Novel aspects of plasma control in ITER

D. Humphreys,[1] G. Ambrosino,[2] P. de Vries,[3] F. Felici,[4] S. H. Kim,[3] G. Jackson,[1] A. Kallenbach,[5] E. Kolemen,[6] J. Lister,[7] D. Moreau,[8] A. Pironti,[2] G. Raupp,[5] O. Sauter,[7] E. Schuster,[9] J. Snipes,[3] W. Treutterer,[5] M. Walker,[1] A. Welander,[1] A. Winter,[3] and L. Zabeo[3]

[1]*General Atomics P.O. Box 85608, San Diego, California 92186-5608, USA*
[2]*CREATE/University of Naples Federico II, Napoli, Italy*
[3]*ITER Organization, St. Paul Lez durance Cedex, France*
[4]*Eindhoven University of Technology, Eindhoven, The Netherlands*
[5]*Max-Planck Institut für Plasmaphysik, Garching, Germany*
[6]*Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543-0451, USA*
[7]*Centre de Recherches en Physique des Plasmas, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland*
[8]*CEA, IRFM, 13108 St. Paul-lez Durance, France*
[9]*Lehigh University, Bethlehem, Pennsylvania, USA*

(Received 7 August 2014; accepted 16 January 2015; published online 12 February 2015)

ITER plasma control design solutions and performance requirements are strongly driven by its nuclear mission, aggressive commissioning constraints, and limited number of operational discharges. In addition, high plasma energy content, heat fluxes, neutron fluxes, and very long pulse operation place novel demands on control performance in many areas ranging from plasma boundary and divertor regulation to plasma kinetics and stability control. Both commissioning and experimental operations schedules provide limited time for tuning of control algorithms relative to operating devices. Although many aspects of the control solutions required by ITER have been well-demonstrated in present devices and even designed satisfactorily for ITER application, many elements unique to ITER including various crucial integration issues are presently under development. We describe selected novel aspects of plasma control in ITER, identifying unique parts of the control problem and highlighting some key areas of research remaining. Novel control areas described include control physics understanding (e.g., current profile regulation, tearing mode (TM) suppression), control mathematics (e.g., algorithmic and simulation approaches to high confidence robust performance), and integration solutions (e.g., methods for management of highly subscribed control resources). We identify unique aspects of the ITER TM suppression scheme, which will pulse gyrotrons to drive current within a magnetic island, and turn the drive off following suppression in order to minimize use of auxiliary power and maximize fusion gain. The potential role of active current profile control and approaches to design in ITER are discussed. Issues and approaches to fault handling algorithms are described, along with novel aspects of actuator sharing in ITER. © *2015 AIP Publishing LLC*. [http://dx.doi.org/10.1063/1.4907901]

I. INTRODUCTION

Tokamaks rely on sophisticated plasma control systems (PCSs) and algorithms to drive a large array of actuators (power supplies for superconducting and resistive coils, heating and current drive systems, fueling systems, etc.) in a coordinated way to produce the scenarios required to satisfy experimental physics and operational goals. A similarly large set of diagnostics provides real-time measurements from which control algorithms determine the necessary commands to these actuators. The ITER PCS will make use of about 45 different diagnostic systems and 20 different actuator systems to determine relevant plasma conditions and produce the necessary coil current waveforms, heating effects, fueling, etc.[1] Each of these actuator and diagnostic systems may in turn consist of several to dozens of components (e.g., 12 superconducting coils and 24 gyrotrons). Using these systems, the ITER PCS must regulate dozens to hundreds of quantities. In addition to actively regulated quantities, ITER must monitor hundreds to thousands of diagnostics of plasma and system states in order to respond appropriately to predicted or detected faults. The number of controlled quantities, actuators, and diagnostics is comparable to those found in operating commercial aircraft, but is smaller than, for example, a typical process-controlled chemical plant.[2] Although the number of actuator systems and the number of basic and physics-supporting diagnostics in ITER is comparable to that in many operating tokamaks (e.g., 18 PF coils, 14 independent heating systems, and hundreds of measurements in DIII-D[3]), the number, complexity, and required reliability of measurements to be monitored for fault responses exceed that in any operating device. The principal challenges for ITER control arise from the complexity and uncertainty of the plasma responses and interactions among different controlled quantities and systems, and the need for high confidence robust performance to be established prior to execution of a given experiment. Experience on operating devices suggests that while many of the control solutions

needed by ITER are well established, many others still present various kinds of challenges, ranging from lack of physics knowledge, to incomplete ITER-specific and integrated mathematics solutions. These challenges are further complicated by the novelty of burning plasmas, which will only have been glimpsed experimentally prior to their production in ITER. Because of the high cost per discharge and ambitious research schedule, once ITER begins operating the experimental time available to develop the required control solutions will be highly constrained, in contrast to present operating devices.

Compared with existing devices, the ITER plasma control development plan must rapidly achieve high levels of control performance and reliability including nuclear licensing requirements. Existing tokamaks have evolved their control capabilities gradually over many years to reach a point comparable to the performance envisioned for the ITER H/He non-activated phase target (the earliest shaped plasmas envisioned in ITER). In ITER, ~90 days of 2-shift operation, or ~1200 discharges, are envisioned for commissioning of target control capabilities for the H/He operating phase, which precedes the deuterium and deuterium-tritium research phases.[4] The H/He phase as a whole will include 10 000–12 000 pulses, during which further control refinement can be accomplished. By comparison, the relatively new superconducting EAST tokamak achieved equilibrium control, perhaps roughly comparable to the H/He campaign target after many thousands of discharges (with much lower requirements for reliability and exception handling (EH) capability).[5] DIII-D was able to exploit experience with earlier Doublet III development, yet similarly reached a roughly comparable level of performance (but not reliability) in equilibrium control to the H/He campaign target after several thousand discharges.[6] It is noteworthy that such development times have often been required for separate individual control areas, although multiple algorithms are often developed and reduced to practice simultaneously, and also in concert with ongoing physics studies. Rapid deployment and commissioning of control algorithms is, of course, often accomplished in mature tokamaks. For example, after more than two decades of operation and using model-based design and testing techniques, JET implemented its Extreme Shape Controller (XSC), and required "limited" commissioning time with minimal impact on operations to make the new algorithm available for experimental use.[7] Initial commissioning of tearing mode (TM) suppression algorithms on DIII-D was accomplished in dozens of discharges using similar model-based design and test methods, and may have achieved H/He campaign target levels of performance (although not reliability) after experimental use and ongoing development over hundreds of discharges.[8] Thus, while ITER control deployment will require high reliability and nuclear licensing, in contrast to presently operating devices, there is high confidence that the mature design tools and preparation approaches under development or available now on existing devices will enable efficient and rapid deployment of reliable control for ITER within the planned commissioning time.

The tokamak control process encompasses many elements, each touching a wide range of fields in fusion science. Figure 1 illustrates the key elements of control development, and some of their interrelationships. Production of a *physics scenario* (defined here as a time-ordered set of target equilibria and time sequences of other plasma states that define an intended discharge history over time) requires definition of a *control scenario* (defined here as the actuator waveforms needed to produce the physics scenario, as well as the definition of control algorithms used at specific times during the discharge). Together, we refer to the combination of physics and control scenarios as the *discharge scenario*. Some variability is usually expected within a given scenario definition, and the degree of specificity of definition varies significantly for different applications. For example, the ITER Baseline Scenario (IBS) includes snapshot equilibria for a range of specified transition times from limited to diverted plasma during startup, nominal transition to H-mode and burn, expected burn (>400 s) and flattop duration, and a current rampdown with gradual reduction in plasma elongation while preserving the divertor configuration for 80 or more seconds. Construction of a control scenario requires definition of control *schemes*, which are descriptions of the set of variables to be regulated to accomplish the necessary control goals, as well as the set of actuators and diagnostics to be used. For example, regulation of the plasma shape and position can be accomplished through control of the current centroid and shape parameters such as elongation, triangularity, and squareness. Alternatively, the plasma boundary can be controlled through regulation of the position of points on the plasma boundary, or spacing of reference gaps between the first wall and plasma separatrix. The latter approach is the fiducial approach for control of the ITER boundary and divertor leg geometry. While axisymmetric magnetic control always makes use of the PF and in-vessel coil power supplies, other types of control may offer different actuator options. For example, current and pressure profile control can be accomplished by a mix of regulation of ohmic coil current, neutral beams (NBs), radio frequency (RF) sources, fueling, and impurity injection, all varying widely among operating devices.

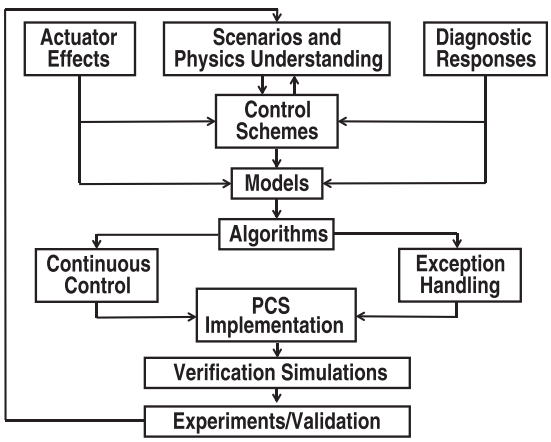

FIG. 1. General elements of the tokamak control process include definition of scenarios, schemes, and control algorithms. Both continuous control and EH functions require design of algorithms. Simulations must verify feasibility and performance prior to use in experiments.

A physics scenario can be nominally produced through pure open loop drive of the designated actuators. However, even in a plasma/tokamak system that is passively stable in all degrees of freedom, purely open loop drive is virtually always insufficient to accurately produce a desired trajectory. Variability in plasma and machine conditions will result in steady state or drifting errors, while disturbances will typically move the scenario away from the target or even produce instability. Thus, even for nominally passively stable plasma states, feedback regulation is essential.

The degree of high performance required of ITER control, along with limitations on control tuning during operation, leads to the requirement that all control algorithms in ITER must be based on *control-level models*. Such models describe appropriate plasma and system responses with sufficient accuracy to enable practical design of feedback algorithms. Often (but not always) the required level of accuracy for plasma models is well below that associated with good physics understanding, although control performance is improved by more accurate models. The need for plasma response and machine system models to design controllers demands significant and specific physics understanding. In some sense, this precisely determines the understanding needed to achieve the goal of operational fusion, and can thus be used to help drive priorities in physics research.

Two different classes of control algorithm are required for ITER, both based on control-level models. *Continuous control* algorithms determine the feedback regulation needed to produce and maintain a desired nominal scenario. These algorithms are in continuous calculation mode, although in some cases the resulting actuator commands may be intermittent. For example, while plasma shaping control will generate continuously varying commands to PF coil power supplies, "continuous" TM control must minimize the use of gyrotron power (ideally maintaining near zero) except for times when a mode is in imminent danger of growth, or is actively growing. *EH control* includes both scenarios and algorithms to provide effective responses to off-normal and fault events. Such events which require a change of some sort in the control action are termed "exceptions." The EH system will monitor many relevant plasma and machine states and evaluate the results of real-time forecasting calculations, to determine the need for asynchronous intervention in some aspect of control. Such interventions will vary from replacement of faulty diagnostic signals, to modification of control gains, to triggering of new scenarios, in order to manage many different off-normal conditions. The EH system has a hierarchy of goals, from maintaining a discharge in order to maximize physics productivity, to preventing disruptions, including executing different types of controlled rapid shutdown in the event of unrecoverable faults.

Following *PCS implementation* of the functional algorithms (including EH) and a specified scenario, a critical and novel requirement of ITER operation is the need to verify both the implementation and performance of control algorithms prior to use. This requirement is driven by the high reliability standards implied by physics productivity demands and low disruptivity constraints. All algorithms must be verified in simulation to confirm correct implementation, and dynamic performance must be confirmed under challenges by expected noise, disturbances, and key potential off-normal events. Finally, each discharge scenario must itself be verified in simulation prior to experimental execution, to ensure consistency with machine limits and expected plasma evolution, and confirm robustness to non-ideal effects and key potential off-normal events.

One of the important choices that must be made in constructing control algorithms is the degree of integration of multiple control goals under a single algorithm. Although many of the fundamental control requirements of ITER can be treated independently of each other where there are weak physics couplings (for example, plasma boundary control and density), many others involve strongly coupled physics and may be optimally regulated by integrating their mathematical algorithms. Where control goals are not explicitly integrated in a single algorithm, shared actuator management (SAM) becomes an important challenge for ITER.

In what follows, we discuss selected aspects of ITER control in order to identify key novel characteristics with respect to presently operating devices. These novel aspects range from differences in physics phenomena and control schemes that must be addressed by control algorithms, to new approaches to control mathematics and simulation to support both operational requirements and the fundamental physics mission. In the area of control physics, TM control and profile control are discussed in substantial detail as key examples illustrating many novel aspects of control in ITER. Similarly, EH and actuator management are emphasized as areas of unique challenge in control mathematics for ITER. Note that the present work is not intended to provide a comprehensive overview of ITER control issues. Many topics known to be of potential importance are omitted in the interest of space, such as sawtooth regulation and disruption mitigation-related control including runaway electron channel position regulation. Omissions should not be construed as implying that such topics are unimportant to ITER. Section II provides an overview of ITER plasma control elements, including key actuators. Section III discusses selected control physics topics, while Sec. IV describes key areas of control mathematics and computational solutions required for ITER control. Concluding remarks are provided in Sec. V.

II. OVERVIEW OF ITER PLASMA CONTROL ELEMENTS

The principal physical control actuators in ITER are summarized in Figure 2. These actuators enable all control functions, including regulation of plasma equilibrium shape and position, plasma kinetic characteristics, current profile control for both nominal scenario and avoidance of stability boundaries, and active control of MHD instabilities. The figure illustrates the high level of actuator and mission sharing (see Sec. IV E) in ITER: all PF coils contribute to both ohmic flux and shaping at the same time; multiple fueling, impurity, and heating sources contribute to kinetic control; the electron cyclotron (EC) system will provide both MHD control and

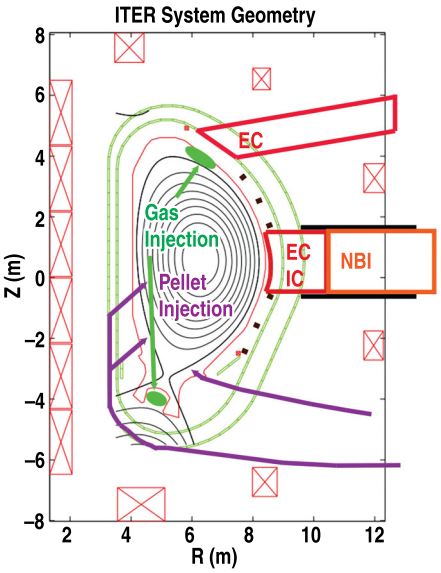

FIG. 2. ITER geometry and principal physical control actuators.

heating. Machine operations and protection functions provided under ITER control include wall conditioning, tritium removal, error field control, and disruption mitigation. Related to and significantly coupled with scenario kinetic regulation control goals are the explicit control of divertor heat flux and plasma burn state. Table I summarizes selected key control categories along with the principal control goals and actuators.

Principal heating and current drive systems in the ITER baseline include 33 MW of 1 MeV negative NBs, 20 MW of EC heating at 170 GHz, and 20 MW of ion cyclotron (IC) heating between 40 and 55 MHz (all values correspond to power coupled to the plasma). A potential upgrade would add 20 MW of lower hybrid (LH) heating and current drive at 5 GHz. The EC system will share its resources between an equatorial launcher (EL) system for plasma initiation and burnthrough, and an upper launcher system primarily for TM suppression. Fueling and impurity gas injection systems will provide H, D, T, and He^4 for plasma fueling, and N, Ne, Ar, and He^3 for impurities (with maximum throughputs of 100 Pa-m^3/s for H, D, and He^4, and 10 Pa-m^3/s for T, N, He^3, Ne, and Ar). Cryogenic pellet injection systems provide maximum throughputs of 120 Pa-m^3/s for H and D, 111 Pa-m^3/s for T, and 10 Pa-m^3/s for impurity species. Pellet injection velocities are in the range of 100–300 m/s.

Certain diagnostics are of particular interest for control discussions below. Both Motional Stark Effect and polarimeter/interferometer systems will provide current profile measurements to constrain plasma equilibrium reconstructions

TABLE I. Summary of selected control functions in ITER, along with principal control goals and actuators. (PF = poloidal field, CS = central solenoid, VS3 = vertical stability coil circuit #3, NBI = neutral beam injection, ICRH = ion cyclotron resonant heating, RMP = resonant magnetic perturbation, ECH/ECCD = electron cyclotron heating/current drive, EFC = error field correction).

Control Category	Principal Control Goals/Control Quantities	Actuators
Plasma equilibrium	Boundary, position	PF coils
	Divertor magnetic configuration	
Plasma current	Magnitude of plasma current	CS coils
	Internal inductance	
Vertical stability	Vertical stabilization	PF coils, in-vessel VS3 coils
	Vertical position (partial)	
Kinetics	Core electron density	Fueling pellets, gas, NBI, ICRH
	Stored energy/beta	
Burn state	Fusion gain	Fueling pellets, gas, NBI, ICRH, in-vessel RMP coils
	Fusion alpha power	
	Confinement	
Divertor	Target heat flux	Impurity gas injection, fueling
	Divertor radiation	
	Degree of detachment, electron temp.	
Current profile	Internal inductance	CS coils, ECH/ECCD
	Q profile	
	Proximity to MHD control boundaries	
ELM control	ELM frequency, amplitude	RMP coils, pacing pellets
	ELM stability	
Sawtooth control	Sawtooth stability	ECH, ICRH
	Sawtooth frequency	
TM control	TM stability	ECH/ECCD, in-vessel RMP
	TM island size	coils, NBI, ICRH, EFC coils
	Mode rotation	
Fast particles	Stabilize Alfven Eigenmodes	ECH/ECCD, ICRH, NBI, in-vessel RMP coils
	Regulate fast particle confinement	
Error field	Error field correction	Error field correction coils, in-vessel RMP coils
	Rotation	
Disruption mitigation	Rapid uncontrolled shutdown	CS/PF coils, VS3 coils, Disruption Mitigation System
	Mitigate disruption effects	(impurity gas, shattered pellet injection)

and enable q-profile determination. Fast magnetic and ECE will provide measurements contributing to determination of TM island geometry and phase. Line averaged electron density will be measured by polarimeter/interferometer, while local measurements will be available from Thomson Scattering and microwave reflectometry. Divertor characteristics will be measured by a large array of diagnostics, including bolometers, IR/visible cameras, pressure gauges, and Langmuir probes.

