GA–A24760

EMERGING APPLICATIONS IN TOKAMAK PLASMA CONTROL

by

M.L. WALKER, D.A. HUMPHREYS, D. MAZON, D. MOREAU,
M. OKABAYASHI, T.H. OSBORNE, E. SCHUSTER

AUGUST 2004

GENERAL ATOMICS

DISCLAIMER

GA–A24760

EMERGING APPLICATIONS IN TOKAMAK PLASMA CONTROL

by
M.L. WALKER, D.A. HUMPHREYS, D. MAZON,* D. MOREAU,†
M. OKABAYASHI,‡ T.H. OSBORNE, E. SCHUSTER§

This is a preprint of a paper submitted
to IEEE Control Systems Magazine.

*EURATOM-CEA Association, DSM-DRFC, CEA-Cadarache, St Paul lez Durance, France
†EFDA-JET CSU, Culham Science Centre, Abingdon, United Kingdom
‡Princeton Plasma Physics Laboratory, Princeton, New Jersey
§Lehigh University, Bethlehem, Pennsylvania

Work supported by
the U.S. Department of Energy
under Contract Nos. DE-FC02-04ER54698,
DE-AC02-76CH03073 and Grant No.
DE-FG02-92ER54141 and performed under
the European Fusion Development Agreement (EFDA)

GA PROJECT 30200
AUGUST 2004

GENERAL ATOMICS

TABLE OF CONTENTS

LIST OF SIDEBARS

LIST OF FIGURES

LIST OF TABLES

INTRODUCTION

Many of the articles in this special issue focus on the problem of control of plasma axisymmetric shape and position. It is the best understood of tokamak control problems and, therefore, has the most results to discuss. There are also a large number of other tokamak control problems that are not nearly as mature in control development, but are equally important to continued progress toward the goal of producing energy from fusion. Work on these problems is active and growing rapidly. Areas of active development include work on basic physics understanding, developing actuators and sensors for control, development of control models including experiments to characterize actuators and sensors, experimental use of initial *ad hoc* controls, and development (and even a few deployments) of sophisticated control algorithms. In this article we provide a brief introduction to several of these control problems. For those problems that are discussed, we provide examples of progress that has been made at different tokamak devices around the world. In almost all cases we are unable, because of space considerations, to provide a complete accounting of all work on a particular class of control problems.

A practical tokamak fusion reactor must operate at high temperature, high pressure, and high current. These high performance tokamak plasmas are susceptible to a large number of instabilities, some of them posing a risk to the device itself. One example seen earlier in this issue is the so-called vertical instability due to noncircularity of the plasma. The vertical instability is axisymmetric, that is the plasma motion is the same at all toroidal angles, and is characterized by a primarily vertical displacement. Increases in growth rates of the vertical instability and of other plasma instabilities correspond to increases in plasma pressure. Consequently, the most attractive operational regimes from the point of view of a fusion power reactor also tend to be those that are nearest to instability.

Increasing temperature, pressure, and current cause several non-axisymmetric instabilities that must be stabilized. Most, but perhaps not all, will require feedback control. In this article, we will describe the control problems presented by instabilities known as resistive wall modes (RWM), neoclassical tearing modes (NTM), and edge localized modes (ELM). The objective of stabilization is to prevent loss of the plasma while retaining high performance; pushing the plasma into higher performance regimes is often what triggers a particular instability. Models upon which stabilization algorithms are based all start with ideal magnetohydrodynamic (MHD) theory (sidebar S3 of Attachment, GA–A24794), but each requires an extension to that basic theory to correctly account for all important effects.

There is also a collection of control problems that might collectively be called control of internal plasma parameters: control of the current, temperature, and density profiles (sidebar S4 of the Attachment, GA–A24794), and control of transport (the continuous flow of particles, heat, and/or current) in the plasma interior. One particular combination of this class of problems will also be described.

There are also a number of *off-normal* events that can occur which must be "handled" rather than controlled in the sense of feedback. Off-normal events are events that are likely to occur due to occasional loss of control in future devices (and that routinely occur in present-day devices). A description of several of these events will be provided. Efforts have been made on various experimental devices to define and implement appropriate responses to some types of off-normal events. A few of these methods will be described.

A Guide for the Reader (sidebar)

This paper covers a number of separate tokamak control problems, each of which requires some background understanding of a variety of basic plasma physics topics. We rely heavily on the introductory paper (Attachment General Atomics Report GA–A24794) and collection of numbered sidebars in this special issue to provide this background. In most cases, a reference to the introduction paper (see Attachment) or appropriate sidebar is included when it is thought that a term or concept might be unfamiliar to some readers. For further assistance, at the end of the introduction paper (Attachment) there is a list of many plasma physics terms and the locations where each is defined. Each section of this paper discusses a completely separate control problem and all are written so that they can be read more or less independently of the other sections. The separate topics are ordered approximately in order of the level of physics background needed to understand them.

SUPPRESSION OF THE NEOCLASSICAL TEARING MODE

Increasing beta (sidebar S2 of the Attachment) in a resistive plasma can make the nested magnetic surface topology required by ideal MHD (sidebar S3 of the Attachment) unstable, resulting in a *tearing* and *reconnection* of the flux surfaces. When this reconnection occurs, a structure called a *magnetic island* is formed (Fig. 1). The instability known as a *neoclassical tearing mode* (NTM) drives an island to a saturated size, which then persists stably in the plasma [1]. The island actually winds helically around the tokamak with a helicity given by the value of the factor q (sidebar S4 of the Attachment) on the surface where it forms. The NTM forms on flux surfaces whose safety factor (q) is rational, the most important of which are the q=3/2=1.5 or the q=2/1=2.0 surfaces. The 2/1 NTM often produces a plasma-terminating disruption (sidebar S14) by triggering an ideal MHD mode, which can grow to be comparable in size to the plasma cross section, while the 3/2 NTM usually remains small enough to merely degrade confinement. The presence of the island can degrade confinement because it effectively connects hotter inner regions of the plasma more directly to colder outer regions with "short circuiting" magnetic field lines, allowing heat to leak out of the plasma core faster than it would without the island. The resulting flattening in temperature, pressure, and current profiles across the island (Fig. 2) corresponds to a steady-state lowering of total plasma internal energy and therefore an overall colder, less efficient plasma.

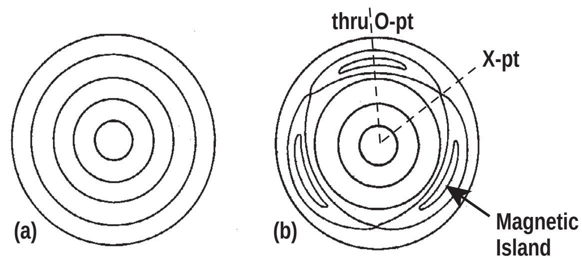

Figure 1. Illustration of magnetic islands topology in a circular cross-section plasma. Perfectly conducting ideal MHD plasmas (sidebar S3 of the Attachment) require nested flux surfaces (a), while resistive plasmas can produce tearing and *reconnection* (hence the name *tearing mode*) of flux surfaces, resulting in magnetic islands (b). Localized current and pressure profiles (sidebar S5 of the Attachment) are flattened across an island (whose center is the *O-point*). The resulting connection between inner and outer island surfaces (joined at the *X-point*) allows heat to leak out of the plasma core faster than it would without the island, thus degrading confinement. (b) shows the island topology corresponding to a 3/2 NTM [periodicity of $m = 3$ (sidebar S4 of the Attachment) in the poloidal cross section shown, and a periodicity of $n = 2$ in the toroidal direction (not illustrated)].

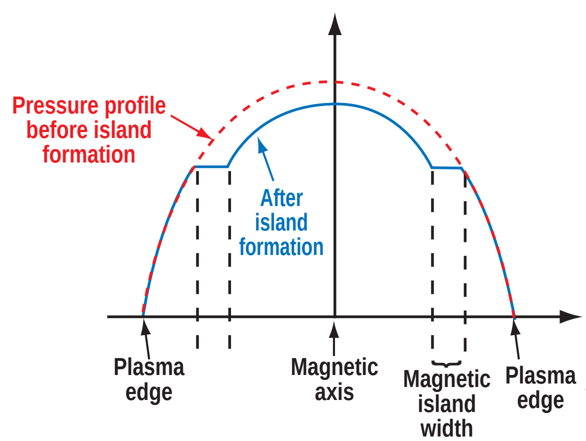

Figure 2. Flattening of the pressure profile caused by the NTM. Islands driven by the NTM are responsible for a lowering of temperature, pressure, and current both on and inside the island.

