, Walker M L, Lao L L, St. John H E, Humphreys D A and Leuer J A 1998 *Nucl. Fusion* **38** 1055
[5] Humphreys D A, Ferron J R, La Haye R J, Luce T C, Petty C C, Prater R and Welander A S 2006 *Phys. Plasmas* **13** 056113
[6] Poli E *et al* 2001 *Comput. Phys. Commun.* **136** 90
[7] Kolemen E *et al* 2013 State-of-the-art neoclassical tearing mode control in DIII-D using real-time steerable electron cyclotron current drive launchers, in preparation
[8] Prater R 2007 *Nucl. Fusion* **47** 371
[9] La Haye R J, Günter S, Humphreys D A, Lohr J, Luce T C, Maraschek M E, Petty C C, Prater R, Scoville J T and Strait E J 2002 *Phys. Plasmas* **9** 2051–60
[10] Park Y S and Welander A S 2006 *Plasma Phys. Control. Fusion* **48** 1447
[11] Bongers W A *et al* 2009 *Fusion Sci. Technol.* **55** 188
[12] Volpe F A G, Austin M E, La Haye R J, Lohr J, Prater R and Strait E J 2009 *Phys. Plasmas* **16** 102502
[13] Kolemen E, Ellis R A, La Haye R J, Lohr J, Noraky S, Penaflor B G and Welander A S 2012 *Bull. Am. Phys. Soc.* **57** 350
[14] La Haye R J, Isayama A and Maraschek M 2009 *Nucl. Fusion* **49** 045005
[15] Federico, Felici F, Goodman T P, Sauter O, Canal G, Coda S, Duval B P and Rossel J X 2012 *Nucl. Fusion* **52** 074001
[16] Chapman I T *et al* 2013 *Nucl. Fusion* **53** 066001