Further details regarding ITER control resources can be found in Ref. 9.

It should be noted that control requirements will vary substantially across the evolution of ITER. First Plasma operations are anticipated to produce a limited plasma configuration with plasma current of ~1 MA or less and very short pulse duration of perhaps several seconds. These conditions, intended to support initial commissioning of many machine systems, will require far lower control resources and algorithm capabilities. EH capability and reliability requirements are also greatly reduced, since the ability of such low power plasmas to damage machine components is very limited.

As noted previously, the ITER Research Plan[4] foresees an initial phase of non-active operation using H/He plasmas following First Plasma during which much of the plasma control capability will be commissioned and applied. For the purposes of defining the key goals which must be achieved during the non-active phase in order to progress to D and DT operation, the non-active phase has been further subdivided into operation at flat-top currents of (i) $\leq$3.5 MA, (ii) $\leq$7.5 MA, (iii) $\leq$15 MA. For the first phase in this sequence, which leads to the achievement of divertor operation at a plasma current of ~3.5 MA, the full equilibrium reconstruction and magnetic equilibrium control capability, including fast feedback control of the vertical position, would be commissioned. Density feedback control would also be necessary. These two control capabilities provide the basic control functionality for much of the present tokamak research program. Commissioning of some aspects of control for machine protection would also be undertaken, including first wall heat flux control (using measurements from in-vessel viewing cameras), disruption detection and mitigation (including runaway electron suppression/mitigation), and error field correction. Application of EH would also be important at this stage in preparation for advancing to higher plasma parameters, as would the demonstration of satisfactory operation of interfaces to machine and safety protection systems. This capability should allow initial diverted operation with highly elongated plasmas to be established for at least several seconds.

The second phase of H/He operation, using currents of up to 7.5 MA (and toroidal fields of either 2.65 T or 5.3 T), would provide the basic plasma scenarios for much of the plasma and plasma control commissioning to be carried out during the non-active phase. These target plasmas would, in particular, be used for much of the H&CD commissioning and lead to high power L- and H-mode operation. Thus, in addition to the extension of the control and protection functionality developed within phase (i) to higher plasma parameters, all elements of control and protection associated with the H&CD systems (e.g., ion cyclotron radio frequency (ICRF) coupling control, NB shine-through protection) would be commissioned as the H&CD systems were brought online. New control functions which would be commissioned during this period would include divertor heat flux control via impurity seeding, ELM control—when sufficiently high powers were attained to permit H-mode access (most likely in helium)—and both divertor and core impurity control. H-mode operation might also provide conditions under which commissioning of beta (stored energy) control and neoclassical tearing mode (NTM) control could be implemented. During this phase, routine exploitation of EH and commissioning of some actuator management algorithms would be essential.

The control functionality developed during phases (i) and (ii) of non-active operation would provide sufficient plasma control capability to support plasma scenario development for phase (iii), the achievement of 15 MA L-mode operation. Although the control and protection functions previously commissioned would need to be adapted as plasma current and associated plasma parameters were increased, no new control functions would be required in progressing from 7.5 to 15 MA.

Following completion of the non-active phase of operation, the Research Plan foresees a short phase of operation in D plasmas in preparation for the transition to full DT operation. Assuming that the commissioning program outlined above for plasma control were fully successful, the required control capability for D operation would already be available (this is, indeed, a major goal of the non-active phase). Control algorithms might need to be retuned to reflect detailed differences between H/He plasmas and D plasmas, particularly in H-mode operation, and some further tests of control robustness would be required as plasma parameters increased. However, further commissioning of plasma control functions would be focused on capabilities required for DT operation. Examples of such DT-specific control include fuel mixture control and burn control, which would be tested using suitable simulation techniques, and possibly in relevant but necessarily limited DD experimental operation. Some capability for current density profile control might also be commissioned if adjustment of the plasma current profile was required to optimize confinement.

The transition to DT operation is likely to occur through a gradual increase in the tritium concentration of initially DT plasmas, leading to increasing levels of fusion power. Virtually, the full scope of the plasma control capability would become available n the course of this phase. In particular, one could expect that a reliable scenario capable of sustaining high fusion power would require the integration of such "advanced" control functions as core impurity control, fuel mixture control, burn control, exhaust power control, and possibly some aspect of current profile control to maintain the quality of plasma confinement necessary for high fusion gain and to limit heat fluxes on plasma facing components. Key aspects of MHD stability control such as error field control, ELM suppression or control, NTM control, sawtooth control, as well as disruption detection, avoidance and mitigation, would also be in routine operation. In this

phase, essentially the full actuator sharing and EH capability would also be required. In the longer term, scenarios exploited for the development of fully non-inductive steady-state operation would require full current profile control and possibly also Alfvén eigenmode control.

This planned evolution of control capability itself places specific requirements on the flexibility and continuous scalability of the control system, which must be built into the design from the earliest stages.

III. SELECTED CONTROL PHYSICS CHALLENGES

A. Control physics

The term "control physics" refers here to the physics understanding required to accomplish the control needed for scenario execution, to derive appropriate schemes to accomplish the required regulation, and to construct the models required to design controllers satisfying given control goals. This process includes identification of the physics phenomena needing to be regulated, as well as quantification of dynamic behavior at the appropriate level of description to define control-level models for algorithm design. In order to achieve and regulate the desired plasma states, the way in which plasma phenomena respond to actuators must also be characterized, along with methods for appropriate measurement of the relevant phenomena in real time. These considerations then lead to development of an effective control scheme. They can also enable derivation of the sequence of events that must be followed in order to accomplish an experimental goal (the physics scenario). This systematic process of scenario construction via modeling and simulation is expected to be followed extensively in ITER, in order to minimize use of machine time to develop discharges.

Because of the significant reliance on model-based design of scenarios and control algorithms, ITER demands a much larger emphasis on physics understanding specific to control requirements. For example, nominal execution of a given discharge scenario depends on specification of the (open loop, or feedforward) trajectories of all relevant actuators, along with descriptions of fiducial target equilibria throughout the scenario. This is the domain of experimental and theoretical scenario physics. However, producing such scenarios accurately, reproducibly, and robustly in the presence of varying wall conditions and disturbances requires active control methods and corresponding model physics understanding often very different from characterizing nominal plasma responses. Plasma dynamics on timescales short compared with discharge and plasma evolution times become much more relevant. Dynamic effects associated with disturbance instabilities and the range of plasma conditions produced by varying wall conditions become important to quantify. The actual size of many kinds of uncertainties, noise, and disturbances must be quantified with some accuracy to enable design of sufficiently robust algorithms.

B. Scenario control

The IBS poses significant control challenges, many of which have been identified in experiments executed on many devices.[10–12] Figure 3 summarizes the evolution of key quantities in a DIII-D experiment simulating the IBS. Points in the scenario of particular challenge include approaches to (and thus the need for avoidance of) control limits. Rampup and rampdown, for example, result in approach to axisymmetric instability limits. High Z_{eff} during rampup can result in an internal inductance (li) near the transition to flattop, which approaches the controllability limit. For this reason, active control of the internal inductance may be required in this phase.[13] The central solenoid (CS) is envisioned to be the principal actuator for this purpose, varying the plasma current ramp rate to regulate li (as was done in DIII-D experiments simulating the IBS). Regulation of li may need to balance constraints for three different purposes: maintenance of controllability for vertical stability, maintenance of stability of kink modes, and minimization of flux consumption to maximize flattop and burn duration. Transition to H-mode improves the equilibrium and vertical controllability, but can also bring PF coils near their current limits, reducing control headroom needed for management of disturbances and ongoing variability in plasma resistivity. The rise in internal inductance during rampdown also poses a challenge for vertical control. With careful reduction of plasma elongation and regulation of the divertor configuration, experimental simulations have successfully reduced plasma current to below 1 MA ITER equivalent before any loss of control.[14]

Transition to H-mode also entails the onset of ELMs, which must be immediately suppressed to minimize divertor erosion. ELM control is not discussed in detail here, but is a significant requirement of ITER, represented by a large and ongoing field of study at devices and institutions worldwide (e.g., Refs. 15–17). Several promising approaches among these studies include control of non-axisymmetric fields with internal coils. However, in general, these experimental studies have used open loop ELM suppression, i.e., preprogrammed coil current waveforms. Providing the necessary level of robustness will require active tracking of variation in

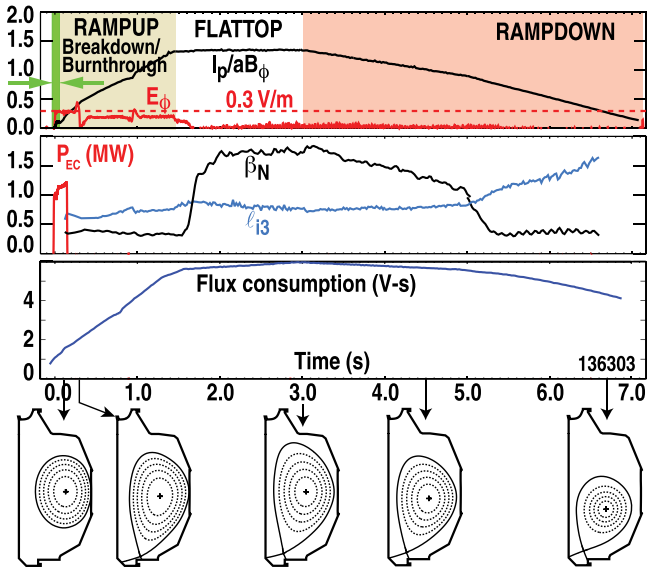

FIG. 3. Summary of IBS as executed in DIII-D.

plasma conditions, implying new control algorithms based on equilibrium, pedestal, and other plasma conditions yet to be determined.

Another unique characteristic of ITER scenarios is the potential need to anticipate certain disturbances that are intrinsically part of the discharge evolution (or may happen unexpectedly), and compensate for their impact on various control categories. For example, transition from H- to L-mode is accompanied by a large major radial shift inboard, which can result in transient wall contact. Studies have shown that anticipatory control programming can begin changing PF coil currents appropriately to minimize this wall contact while satisfying plasma boundary and divertor constraints.[18] Alternatively, optimized controllers may provide a similar function with reduced need for pre-emptive feedforward action.[19]

Another key challenge to robust operation of the IBS is the potential for TMs, which degrade confinement and are potentially disruptive (see also Sec. III D). Analysis of IBS experiments in DIII-D shows that at low torque these scenarios experience a high incidence of TM onset.[20] A key element of robust scenario execution is to monitor and maintain proximity to stability limits at the optimal point consistent with the necessary confinement while preventing onset of potentially disruptive modes (e.g., Ref. 21). These control goals, along with the need to follow a prescribed evolution in the target equilibria, imply a strong need for profile control in ITER scenarios (both IBS and steady state).

C. Current profile control

Although various kinds of profiles (current density, safety factor, rotation, density, pressure, radiation, etc.) may prove to be necessary to control in ITER to some degree, the principal target expected to require active regulation is the current density profile, a strong determiner of both MHD stability and confinement. For the purposes of the present discussion, "profile control" will refer to current density profile control, or its effective equivalents (at fixed toroidal field) safety factor, iota, poloidal flux, or similar basis representations. The principal actuators for accomplishing profile control in ITER scenarios include NBs, ECH/ECCD (both equatorial and upper launch systems), fueling sources, impurity injection, regulation of the CS to affect the total plasma current and internal inductance, potential use of nonaxisymmetric coils to affect confinement, and LHCD if made available in ITER.

It is sometimes proposed that feedback regulation of a scenario profile is unnecessary, since in most cases profile evolution resulting from action of all heating and current drive actuators represents a dynamically stable process. For example, when the (axisymmetric) internal resistive plasma response is linearized, all eigenvalues are negative; the system is purely dissipative. In such cases, and in the absence of noise, disturbances, and uncertainty in plasma conditions, preprogrammed open loop drive by (well-understood) actuators will produce a nominally stable desired target scenario. However, experimental plasmas operate in quite nonideal conditions including significant noise, disturbances, and varying impurity content, so that undesirable drift, fluctuations, and steady state errors will naturally result without feedback. In addition, significant nonlinearities in the plasma response, as well as sensitivity to nonaxisymmetric instabilities, further limit the reliability of pure open loop drive. Although some degree of active profile control is likely to be useful in ITER and remains an explicit requirement of the ITER control system, this capability is classed as an "advanced control" requirement, and is of primary relevance to the high beta operation phases of ITER.[1,9] Basic plasma control categories such as equilibrium, vertical stability, and density regulation are higher priority and will be available earlier in ITER operations.

Profile control in ITER has two fundamental goals: providing robust tracking of desired profile evolution for specified scenarios, and regulating proximity to controllability and/or stability boundaries to prevent loss of control and disruption. Both of these goals are expected to require control of the current or safety factor profile. However, the nature of the control is likely very different in the two cases. Tracking of the desired scenario profile requires reasonable control authority exercised on q values themselves across most of the q-profile. Regulation of proximity to the vertical controllability boundary will likely require simple internal inductance control (e.g., during rampup and rampdown). Regulation of proximity to $n \neq 0$ MHD controllability boundaries (beyond which active mode stabilization techniques are calculated to be insufficiently robust) likely requires more specific control of localized regions such as the location and value of the minimum q, localized profile gradients, or more complex parameters (which may include global characteristics) exemplified by the classical tearing stability parameter Δ'. It remains a challenge to identify sensitive parameters such as Δ' in real-time in presently operating devices. The challenge may prove even greater in ITER with more limited profile measurement resolution. Techniques such as active MHD stability probing may prove important in this case, provided an appropriate profile control response can be identified for a given stability response signal. Active MHD spectroscopy with applied fields can also provide information on the q-profile itself.

Both profile control missions will benefit significantly from advances in predictive modeling to satisfy ITER requirements for physics-based control algorithms. However, identification of the important and effective control parameters for regulation of proximity to controllability and stability boundaries is a relatively new area of study and remains a critical topic of control physics research.[20,22]

Most schemes under consideration for current profile control in ITER (and experimentally studied to date) involve regulation at multiple discrete points of related quantities such as safety factor (q), rotational transform (iota), poloidal flux (ψ), or the radial poloidal flux gradient (sometimes denoted by θ). Each of these choices has advantages and disadvantages, and the basis variables to be used for profile control in ITER have not been selected (and indeed may vary depending on the experimental goals of a particular operating scenario). For example, regulation of the safety factor enables the controller to act directly on qmin, q95, and

other features often used to specify a target scenario. However, the safety factor in a sense depends inversely on the current density, so that the model response to current drive actuators is highly nonlinear and sensitive to the operating point. By contrast, the response of iota $\equiv 2\pi/q$ depends more proportionally on the local current drive, and so tends to be naturally linear. However, MHD stability characteristics themselves tend to be directly related to q and not to iota. The poloidal flux profile is not simply related to any stability metric. However, its role as the fundamental dependent variable in toroidal equilibrium and in magnetic diffusion models helps simplify the relationship between models and controllers.

The physical dynamics involved in current profile control are typically described by resistive (axisymmetric) magnetic diffusion within a fixed-boundary plasma, expressed in one form by[23,24]

$$\frac{\partial \psi}{\partial t} = \frac{\eta(T_e)}{\mu_0 \rho_b^2 \hat{F}^2} \frac{1}{\hat{\rho}} \frac{\partial}{\partial \hat{\rho}} \left(\hat{\rho} \hat{F} \hat{G} \hat{H} \frac{\partial \psi}{\partial \hat{\rho}} \right) + R_0 \hat{H} \eta(T_e) [j_{aux} + j_{bs}] \quad (1)$$

with boundary conditions

$$\left. \frac{\partial \psi}{\partial \hat{\rho}} \right|_{\hat{\rho}=0} = 0 \quad \left. \frac{\partial \psi}{\partial \hat{\rho}} \right|_{\hat{\rho}=1} = -k_{ip} I_p(t),$$

where ψ is the poloidal stream function which is closely related to the poloidal flux Ψ ($\Psi = 2\pi\psi$), t is the time, η is the plasma resistivity (dependent on the electron temperature T_e), μ_0 is the vacuum magnetic permeability, j_{aux} is the noninductive current density driven by auxiliary sources, j_{bs} is the noninductive current density driven by the bootstrap effect, k_{ip} is a geometric factor and I_p is the total plasma current. The normalized spatial coordinate $\hat{\rho} = \rho/\rho_b$ is used to index the magnetic flux surfaces, where ρ is the mean effective minor radius of a magnetic flux surface, i.e., $\Phi(\rho) = \pi B_{\phi,0} \rho^2$, Φ is the toroidal magnetic flux, $B_{\phi,0}$ is the vacuum toroidal magnetic field at the geometric major radius R_0 of the tokamak, and ρ_b is the mean effective minor radius of the last closed magnetic flux surface. The parameters $\hat{F}, \hat{G}, \hat{H}$ are geometric spatial factors pertaining to the magnetic configuration of a particular MHD equilibrium. For current-profile control design purposes, the magnetic diffusion equation (Eq. (1)) needs to be closed by empirical correlations obtained from physical observations and experimental data for the electron temperature, the plasma resistivity, and the non-inductive current drive.

It is important to note that such a "physics based" modeling approach does not require a fixed plasma boundary, although this simplifying choice is often made. The spatial geometric factors F, G, H, and k_{ip} can indeed be functions of time, and can even be updated in real time based on measured or inferred plasma conditions.