S14. Plasma Disruptions (sidebar)

Disruptions are rapid events in which a large fraction of the plasma thermal energy is lost due to the uncontrolled growth of some large-scale plasma instability. These large-scale instabilities take place over a large portion of the plasma cross section. In most cases disruptions are unrecoverable, expelling nearly 100% of the thermal energy and leading to complete termination of the plasma current. More rarely, disruptive instabilities can fail to expel all of the thermal energy and may even allow the discharge to recover (a *minor disruption*). Two broad classes of events account for the vast majority of plasma-terminating disruptions in tokamaks: major disruptions and vertical displacement events (VDEs). VDEs are unique to plasmas that are intrinsically unstable to vertical displacements (the *vertical instability*; see sidebar S6 of the Attachment). This only exists in plasmas that are sufficiently vertically elongated (have cross sections that are higher than they are wide). A VDE is characterized by a loss of vertical position control, which allows the vertical instability to grow. The resulting vertical drift causes the plasma to strike a nearby surface, which erodes away the plasma as it continues to move into the wall. This erosion eventually destabilizes a large-scale instability, which expels all of the thermal energy. Following this *thermal quench* the plasma is always too cold and resistive for the ohmic transformer (sidebar S13) to sustain the plasma current, which decays away and terminates the plasma. By contrast, a major disruption is characterized by a sudden thermal quench triggered by a large-scale instability before any loss of position occurs. In this case the plasma becomes cold and resistive when more or less centered in the tokamak, and much of the plasma current can decay away before any subsequent loss of position causes the plasma to strike the wall.

Disruptions are undesirable in reactors due to a host of potentially damaging effects. The thermal quench can apply a large heat load to the first wall or divertor (sidebar S9 of the Attachment), possibly melting protective surfaces. Large currents can be driven by the disruption in both plasma facing components and conducting structures, resulting in potentially damaging electromagnetic forces. The very high voltages produced in the cold post-thermal quench plasma can accelerate electrons into a relativistic beam. This million electron volt (MeV) scale electron beam can destroy many types of plasma facing components as the field lines they travel on intercept the first wall, much as an electron beam lithograph can etch features into semiconductor substrates.

Fortunately, many of these damaging disruption effects can be mitigated by taking corrective action as a stability limit is being approached, or even after the disruption is underway. Furthermore, understanding of the instabilities whose growth lead to disruptions has advanced to the point that many of these disruptive stability limits are predictable. However, it remains a challenge to be able to predict an impending disruption in real time early enough and with sufficient reliability to execute the necessary corrective or mitigating action.

S13. Plasma Heating and Current Drive

There are several methods in use at experimental devices for heating and driving current in the plasma. Those methods that have reached a maturity that makes them suitable for everyday experimental use include ohmic heating (OH) and current drive, neutral beam injection (NBI), and various forms of radio-frequency (RF) heating and current drive. All of these methods can actually accomplish both, but they are often configured with a primary purpose of either heating or current drive.

Ohmic current drive operates through a transformer action. Continuously changing current in one or more poloidal field coils produces a changing poloidal flux Ψ, known as the *ohmic flux*, at the plasma. The derivative of this flux defines an induced voltage $V = d\Psi / dt$ known as the plasma *loop voltage*, which drives current in the plasma according to $LdI / dt + RI = V$, just as in a standard transformer. Here the values L and R represent the plasma bulk equivalent self-inductance and resistance. Resistive losses in the plasma are responsible for the heating effect, hence the origin of the term *ohmic*.

Radio-frequency heating is a process by which electromagnetic waves are transmitted into the plasma and a portion of their energy is absorbed by the plasma. Wave energy is coupled to the plasma particles primarily through resonant absorption. Resonant absorption occurs when the wave frequency is the same as a particle cyclotron frequency (see Fig. 9 of the Attachment), so that the wave’s electric field increases the perpendicular velocity of the resonant particles. The direction of propagation of the RF waves determines whether this absorption results principally in heating the plasma or in a combination of heating and current drive. Wave propagation with a component parallel to the field lines can produce signficant current drive, while perpendicular acceleration principally produces heating. Typical ion cyclotron frequencies (ICH, ICCD) lie between 30 and 120 Mhz, lower hybrid resonant frequencies (for a plasma mode that is a hybrid of electron and ion cyclotron motions) lie between 1 and 8 GHz, and electron cyclotron resonant frequencies (ECH, ECCD) lie between 70 and 200 GHz. Current can also be driven by coupling to fast magnetosonic waves (fast wave current drive, FWCD) or ion Bernstein waves (IBW) in the plasma [7].

Neutral beam injection (NBI) is the process by which neutral hydrogen or deuterium atoms are injected into the plasma at high speed, and then become ionized through collisions with plasma particles. The resulting ions and electrons then become part of the plasma. The kinetic energy carried by the originally neutral atoms is transferred to the plasma by both the initial and subsequent collisions, resulting in an increase in thermal energy (temperature) of the plasma. Because the high-energy beam ions collide primarily with the plasma thermal electrons, NBI can produce significant current drive when the beam is injected tangentially (i.e., in the toroidal direction). Tangentially injected neutrals also transfer their momentum to the plasma, thereby increasing the speed of plasma fluid rotation (see previous section on RWM). Perpendicular injection of neutral beams produces only heating.

Another interaction between heating and current drive is produced by the *resistivity* of the plasma. In the same way that resistance relates current and voltage in a resistor, resistivity relates local current density (electrical current flow per square meter of plasma cross section) to local electric field in the plasma. The resistivity varies in both spatial distribution and in time (at a given location). The change in resistivity is primarily determined by the local temperature, with increasing temperatures responsible for decreasing resistivity. Thus, (local) heating of the plasma will decrease the (local) resistivity, which in turn will tend to increase the (local) current flow. Thus, while NBI drives significant non-inductive current when injected tangentially, its primary effect is as a method of plasma heating, which has the collateral effect of broadening the current profile and increasing the plasma current for the same loop voltage. In a similar manner, electron cyclotron heating (ECH), when used for heating the plasma, can also have the effect of increasing the local current in the (highly localized) deposition region.

In most cases, the NTM requires a triggering instability such as an ELM (see previous section) to produce a *seed island*, which will grow if the NTM is unstable [2]. One way of controlling the NTM is, therefore, to ensure that such triggering instabilities do not occur. Unfortunately, experimental evidence suggests that there are also seedless tearing modes, or NTMs triggered by background plasma turbulence, which is difficult or impossible to completely eliminate [3]. For this reason, relying on eliminating events that may produce a seed island may not be a sufficiently robust approach to NTM control. An alternate method for controlling NTMs using auxiliary current drive has been successfully demonstrated on several tokamaks. This approach replaces the current lost in the process of flattening the profiles across the island, thereby shrinking the island, restoring the nested flux surface magnetic topology [Fig. 1(a)], and stabilizing the mode [4]. As an alternative to this local control of the current profile (sidebar S5 of the Attachment), global control of the current profile may also prevent NTMs. (See section on profile control.)

Current can be driven at the flux surface containing islands using a variety of methods. One method that can produce highly localized current drive and has been successful in stabilizing the NTM is electron cyclotron current drive (ECCD; see sidebar S13) [5]. This technique drives current in regions typically a few centimeters wide by injecting microwave frequency electromagnetic waves that resonate with the cyclotron orbits (Fig. 9 of the Attachment) of the current-carrying plasma electrons. In tokamaks such waves are typically produced by high power gyrotrons, similar to the wave generators used in satellite communications. The electron cyclotron frequency depends primarily on the toroidal magnetic field, so the current is driven where the injected wave path intersects the major radial location at which the waves are resonant. Figure 3 illustrates this geometry in a DIII–D discharge. The injection chord is typically not straight owing to refractive effects of the plasma. For large tokamaks operating today (e.g., JET, DIII–D, ASDEX-U) injected power on the order of a few megawatts is required to produce the tens of kiloamps of plasma current needed to stabilize the mode [6]. Because islands are typically a few centimeters wide themselves, the alignment of island and deposition locations must be accomplished with accuracy on the order of a centimeter as well. The deposition need only lie on the flux surface containing the NTM islands, rather than directly threading the center of each helical island itself. However, actually driving current within each island (and not outside) would serve to reduce the power required to stabilize the mode.