The development of a physics-based nonlinear partial-differential-equation (PDE) model (e.g., Eq. (1)) capturing the spatially distributed nonlinear dynamics of the system relevant for control design enables not only the design of feedback controllers for regulation or tracking but also (1) the design of optimal feedforward controllers for a systematic model-based approach to scenario planning, (2) the design of state estimators for a reliable real-time reconstruction of the plasma internal profiles based on limited and noisy diagnostics, and (3) the development of a fast (potentially real-time), control-oriented, predictive simulation code for closed-loop performance evaluation of the developed controllers before experimental implementation.[25,26] It is important to note that a linear ordinary-differential-equation (ODE) response model can always be obtained from the nonlinear PDE response model through spatial discretization and linearization around either a reference plasma state (time-invariant linear model) or a plasma state trajectory (time-varying linear model) without the need of dedicated system identification experiments. Moreover, the linearization can be easily carried out in real time around the current plasma state, which results in a linear time-varying "adaptive" (non-static) response model.

Substantial experimental and theoretical research has been done to address the physics understanding and models needed, and then to explore candidate approaches to algorithm design for the profile-tracking mission. This research has employed a wide variety of methods, including black or gray box system identification,[27–29] largely physics-based but simplified transport equations,[23,24,30,31] or adaptive techniques.[32] Because of the complexity of the physics involved in current profile dynamics, and the difficulty in predicting plasma conditions, much of this research has made use of black or gray box approaches to modeling. A black box model is typically derived by identifying fundamental system dynamics without exploiting knowledge of the underlying physics equations. A common approach is to use system identification techniques[33] to determine parameter values for matrices in a linear state space model of the form

$$\begin{aligned} \frac{dx_{ID}}{dt} &= A_{ID} x_{ID} + B_{ID} u \\ y &= C_{ID} x_{ID} + D_{ID} u, \end{aligned} \quad (2)$$

where A_{ID}, B_{ID}, C_{ID}, and D_{ID} are matrices intended to represent the dynamics of the modeled system, x_{ID} is a state vector for the identified system model, u is the input vector, and y is the output vector for the modeled system. The order of the matrices used in the model is determined by balancing considerations of the degrees of freedom actually contained in the physical system with analysis or expectations for how many of these degrees of freedom are important to the control problem. The form of the model (Eq. (2)) is a generic linear state space, and therefore does not (necessarily) reflect the actual structure of the underlying physics equations. A gray box model makes some use of knowledge of the underlying physics or dynamic equations, but identifies some or all of the relevant parameters from experimental data (often statistical). A white box, or physics-based, model implies derivation from first principles, including use of a version of the actual dynamic equations and determination of most or all of the relevant equation parameters independent of experimental data.

Figure 4 illustrates the application of a q-profile controller in JET based on a black box model, identified from extensive experimental data.[27] The controller drives LHCD, NBI,

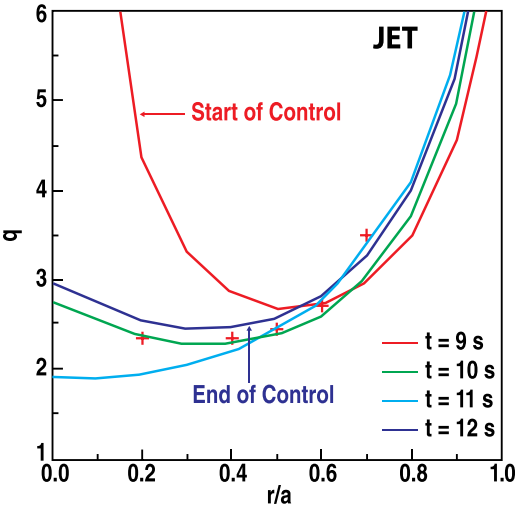

FIG. 4. Active control of the q-profile in JET using LHCD, NBI, and ICRH (pulse #58474, BT = 3 T, Ip = 1.8/1.5 MA). The q-profile control target is indicated by red crosses. Adapted from Fig. 6b, courtesy of the International Atomic Energy Agency, Institute of Publishing and D. Moreau *et al.*, Nucl. Fusion **43**, 870 (2003).

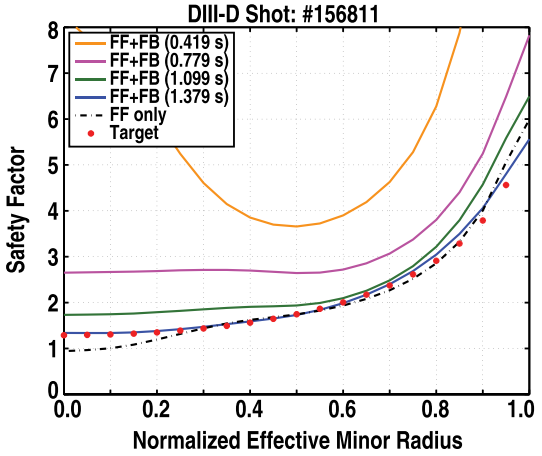

FIG. 5. Active control of the safety factor profile in DIII-D using a multivariable controller designed based on a physics-based model. Both off-axis (co-injection) and on-axis (co-injection and counter-injection) NBI is used in combination with plasma current regulation.

and ICRH actuators to regulate five points in the q-profile, specified at particular values of r/a. The model construction makes use of SVD methods to fit system dynamics in a least squares sense and describes the response of three principal component vectors in the five-point space representing the q-profile. Dynamic variation in the q-profile under feedback control can be seen beginning with the initial highly reversed-shear pre-control state (9.0 s), passing near the target state (10 s), overshooting (11 s), and settling near the target state at the end of control (12 s). Because such approaches use a model that can be very accurate for a specific target plasma, the dynamic control and tracking can be similarly accurate for that plasma (although not necessarily for even slightly different target parameters). Such solutions could be used in ITER given enough operating time to build up a sufficient data set for reliable system identification prior to controller use. A related approach which blends the strengths of physics (equation) based design and system identification-based design is to identify a system response model from data generated by a purely predictive simulation rather than experimental data.[27]

Figure 5 illustrates the application in DIII-D of optimized feedforward (open-loop) actuator trajectories modified by feedback (closed-loop) regulation. The feedforward control law is obtained by embedding the physics-based model (Eq. (1)) in a nonlinear, constrained optimization algorithm.[26] The goal of the optimization algorithm is to specify the actuator trajectories necessary to steer the plasma through the tokamak operating space to a specified target scenario, subject to the plasma dynamics described by the physics-based model (Eq. (1)), while attempting to avoid MHD stability limits and respecting actuator constraints. The optimization algorithm can be utilized in multiple ways for the development of target scenarios. Two applications that have been employed in DIII-D are to improve the reproducibility of plasma startup conditions to reach a target current profile at the end of the plasma current ramp-up, and to optimize the formation of high-beta, time-stationary plasma conditions characteristic of steady-state scenarios. Figure 5 illustrates the results of the former application. While the optimized actuator trajectories are able to drive the q profile close to the desired target, the matching is affected by the sensitivity of this approach to model uncertainties, perturbations in the initial conditions, plasma disturbances, and unaccounted actuator constraints. To add robustness to the control scheme, a feedback control law must be mounted on top of the feedforward control law.[26,34] The matching of the target profile is systematically improved by feedback controller through a tighter regulation of both q0 and q95.

A key challenge for ITER is reliably predicting plasma conditions prior to design of the relevant controllers. Real-time identification, or adaptive, methods may address this challenge by updating an algorithm dynamically based on varying plasma characteristics identified from measurements. A computationally attractive variant of this approach seeks to identify linearized profile response functions in real-time and employ these to determine the appropriate control action. An example of this approach applied to regulation of the q-profile in ITER is illustrated in Fig. 6. In this approach, linearized profile response functions are obtained by directly simplifying the underlying physics with assumptions adequate for feedback control. The required control action is directly computed by inverting the response matrix for multiple target profiles and multiple actuators. The saturation and quantization of the actuator powers are also taken into account. Work is presently underway to apply and evaluate such an approach experimentally in the KSTAR tokamak.[35]

Despite the application of active control systems to achieve optimal profiles and regulate the degree of controllability or stability, ITER may operate beyond the robust metastability limit for TMs in various scenarios. In this case, active TM stabilization is essential.

D. Active Suppression of TMs

TM control in ITER encompasses many aspects of novel control in ITER, including model-based design, significant

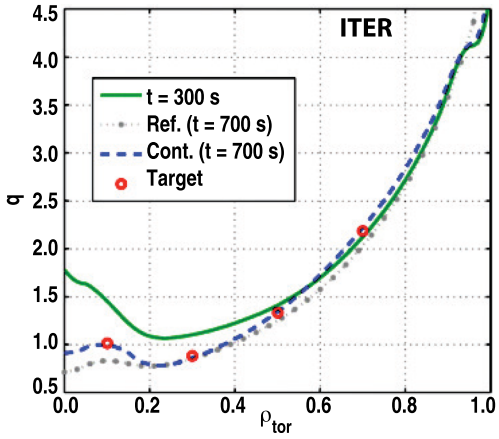

FIG. 6. Active control of the ITER q-profile using an adaptive algorithm that identifies profile response characteristics in real time. A combination of EC, LH, and NB heating power is used. LH power is deposited well off-axis, while EC and NB power is centrally deposited. The gray dotted trace shows the q-profile without control (ref), the control target is denoted by red circles, the initial condition at start of control (t = 300 s) is shown in solid green, and the steady controlled profile agreeing well with the target is shown achieved at t = 700 s in dashed blue [Reprinted with permission from Kim *et al.*, Nucl. Fusion **52**, 74 002 (2012). Copyright 2010 International Atomic Energy Agency, Institute of Publishing].

control physics needs, scenario (profile) regulation to maintain distance from controllability limits, continuous control for active suppression, actuator management (see Sec. IV E), and EH responses when control sufficiency is threatened or degraded. In this section, we therefore discuss TM in some depth.

TMs are resistive instabilities that can become metastable to fluctuations or disturbance-driven seed islands under certain conditions (e.g., sufficiently high plasma beta). This metastability implies that seed islands (produced, e.g., by sawteeth or ELMs) above the "marginal island" size will cause growth of resistive islands to a large saturated state, which in turn can slow or stop any plasma rotation, leading to even larger island growth, a drop in confinement, and even disruption.[36] Operation in the IBS at high beta, particularly at the low rotation levels frequently predicted for ITER, may result in metastability to TMs in a regime such that disturbance-driven seeds exceed the marginal island size. Frequent triggering of TMs has been observed in IBS discharges in DIII-D, with increased incidence in low torque experiments.[20] Since such ITER scenarios therefore operate above the nominal metastability limit, TMs must be actively suppressed in order to maintain performance and prevent disruption. Because this control requirement is an integral part of the scenario operation, the TM control falls into the "continuous" control category, and does not constitute an exception. The continuous TM control scheme envisioned for ITER under these conditions makes use of ECCD to drive current at the relevant rational surface and in the corresponding magnetic island.[37–39] This stabilization process has been found experimentally to be well-described by the Modified Rutherford Equation (MRE),[40] one form of which is given by

$$\frac{\tau_R}{r}\frac{dw}{dt} = \Delta'_0 r + \delta\Delta' r + a_2 \frac{J_{bs}}{J_{||}}\frac{L_q}{w}\left[1 - \frac{w_{marg}^2}{3w^2} - K_1\frac{J_{ec}}{J_{bs}}\right], \quad (3)$$

where w is the island width, τ_R is the island resistive diffusion time, r is the minor radius of the rational surface, $\Delta'_0 r$ is the classical stability index, $\delta\Delta'_0 r$ represents the change in $\Delta'_0 r$ resulting from the off-axis current drive, J_{bs} is the bootstrap current density, $J_{||}$ is the parallel equilibrium current density, L_q is the safety factor shear length, w_{marg} is the "marginal" island size (below which the mode cannot sustain itself, and above which the mode will grow to saturation), and J_{ec} is the EC driven current density. Several of the parameters in Eq. (3) are difficult to quantify from measurements, and this poses a challenge to real-time monitoring and prediction of the controllability boundary, as well as for control robustness in ITER.

The final term of Eq. (3) describes the effect of current driven at the rational q-surface (q = m/n) corresponding to the m/n structure of the mode. In equilibria with monotonic magnetic shear and plasma beta values characteristic of the IBS, island growth is metastable to sufficiently large seed islands (i.e., a sufficiently large seed island produced by some perturbation will grow to saturation). The effectiveness parameter, K_1, represents the degree of current drive alignment with the rational surface, the relative size of the deposition spot, and the balance of current driven within the island (near the island O-point) relative to that driven at the island X-point. Current driven within the island is stabilizing, while current driven at or near the island X-point is generally destabilizing. However, for sufficiently narrow deposition relative to island size, the stabilizing effect of current driven near the O-point is larger than the destabilizing effect of equal current density driven near the X-point. Thus, continuous and equal current deposition across the entire island O-point and X-point regions produces a net stabilizing effect.

Figure 7 shows the rate of island growth, dw/dt, as a function of island size as represented by the MRE for the IBS during the high beta burn state. The dw/dt = 0 point with positive derivative occurs at the "marginal island size" w_marg (~2 cm in the IBS). Seed islands greater than this size will grow to the saturated size indicated by w = w_sat.

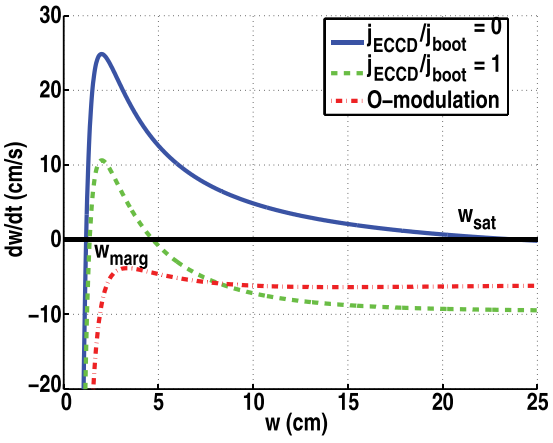

FIG. 7. MRE description of TM island rate of growth as a function of island size. The solid-blue curve shows the case for zero ECCD current drive, while the dashed-green curve represents equal ECCD and bootstrap current density, and the red dashed-dotted curve shows complete stabilization with O-point modulation and equal ECCD and bootstrap current density.

As illustrated in the figure, at the highest beta expected in the Q = 10 IBS, the saturated the 2/1 mode island size is predicted to be on the order of 20–25 cm (assuming the beta is maintained, and the mode does not lock). However, even if initially rotating in ITER plasmas, these modes are expected to lock at island dimensions on the order of 5–10 cm, requiring a strong EH response.

The MRE shows that driving sufficient current density Jeccd inside the magnetic island can make the growth rate negative and stabilize the mode. This general approach has been demonstrated on many devices by driving constant ECCD current at or near the resonant surface, and stabilizing the mode when sufficient current density is driven.[41–45] Preemptive constant current applied to this region has also been shown to provide increased mode stability, maintaining suppression in the presence of disturbances that would otherwise trigger TM with no current applied.[46] Up to 20 MW of 170 GHz ECCD power will be available from the ITER upper launcher system (Fig. 2) for TM suppression (but will be shared with the ELs to accomplish other missions including profile control). Current deposition widths for the upper launcher system are predicted to lie in the range of twice the marginal island width, implying reduced efficiency in stabilizing the mode, with high sensitivity to misalignment. Misalignments as small as 1.7 cm will drop the control effectiveness to near zero.[38] Modulating the ECCD power to synchronize with deposition in the O-point of the island is planned for ITER to increase the stabilization efficiency and compensate somewhat for the wide deposition (Fig. 7, red, dashed-dotted curve).

The TM continuous control algorithm envisioned for ITER includes several components, some elements of which have been demonstrated separately in operating devices:

(1) Active maintenance of profiles away from *controllability* limits (to the degree possible)
(2) Suppression of seeding instabilities: ELM suppression (required primarily to extend divertor lifetime, but also key to preventing TM island seeding); Sawtooth pacing (required to maintain sawtooth amplitude and thus seed island amplitude below the controllability limit);[47] Fast particle instability control (potentially, but presently without a defined control scheme)
(3) Active suppression of TM islands with ECCD, including active tracking to maintain alignment after an island is suppressed, and repetitive re-suppression if islands are re-triggered

The first component in this continuous control process is the use of profile control to maintain controllability within the available system capabilities (see Sec. III C above). The second component, suppression of instabilities with potential for excitation of seed islands, will likely be accomplished via separate algorithms with different control goals and metrics. For example, ELM suppression must be applied effectively and robustly to prevent divertor erosion. Control of fishbones and other energetic particle instabilities that may seed TMs may be necessary to accomplish separately from TM control to minimize fast particle deconfinement and first wall impacts. Sawtooth pacing or amplitude control is more likely to be integrated with continuous TM control due to its limited impact on other aspects of scenario physics (although present experiments with tungsten walls show that it has a significant impact on impurity control[48]).

The third component and principal function of the TM continuous control algorithm will be to perform suppression of islands using ECCD deposition at the resonant surface, as described above. Steerable mirrors on the upper launcher system will enable real-time alignment of the deposition location with the relevant resonant surface. The 2/1 TM is most essential to suppress, owing to its high disruptivity potential, but 3/2 islands can significantly degrade confinement, impacting the Q = 10 performance goal. In fact, the suppression process itself can significantly impact the Q = 10 goal unless the average power used in this control is minimized. Since fusion Q is defined as the ratio of fusion power to auxiliary input power, any auxiliary power used will decrease the Q achieved. The consequences for this are illustrated in Fig. 8.[39] The family of curves shown relate the achieved Q to the ECCD power for variation in the H-mode confinement factor HH between HH = 0.75 and HH = 1.25 (where HH is defined here as the energy confinement time normalized to an assumed undegraded confinement time of 3.7 s, so that $HH = \tau_E/3.7\,s$). Using no ECCD power, an unstabilized saturated 2/1 mode degrades confinement so as to achieve Q ~ 4.7. Use of 10 MW average ECCD power with an undegraded confinement corresponding to HH = 1 yields Q ~ 8.5.