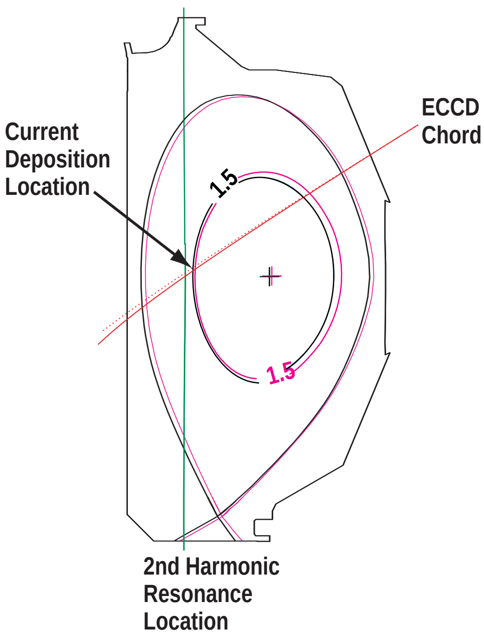

Figure 3. Geometry of current drive to suppress NTM islands. Driving current at the flux surface which contains islands can restore the current lost in island formation. This shrinks and can even eliminate the island, stabilizing the mode. In the case shown, the flux surface (for two different times during the discharge) being affected corresponds to the 3/2 NTM, and is labeled by the safety factor contour value of 1.5. The location where current is actually deposited is slightly offset from the second harmonic resonance due to a Doppler shift.

The timescale for ECCD-driven current to rise in a typical present-day tokamak is tens of milliseconds, comparable to the time required for the island to grow to a saturated state or to respond to the current drive and reduce in size (providing the deposition region is sufficiently well aligned with the island flux surface to produce suppression). Several mechanisms are available to perform the alignment on roughly this same timescale. One approach is to actively vary the wave launcher mirror angle, which in turn varies the angle of the injection wave path (Fig. 3) [5]. This has the advantage of leaving the plasma equilibrium characteristics unchanged. Another approach is to vary the toroidal field. This moves the major radial location of the harmonic resonance (and thus the deposition location) back and forth relative to the island (see Fig. 3). This leaves the plasma shape and position unchanged, allowing divertor pumping (sidebar S9 of the Attachment) and stability characteristics that depend on the shape to be held constant. This approach has been successfully used in stabilizing both 3/2 and 2/1 NTMs [8]. Still another approach involves movement of the plasma position (and thus the island position) relative to the approximately fixed deposition location. Moving the plasma radially can be accomplished with approximately fixed divertor pumping and plasma shape, but moving the plasma vertically tends to significantly affect the divertor configuration (although with little effect on the shape). Nevertheless, while modifying the toroidal field typically requires >100 ms owing to the large L/R time of toroidal field coils, varying the plasma position requires <10 ms. This speed advantage means that plasma position control can produce adequate alignment in a shorter time, and phase lags in the control action and island response are significantly reduced. Launcher angle control can, in principle, be comparable in speed to plasma position control.

The central problem in using these methods to align the current drive deposition and island locations arise from the difficulties in determining the location of the island flux surface and the deposition location in real time. Present-day diagnostics and equilibrium reconstructions produce estimates of the island flux surface location with typical accuracy of ±1 to 1.5 cm, comparable to the control accuracy of the plasma position and toroidal field themselves. In addition, determination of the deposition location requires complex computation that at present cannot be accomplished in realtime. A suite of search and tracking algorithms has been developed on DIII–D to address these difficulties and produce successful, sustained NTM suppression. In particular, the DIII–D NTM control system makes use of three coupled algorithms: the Search and Suppress, Active Tracking, and Target Lock routines. Each of these can affect any of three different island/ECCD alignment control quantities: the plasma major radial position, the toroidal field, or the plasma vertical position. The operation mode that has been most successfully and routinely applied to date is a combined Search/Suppress and Active Tracking mode (Fig. 4). When the control phase is enabled in this mode, the control algorithm fixes the selected control actuator quantity for a specified dwell time to determine whether the degree of alignment is sufficient to suppress the mode. If at the end of this dwell time the (filtered) mode amplitude has been reduced at a sufficiently high rate relative to a specified threshold rate, or has actually fallen below a specified threshold amplitude, the algorithm continues to hold the control quantity fixed.

If the rate of suppression has not exceeded the threshold rate, the algorithm executes a search by incrementing the control quantity by a specified amount, freezing the control quantity for another dwell time, and examining the resulting effect on the mode. This search/dwell/search sequence continues until a specified limit is reached, and the search reverses the sign of the control quantity increment. This process continues until sufficient alignment is detected and the mode is suppressed below the specified amplitude threshold.

Once a sufficient alignment is detected, the control quantity is nominally frozen, and the Active Tracking algorithm is engaged. This algorithm adjusts the control quantity to maintain alignment in the absence of the mode. The required adjustment is determined by either a linear or nonlinear predictor calculation based on magnetic measurements, or on an explicit equilibrium reconstruction including measurement of

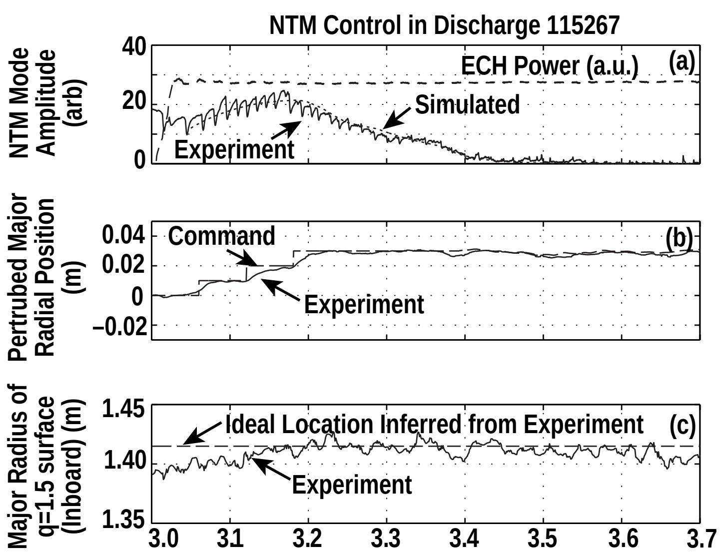

Figure 4. The DIII–D Search and Suppress algorithm with Active Tracking has been successful in suppressing and sustaining stabilization of the NTM. Search and Suppress is engaged at approximately 3.0 s and Active Tracking is engaged at approximately 3.4 s. (a) Comparison of model predicted and experimental amplitude of the NTM mode (related to the island width), and input ECH power (b) change in the plasma major radius to achieve and maintain alignment between NTM island and ECCD deposition location, (c) major radius of the 3/2 NTM island flux surface (dashed line = empirically determined ECCD deposition location).

the internal magnetic topology with motional Stark effect (MSE) sensors [9]. The MSE sensors measure the local ratio of poloidal and toroidal magnetic fields at many points within the plasma. Several predictors are available to the routine, including neural network-based algorithms. These are trained on previous experimental discharge data or artificially generated data to produce an estimate of the position of the relevant q-surface relative to the plasma geometric centroid. If toroidal field control is being used, the required toroidal field correction is derived from the predicted perturbation in the relative major radial position. While neural network predictors have been very successfully applied in DIII–D, flux surface reconstruction based on direct magnetic measurements has provided the most accurate and reliable sustained alignment in recent experimental campaigns.

Design of the parameters governing the Search and Suppress and Active Tracking algorithms was accomplished using accurate, experimentally validated dynamic models of NTM island response to ECCD and plasma response to position commands. Figure 4 shows a comparison of a mode suppression model and the experimental response to variation in the degree of alignment between island and ECCD location (top frame). The model response, based on a simplified version of the modified Rutherford equation that describes island dynamics, shows sufficiently accurate representation of island response dynamics to allow for good control design [3]. In the case shown, the plasma major radius (middle frame) was varied to adjust the alignment of the q=1.5 surface with the ECCD deposition location. The NTM control algorithm is also integrated with a special plasma shape/position regulation scheme that fixes the strike points (Fig. 14 of the Attachment) while varying the major radius to adjust the island/ECCD alignment. This allows for constant divertor pumping (sidebar S9 of the Attachment) while NTM suppression is performed. Following suppression of the mode, where it is assumed that deposition is aligned with the q=3/2 surface, the Active Tracking algorithm is enabled to compensate for variations in the q=1.5 surface due to changes in the current profile and poloidal beta. The Active Tracking action can be seen in the fluctuating major radius perturbations following suppression at $t \approx 3.4$ s.

Figure 4 also summarizes experimental results from use of the Search and Suppress with Active Tracking. The growth of the mode is slowed even when the island and ECCD are misaligned by as much as 1.5 to 2 cm. Adjustment of the major radius by the Search and Suppress produces sufficient alignment to fully suppress the mode within 200 ms. After design and testing of the basic scheme, proper functioning of the algorithm requires specification of various parameters and thresholds to match the dynamic characteristics of the target equilibrium. Dynamic models of NTM response were sufficiently reliable that control system parameters designed using these models produced successful suppression and tracking of the evolution of the profile to maintain island/ECCD alignment the first time this integrated active suppression was attempted experimentally.

An alternative to the Search/Suppress algorithm is the Target Lock algorithm. This control scheme uses the observed response of the mode amplitude to either natural fluctuations or preprogrammed variations in the control quantity to infer the proximity to ideal alignment. An approximate form of the modified Rutherford equation is implemented in this algorithm to provide an estimate of the expected mode decay or growth rate based on the degree of misalignment.

Use of Search and Suppress, Active Tracking, and Target Lock algorithms has enabled full and sustained suppression of both 3/2 and 2/1 NTMs in DIII–D under closed loop control without previous experimental determination of the ideal alignment location. Suppression of the 3/2 NTM has allowed operation at normalized beta values approximately 50% above the value achieved in the presence of the unsuppressed mode (from β_N=2.3 to β_N=3.4) [3]. The duration of increased-beta operation is presently limited only by the length of time that the gyrotrons can inject power into the plasma.

Eventual application of NTM suppression to fusion reactors envisions several possible modes of operation. In one scenario, current drive will be applied in steady state at the relevant flux surfaces to pre-emptively suppress any seed islands that may be triggered by background MHD instabilities or turbulence. Because the current drive source would have to operate constantly in this scenario, this would require a large amount of auxiliary power. This approach would also require constant tracking of the deposition location and target flux surfaces to maintain alignment. Another approach is to detect the presence of NTM islands and enable a suppression system to suppress the mode as rapidly as possible, hopefully before a 2/1 island reaches a potentially disruptive saturated size, or a 3/2 island significantly degrades confinement. This demands a very rapid island acquisition and alignment system to achieve suppression within tens of milliseconds after the onset of island growth. Of course, constant and accurate determination of the flux surface geometry and computation of the expected deposition location would allow the system to be engaged very quickly and could meet this suppression time requirement. Gyrotrons, for example, can be ramped to full power in significantly less than 10 ms.

Future Directions

It is now widely accepted that localized current drive can replace the missing bootstrap current that characterizes the NTM, thereby stabilizing the mode. Several closed loop feedback approaches have demonstrated NTM suppression sustained for several seconds on various tokamaks. However, before current drive suppression can be used effectively in a reactor-grade plasma, several key capabilities must be demonstrated. The principal requirement in a reactor is reliable, simultaneous, and steady state suppression of both the 3/2 mode (which mainly degrades confinement) and the 2/1 mode (which can lead to a disruption). Such a demonstration will require more installed current drive power than any machine presently has, independently steerable launchers, more accurate and reliable realtime reconstruction of internal magnetic surfaces, accurate realtime determination of the current deposition location, and algorithms that can deal with changing equilibrium and machine conditions. Work has begun on installing steerable launchers in several devices, and realtime steering has been demonstrated in JT–60U.

Realtime magnetic surface reconstruction has been demonstrated in DIII–D, but present levels of noise and accuracy have limited the effectiveness of suppression over long periods. Algorithms for detecting the location of optimal alignment and maintaining alignment once the mode is suppressed have been developed, but leave much room for improvement. The field of NTM suppression in tokamaks is rapidly evolving, but would benefit greatly from contributions in the areas of improved control algorithms and approaches, estimation, realtime computation, actuator technology, and diagnostic signal interpretation.

DETECTION, CORRECTION, AND MITIGATION OF OFF-NORMAL EVENTS

An off-normal event is any event that would not normally occur during well-controlled routine steady state operation. There are a large number of such events, some examples of which will be described in this section. The impact of off-normal events in a tokamak can be summarized according to the severity of their consequences:

1. Risk of Personnel Safety. These are events that have the potential to cause harm to operating personnel or to the general public. A categorization of off-normal events from a safety point of view is given in [10].
2. Risk of Equipment Safety. These are events that have the potential to cause damage to either the device or the facility, but do not risk the safety of personnel.
3. Performance Degradation. These are events that can cause performance of tokamak operation to degrade, but do not create risk to either personnel or equipment.

The primary personnel risks at a tokamak facility are due to the high currents and voltages used in confining and heating the plasma, radiation from activated materials or tritium release, and conventional process control risks such as use of toxic chemicals and systems under pressure or with extremely high or low temperatures. Most of the approaches for dealing with these risks are rather conventional and are already in use at the major facilities, so will not be discussed in this section. The highest risks to the device or facility derive primarily from the large energy content in the plasma while the tokamak is operating. Off-normal events that lead to plasma termination such as major disruptions (sidebar S14) can deposit a significant amount of this energy onto plasma facing components (PFC) or create large and potentially damaging mechanical forces.

The handling of off-normal events represents another highly coupled category of control in tokamaks, and includes detection and identification of the off-normal event, determination and execution of corrective action when possible, and execution of a mitigating response when correction or recovery is not possible. In future power producing reactors, the system responsible for off-normal event detection and response must be thoroughly integrated with many other control subsystems in order to provide effective and appropriate action coordinated with other plasma control responses. Such responses must also be well-integrated with the overall safety and supervisory system.

Responses to off-normal events generally depend on the risk associated with the event. A *fast plasma termination* would usually be initiated if an event poses an imminent threat to personnel or to the device or facility. A fast plasma termination can cause structural damage by localized deposition of thermal or magnetic energy, however. To mitigate these effects, termination techniques that distribute the plasma energy more homogeneously are necessary. A *controlled plasma termination* is typically initiated when continued operation could lead to a potential risk situation. In present devices, discharges are typically not terminated for reasons of reduced performance, although this may not be the case with future devices. Usually, efforts are made to recover from the performance reducing event, either in the same discharge or in those that follow.

A fast plasma termination is not always the result of a controlled safety procedure. Sometimes it is the consequence of a control failure. This loss of controllability originates either in the failure of a component or subsystem of the feedback loop, or in the operation of the system at operating points for which the controller was not designed or at operating conditions that trigger other types of instabilities.

In the following sections, we provide descriptions of some methods for handling disruptions, system fault detection and isolation, and performance optimization. These are only a few examples of a much broader collection of problems that must be solved to enable a viable power producing fusion reactor.

Disruption Avoidance and Mitigation

Erosion due to disruption may consume a high percentage of the designed plasma facing components (PFC) lifetime. Sometimes, long-term reconditioning of plasma-facing surfaces after disruptions will be required before normal plasma operation can be resumed. It is desirable to avoid the occurrence of disruptions whenever possible and to reduce the direct and consequential effects of such disruptions. The approaches to disruption avoidance can be divided into two categories [11]: (1) a priori avoidance of the operation conditions that lead to disruption, and (2) active intervention in a discharge scenario after early prediction of disruption onset.

Disruptions can in principle, be avoided during tokamak operation by the use of a *disruption-free* scenario (Appendix sidebar S18) that avoids the various operational limits and conditions that cause disruptions, and by the provision of adequately reliable plasma operation and control systems such that all critical parameters of the prescribed scenario can be reliably obtained and repeated. However, optimizing plasma performance often requires operating near disruption limits.