In order to minimize the average ECCD power used, suppression of tearing islands must be accomplished with repetitive, as-needed application of current drive for brief intervals. One scheme to accomplish this (termed "Catch and Subdue" in its instantiation at DIII-D) entails detection (potentially prediction) of island growth at an early enough point to trigger alignment.[8,49] The scheme includes the following control actions:

(1) Detect growth of island to beyond a specified threshold amplitude
(2) Turn on ECCD power
(3) Align ECCD deposition location with the island, potentially including synchronizing ECCD modulation with the island O-point ("Catch").

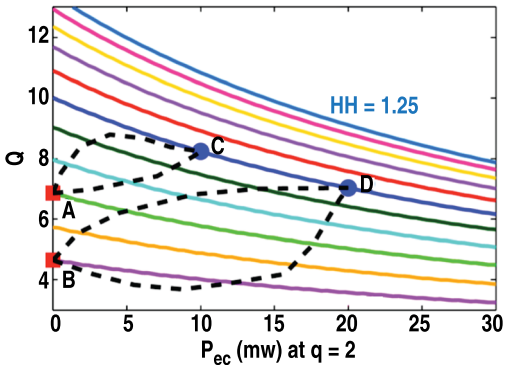

FIG. 8. Effect of ECCD power P_{ec} on fusion performance Q [Reprinted with permission from Sauter *et al.*, Plasma Phys. Controlled Fusion. **52**, 025 002 (2010). Copyright 2010 IOP, Institute of Publishing].

(4) Suppress mode ("Subdue")
(5) Turn off ECCD power
(6) Maintain alignment following mode suppression (i.e., in the absence of a detectable island and with no ECCD turned on)
(7) Repeat from step 1 for subsequent triggering and growth of seed islands

Several approaches to Step 3 (alignment, or "catch") may use direct measurement/reconstruction of the q-profile and calculation of the ECCD deposition location, or apply a convergent search technique based on the island response as an indirect detector of the alignment.[50,51] Which approach (or a combination) is used for ITER will depend on the accuracy and precision provided by profile diagnostics, and the accuracy of real-time calculation of microwave ray paths and deposition location. Step 6 requires sufficiently accurate reconstruction of the q-profile in the absence of an island (although the accuracy of the maintained alignment in this step need not be as great as in Step 3). Even in operating devices, achieving and maintaining good alignment for long periods can be challenging.[51]

A simulation of this scheme applied to an ITER Start-Of-Burn (SOB) plasma is illustrated in Fig. 9. Random seeds are generated with average frequency of 1 Hz, and the launcher mirrors are initially set to be misaligned by $\Delta\rho \sim 0.05$ in normalized minor radius ($\rho \equiv \sqrt{\Phi/\Phi_b}$, where Φ is the toroidal flux within the flux surface and Φ_b is the total toroidal flux within the plasma boundary), corresponding to $\Delta r \sim 10$ cm. Upon detection of island growth beyond the threshold size of 2 cm, the ECCD power is turned on (to a constant value of 12 MW) and alignment is accomplished with dynamic response times based on modeling expected delays and the specified launcher mirror sweep rate corresponding to deposition spot velocity of ~50 cm/s at mid-radius (in this case also assuming direct measurement/calculation of resonant surface and deposition to accomplish alignment). Notice that with the level of noise and drift

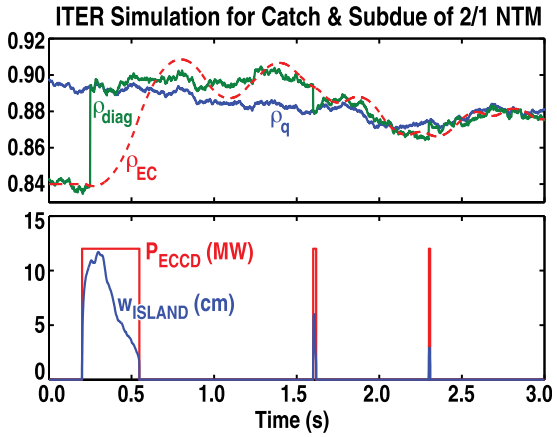

FIG. 9. Simulation of continuous TM control in ITER. (a) locations of actual resonant q-surface (ρ_q), measured q-surface (ρ_{diag}), and ECCD deposition (ρ_{EC}) in normalized radius; (b) ECCD power (P_{ECCD}]) and island width (w_{island}). The first activation of Catch and Subdue suppresses the island and brings the measured location into good alignment. After the mode is suppressed, good alignment is sustained by an algorithm actively tracking the q-surface through a profile reconstruction.

specified in the simulation, the misalignments prior to the 2nd and 3rd seeds are much smaller than the initial value. The islands begin suppressing immediately in these cases, good alignment is achieved more rapidly, and islands are therefore completely suppressed much more quickly (in ~50 ms).

For this kind of control scenario, the resulting average ECCD power P_{av}^{EC} is a simple function of the average seed trigger frequency f_{seed} and the average time T_{EC} required for the ECCD to suppress the mode

$$P_{av}^{EC} = P_0 f_{seed} T_{EC}, \tag{4}$$

where P_0 is the peak (constant) ECCD power. Note that this assumes a suppression time that is less than the time between seeds. If the principal seeds result from sawtooth crashes, the sawtooth period in ITER is predicted to be ~100 s, so that $f_{seed} \sim 0.01\ s^{-1}$.[52] For the ITER control scenario shown in Fig. 9, this average power is ~0.5 MW, with a minimal impact on Q. Interestingly, if a smaller P_0 is used for NTM stabilization (e.g., due to conflicts with other subscribers to the gyrotron resources), it is possible that stabilization will require a longer time and the average power may not be reduced as dramatically (or at all).

Proper function of this continuous control algorithm requires design of key parameters based on the MRE, as well as models of the relevant diagnostic and actuator dynamics. One approach to optimized search algorithms makes direct use of a real-time model of island dynamics based on the MRE, which is effectively inverted to infer the degree of misalignment.[51] The overall algorithm designs thus derived are then tested and optimized in simulations such as that illustrated in Fig. 9.

While this continuous control function is expected—and should be designed to—maintain TM suppression under nominal seeding and plasma regime assumptions, exception events may still result in violation of the robust controllability boundaries. For example, loss of a gyrotron, degradation in profile measurements, or similar resource-denial faults may require an asynchronous change in control action. EH algorithms in this class may need to repurpose other gyrotrons, change the nature of the TM control algorithm itself, change plasma conditions, apply torque to a locking mode, trigger a rapid uncontrolled shutdown, or execute one of various other actions. Many such responses require design and verification of algorithms in the same way as continuous control design, and require effective integration solutions such as actuator management policies for sharing oversubscribed resources (see Sec. IV).

E. Divertor, fast particle, and burn control

Regulation of the kinetic state of the plasma (defined here broadly as the thermal energy and plasma density, in the latter case perhaps local to regions including the core and divertor legs) tightly integrates various control goals and actuators. For example, regulation of the core density is tightly coupled to the simultaneous goal of core radiation regulation,[9] as well as to goals of regulating the burn state

and divertor target heat flux.[53] Similarly, many approaches to regulation of the core energy confinement time will also affect the core particle confinement time and core radiation, in turn affecting both the burn state and divertor heat flux. Regulation of the burn state is of course an entirely unique characteristic of ITER, as the first highly self-heated burning plasma experiment. The need to simultaneously regulate divertor conditions represents a further unique challenge. Although a fully consistent solution for ITER divertor operation with sufficiently low target heat flux has not yet been determined, it is expected that the solution will include active methods to produce sufficient levels of divertor radiation and produce partial detachment of the core plasma from divertor target. Many schemes have been envisioned for regulation of divertor conditions. Controlled variables may include the local radiated power in core and divertor, measured or inferred heat flux to divertor targets, the local electron temperature or density in relevant regions of the divertor legs or x-point, or proxies for the actual degree of detachment such as local ratios between selected impurity radiation lines. Relevant actuators to affect core and divertor conditions include fueling gas and pellets, and impurity gas and pellet injection. Experimental efforts have begun to study potential approaches to regulation of core and divertor states relevant to ITER.[54,55]

An approach which focuses on combined control of core and divertor radiation is illustrated in Fig. 10, showing results of its application in ASDEX Upgrade (AUG). In this scheme, argon gas injection at the outer midplane is used to regulate core radiation, while nitrogen injection in the private flux region is used to regulate divertor target heat flux. Peak divertor heat flux was held below 5 MW/m^2 under active control in a discharge with 23 MW of input power with high P/R, while good confinement ($H_{98(y,2)} = 1$) and high normalized beta ($\beta_N = 3$) were maintained.[54] The algorithm uses the heating power and the main chamber radiation to infer the power flux into the divertor, and a proxy of the divertor temperature to infer the heat flux to the outer target. Figs. 10(e) and 10(f) show good tracking of these quantities (P_{div}, T_{div}) under regulation. Alternatively to argon, krypton has successfully been used as the core radiator in AUG.

Another approach is illustrated by Fig. 11, in which a divertor detachment and radiation control scheme was applied in DIII-D.[55] The algorithm used divertor temperature measurements from real-time Thomson diagnostics and a line ratio measurement to compute the detachment level, along with a real-time bolometer diagnostic to determine core and divertor radiation. This system was used to regulate deuterium and impurity gas injection level to control detachment, power exhaust, and radiation, and study feasibility of the envisioned ITER partial-detachment operation. This control stabilized the detachment front (where the electron temperature drops to a few eV) fixed at the mid distance between the strike point and the X-point throughout the shot.

The fast particle content of ITER self-heated plasmas, including fusion-generated alpha particles, beam injected fast ions, and fast particles accelerated by RF heating systems, implies the possibility of high localized heat loads due to deconfinement of these particles. This deconfinement can result from the action of energetic particle or other plasma instabilities, as well as orbit loss. Experiments on various devices have begun to understand energetic particle modes such as toroidal and reversed shear Alfven Eigenmodes, and to explore possible mechanisms for stabilizing them or otherwise reducing the fast particle deconfinement effects.[56] However, significant research remains to be done to determine whether fast particle deconfinement will pose a problem to ITER, and if so, what control methods might be used to regulate or mitigate the impacts.

The goal of burn control in ITER is to regulate some characteristic of the burn state, such as fusion gain or total alpha power. Burn control serves many purposes, including producing a steady scenario, satisfying the ITER $Q = 10$ milestone, and enabling exploration of the fundamental stability of fusion output power. Principal challenges to burn control include sufficiency of the fueling and heating systems to directly adjust the burn state with short enough delay times, different effectiveness and dynamics of these

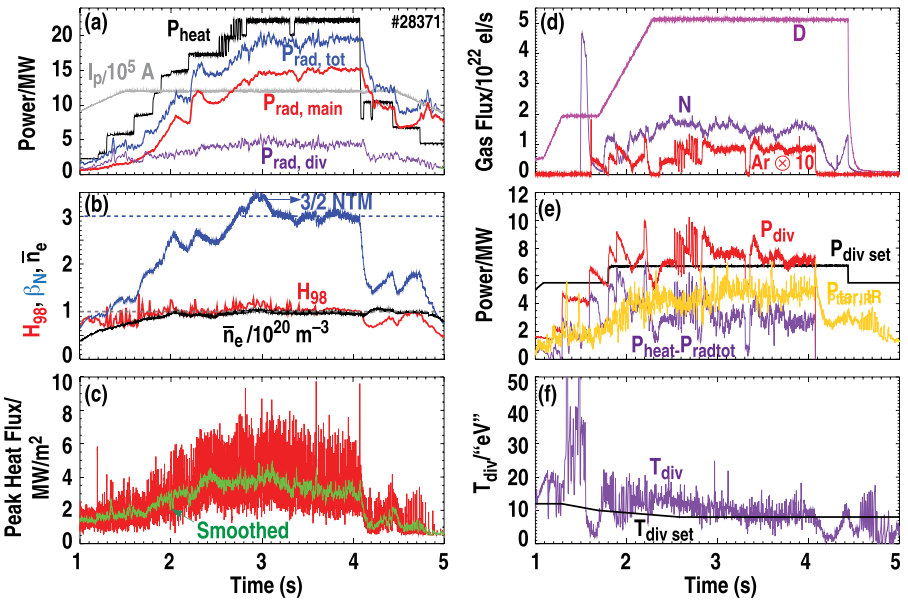

FIG. 10. Time traces of ASDEX-U discharge demonstrating simultaneous control of core radiation (via midplane argon injection), and control of heat flux to the divertor target (via nitrogen injection in the private flux region) [Reprinted with permission from Kallenbach *et al.*, Nucl. Fusion **52**, 122 003 (2012). Copyright 2012 International Atomic Energy Agency, Institute of Publishing].

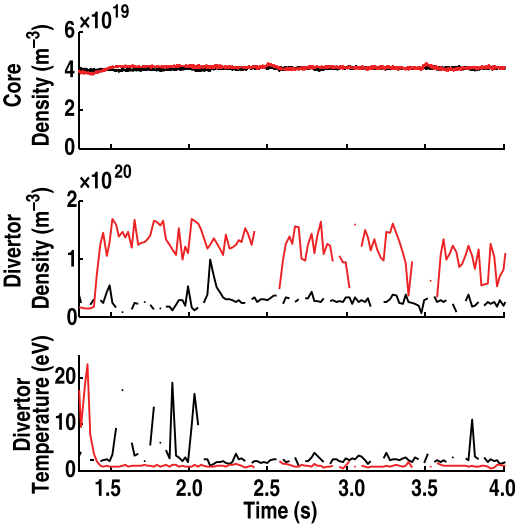

FIG. 11. Illustration of divertor detachment control in DIII-D. Red–detachment feedback control on (#153816). Black–detachment control off (no divertor fueling #153814). Top: Line average core density Middle: Divertor temperature measured by divertor Thomson. Bottom: Divertor density measured by divertor Thomson [Adapted from Fig. 5, Elsevier and Kolemen *et al.*[55]].

actuators depending on whether fusion power is being augmented or reduced, and the complex impacts of many other control functions on the burn state. While several theoretical studies have been done on approaches relevant to burn control in ITER (e.g., Refs. 57–59), these remain relatively untested experimentally. Although operating devices can simulate many aspects of burn control using subsets of their heating systems as proxies for alpha-particle heat deposition, no consistent and effective solution has been demonstrated even in this experimental simulation mode. However, some limited experimental studies have begun, showing promise for use of novel actuators such as the in-vessel RMP coils for modification of edge transport characteristics.[60]

F. EH physics and disruption mitigation control

In order to minimize the impact of off-normal events and faults, while maximizing the physics productivity of discharges, ITER must have a highly reliable integrated system to forecast, detect, and respond to these events. ITER disruption tolerance analysis implies that unmitigated major disruption rates can begin in the range of 20%–30% of all discharges in early H/He operations, and must steadily drop to ~5% of discharges by the time high performance DT operations begin. Correspondingly, disruption prediction success rates can begin in the 60%–75% range in early operations, and must improve to ~95% reliability in DT operation. Vertical displacement events (VDEs) are much more strongly constrained, permitting only 1–10 unmitigated events at the highest current and stored energy (depending on resulting halo current fraction and forces).[61] Approaches will be required to reduce the rate of disruption to very low levels, and to mitigate damaging effects where prediction, detection, and response is successful. High reliability control that includes effective avoidance of stability boundaries is expected to largely accomplish the first requirement. However, both satisfying that requirement, and fulfilling the ability to apply effective damage mitigation strategies, depend on effective responses to exceptions (off-normal events requiring a change in control action). The responses which implement the change in control action fall in the category of "EH." While the general requirement of recognizing exceptions and responding to them is not unique to ITER, the level of complexity, effectiveness, and reliability required for ITER far exceeds that of any fault response system on operating devices.

Examples of possible exceptions include failure of a magnetic probe, loss of a gyrotron, growth of a TM beyond some threshold island size, or *prediction* that a power supply limit will be reached in 5 s with the present planned scenario trajectory. Thus, exceptions can be faults that are either detected or predicted, and can include system faults or plasma events of many kinds. The definition of an exception includes identification of a signal or combination of signals that reflect the event, as well as specification of the context under which the exception is active. For example, loss of a gyrotron may be reported by an availability signal provided by the ECH system, but this may only constitute an exception if the gyrotron is presently (or in a specified future period) subscribed for use in the scenario.

EH begins with the definition of relevant exceptions, coupled with response policies. Response policies may consist of simple actions (e.g., setting of a status flag or sending of a signal to the CIS), scenarios (sequences of steps to take and command waveforms to generate), and algorithms needed to execute those scenarios (see also Sec. IV D). The EH control process consists of either detection or prediction of a defined event, followed by applying the control needed to produce the desired response. Because these response policies have all of the characteristics of any other control function, significant control physics understanding is needed before effective ITER solutions are possible. For example, in order to identify or predict the crossing of a controllability or stability boundary to trigger an exception response, sufficient understanding of the parameters governing these phenomena must be achieved, enabling quantitative mapping of relevant parameters to stability boundaries. This understanding must also facilitate reliable, rapidly executed computational tools that can assess stability boundary proximity in real time based on real-time equilibrium reconstruction data, MHD spectroscopy, and other available measurements. Similarly, a real-time computational process that can assess the risk of uncontrolled growth of a potentially *disruptive* instability is critical to enable timely and effective execution of disruption mitigation responses (see also Sec. IV D). For example, prediction of potentially disruptive growth of a TM could trigger an exception response in which rotating magnetic fields might be applied by the in-vessel coil array to apply torque and entrain the plasma.[62] The level of physics understanding, accuracy of predictive algorithms, and complexity of EH control algorithms that may be needed in order to guarantee sufficiently effective responses is presently unique to ITER.