Disruption-free scenarios are often based upon conservative plasma operation parameters that do not press close to known disruption-initiating limits or plasma control limits. A priori disruption avoidance procedures can in principle be extended to operation scenarios that come closer to several operational limits. In these cases, observation of the limits involved and provision of real-time disruption prediction or onset warning capability become important. Using a warning indicator of a potentially impending onset of disruption to effect feedback controlled intervention can lead to reliable operation near a limit that can initiate disruption. Basing disruption prediction on single parameter proximity or the confluence of several single parameter limits may not necessarily provide complete certainty for disruption avoidance, or conversely, may unduly restrict the accessible operation domain. A possible improvement can be made by implementation of a neural network disruption predictor, wherein multiple disruption-related indicators or diagnostic signals are combined via a neural network to provide a composite impending disruption warning indicator that is more robust and reliable than simple single- or multiple-parameter indicators. For example, after training, neural networks were able to successfully predict disruptions in DIII–D [12] and classify disruptions in JET [13]. Enhanced predicting capabilities (85%) were also achieved in ASDEX-U using a neural network disruption predictor [14]. More complex systems [15] are capable of disruption prediction with a probability of 95%.

Ideally, an impending disruptive MHD instability would be detected and avoided by modifying the target equilibrium, heating, density, or other operating parameters. Should such avoidance not be possible, mitigation of its effects is required. PFC damage is primarily due to an excessive surface temperature rise leading to melting or ablation, and this situation cannot be ameliorated through improved heat removal capability on the PFC heat sink. Effective thermal mitigation approaches involve maximizing the time over which the energy is released or expanding as much as possible the region over which the energy is deposited. This goal can be achieved by the fast injection of impurities. UV line radiation from the injected impurities distributes the plasma energy more uniformly on the first wall, reducing the thermal load to the divertor (sidebar S9 of the Attachment). The most effective injection methods are solid pellet injection [16] and intense gas puffing [17].

Fault Detection and Isolation

During operation of a tokamak, hundreds of subsystems must operate correctly and simultaneously for a completely successful plasma discharge. On many devices, verifying proper operation of the subsystems most prone to failure is done manually by human operators after the discharge. Because of the tedious nature of this task and the large number of systems, inoperative or malfunctioning systems can sometimes remain undetected until several experimental discharges have passed. Occasionally, problems will not be detected until days later if the failure does not prevent operation. Due to the increasing complexity of tokamak experimental devices, some efforts have been made toward development of automated fault detection systems.

Fault detection and isolation (FDI) techniques have been under development for the last three decades. In the course of this development an FDI approach based on analytical redundancy has emerged in place of the more traditional FDI approach based on hardware redundancy [18]. In the latter approach, redundant physical subsystems (multiple sensors, for example) are built and their output signals are compared for consistency. In the event of failure, a subsystem backup is switched in. Higher costs, additional space requirements, and complexity make this approach often unattractive. In the former approach, the inherent redundancy contained in the static and dynamic relationships among the inputs and outputs of the system is exploited for fault detection and isolation. The measurements of the system inputs and outputs are processed analytically to estimate the value of a desired variable. This estimate can be generated using either quantitative or qualitative models. For quantitative (mathematical) models, the model-based predictors are state estimation, parameter identification, and parity space. For qualitative (non-mathematical) models, the prediction is based on decision-table-based methods, knowledge-based expert systems, and neural-network-based methods. The estimated variable is then compared with the measured value of the variable to generate a residual. Deviation of this residual from the normal behavior significant enough to indicate a failure must then be detected by methods of change detection such as Bayes decision and hypothesis testing. Each type of fault is characterized by a specific combination of symptoms. Classification methods such as fuzzy clustering, artificial neural networks, and geometrical distance are used to determine the type of fault [19]. If more information about the relations between symptoms and faults is available in the form of diagnostic models, methods of reasoning can be applied. The reasoning strategies for fault diagnosis are probabilistic reasoning, rule-based reasoning, sign directed graph, fault symptom tree, and fuzzy logic. Due to the complexity of tokamak systems and the large number of variables to be monitored, current efforts toward fault detection and isolation in tokamaks are based primarily on qualitative approaches.

At JET, an automatic modular sensor fault detection and classification (SFDC) [20] system has been built for the sensors measuring the vertical mechanical stresses on the supports of the vacuum vessel of the tokamak. Experts are interested in the reliability of these measurements during specific time windows, corresponding to the occurrence of disruptions. During the usual operational life of the tokamak, in fact, mechanical stresses are weak and do not need to be monitored. During a disruption, on the contrary, fast dynamic vertical displacement events (VDEs) (sidebar S14) occur, causing an impulsive force and mechanical oscillations of the vacuum vessel, which must be monitored to assure the mechanical integrity of the machine. One of the actions related to the mechanical monitoring of the stresses is to suspend the experimental campaign when more than a fixed number of VDEs trespassing a certain stress threshold occurs in a day. The reliability of the measurements is therefore very important to avoid both unmotivated suspension of the experiments and dangerous experiments carried out above the operational limits. The strategy used is based on a modular system that consists of two stages. The first stage consists of a multilayer perceptron (MLP) neural network that has been trained to predict some features of the considered signals on the basis of a selected set of inputs [21]. The predicted features provided by the neural model can be compared with the actual feature values, in order to raise alarms indicating sensor faults if the corresponding residuals are too high. This task is part of the fault detection phase, which consists of revealing the presence of a fault. The second stage focuses on fault classification, which is accomplished by a fuzzy inference system (FIS). In this case, the fault classification rules were established on the basis of manual fault classification previously performed by experts. The tuning of the membership functions was set by trial and error, taking sensor accuracy and disturbance level into account. As mentioned above, validation of stress measurements is only relevant during disruptions. The occurrence of a disruption is notified by a suitable flag provided by an automatic disruption detection tool already in existence at JET. It is possible, in this way, to select the experiments in which the sensor validation should be performed.

At present, an expert-system-based fault detection system is used routinely during DIII–D operations and has led to an increase in tokamak productivity. The Fault Identification and Communication System (FICS) [22] executes automatically after every plasma discharge to check dozens of device subsystems for proper operation and to communicate the test results to the tokamak operator. The primary purposes of FICS are fault detection and fault prediction. Fault detection refers to determining which systems were not working properly during an experiment, even if they do not cause the loss of the experimental discharge (shot). Fault prediction refers to determining which systems look like they are having trouble and may cause a future fault. This includes detecting programming errors, that is, determining whether the operator specified setup for a shot is self-consistent. The core of FICS uses the public domain software package called CLIPS (C language integrated production system), originally developed by NASA [23]. CLIPS, a computer language designed for implementing expert systems, provides two powerful capabilities not provided by conventional programming languages. *Chaining* provides the ability to emulate a human chain of reasoning in software. *Data driven* execution enables the expert defined rules to activate as soon as a knowledge or data source becomes available. The CLIPS shell performs the inferences, executing rules of the form “if A, then B”. CLIPS in its simplest form consists of facts and rules. Rules are executed when specified facts are asserted (A in the clause “if A, then B”). The consequences of a rule execution are to assert other facts (B), which can then execute other rules, and so on. A very important side effect of rule execution is the ability to activate functions that extract and manipulate data and return the results of those manipulations to the expert system. This information can then also be used to assert more facts to drive other rules, and so on. The order of rule execution can be influenced by a priority value assigned to each rule. Rules are executed according to their relative priority and according to when their data becomes available. Since the program is data driven, each rule executes if and when the data shows up. If a particular piece of data is not available, tests that require that data do not execute and an alarm is raised saying the needed data was not acquired.

The success of any fault detection system is linked to the availability of adequate measured data. Present devices already acquire a huge amount of experimental data, sometimes several Gigabytes per discharge of only a few seconds duration. Steady state operation will put significant demands on data storage, even with relatively slow rates of data acquisition. On the other hand, the data acquisition system must be able to capture rapid and unpredictable changes for use in fault detection and identification. For this purpose, new data acquisition methods were developed for the TRIAM–1M [24] and JET [25] tokamaks in order to combine coarse data from quiescent steady-state phases with fine data from rapid and unpredictable transitions.

Performance Optimization through Event Handling

Catastrophic loss of plasma or failure of tokamak systems are the most obvious off-normal events. Other events are more subtle in that they simply degrade performance. Thus methods to handle and correct these events can be thought of as performance optimization. A good example of this is the detection and handling of the H–mode to L–mode back-transition (sidebar S11) that occurs in ASDEX-U [26]. The performance control in ASDEX-U is dedicated to the control of plasma characteristics like confinement or radiative behavior, which can be relatively well separated from those related to the plasma position and shape control. Several process controllers, usually simple single-variable PI controllers, have been implemented to control different characteristics of the plasma such as various forms of density, temperature, and pressure as well as fueling mixtures and power flows. These basic process controllers are combined to define so-called control recipes that are switched on and off during the discharge. Activation of the control recipes at preset times during the discharge to optimize some plasma characteristic is not practical because it is difficult, if not impossible, to predict the conditions of the plasma. A real-time algorithm for plasma regime recognition was therefore developed to be used as the trigger mechanism for the control recipes. The first version implemented identifies five different confinement regimes (sidebar S11): Ohmic Phase (OH), standard L–Mode (L), standard H–Mode (H), highly radiating L–Mode (HRL), highly radiating H–mode (HRH). This algorithm allows the switching of control recipes to recover plasma performance. For example, the working point for best HRH plasma performance is close to the H–mode to L–Mode back transition, so occasional back transitions occur. When they do, the regime recognition algorithm is used to dynamically switch to another control recipe to restore the desired HRH mode.

S11. Confinement Modes (sidebar)

The *energy confinement time,* τ_E (sidebar S2 of the Attachment), is an important measure of the ability of a plasma to retain energy and thereby support continuing fusion reactions. In the tokamak experimental community, plasmas are typically categorized as one of four different *confinement modes* when describing their confinement properties. The first of these is the ohmic plasma, which is heated only by the ohmic transformer action (sidebar S13). The three remaining modes all use methods of auxiliary heating as well as ohmic heating. In order of increased confinement, these are the *low confinement mode* or *L–mode*, the *high confinement mode* or *H–mode*, and several versions of very high confinement modes. These modes are primarily distinguished by the shape of their temperature, pressure, and density profiles (sidebar S5 of the Attachment) with the higher confinement modes exhibiting steeper gradients there. The terms ohmic plasma, L–mode, and H–mode have become more or less accepted terminology, while the term describing the highest confinement regime varies between institutions.

Even though higher confinement regimes have been identified, H–mode [28] is the target regime for the advanced tokamak (AT) concept, because the higher confinement modes have only been produced transiently, always being terminated by a severe MHD instability. The AT requirement for active control has driven much of the recent research in methods of control for tokamak quantities.

Future Directions

Methods for ensuring personnel safety have been largely systematized in existing tokamak devices. Anticipated new dangers in next generation devices include the use of larger amounts of radioactive tritium, greater neutron production within the device, and the resulting neutron activated plasma facing materials. It is not expected that this will present any major difficulties because new facilities will be able to draw on many decades of nuclear industry experience. In addition, various types of loss of cryogenic (for superconducting coils) coolant accidents are envisioned that may or may not present a danger to personnel, but will certainly endanger the device.

A number of new methods will be needed to deal with other types of off-normal events. Many of the present methods of fault detection and diagnosis for these events execute primarily off-line and between plasma pulses. These methods will need to be converted to on-line detection algorithms when steady state devices are put into use. A great deal of the required knowledge for how to detect faulty tokamak systems now resides only in the minds of experienced operators. This knowledge will need to be captured and models of various tokamak support systems will need to be developed in order to incorporate them into on-line fault diagnosis software. Present day devices typically do not maintain the functional models of tokamak support systems that would be necessary for fault detection.

For those faults that can lead to device safety issues, methods for controlled shutdown of the plasma pulse must be further developed and must be expanded to include a greater number of faults. In particular, reliable and safe on-line methods for remediation of large energy events must be implemented and proven. For example, mitigation of plasma disruptions through massive gas puffing is now being implemented and tested on several experimental devices but, even at DIII–D where these methods were pioneered, they are not yet a part of routine operation.

EDGE LOCALIZED MODES

A defining feature of high energy confinement (H–mode; see sidebar S11) tokamak plasmas is the existence of a region of significantly reduced thermal and particle diffusion near the plasma boundary, called the *edge transport barrier* or ETB (sidebar S12). Although the high performance of H–mode plasmas results from high pedestal temperatures and densities (Fig. 6) obtainable with an ETB, the very low transport in the edge region generally leads to a rise in the gradient of the pressure distribution near the edge, which triggers an instability localized to the plasma edge known as the *edge localized mode* (ELM) [27].

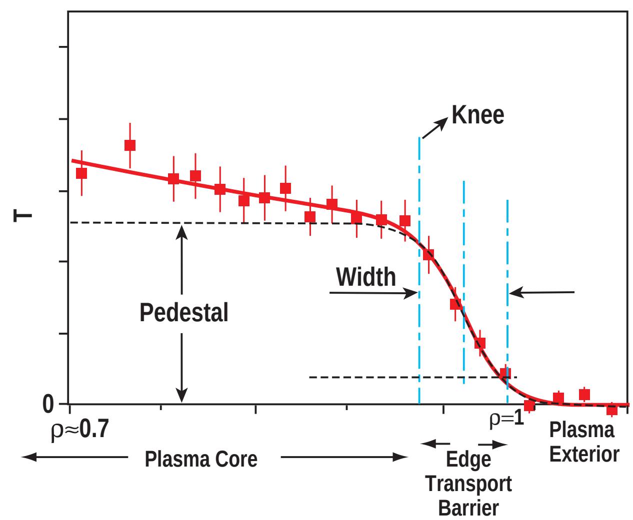

Figure 5. Illustration of definitions relevant to ELMs. The solid red curve is fitted to measured electron temperature data. Lines through the data points indicate plus or minus one standard deviation (estimated) from measurements. The *edge transport barrier* (ETB) is a region at the edge of the plasma that is a barrier to the *transport* (diffusion) of heat and particles out of the plasma. The ETB and other forms of transport barrier are characteristic of high energy confinement mode (H–mode) plasmas, since they tend to prevent heat from escaping the plasma. Width of the ETB is defined to be the width of the steep gradient region in the electron temperature profile. This edge region is defined to be the region between the knee of the fitting function and the plasma last closed flux surface. The *pedestal* in temperature coincides with the plasma interior region. (Unfortunately, terminology in the literature is inconsistent, with "pedestal" sometimes being used to refer to the edge region.) The pressure profile in H–mode plasmas takes a similar form, so in discussions of edge transport barriers, temperature and pressure are often used interchangeably. However, pressure can sometimes have a narrower steep gradient region. The close proximity of the high pressure region inside the ETB and the low pressure region outside the ETB is the source of the ELM instability.

S12. Edge Transport Barriers (sidebar)

Discussion of the ELM instability centers on transport and MHD behavior in the *edge transport barrier* of the plasma (Fig. 5). In contrast to global modes such as the RWM (see second section of this article), the ELM is primarily localized to this edge region. The height of the temperature pedestal (Fig. 5) acts almost as a multiplier for temperatures inside the plasma. The greater the pedestal height the greater the total energy content of the plasma. Inside the knee, the temperature is approximately 10 times that of the sun's center (1000 eV), while exterior to the plasma it is about 10 eV. The source of the ELM instability is the resulting large pressure gradient in the edge region. During an instability, a significant amount of the pedestal energy (up to 25%) can leave the plasma and be deposited on plasma facing components.

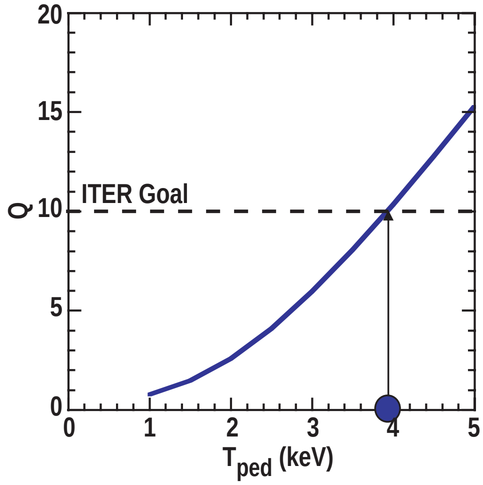

Figure 6. Plot of projected dependence of Q on pedestal temperature. Q is the ratio of fusion power output to additional heating power into the plasma and is a measure of efficiency of fusion power production. This plot assumes fixed available input power (40 MW) and constant density across the pedestal (Fig. 5). The pedestal temperature must be maintained at approximately 4 keV (4 thousand electron volts) in order to sustain the ITER minimum value of Q = 10.

ELMs are most often observed as bursts of light (measured using photodiodes) from excited hydrogen or deuterium atoms [Fig. 7(a)] in regions where the ELM power flux reaches the vessel wall. The magnetohydrodynamic instability that is believed to be responsible for the larger type of ELMs in Fig. 7(c) perturbs the magnetic field and can also be observed on magnetic probes mounted on the vessel wall [Fig. 8(b)].