Disruptions pose a particular challenge to ITER owing to the relatively small number of high force and thermal load events for which the machine is designed. Prediction of an

unavoidable high-risk disruptive state with sufficient look-ahead time can greatly enhance the effectiveness of mitigation action by enabling preparatory and post-mitigation control responses. For example, given several seconds of warning, an ITER plasma could be reduced in elongation and limited at an appropriate position to prepare for massive impurity injection. Such preparatory actions could minimize vertical motion of both thermal plasma and a runaway current channel following the thermal quench, and could enable sustained control of a runaway channel while it is gradually damped by carefully regulated impurity injection or other means. Scenarios such as these involving some control action before, during, and following a (mitigated) disruption have been explored in existing devices (e.g., Refs. 63 and 64), but much of the relevant physics is not yet well understood, and no full control solution (particularly for preparatory action before disruption, and runaway management following the thermal quench) has yet been identified.

IV. SELECTED CONTROL MATHEMATICS CHALLENGES

A. Control mathematics

The term "control mathematics" refers here to the field and activities concerned with determining the mathematical algorithms needed to meet control performance requirements, including the logical structure, decision elements, architecture, and integration issues. Typically, the control mathematics design and analysis effort takes a strong role once sufficient physics understanding and control-level models have been developed. In the control development process outlined in Fig. 1, there is some overlap in control physics and mathematics efforts at the point of scheme development, since the choice of scheme characteristics must be informed by both physics constraints and performance requirements. Control mathematics must inform how controlled variables are to be regulated, and how best to achieve the control performance specifications. Algorithmic solutions must be identified to provide the required performance, including robustness to noise, disturbances, and uncertainties. It should be noted that particular control design techniques are not dealt with in detail here. The applications of specific techniques themselves typically do not significantly distinguish ITER control from that of other fusion devices, or from other fields of control beyond fusion.

Mathematical control approaches apply to both continuous control (often resulting in linear multivariable control algorithms) and EH control (often resulting in nonlinear or logic-based algorithms). ITER will require systematic application of control design techniques with control level models to produce quantified levels of reliability in order to satisfy its nuclear regulation requirements. This approach will also maximize scientific output of the limited number of discharges and greatly help elucidate understanding of relevant physics phenomena by enabling isolation of specific variables.

B. Robust control and operation

Perhaps, the largest gap in control accomplished in today's tokamaks and the requirements of ITER lies in the robustness of sustained control demanded by ITER. Varying machine conditions, operator error in setting up discharges, insufficient actuator and diagnostic capability, physics model uncertainty, and hardware faults are among the effects that can challenge the robustness of control performance shot-to-shot or within a discharge. Experimental programs today recognize that the effort required to achieve ITER levels of robustness would require reducing the level of resources applied to advancing physics understanding, and therefore do not pursue high robustness in their own operations. ITER limits on total number of discharges, actuator capability, etc., require higher knowledge output per shot, and therefore require greater reliability in operation. More stringent design limits on disruption frequency, wall loads, divertor erosion, etc., also require higher control robustness than present devices.

Achievement of quantifiably robust operational control requires systematic design and verification as described in Fig. 1, using models with quantified levels of uncertainty, along with quantified disturbances and noise levels. The need for physics characterization of these quantities is thus a relatively unique requirement of ITER. Quantifying and guaranteeing performance for a given control category requires specification of a relevant metric. For example, vertical control performance in ITER can be characterized by the maximum vertical displacement (denoted ΔZ_{MAX}) that the control system is capable of restoring. Once a theoretical performance metric such as ΔZ_{MAX} is identified, the performance requirement for the metric must be specified. Often such requirements arise from complex characteristics of the control that depend on engineering details of a given machine. In the case of vertical control, the ΔZ_{MAX} capability of the system is set by the level and spectrum of noise and amplitude of disturbances expected to occur. In the absence of a predictive capability for the relevant ITER noise environment, statistical studies of operating machines have been used to establish an empirical requirement of $\Delta Z_{MAX} > 0.1a$ for sufficiently robust control to prevent noise and disturbance-driven VDEs.[65] The ITER in-vessel vertical control coil (the VS3 circuit) has been designed to satisfy this robustness criterion. A more general robustness requirement takes the form

$$\frac{\Delta Z_{MAX}}{\Delta Z_{PERT}} > C_{Zrob}, \tag{5}$$

where ΔZ_{PERT} represents the maximum perturbation amplitude produced by noise and disturbances, including such phenomena as rapid changes in confinement, transient locked modes, and ELMs. C_{Zrob} is a robustness coefficient for vertical control, defining the degree of robustness needed beyond marginal controllability. While typically $C_{Zrob} \sim 2$–3 in operating devices for sufficient (empirical) robustness, each of these quantities require further study to fully quantify ITER requirements in terms of perturbation sources.

ITER will require identification and quantification of performance requirements for such metrics in many control categories, particularly those involving stabilization of potentially disruptive instabilities, such as the TM control discussed previously in Sec. III D. Both identification of physics-based metrics and quantification of robust performance requirements require significant ongoing and future research. In the case of TMs, a key metric for robust stability is the amplitude of seed island perturbation produced by disturbances such as ELMs, sawteeth, and fast particle instabilities. Island dynamics governed by the MRE imply that a useful criterion for robustness to seed island perturbations may take the form of

$$\frac{\Delta W_{MARG}}{\Delta W_{PERT}} > C_{Wrob}, \tag{6}$$

where ΔW_{MARG} is the marginal island size (the island size above which the island growth rate becomes positive), ΔW_{PERT} is the maximum seed island amplitude produced by the most perturbing disturbance or noise effect, and C_{Wrob} is a robustness coefficient for TM control, defining the degree of robustness required beyond marginal controllability. Further research remains to quantify the coupling between various sources of seed islands and the amplitude of the seed produced, as well as to quantify the robustness coefficient needed for non-disruptive ITER operation. Research is also needed to relate the physics of "seedless" island triggering to corresponding controllability or stability robustness metrics.[20,66]

A powerful tool for mapping performance metric values into operating regimes is the Control Operating Space (COS), an example of which is shown in Fig. 12. The COS plots the amplitude of a relevant robustness metric (in this case ΔZ_{MAX}) against a quantity reflecting the degree of instability and/or operating regime (in this case parameterized by the internal inductance, which correlates with the vertical growth rate at fixed elongation, beta, and plasma current). This form of the COS thus shows the ΔZ_{MAX} *capability* required by various operating regimes mapped to ranges of internal inductance or growth rate. Higher capabilities in

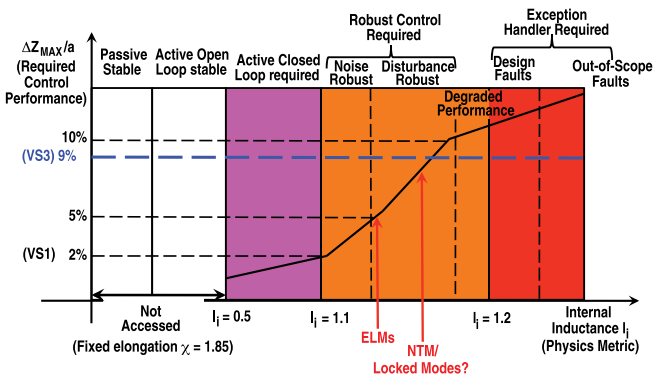

FIG. 12. The COS quantifies requirements for robust performance in terms of a control performance metric (here exemplified by ΔZ_{MAX}) as a function of physics operating parameters (here exemplified by internal inductance), and identifies regions in operating space by degree of stability or robustness expected.

ΔZ_{MAX} required by disturbances (e.g., ELM's and NTM growth) that apply a perturbation to the vertical position are indicated notionally. The COS enables identification and specification of operating regimes requiring different levels of robustness, as well as different approaches to control. ITER will generally design for a sufficient level of robustness for a range of expected disturbances and simple faults (provided by Continuous Control algorithms), as well as more unlikely, unexpected, and complex faults (which may trigger EH responses), while providing some additional margin for completely unforeseen, "out-of-design" faults (largely provided by EH algorithms).

Another important way the COS enables robust control performance is through real-time monitoring and regulation of controllability metrics (i.e., the present operating point in the COS). In many cases, specific algorithms can be run to maintain sufficient distance from a controllability boundary to ensure low likelihood of loss of control (in the absence of a major fault). For example, real-time assessment of Eqs. (5) and (6) would allow identification of the relative degree of controllability for vertical stability and TMs, respectively, to be compared with appropriate administrative robustness limits. Specific algorithms can be designed and operated to regulate the distance from controllability boundaries as a control goal by adjusting selected profile characteristics, plasma beta, etc. Such control goals are distinct from nominal scenario and active stability control and can often run in parallel with these algorithms. For example, nominal scenario profile control may seek to match a target q-profile on average, while a controllability regulation algorithm may seek to simultaneously maintain a local current density gradient or a global stability parameter such as Δ' beyond a specified level. In addition to these continuous control functions, real-time prediction of the trajectory in the COS can enable early asynchronous (EH) action to be taken to prevent impending loss of control (see Sec. IV D).

An essential and highly unique element of ITER robust control procedures is the required use of discharge verification tools in routine experimental operation. Shown as the final step before an experiment in the systematic design process of Fig. 1, each discharge must be validated and verified prior to experimental execution. In this process, simulations will be used to confirm that all administrative and performance limits will be avoided under the given pulse schedule, and that all programmed control characteristics will provide adequate performance even in the presence of expected noise, disturbances, and selected exceptions. The exact forms of the simulations used to develop discharges, and to certify them as verified prior to execution, have yet to be specified. However, the general architecture and functional requirements of such simulations has been determined in ITER design activities.[67] Figure 13 shows the functional elements of such a simulation. The core of the system is a tokamak plant simulator, which can connect to the actual ITER PCS hardware (or a software simulation thereof) to perform and verify a simulated discharge. Modules are also provided to enable simulation of off-normal events that are likely to challenge the control performance. Verification of control performance will require demonstrating both robust

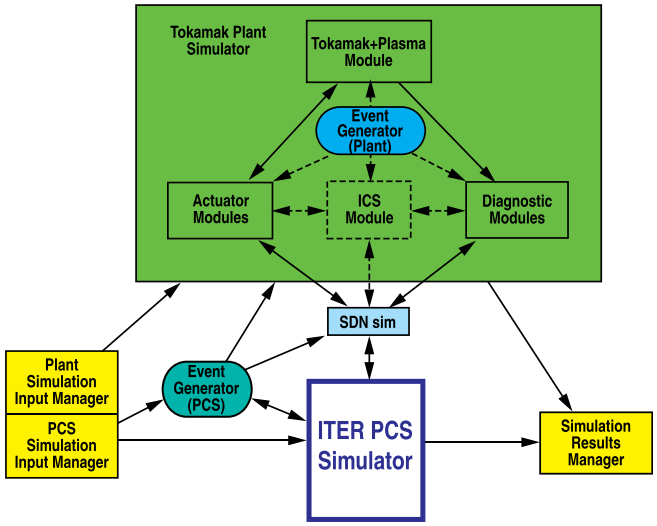

FIG. 13. Discharge verification and control performance qualification is accomplished using a simulation that can connect to and test the actual ITER PCS. Such simulations will verify control performance by triggering key relevant events that challenge both continuous and EH control.

continuous control in the presence of disturbance events, and effective responses to likely exceptions.

C. Model-based control: Design and state estimation

A critical and significantly novel aspect of ITER plasma control is its extensive reliance on model-based design in order to achieve the necessary high level of confidence in control performance (see the third design step in Fig. 1). While scenarios and control represent the "observables" of fusion science, the tangible outputs that enable understanding to become reality, the principal way in which physics understanding interfaces with control, is through control-level models. Sufficiently accurate predictive models enable design of control algorithms that can guarantee quantifiable levels of performance under defined conditions. Fundamentally, a model appropriate for control design describes a plasma or system response to inputs such as voltage applied to conductors, or current drive applied to a plasma, sufficient to enable determination of dynamic algorithms to guarantee the required level of performance. Such models often include internal dynamics that are either self-driven or can be excited by actuator inputs. A key implication of ITER control requirements is that all control areas require models that are both *sufficiently* accurate, and include the dynamic responses needed to produce satisfactory control performance. These are characteristics that often distinguish ordinary physics understanding goals from control physics goals. Sufficient accuracy implies that models should be as accurate (and as complex) as needed, and no more so. In most cases control performance is a weak function of model accuracy, since a key feature of feedback loops is their ability to *compensate* for inaccuracies.

Models used for plasma control design thus range from first order differential equations in one or many variables, to highly nonlinear multivariable systems, to mixes of discrete, continuous, and logical mathematical operations. An extremely common and powerful representation used for a wide range of linear multivariable models is the state space representation, already encountered in Eq. (2) above. A rich array of design approaches has been developed in the field of multivariable control for linear first order matrix equation models of this form, making it a highly desirable representation for plasma control models as well.

Black, gray, and white box models have already been discussed in the context of current profile control (Sec. III C). Each has virtues and issues for different applications, so that the appropriate model approach must be carefully determined for a given control problem. A completely black box model, identified through some training or fitting process based entirely on experimental data, avoids the need for actual understanding of the relevant control physics processes. However, this comes at the cost of requiring enough *a priori* experimental data to calculate and verify a sufficiently accurate model. Potential limitations on ITER discharge time available for both collecting data and testing controllers may limit the usefulness of such approaches. Fully physics-based ("white box") models are therefore preferable where sufficient understanding exists and sufficient accuracy can be provided. "Gray box" solutions intermediate between these extremes may find substantial application in ITER by determining key parameters through limited experimental experience or even real-time measurement.

One possible approach to combining the virtues of physics-based and black box models is to produce training data with a fully physics-based simulation.[29] The "black box" models thus produced can be structured in state space or other convenient form and have arbitrary internal states as well as dimensionality driven by the simulation data, while still being based on a physics model. Such models could be produced without any experimental data, and thus satisfy ITER limitations on discharge time for control development.

Model-based design has been used successfully in many demonstrations on existing devices, and is used routinely in some areas of plasma control. Axisymmetric magnetic (equilibrium) control and vertical stability are areas in which the approach has been well-developed and applied effectively in many devices. Figure 14 illustrates the typical plasma response model used for design of such control algorithms: a linear equilibrium response to variation in an outboard PF coil circuit in the JET tokamak.[68] The active coils in this circuit (shown in red and green) apply a radial field, resulting in a nonrigid but largely vertical displacement of the plasma (indicated by the orange band). Linear axisymmetric plasma response models such as this exemplify a large and powerful class of such models in which static (memoryless) plasma response functions are coupled to conductor circuit dynamics (with memory) to produce accurate dynamic system descriptions. The complete model describing the response of the plasma to all PF coil and passive conductor currents enabled design of both the JET XSC, responsible for dynamic regulation of the plasma boundary, and JET vertical stability controllers.[69] Similarly, the highly nonlinear MRE (Sec. III D; Eq. (3)) and plasma transport equations used for profile control (Sec. III C) have been applied in design of many control schemes to produce robust algorithms. Although such physics phenomena are frequently difficult to describe in

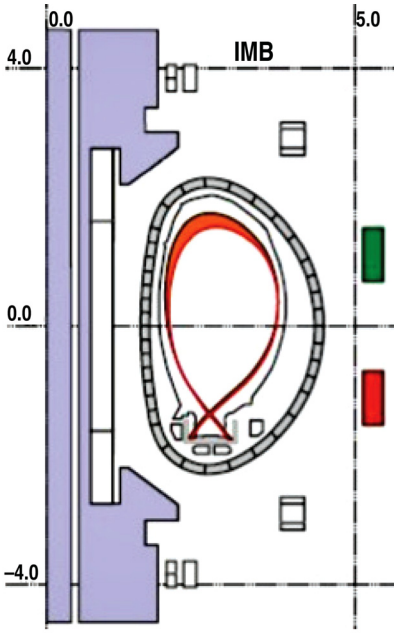

FIG. 14. The JET XSC is an example of mature model-based algorithm design, making use of relatively simple linear axisymmetric plasma response models calculated from MHD force balance [Reprinted with permission from Sartori *et al*., IEEE Control Syst. Mag. **26**(2), 68 (2006)]. Models for control design should be as complex as needed to be sufficiently accurate, but no more so.

linearized form while retaining sufficient accuracy for algorithm design, they can often be described by appropriate nonlinear and complex but still control-level descriptions.

Since control models describe responses of plasmas or other systems to inputs (including internal, self-driven dynamics), in a sense the goal of control design is to "invert" this mapping to determine the inputs required to produce the desired response. Advanced control science provides a rich array of design solutions to accomplish this "inversion" while satisfying sophisticated performance goals, even for nonlinear and complex models. Some examples of broad design approaches that have been effectively applied to various plasma control problems include

- Linear Quadratic Gaussian design,[70,71] in which the controller typically includes a linear state observer and a feedback regulator based on the states, and a cost functional is minimized by the control action. Typical cost functionals involve parameters such as measured errors, actuator power, or closed loop response to noise and disturbances;
- Multivariable pole placement[72,73] or loop shaping[74,75] design, in which desired dynamic characteristics of the closed loop system are explicitly built into the design process;
- Robust techniques (e.g., H-infinity optimal, mu-synthesis[74,76]), in which characterized plant uncertainties are explicitly included in the design process, and controllers are guaranteed to perform robustly to the characterized level of uncertainty;
- Extremum Seeking,[77,78] in which the controller is designed to follow a trajectory to find a specified extremum in relevant physics or other performance parameters. It should be noted that a feature of the Extremum Seeking approach is that it does not in principle require a system model for its basic design, although such a model is still useful to confirm adequate performance in simulation prior to deployment.
- Model predictive control,[79,80] in which a model of the controlled system or process is used in real time to predict its future evolution and calculate an optimized control response. The optimization is carried out over a finite time period into the future, and can readily include constraints.

The extensive use of model-based control is a strong driver for control physics research, which must be informed by the needs of control, just as control solutions are in turn informed by physics understanding. Control solutions and implementations are not only demanding in terms of mathematical sophistication and understanding. They require specific and sufficiently accurate physics understanding to enable transformation from knowledge to application.

Examples of areas requiring significant and specific research to provide sufficient models for ITER control design include TM dynamics; divertor radiation and detachment responses to core, divertor, and scrapeoff layer conditions; ELM stability and suppression response to applied 3D fields; sawtooth stability and triggering response to applied RF heating; fast particle instabilities and responses to RF and applied nonaxisymmetric fields; plasma responses to all relevant heating and current drive systems, including the effects of intrinsic rotation. Each of these areas requires determination of associated disturbance effects and couplings to other control physics. For example, the amplitude of seed islands driven by sawteeth remains a critical specification to inform TM control algorithms. The ability of the in-vessel nonaxisymmetric coils to apply torque to the plasma and produce rotation, as well as the effect of low rotation on TM stability, is similarly critical to inform design of EH algorithms that might seek to spin up locked modes and minimize disruptivity.

The use of control models is not restricted to controller design. Another important application is that of real-time (plasma and tokamak system) state estimation, in which real-time diagnostic measurements are combined with a-priori knowledge about the expected plasma/system evolution from a model. This estimation function, often called a "state observer," can be accomplished using linear algorithms such as the Kalman filter and its variants, or with more complex, nonlinear models.[81] Given a previous state estimate and present diagnostic measurements, a Kalman filter uses a dynamic model of the system response to provide a real-time estimation of the present state and (in discrete form) a one-step-ahead prediction of the next state and predicted diagnostic measurements. When the true diagnostic measurements become available, the difference between predicted and true measurements is added as an *a posteriori* update to the model-based state estimate. A key advantage of using a state observer is that control algorithms no longer rely on individual diagnostic measurements to derive the present value of the controlled variables. Instead, all available information is distilled by the model-based observer into one unique state estimate, which distributes the relevant quantities to each

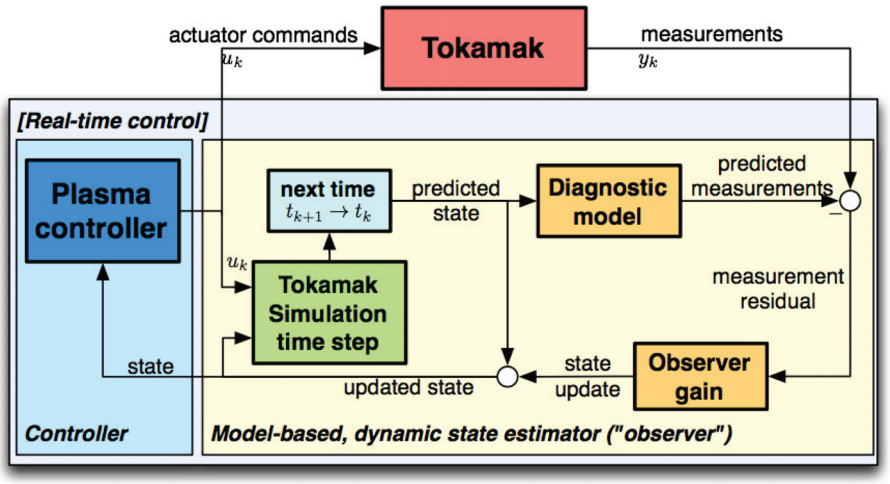

FIG. 15. Illustration of a complex model-based real-time state observer. A tokamak simulator executes in real-time in parallel to the physical evolution of the plasma. Differences between the expected and true diagnostic signals are fed back as a correction to the state estimate. Unexpectedly large values of the measurement residual can be used to detect real-time diagnostic faults.

control algorithm. Faults in diagnostics can be detected and handled in a variety of ways, either by excluding the diagnostic from the state reconstruction, switching to a backup signal, or signaling an exception to the event handling system. This helps alleviate reliability requirements on real-time diagnostics, since the state estimate, and hence the control, can be made more robust to diagnostic failures.

Implementation of a state observer can include a real-time capable simulator of relatively complex plasma dynamics of interest. Figure 15 illustrates the use of such a real-time simulator providing state estimation as part of the control loop. An example of such an application is the RAPTOR code, used as such a real-time simulator for plasma profile evolution on TCV and ASDEX-Upgrade.[30,82] The code receives real-time information from equilibrium reconstruction as well as several profile diagnostics, and provides a consistent real-time estimate of the q profile as well as various kinetic profiles (see also Sec. IV D following). This use of control models for state estimation also finds important application in exception forecasting and handling in ITER.

D. Exception forecasting, detection, and handling

Exceptions are off-normal events that require some change in control action in order to maximize physics productivity of a discharge or to prevent disruptions and other adverse impacts to ITER. An effective EH system is thus necessary to successful ITER operation. Such events include faults ranging from signal degradation in a diagnostic, to loss of a heating resource, to growth of a potentially disruptive instability. An exception can also result from a predicted impending event.

An exception to be responded to, and the form of the response are defined by an EH policy, which includes:[83]

- An exception identification rule and *detection* algorithm that define the exception. Specific measurements need to be monitored for detection or prediction (e.g., wall or divertor target temperature, diagnostic signal quality value, predicted profile stability)
- A *decision* algorithm defining the condition that triggers a response. Apart from the exception occurrence, this condition also comprises the context under which it is relevant (e.g., the present discharge phase, plasma confinement regime, stored energy level, or value of Ip)
- A *response* algorithm or scenario, possibly including specified signals to alert other subsystems (e.g., change of control gains, execution of pre-calculated rapid shutdown scenario, signals to CIS)

Thus, design of the EH system will require identifying relevant contexts, or machine and plasma states for key exceptions, along with the parameters and conditions that define each exception. Finite state machines may be used to track the relevant machine and plasma states, as well as previous EH decisions (Fig. 16). Normal states correspond to those programmed in the nominal pulse schedule. Alternate states correspond to those an exception response policy could transition to if nominal states become unachievable or pose too high a risk. Recovery states represent temporary states the plasma may be moved to with potential to return to the nominal pulse schedule. However, it is important to note that a simple enumeration of all possible relevant states and transitions that may occur can rapidly lead to the explosion illustrated in Fig. 17, which shows a series of possible states and decision points responding to various exception examples, such as NTM onset or loss of vertical controllability. One possible way to avoid this is to define policies with efficient and general decision algorithms that collapse a large space of conditions and responses such as that shown in the figure to a small set of (more complex) functions. Thus, many different exceptions would lead to a much smaller set

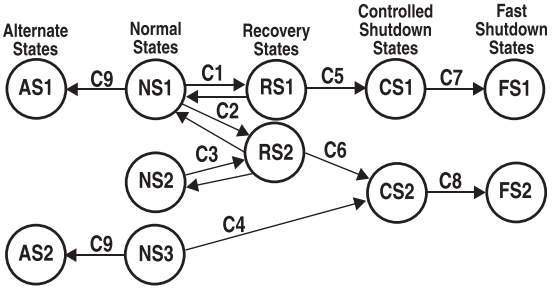

FIG. 16. Candidate finite state machine structure as part of an integrated pulse schedule and EH system.

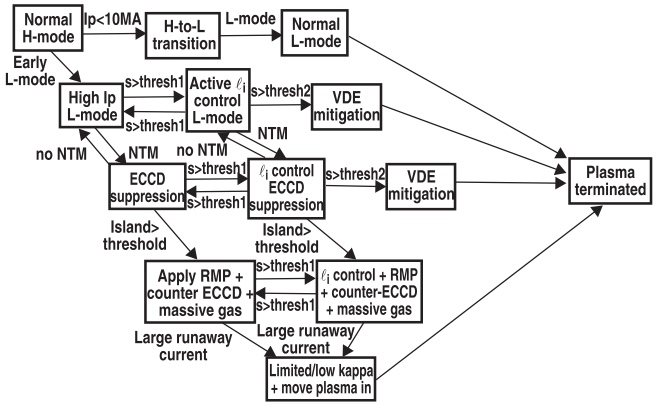

FIG. 17. What EH should not look like: poor choices in exception definition and finely detailed enumeration of states can lead to an intractable explosion of branching options.

of responses, each of which could be tractably validated under varying simulated conditions. An obvious example is the small set of plasma shutdown scenarios. Rapid controlled shutdown will likely be defined as a regulated rampdown of all coils, executed as rapidly as possible within the constraints of removing plasma thermal and magnetic energy without disruption. The specific trajectories of PF coil currents, fueling gas, impurity injection, heating, etc., required to execute the shutdown must be calculated in real time as a function of the discharge phase and conditions from which such a shutdown is commanded. However, a single EH response function incorporating this real-time calculation may be specified and verified in simulation.

Because EH fundamentally produces a change in control response, the EH system includes both scenarios and model-based algorithm designs, in the same way as continuous control. Specialized and unique control algorithms will be required for certain instances (e.g., position regulation with controlled deconfinement or damping of a runaway beam), while some exceptions may use continuous control algorithms in a different way (e.g., changing control gains in the same algorithm used for continuous control).

Although a detailed architecture and algorithmic solutions for most policies in the EH system have not yet been determined, a functional architecture in the PCS has been designed. As illustrated in Fig. 18, many event detection and response policy execution functions will reside in a Pulse Supervision layer. However, the philosophy of this design is to handle exceptions with local control whenever possible, only escalating to the supervisory level when the exception cannot be handled locally, or when local actions have been taken but failed to resolve it.[83] Thus, EH functions will also reside in a distributed form throughout the ITER PCS and plant systems. For example, failure of some sensors (e.g., magnetic probes) may be managed locally at the responsible plant system through replacement by redundant signals. Similarly, failure of some actuators, e.g., gyrotrons, may be managed locally at the gyrotron plant system through a similar replacement with available tubes.

The EH system will rely on several support functions to provide real-time complex analysis of diagnostic measurements. The global service functions of real-time equilibrium

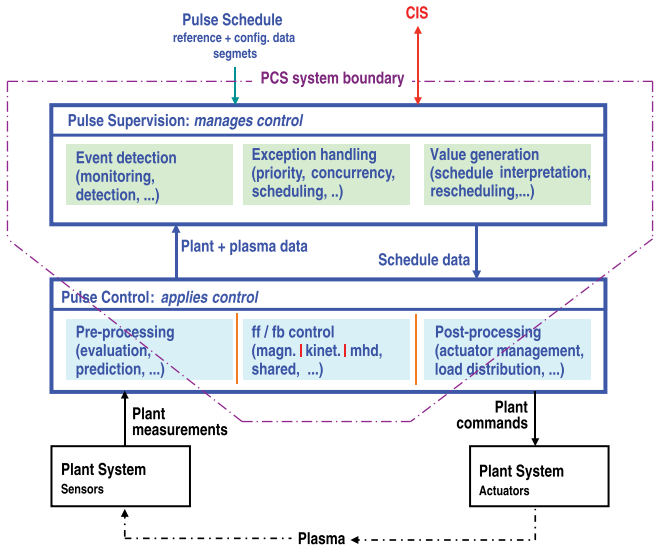

FIG. 18. ITER PCS architecture highlighting the Pulse Supervision layer and EH components. EH functions will also reside in a distributed form throughout the ITER PCS and plant systems.

reconstruction and plasma state observer (see Sec. IV C) provided to the entire PCS are one example of this. A key function of the EH system will involve monitoring key controllability metrics to enable EH response functions to prevent a fault-driven violation of a control boundary (see Sec. IV B, e.g., Eqs. (5) and (6)). Support functions will also be provided to do real-time calculation of these metrics.

A facility unique to the EH system is the plasma and system state forecasting system (FS), which will enable sufficiently early prediction of an impending violation of a controllability boundary, or an unavoidable disruptive state. Sufficient lead-time must be provided by this forecasting function to enable effective action to prevent or mitigate disruptions. For example, triggering of massive gas injection is estimated to require a lead-time of ~10 ms in ITER to allow time for impurity gas to reach the plasma and produce a thermal quench.[84] If modifications to the equilibrium are also required to prepare for such a rapid shutdown (e.g., reducing elongation, or positioning the plasma near the vertical displacement neutral point[85] or at an armored limiter surface), the necessary lead-time lengthens to several 100 ms. The longest look-ahead time expected for predictive simulation is related to resistive evolution of either current profile regions local to a resonant surface, or a tearing island, likely whichever is longer. This ranges from several 100s of ms to perhaps a few seconds in ITER high performance discharges. This implies the need to apply sophisticated prediction algorithms to monitored tokamak systems and plasma evolution data in order to forecast impending threats to sustained operation. Forecasting capabilities of the ITER PCS will likely include direct projection based on algorithms such as neural nets or extrapolation in time with a linear response model, and actual faster-than-real-time simulation (FRTS). Plasma state, predicted performance and health for plant systems, proximity to stability and controllability boundaries, and disruption probability are among the important outputs of the system. It will accept inputs including all relevant signals for

monitoring key hardware systems, present equilibrium reconstructions, and plasma measurements relevant to determining stability. The FS is envisioned as being implemented as a general Support Function with functionality illustrated in Fig. 19.

The FS is a novel feature of the ITER PCS, which distinguishes it from present-day PCSs. It is an area where some R&D is still required to fully finalize the functional specification. At the present stage, only generic functional blocks have been identified where the detailed scope still needs to be defined. The final scope will also depend on the availability of suitably validated faster than real-time control-level physics codes. Note that the FS is not presently envisioned to evolve FRTS simulations by simulating the real-time feedback operations of the PCS in detail. In general, reasonable following of pulse schedule waveforms will likely be assumed with some account taken for simple dynamic lags and errors due to control action. However, these simplified control loops are essential to forecasting several key plasma phenomena, including kinetic quantities and profile evolution. For example, the effects of resource limitations must be taken into account in FRTS execution to predict the consequences of such asynchronous events.

Experimental progress toward this FRTS/FS functionality includes the RAPTOR code, presently used for real-time plasma state estimation at ASDEX-Upgrade and TCV (see Sec. IV C).[30] The use of RAPTOR and its role in the TCV control loop is shown in Fig. 20. The algorithm uses concurrent diagnostic measurements to constrain a real-time simulation of current profile evolution and can modify modeled plasma characteristics to match experimental conditions in real-time as well. Given the short time scales of pulses in these devices, RAPTOR currently runs in real-time. However, the present implementation of RAPTOR is already capable of FRTS calculation on ITER pulse timescales.

In addition to forecasting simulation, the ITER FS must include real-time analysis of both present and projected profiles to determine evolution of proximity to stability and controllability boundaries. The key challenge in ITER will be achieving sufficiently accurate and high-resolution equilibrium reconstruction and profile projection for accurate stability calculation, while providing such accurate converged stability and controllability computations with sufficient look-ahead time to enable effective EH control action.

E. Actuator management

Driven primarily by the large cost of the device and consequent constraints on resources, actuators in ITER (Figure 2) must be actively managed in an extremely efficient way.[9] While all tokamaks must manage actuators in the basic sense of distributing commands to multiple systems for a common goal, the degree to which ITER must perform SAM is significantly different from operating devices. The need for a separate SAM function arises when multiple control algorithms that are not designed in a fully integrated way may seek to command the same resource at the same (or nearly the same) time.

Two broad kinds of actuator resource sharing can be identified in control of ITER: *Simultaneous Multiple Mission (SMM) sharing* refers to simultaneous use of a common resource by two or more control algorithms. Examples of SMM include simultaneous use of PF coils for shaping, vertical stability, and divertor regulation, or simultaneous use of RMP coils for ELM suppression and trimming of residual error field components (i.e., spectral content not dealt with by the superconducting error field correction coils). *Repurposing* (RP) *sharing* refers to rapid hand-off of actuator use from one algorithm to another, perhaps rapidly enough to challenge limits on RP speed. The use of two or more actuators for a *common* control goal, even if involving more than one control algorithm, is not considered actuator sharing and will not be discussed here. There are multiple types of RP. A very common case involves RP of actuators when switching from one scenario or one identified shot phase to another in the pulse schedule. For example, after plasma breakdown, the PF coils will be converted from an objective of producing a good field null with large loop voltage to the role of providing boundary control and regulating

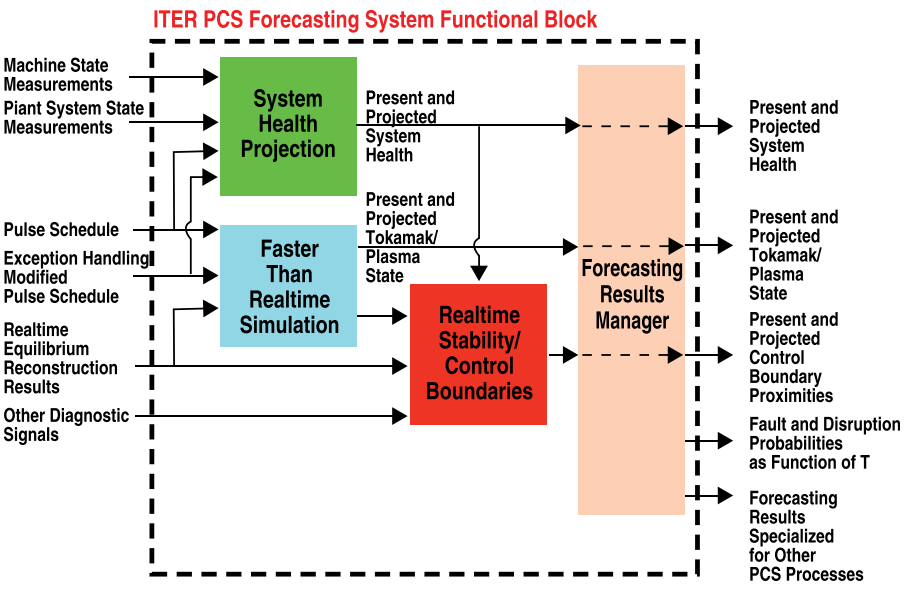

FIG. 19. Principal components of the ITER FS. The FS projects system evolution forward in time with FRTS, and also explicitly evaluates proximity to stability and controllability boundaries.