The physics processes of the ELM are not completely understood although significant progress has been made. Theories based on ideal MHD (sidebar S3 of the Attachment) seem to be in good agreement with the early linear phases of the mode growth. The later nonlinear growth of the mode is an active area of research. Figure 9 shows the mode structure in the nonlinear growth phase as predicted by one proposed model [30]. This figure is consistent with the intuitive picture of ELMs — the nested flux surfaces characterizing ideal MHD are broken by the instability and, consequently, particles and heat can be removed through the plasma edge.

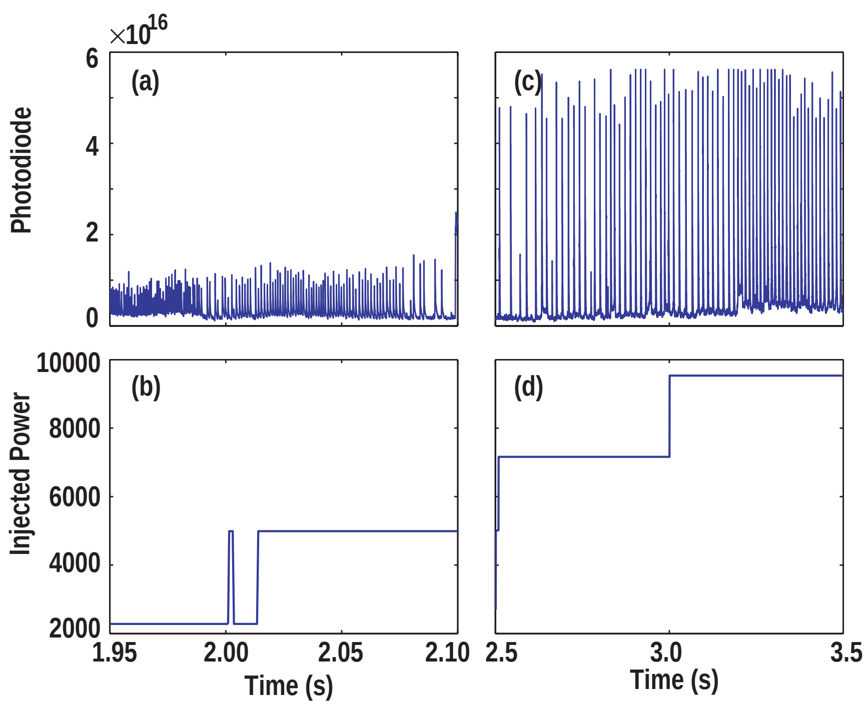

Figure 7. Illustration of different ELM types' dependence on injected power. (a) Detected light emission for Type III ELMs, (b) neutral beam power injected into the plasma simultaneous with measurement (a), (c) detected light emission for Type I ELMs, (d) neutral beam power injected into plasma simultaneous with (c). One characteristic which distinguishes Type I and Type III ELMs is their response to increases in injected power. The frequency of Type I ELMs increases with increasing power while the frequency of Type III decreases. Another characteristic is the typically much larger peak amplitude of Type I ELM light emissions.

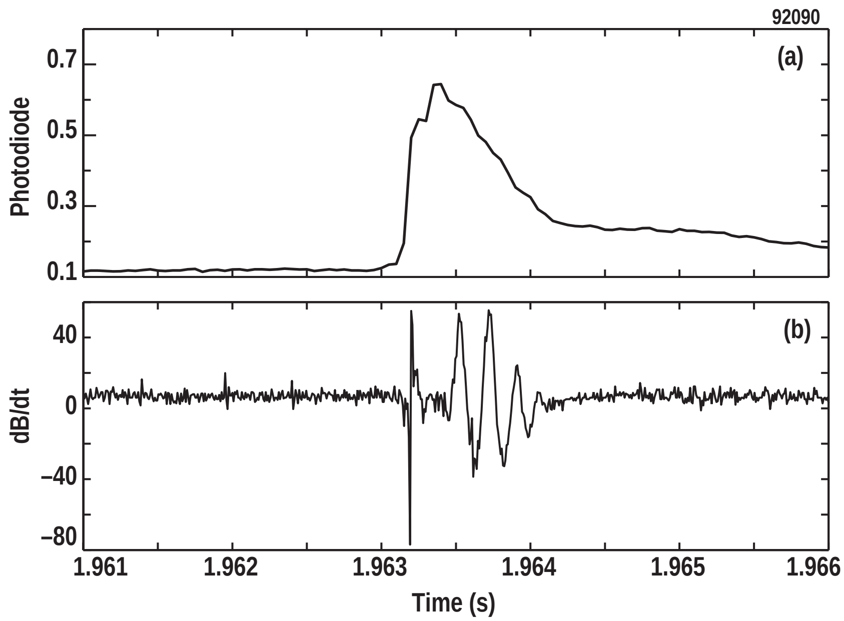

Figure 8. Illustration of the magnetic character of an ELM. (a) Expansion of a single instability similar to those shown in Fig. 7(c); (b) measurement of the derivative of magnetic field at the outboard midplane showing magnetic behavior simultaneous with the light emission in (a).

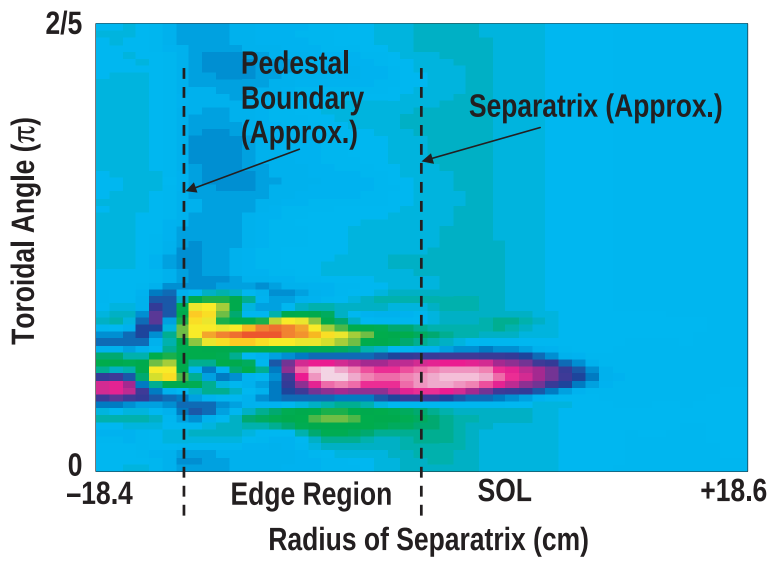

Figure 9. Intensity plot of perturbed density at the plasma outer midplane during the later nonlinear phase of simulation of the growth of the edge localized mode. The early phase of the mode growth is linear and approximately represented by ideal MHD. The local nature of the mode growth is illustrated with the “finger” of plasma radiating out from the plasma edge and toward the vacuum vessel wall. It also extends along the magnetic field (into and out of the page). Large transport through the walls of the finger or possibly the breaking off (magnetic reconnection) of the finger are possible mechanisms for the ELM energy loss.

The three ELM types that are observed in most tokamaks [27] are summarized in Table 1. They are experimentally distinguished by the dependence of the ELM frequency on heating power, df_{ELM}/dP (Fig. 7), the density of the plasma, and the shaping applied to the plasma cross section.

Table 1. Experimental characteristics of different ELM types. The ELM type with the widest operational range, Type I [Fig. 7(c)], allows large pressure in the ETB (the pedestal energy W_{PED} = pressure on inside edge of the ETB times the plasma volume) but the energy loss at each ELM, ΔW_{ELM}, is also large. (The symbol W is widely used in plasma physics to denote quantities of energy. Note that pedestal energy is not the same as total stored energy.) The Type II ELM has low ELM energy loss and high ETB pressure but is observed only in a limited range of plasma shaping and density that may not be applicable to a tokamak reactor. The Type III ELM [Fig. 7(a)] has a wider operational range than the Type II and low ELM energy loss but also reduced pressure in the ETB. The pedestal energy W_{PED} and density shown here are values relative to the conditions under which Type I ELMs occur.