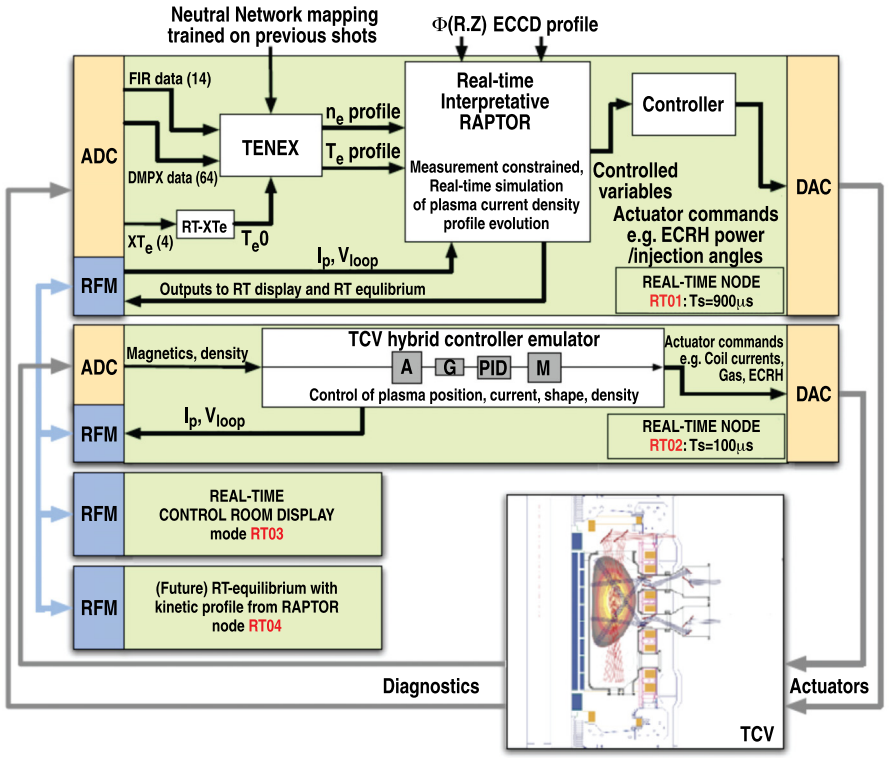

FIG. 20. Implementation of the real-time RAPTOR code in the distributed TCV digital control system. Although RAPTOR adapts to real-time measurements and operates as a concurrent-with-real-time simulation, the algorithm is extendable to FRTS execution using real-time measurements to project future kinetic evolution [Reprinted with permission from Felici *et al.*, Nucl. Fusion **51**, 083 052 (2011). Copyright 2011 The International Atomic Energy Agency, Institute of Publishing].

I_p, the EC system will transition from pre-ionization to heating, and the fuelling and pumping systems will transition from neutral particle control to plasma density control. This is a *planned* RP corresponding to a planned change in scenario triggered by a Pulse Schedule action, even if the time of this transition might not be known in advance. There is a type of *unplanned* but *expected* RP in which an event is expected to occur occasionally, e.g., a growing TM expected in a high beta plasma regime, at which point EC is repurposed until the NTM is suppressed, then returned to the previous duty (part of "Catch and Subdue" scheme, see Sec. III D). Then, there is a RP that is both *unplanned* and *unexpected*, e.g., dealing with a system fault or impending disruption. The unplanned but expected case is typically considered part of the "continuous" pulse control mission, while the unplanned and unexpected is considered part of EH. Because the nature of RP depends critically on the timescales involved, models used to design and simulate these functions must contain good descriptions of RP dynamic characteristics and speed limitations. Similarly, real-time decision-making algorithms to accomplish RP must incorporate real-time knowledge of the consequences of such limitations. Thus, if an EC launcher mirror will be incapable of reaching a repurposed target in time to affect a potentially disruptive TM, the RP will likely be aborted and an exception generated to otherwise deal with the state (e.g., discharge termination).

Resolution of potential conflict in demands on actuator resources generally requires three operations: assessment of resource sharing criteria (which must be supplied to the SAM functional block), application of sharing policies to the criteria to determine the appropriate sharing resolution algorithm, and application of sharing resolution algorithms to the actuator command requests. Resolution of the competition for resources can be based on many different criteria, including the present plasma state and related control goals, present actuator status, or present control operating state (e.g., operating well within the nominal scenario, approaching control limits, operating in a recovery state or an alternate scenario state having already responded to an off-normal event, operating in a pre-disruptive or post-disruptive state, etc.). Sharing policies broadly describe how the criteria map onto selection of the resolution algorithm(s). Examples of such policies include a simple fixed priority ranking of control algorithms or control goals, conditions under which commands are simply summed, or a nonlinear operation is applied, or conditions under which a different priority ranking is assigned to competing control algorithms or goals. Examples of sharing resolution algorithms include simple selection of a single control goal to own a resource, summation of commands, or use of a complex nonlinear or logic function to resolve the conflict. Establishment of resource sharing criteria, policies, and algorithms represents a significant research and design effort, which in turn relies on specification of detailed ITER operation constraints and objectives.

Forecasting may play an important role in providing criteria for SAM. While both SMM and RP resolutions may benefit from forecasted evaluation of criteria, resolution of potential conflict in RP sharing cases may particularly require some level of forecasting capability. In such cases, a faster-than-real-time projection of control scenario evolution could identify an upcoming need for RP, for example, producing a revised priority in a resolution policy that has flexible priorities. Aside from this temporal dimension to the decision process, RP sharing decisions can also be based largely on the same kinds of criteria as SMM sharing decisions.

An example of the unique level of actuator management required by ITER is use of the EC system, responsible for several scenario operation and robustness missions. It has been specified that the EL will be used primarily for profile control using broad midplane heating and current drive deposition. The four upper port launchers (UL) are primarily dedicated to TM (TM) suppression.[38] However, these actuators are shared in two crucial ways: the profile control accomplished by the EL may be a key element in robustly preventing onset of TM, and both systems will share power supplies, which require 2–3 s to repurpose from EL to UL. The UL themselves may be shared for suppression of both 3/2 and 2/1 TM. This may require RP mirror motion to shift deposition between the corresponding resonant surfaces. It may also require SMM application of a subset of the launchers to the two surfaces, either pre-emptively or following detection of island growth. Under perfect alignment conditions, 10 MW out of the available 20 MW will enable suppression of both modes.[39,46] However, non-ideal control conditions may erode this apparent margin, so that further analysis is needed to identify realistic requirements. Management of the sharing of EC resources may thus be accomplished through a mix of prioritization and more complex decision algorithms. Prior to island triggering, competition between maintaining distance from the triggering threshold and other profile regulation goals would have to be balanced depending on scenario conditions and physics goals. The presence of a sufficiently threatening TM island will likely invoke simple prioritization of UL actuators in favor of TM stabilization over scenario goals. Other actuator systems that will likely require sharing by multiple control algorithms include the gas/pellet injection system (D/T fuelling for density, divertor, and burn control, and impurity injection for divertor and burn control as well as disruption mitigation), and the ICRF system (for bulk heating and sawtooth control).

V. CONCLUSIONS

Many aspects of ITER plasma control are significantly novel relative to presently operating devices. Although many of the fundamental schemes for scenario and stability control will be familiar from present tokamaks, the specific forms required by ITER, as well as the high level of performance and high level of integration demanded, will be unique. For example, although operating devices have explored various approaches to TM stabilization, an integrated and reliably effective ITER solution has not yet been demonstrated. While extensive research has been done on profile control solutions, many of the physics and algorithmic elements remain unsolved for ITER. In addition to specific new control physics solutions, ITER control will require a large step in the degree of robustness and confidence established in control performance prior to use. These requirements will drive a unique reliance on model-based algorithm design, as well as unprecedented levels of verification before experimental application. New integration solutions will be needed for managing actuators and handling off-normal and fault events that require a change in control action.

A significant amount of control physics understanding remains in order to complete such a robust ITER solution. Examples include:

- Identification of key profile characteristics determining TM stability
- Size of effective TM seeds produced by disturbances such as sawteeth and ELMs
- Quantification of controllability/robustness metrics corresponding to seedless TM triggering
- Effective control methods for burn control
- Sufficient look-ahead predictive capability for profile evolution using real-time data
- Reliable real-time calculation of proximity to stability and controllability boundaries
- Physics-based models for many relevant control functions, including noise, disturbance, and robustness specifications
- Scenarios for EH response to minimize or prevent disruptions

Equally as challenging as the control physics needs for ITER, a significant amount of control mathematics research remains to determine effective algorithmic and architectural solutions for ITER. Examples include:

- Specification of physics model requirements for robust control design
- Determination of algorithmic approaches capable of providing quantified robust performance in continuous control
- Design of algorithms and scalable architecture for effective EH responses
- Design of effective actuator management solutions to enable prioritized and effective sharing and ensure robust operation

These issues motivate a substantial R&D program within magnetic fusion research to provide a robust basis for plasma control in ITER. While the challenges remain great, the strong control research programs now underway at many institutions and devices worldwide, including targeted physics understanding and mathematical control research efforts, are aggressively engaged in addressing these issues to ensure the success of ITER.

ACKNOWLEDGMENTS

This material is based on work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award DE-FC02-04ER54698 and DE-AC02-09CH11466. DIII-D data shown in this paper can be obtained in digital format by following the links at https://fusion.gat.com/global/D3D_DMP.

[1]J. A. Snipes, S. Bremond, D. J. Campbell, T. Casper, D. Douai, Y. Gribov, D. Humphreys, J. Lister, A. Loarte, R. Pitts, M. Sugihara, A. Winter, and L. Zabeo, "Physics of the conceptual design of the ITER plasma control system," Fusion Eng. Des. **89**, 507 (2014).
[2]B. A. Ogunnaike and W. H. Ray, *Process Dynamics, Modelling, and Control* (Oxford University Press, Oxford, 1994).

[3]D. N. Hill and DIII-D Team, "DIII-D research towards resolving key issues for ITER and Steady state tokamaks," Nucl. Fusion **53**, 104001 (2013).
[4]D. Campbell, "Challenges in Burning Plasma Physics: the ITER Research Plan," in *Proceedings of 24th IAEA Fusion Energy Conference, San Diego, USA, October 8–13* (International Atomic Energy Agency, Vienna, 2012), IAEA-CN-197, ITR/P1-18.
[5]Q. P. Yuan, B. J. Xiao, Z. P. Luo, M. L. Walker, A. S. Welander, A. W. Hyatt, J. P. Qian, R. R. Zhang, D. A. Humphreys, J. A. Leuer, R. D. Johnson, B. G. Penaflor, and D. Mueller, "Plasma current, position, and shape feedback control on EAST," Nucl. Fusion **53**, 043009 (2013).
[6]S. Kinoshita, *et al.*, "Independent control of gaps in single-null divertor discharges on the DIII–D tokamak," General Atomics Report No. GA-A19584 (1989).
[7]R. Albanese, G. Ambrosino, M. Ariola, A. Cenedese, F. Crisanti, G. De Tommasi, M. Mattei, F. Piccolo, A. Pironti, F. Sartori, and F. Villone, "Design, implementation, and test of the XSC extreme shape controller in JET," Fusion Eng. Des. **74**, 627 (2005).
[8]E. Kolemen, A. Welander, R. La Haye, N. Eidietis, D. Humphreys, J. Lohr, S. Noraky, B. Penaflor, R. Prater, and F. Turco, "State-of-the-art neoclassical tearing mode control in DIII-D using real-time steerable electron cyclotron current drive launchers," Nucl. Fusion **54**, 073020 (2014).
[9]J. A. Snipes, D. Beltran, T. Casper, Y. Gribov, A. Isayama, J. Lister, S. Simrock, G. Vayakis, A. Winter, Y. Yang, and L. Zabeo, "Actuator and Diagnostic Requirements of the ITER Plasma Control System," Fusion Eng. Des. **87**, 1900 (2012).
[10]E. J. Doyle, J. C. DeBoo, J. R. Ferron, G. L. Jackson, T. C. Luce, M. Murakami, T. H. Osborne, J.-M. Park, P. A. Politzer, H. Reimerdes, R. V. Budny, T. A. Casper, C. D. Challis, R. J. Groebner, C. T. Holcomb, A. W. Hyatt, R. J. La Haye, G. R. McKee, T. W. Petrie, C. C. Petty, T. L. Rhodes, M. W. Shafer, P. B. Snyder, E. J. Strait, M. R. Wade, G. Wang, W. P. West, and L. Zeng, "Demonstration of ITER operational scenarios on DIII-D," Nucl. Fusion **50**, 075005 (2010).
[11]A. C. C. Sips, T. A. Casper, E. J. Doyle, G. Giruzzi, Y. Gribov, J. Hobirk, G. M. D. Hogeweij, L. D. Horton, A. E. Hubbard, I. Hutchinson, S. Ide, A. Isayama, F. Imbeaux, G. L. Jackson, Y. Kamada, C. Kessel, F. Kochl, P. Lomas, X. Litaudon, T. C. Luce, E. Marmar, M. Mattei, I. Nunes, N. Oyama, V. Parail, A. Portone, G. Saibene, R. Sartori, J. K. Stober, T. Suzuki, S. M. Wolfe, C-Mod Team, ASDEX Upgrade Team, DIII-D team, and JET EFDA Contributors, "Experimental studies of ITER demonstration discharges," Nucl. Fusion **49**(8), 085015–085025 (2009).
[12]J. Schweinzer, V. Bobkov, A. Burckhart, R. Dux, C. Fuchs, A. Kallenbach, J. Hobirk, C. Hopf, P. T. Lang, A. Mlynek, Th. Pütterich, F. Ryter, G. Tardini, J. Stober, and ASDEX Upgrade team, "Demonstration of the ITER baseline scenario on ASDEX Upgrade," in *Proceedings of the 40th EPS Conference on Plasma Physics, Spain* (2013), P2.134.
[13]G. L. Jackson, T. A. Casper, T. C. Luce, D. A. Humphreys, J. R. Ferron, A. W. Hyatt, T. W. Petrie, and W. P. West, "Simulating the ITER Plasma Startup Scenario in the DIII-D Tokamak," Nucl. Fusion **49**, 115027 (2009).
[14]P. A. Politzer, G. L. Jackson, D. A. Humphreys, T. C. Luce, A. W. Hyatt, and J. A. Leuer, "Experimental simulation of ITER rampdown in DIII-D," Nucl. Fusion **50**, 035011 (2010).
[15]T. E. Evans, R. A. Moyer, P. R. Thomas, J. G. Watkins, T. H. Osborne, J. A. Boedo, E. J. Doyle, M. E. Fenstermacher, K. H. Finken, R. J. Groebner, M. Groth, J. H. Harris, R. J. LaHaye, C. J. Lasnier, S. Masuzaki, N. Ohyabu, D. Pretty, T. L. Rhodes, H. Reimerdes, D. L. Rudakov, M. J. Schaffer, G. Wang, and L. Zeng, "Suppression of large edge-localized modes in high-confinement DIII-D plasmas with a stochastic magnetic boundary," Phys. Rev. Lett. **92**, 235003 (2004).
[16]Y. Liang *et al.*, "Active control of type-I edge-localized modes with n=1 perturbation fields in the JET tokamak," Phys. Rev. Lett. **98**, 265004 (2007).
[17]W. Suttrop *et al.*, "First observation of edge localized modes mitigation with resonant and nonresonant magnetic perturbations in ASDEX upgrade," Phys. Rev. Lett. **106**, 225004 (2011).
[18]M. Cavinato, G. Ambrosino, R. Ambrosino, Y. Gribov, F. Koechl, M. Mattei, V. Parail, A. Pironti, G. Saibene, R. Sartori, and L. Zabeo, "ITER plasma position control system and scenario optimization," in Proceedings of the 24th IAEA Fusion Energy Conference, San Diego, CA (2012), ITR/P1-16.
[19]V. Parail, R. Albanese, R. Ambrosino, J. –F. Artaud, K. Besseghir, M. Cavinato, G. Corrigan, J. Garcia, L. Garzotti, Y. Gribov, F. Imbeaux, F. Koechl, C. V. Labate, J. Lister, X. Litaudon, A. Loarte, P. Maget, M. Mattei, D. McDonald, E. Nardon, G. Saibene, R. Sartori, and J. Urban, "Self-consistent simulation of plasma scenarios for ITER using a combination of 1.5D transport codes and free boundary equilibrium codes," in Proceedings of the 24th IAEA Fusion Energy Conference, San Diego, CA (2012), ITR/P1-10.
[20]F. Turco and T. C. Luce, "Impact of the current profile evolution on tearing stability of ITER demonstration discharges in DIII-D," Nucl. Fusion **50**, 095010 (2010).
[21]J. M. Hanson, H. Reimerdes, M. J. Lanctot, Y. In, R. J. La Haye, G. L., Jackson, G. A., Navratil, M. Okabayashi, P. E., Sieck, and E. J., Strait, "Feedback control of the proximity to marginal RWM stability using active MHD spectroscopy," Nucl. Fusion **52**, 013003 (2012).
[22]G. L. Jackson, T. C. Luce, R. J. Buttery, A. W. Hyatt, J. R. Ferron, R. J. La Haye, P. A. Politzer, W. M. Solomon, F. Turco, and E. J. Doyle, "Stability boundaries and development of the ITER baseline scenario," Nucl. Fusion **55**, 023004 (2015).
[23]Y. Ou, T. C. Luce, E. Schuster, J. R. Ferron, M. L. Walker, C. Xu, and D. A. Humphreys, "Towards model-based current profile control at DIII-D," Fusion Eng. Des. **82**, 1153–1160 (2007).
[24]J. E. Barton, W. Shi, K. Besseghir, J. Lister, A. Kritz, E. Schuster, T. C. Luce, M. L. Walker, D. A. Humphreys, and J. R. Ferron, "Physics-based control-oriented modeling of the safety factor profile dynamics in high performance tokamak plasmas," in Proceedings of the 52nd IEEE Conference on Decision and Control, Florence, Italy, December 10–13, 2013.
[25]F. A. A. Felici and O. Sauter, "Non-linear model-based optimization of actuator trajectories for tokamak plasma profile control," Plasma Phys. Controlled Fusion **54**, 025002 (2012).
[26]J. Barton, W. Shi, E. Schuster, T. C. Luce, J. R. Ferron, M. L. Walker, D. A. Humphreys, F. Turco, R. D. Johnson, and B. G. Penaflor, "Nonlinear physics-model-based actuator trajectory optimization for advanced scenario planning in the DIII-D tokamak," in Proceedings of the 19th IFAC World Congress, Cape Town, South Africa, August 24–29, 2014.
[27]D. Moreau, F. Crisanti, X. Litaudon, D. Mazon, P. De Vries, R. Felton, E. Joffrin, L. Laborde, M. Lennholm, A. Murari, V. Pericoli-Ridolfini, M. Riva, T. Tala, G. Tresset, L. Zabeo, K. D. Zastrow, and Contributors to the EFDA-JET Work Programme, "Real-time control of the q-profile in JET for steady state advanced tokamak operation," Nucl. Fusion **43**, 870 (2003).
[28]D. Moreau, D. Mazon, M. Ariola, G. De Tommasi, L. Laborde, F. Piccolo, F. Sartori, T. Tala, L. Zabeo, A. Boboc, E. Bouvier, M. Brix, J. Brzozowski, C. D. Challis, V. Cocilovo, V. Cordoliani, F. Crisanti, E. De La Luna, R. Felton, N. Hawkes, R. King, X. Litaudon, T. Loarer, J. Mailloux, M. Mayoral, I. Nunes, E. Surrey, O. Zimmerman, and JET EFDA Contributors, "A two time scale dynamic model approach for magnetic and kinetic profile control in advanced tokamak scenarios in JET," Nucl. Fusion **48**, 106001 (2008).
[29]D. Moreau, M. L. Walker, J. R. Ferron, F. Liu, E. Schuster, J. E. Barton, M. D. Boyer, K. H. Burrell, S. M. Flanagan, P. Gohil, R. J. Groebner, C. T. Holcomb, D. A. Humphreys, A. W. Hyatt, R. D. Johnson, R. J. La Haye, J. Lohr, T. C. Luce, J. M. Park, B. G. Penaflor, W. Shi, F. Turco, W. Wehner, and ITPA-IOS Group Members and Experts, "Integrated magnetic and kinetic control of advanced tokamak plasmas on DIII&unknown_hyphen;D based on data-driven models," Nucl. Fusion **53**, 063020 (2013).
[30]F. A. A. Felici, O., Sauter, S. Coda, B. P. Duval, T. P. Goodman, J.-M. Moret, and J. I. Paley, "Real-time physics-model-based simulation of the current density profile in tokamak plasmas," Nucl. Fusion **51**, 083052 (2011).
[31]E. Witrant, E. Joffrin, S. Bremond, G. Giruzzi, D. Mazon, O. Barana, and P. Moreau, "A control-oriented model of the current profile in tokamak plasma," Plasma Phys. Controlled Fusion **49**, 1075 (2007).
[32]S. H. Kim and J. B. Lister, "A potentially robust plasma profile control approach for ITER using real-time estimation of linearized profile response models," Nucl. Fusion **52**, 074002 (2012).
[33]K. Keesman, *System Identification: An Introduction, Advanced Textbooks in Control and Signal Processing* (Springer-Verlag, London, 2011).
[34]M. D. Boyer, J. Barton, E. Schuster, M. L. Walker, T. C. Luce, J. R. Ferron, B. G. Penaflor, R. D. Johnson, and D. A. Humphreys, "Backstepping control of the toroidal plasma current profile in the DIII-D tokamak," IEEE Trans. Control Syst. Technol. **22**, 1725 (2014).
[35]S. H. Kim, private communication (24 June 2014).
[36]R. J. Buttery, S. Gunter, G. Giruzzi, T. C. Hender, D. Howell, G. Huysmans, R. J. La Haye, M. Maraschek, H. Reimerdes, O. Sauter, C. D.

Warrick, H. R. Wilson, and H. Zohm, "Neoclassical tearing modes," Plasma Phys. Controlled Fusion **42**, B61 (2000).
[37]R. J. La Haye *et al.*, "Requirements for alignment of electron cyclotron current drive for neoclassical tearing mode stabilization in ITER," Nucl. Fusion **48**, 054004 (2008).
[38]R. J. La Haye *et al.*, "Prospects for stabilization of neoclassical tearing modes by electron cyclotron current drive in ITER," Nucl. Fusion **49**, 045005 (2009).
[39]O. Sauter, M. A. Henderson, G. Ramponi, H. Zohm, and C. Zucca, "On the requirements to control neoclassical tearing modes in burning plasmas," Plasma Phys. Controlled Fusion **52**, 025002 (2010).
[40]C. C. Hegna and J. D. Callen, Phys. Plasmas **4**, 2940 (1997).
[41]M. Maraschek, "Control of neoclassical tearing modes," Nucl. Fusion **52**, 074007 (2012).
[42]A. Isayama, Y. Kamada, N. Hayashi, T. Suzuki, T. Oikawa, T. Fujita, T. Fukuda, S. Ide, H. Takenaga, K. Ushigusa, T. Ozeki, Y. Ikeda, N. Umeda, H. Yamada, M. Isobe, Y. Narushima, K. Ikeda, S. Sakakibara, K. Yamazaki, K. Nagasaki, and Jt-60 Team, "Achievement of high fusion triple product, steady-state sustainment and real-time NTM stabilization in high-Betap ELMy H-mode discharges in JT-60U," Nucl. Fusion **43**, 1272 (2003).
[43]G. Gantenbein, H. Zohm, G. Giruzzi, S. Gunter, F. Leuterer, M. Maraschek, J. Meskat, Q. Yu, ASDEX Upgrade Team, and ECRH-Group (AUG), "Complete suppression of neoclassical tearing modes with current drive at the electron-cyclotron-resonance frequency in ASDEX upgrade tokamak," Phys. Rev. Lett. **85**, 1242 (2000).
[44]F. A. A. Felici, T. P. Goodman, O. Sauter, G. Canal, S. Coda, B. P. Duval, and J. X. Rossel, "Integrated real-time control of MHD instabilities using multi-beam ECRH/ECCD systems on TCV," Nucl. Fusion **52**, 074001 (2012).
[45]R. Prater, R. J. Lahaye, J. Lohr, T. C. Luce, C. C. Petty, J. R. Ferron, D. A. Humphreys, E. J. Strait, F. W. Perkins, and R. W. Harvey, "Discharge improvement through control of neoclassical tearing modes by localized ECCD in DIII-D," Nucl. Fusion **43**, 1128 (2003).
[46]R. J. La Haye, "Neoclassical tearing modes and their control," Phys. Plasmas **13**, 055501 (2006).
[47]T. P. Goodman, F. Felici, O. Sauter, and J. P. Graves, Phys. Rev. Lett. **106**, 245002 (2011).
[48]T. Pütterich, R. Dux, R. Neu, M. Bernert, M. N. A. Beurskens, V. Bobkov, S. Brezinsek, C. Challis, J. W. Coenen, I. Coffey, A. Czarnecka, C. Giroud, A. Gude, E. Joffrin, A. Kallenbach, M. Lehnen, E. Lerche, E. de la Luna, S. Marsen, G. Matthews, M.-L. Mayoral, R. M. McDermott, A. Meigs, M. Sertoli, G. van Rooij, J. Schweinzer, ASDEX Upgrade Team, and JET EFDA Contributors, "Taming tungsten in JET and ASDEX Upgrade," in Proceedings of the 40th EPS Conference on Plasma Physics, Espoo, Finland, I2.106 (July 2013).
[49]A. S. Welander, E. Kolemen, R. J. La Haye, N. W. Eidietis, D. A. Humphreys, J. Lohr, S. Noraky, B. G. Penaflor, R. Prater, and F. Turco, "Advanced control of neoclassical tearing modes in DIII-D with real-time steering of the electron cyclotron current drive," Plasma Phys. Controlled Fusion **55**, 124033 (2013).
[50]W. Wehner and E. Schuster, "Control-oriented modelling for neoclassical tearing mode stabilization via minimum-seeking techniques," Nucl. Fusion **52**, 074003 (2012).
[51]D. A. Humphreys, J. R. Ferron, R. J. La Haye, T. C. Luce, C. C. Petty, R. Prater, and A. S. Welander, "Active control for stabilization of neoclassical tearing modes," Phys. Plasmas **13**, 56113 (2006).
[52]I. T. Chapman, R. J. La Haye, R. J. Buttery, W. W. Heidbrink, G. L. Jackson, C. M. Muscatello, C. C. Petty, R. I. Pinsker, B. J. Tobias, and F. Turco, "Sawtooth control using ECCD in ITER demonstration plasmas in DIII-D," Nucl. Fusion **52**, 063006 (2012).
[53]A. W. Leonard, M. A. Mahdavi, C. J. Lasnier, T. W. Petrie, and P. C. Stangeby, "Scaling radiative divertor solutions to high power in DIII-D," Nucl. Fusion **52**, 063015 (2012).
[54]A. Kallenbach, M. Bernert, T. Eich, J. C. Fuchs, L. Giannone, A. Herrmann, J. Schweinzer, W. Treutterer, and ASDEX Upgrade Team, "Optimized tokamak power exhaust with double radiative feedback in ASDEX Upgrade," Nucl. Fusion **52**, 122003 (2012).
[55]E. Kolemen, S. L. Allen, B. D. Bray, M. E. Fenstermacher, D. A. Humphreys, A. W. Hyatt, C. J. Lasnier, A. W. Leonard, M. A. Makowski, A. G. McLean, R. Maingi, R. Nazikian, T. W. Petrie, V. A. Soukhanovskii, and E. A. Unterberg, "Heat flux management via advanced magnetic divertor configurations and divertor detachment," J. Nucl. Mater. (published online).
[56]M. A. Van Zeeland *et al.*, "Measurements, modelling and electron cyclotron heating modification of Alfvén eigenmode activity in DIII-D," Nucl. Fusion **49**, 065003 (2009).
[57]S. W. Haney, L. J. Perkins, J. Mandrekas, and W. M. Stacey, Jr., "Active control of burn conditions for the International Thermonuclear Experimental Reactor," Fusion Tech. **18**, 606 (1990).
[58]M. Firestone, C. E. Kessel, J. Morrow-Jones, S. C. Jardin, and T. K. Mau, "Comprehensive control of a tokamak reactor," in Proceedings of the 14th IEEE/NPSS Symposium on Fusion Eng. (1991), Vol. 1, p. 209.
[59]E. Schuster, M. Krstic, and G. Tynan, "Burn control in fusion reactors via nonlinear stabilization techniques," Fusion Sci. Technol. **43**(1), 18–37 (2003).
[60]R. J. Hawryluk, N. Eidietis, B. A. Grierson, A. Hyatt, E. Kolemen, R. Nazikian, C. Paz-Soldan, W. Solomon, and S. Wolfe, "Control of plasma stored energy for burn control using DIII-D in-vessel coils," Nucl. Fusion (submitted).
[61]M. Sugihara, S. Putvinski, D. J. Campbell, S. Carpentier-Chouchana, F. Escourbiac, S. Gerasimov, Yu. Gribov, T. C. Hender, T. Hirai, K. Ioki, R. Khayrutdinov, H. Labidi, V. Lukash, S. Maruyama, M. Merola, R. Mitteau, S. Miyamoto, J. Morris, G. Pautasso, R. A. Pitts, R. Raffray, V. Riccardo, R. Roccella, G. Sannazzaro, T. Schioler, J. Snipes, and R. Yoshino, "Disruption impacts and their mitigation target values for ITER operation and machine protection," in Proceedings of 24th IAEA Fusion Energy Conference, San Diego, USA, October (2012), ITR/P1-14.
[62]F. Volpe, R. J. La Haye, R. Prater, and E. J. Strait, "Locked mode NTM control on DIII-D by ECCD and magnetic perturbations," in Proceedings of 4th IAEA Technical Meeting on ECRH Physics and Technology for ITER, Vienna, Austria (2007).
[63]N. Commaux, L. R. Baylor, S. K. Combs, N. W. Eidietis, T. E. Evans, C. R. Foust, E. M. Hollmann, D. A. Humphreys, V. A. Izzo, A. N. James, T. C. Jernigan, S. J. Meitner, P. B. Parks, J. C. Wesley, and J. H. Yu, "Novel rapid shutdown strategies for runaway electron suppression in DIII-D," Nucl. Fusion **51**, 103001 (2011).
[64]V. Riccardo, G. Arnoux, P. Cahyna, T. C. Hender, A. Huber, S. Jachmich, V. Kiptily, R. Koslowski, L. Krlin, M. Lehnen, A. Loarte, E. Nardon, R. Paprok, D. Tskhakaya, Sr., and JET-EFDA Contributors, "JET disruption studies in support of ITER," Nucl. Fusion **52**, 124018 (2010).
[65]D. A. Humphreys, T. A. Casper, N. Eidietis, M. Ferrara, D. A. Gates, I. H. Hutchinson, G. L. Jackson, E. Kolemen, J. A. Leuer, J. Lister, L. L. Lodestro, W. H. Meyer, L. D. Pearlstein, F. Sartori, M. L. Walker, A. S. Welander, and S. M. Wolfe, "Experimental vertical stability studies for ITER performance and design guidance," Nucl. Fusion **49**, 115003 (2009).
[66]D. P. Brennan, R. J. La Haye, A. D. Turnbull, M. Chu, T. H. Jensen, L. L. Lao, T. C. Luce, P. A. Politzer, E. J. Strait, S. E. Kruger, and D. D. Schnack, "A mechanism for tearing onset near ideal stability boundaries," Phys. Plasmas **10**, 1643 (2003).
[67]A. Winter *et al.*, "Design progress of the ITER plasma control system," in Proceedings of the 9th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research, Hefei, China (2013), O1-2.
[68]F. Sartori, G. De Tommasi, and F. Piccolo, "The Joint European Torus-Plasma position and shape control in the world's largest tokamak," IEEE Control Syst. Mag. **26**(2), 64–78 (2006).
[69]G. Ambrosino, R. Albanese, M. Ariola, A. Cenedese, F. Crisanti, G. De Tommasi, M. Mattei, F. Piccolo, A. Pironti, F. Sartori, F. Villone, and JET-EFDA Contributors, "XSC plasma control: Tool development for the session leader," Fusion Eng. Des. **74**, 521 (2005).
[70]M. Athans, "The role and use of the stochastic Linear-Quadratic-Gaussian problem in control system design," IEEE Trans. Automatic Control **AC-16**, 529 (1971).
[71]C. E. Kessel, M. A. Firestone, and R. W. Conn, "Linear optimal control of tokamak fusion devices," Fusion Technol. **17**, 391 (1990).
[72]J. M. Maciejowski, *Multivariable Feedback Design* (Addison-Wesley, Wokingham, England, 1989).
[73]D. A. Humphreys and I. H. Hutchinson, "Axisymmetric magnetic control design in tokamaks using perturbed equilibrium plasma response modeling," Fusion Technol. **23**, 167 (1993).
[74]S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd ed. (Wiley, 2005).
[75]M. L. Walker, J. R. Ferron, D. A. Humphreys, R. D. Johnson, J. A. Leuer, B. G. Penaflor, D. A. Piglowski, M. Ariola, A. Pironti, and E. Schuster, "Next-generation plasma control in the DIII-D tokamak," Fusion Eng. Des. **66-68**, 749 (2003).

[76]G. Ambrosino, M. Ariola, Y. Mitrishkin, and A. Pironti, "Plasma current and shape control in tokamaks using H-infinity and mu-synthesis," in Proceedings of the 36th Conference on Decision and Control (1997), 3697.
[77]K. Ariyur and M. Krstic, *Real Time Optimization by Extremum Seeking Control* (Wiley, 2003).
[78]Y. Ou, C. Xu, E. Schuster, T. C. Luce, J. R. Ferron, M. L. Walker, and D. A. Humphreys, "Design and simulation of extremum-seeking open-loop optimal control of current profile in the DIII-D tokamak," Plasma Phys. Controlled Fusion **50**, 115001 (2008).
[79]C. Garcia, D. M. Prett, and M. Morari, "Model predictive control: theory and practice—a survey," Automatica **25**, 335 (1989).
[80]M. Mattei, C. V. Labate, and D. Famularo, "A constrained control strategy for the shape control in thermonuclear fusion tokamaks," Automatica **49**, 169 (2013).
[81]D. Simon, *Optimal State Estimation* (Wiley and Sons, Hoboken, NJ, 2006).
[82]F. Felici, P. Geelen, L. Giannone, D. Kim, E. Maljaars, P. Piovesan, C. Piron, C. Rapson, O. Sauter, A. Teplukhina, W. Treutterer, L. Barrera, E. Fable, M. Reich, and G. Tardini, "First results of real-time plasma state reconstruction using a model-based dynamic observer on ASDEX-Upgrade," in Proceedings of the 41st EPS Plasma Physics Controlled Fusion Conference, Berlin, Germany, June 2014, P2.002.
[83]E. H. Raupp, V. Mertens, G. Neu, W. Treutterer, D. Zasche, and Th. Zehetbauer, "Real-time exception handling—Use cases and response requirements," Fusion Eng. Des. **87**, 1891 (2012).
[84]M. Sugihara, M. Shimada, H. Fujieda, Yu. Gribov, K. Ioki, Y. Kawano, R. R. Khayrutdinov, V. Lukash, and J. Ohmori, "Disruption scenarios, their mitigation and operation window in ITER," Nucl. Fusion **47**, 337 (2007).
[85]Y. Nakamura, R. Yoshino, Y. Neyatani, T. Tsunematsu, and M. Azumi, "Mechanism of vertical displacement events in JT-60U disruptive discharges," Nucl. Fusion **36**, 643 (1996).