Type	W_{PED}	ΔW_{ELM}	df_{ELM}/dP	Density	Shaping
I	1	0.05–0.25 W_{PED}	> 0	1	Any
II	1	< 0.01 W_{PED}	< 0	» 1	Strong
III	< 2/3	< 0.01 W_{PED}	< 0	» 1 or « 1	Any

Each ELM causes a collapse of the ETB and can cause a loss on the order of 5% of total plasma stored energy on a very short time scale. Although the ETB forms again following the ELM and the energy confinement of H–mode discharges with ELMs is still much superior to discharges with no ETB, the very large power loss during the ELM (in a tokamak reactor this is projected to be tens of gigawatts) creates severe difficulties in the design of the tokamak power exhaust handling structures (Fig. 10). Work on ELMs has therefore focused on developing techniques for either (1) reducing the ELM power loss (either by reducing the total ELM energy loss or extending the time over which the energy is lost), or (2) enhancing the particle transport in the ETB or possibly raising the ELM instability threshold, to keep the pressure gradient below the critical level that triggers the ELM. However these approaches must also maintain a high quality ETB for good overall confinement.

The second of the two ELM control techniques above would seem to be solved by operating with the so called “ELM-free” discharges sometimes produced in tokamaks. In ELM-free discharges, the plasma has a high critical (threshold) pressure gradient for the ELM instability and input power is kept sufficiently low that this critical pressure gradient is not reached. However this results in another difficulty. Impurities are continuously produced by plasma interaction with device components that face the plasma or as products of fusion reactions. Impurities entering the plasma edge are generally ionized in or near the ETB. They are then transported inward to the plasma core (by diffusion and by another mechanism that will not be discussed here) where they strongly interfere with fusion power production. ELMs act to reduce the source of these impurities by removing them periodically in their region of ionization. Thus ELM-free discharges typically exhibit impurity accumulation and are not considered to be viable solutions. Any technique that solves the ELM problem by eliminating the ELMs must also provide an alternate mechanism for reducing the impurities in the plasma. In addition to impurity removal, ELMs provide a mechanism for density control. If ELMs are eliminated, an alternate method for density control is also needed.

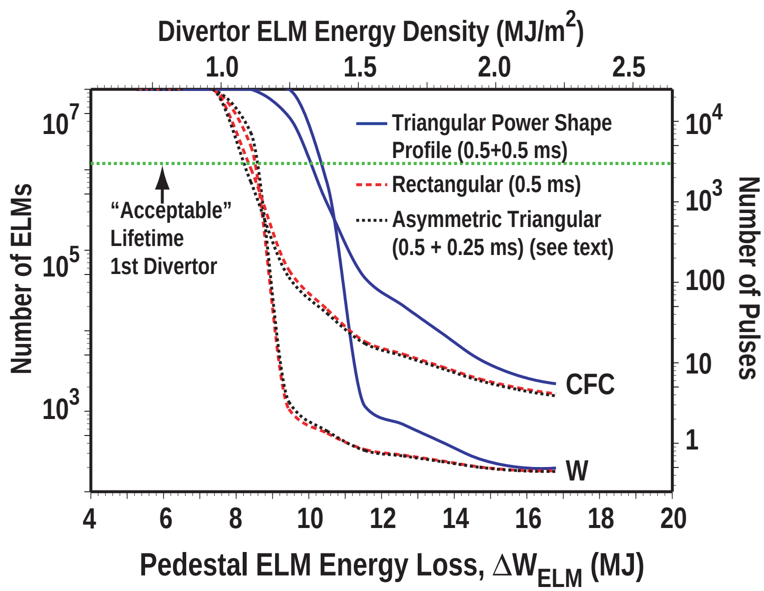

Figure 10. Expected erosion lifetime of ITER divertor plasma facing components (expressed in terms of number of ELMs or corresponding ITER full power pulses) as a function of ELM energy loss from the pedestal, ΔW_{ELM} (see Table 1). Curves are shown for two possible materials — carbon fiber composites (CFC) and tungsten (W) — and for three different approximations to the power signal during an ELM. The lifetime of the ITER divertor (sidebar S9 of the Attachment) drops quickly as the energy lost per ELM increases. Uncertainties in extrapolating expected ELM energies from present devices make it difficult to know precisely what to expect in the ITER device. (Reproduced from Ref. 29 by permission of G. Federici and the Institute of Physics.)

Experimental Approaches to Control

Several promising approaches for reducing or eliminating ELMs have been explored, but none yet presents a clear solution path. An ideal method would provide transport of particles out of the plasma, but without the associated transport of heat that degrades confinement. None of the currently proposed control methods involves feedback.

Several experiments have shown that regular injection of deuterium fuel or impurities in the form of frozen gas pellets or by gas puffing can trigger ELMs and/or change their characteristics. For example, experiments at the ASDEX-Upgrade (ASDEX-U) tokamak [31] have shown that repetitive deuterium pellet injection can trigger more frequent, smaller ELMs. Energy losses are thus spread over longer times, with smaller peak losses. There is some resulting loss of confinement, but not as severe as with naturally occurring ELMs. Using this method, only ELMs initiated by pellets occur. Gas fueling has also been shown to produce ELMs of smaller size. In cases where the Type I ELM size is reduced with increased gas fueling, the reduced ELM size is believed to be caused by a narrowing of the steep pressure gradient region near the separatrix relative to the ETB width (see Fig. 5). (The narrower pressure gradient region enables instabilities to be triggered more easily, allowing less time for energy to build up between ELMs.) In present experiments this effect also results in reduced core stored energy. However, in a reactor scale tokamak, it is speculated that a confluence of factors enabled by the expected much higher pedestal temperature may mitigate this reduction in performance.

One class of experiments has shown that Type I ELMs can be converted to Type II or III by means of oscillation of the plasma position. At TCV [32], vertical position oscillations of a few millimeters induced higher frequency ELMs, some apparently locking to the oscillation frequency. If this is done on a frequency greater than the natural ELM frequency, the ELM size is reduced. The cause of the ELM triggering is unclear at present although it is speculated that the vertical motion induces current in the plasma edge triggering the instability. It is unknown whether this technique will be compatible with good performance in a reactor scale tokamak.

Experimentalists at the Alcator C–Mod tokamak have discovered an ELM-free H–mode regime called enhanced D-alpha (EDA) [33] having enhanced particle transport without an accompanying increased energy transport. This regime is characterized by the presence of high frequency (>100 kHz) fluctuations in the edge that seem to provide the necessary mechanism for enhanced particle transport for density control and for removal of impurities. The EDA has the desired characteristics, but other devices have so far been unable to reproduce it (except, perhaps, for the JFT-2M tokamak in Japan [34]). In addition, the continuous edge instability is thought to be associated with high edge resistivity (sidebar S13) and thus may not occur in a reactor scale tokamak.

Experiments at DIII–D have demonstrated a new regime known as Quiescent H–mode (QH–mode) [35] that allows ELM-free operation, with a key feature being the presence of an edge electromagnetic oscillation known as the edge harmonic oscillation (EHO) [36]. The EHO enhances the particle transport through the edge without significantly increasing the thermal transport. The QH–mode operation has also been demonstrated at ASDEX-U [37] and QH–mode periods have been seen in discharges in JET and JT60–U. Although QH–mode occurs in a more reactor-relevant regime than EDA, it appears to require some level of toroidal rotation of the edge plasma in the direction opposite to the direction of the plasma current. The physics of the QH–mode is not understood presently and its applicability to a reactor scale tokamak is unclear.

Evans, et al. [38] has demonstrated a method for suppression of large ELMs in high confinement plasmas by creating a *stochastic magnetic boundary*. A stochastic boundary refers to a randomization of magnetic flux contours at the plasma edge in place of the nested contours characteristic of ideal MHD (sidebar S3 of the Attachment). Experiments at DIII–D and subsequent analysis show that imposition of a non-axisymmetric field can randomize the flux at the plasma edge and provide a means for steady state transport of particles out of the plasma, in contrast to the impulsive transport of ELMs. ELMs are reduced or eliminated while maintaining a high H–mode pedestal (Fig. 5). Present stochastic boundary experiments used steady state non-axisymmetric magnetic perturbations, which have the side effect of slowing plasma rotation, and could result in destabilizating the RWM instability. (See section on stabilization of the RWM.) Proposed methods would use an oscillating field perturbation, which may not have this problem. This approach is promising for ITER, because it is practical to implement, but the physics is not well understood yet. Although the method was successful in experiments with an ITER shape scaled to fit the DIII–D vacuum vessel, extrapolation to the complete set of ITER plasma parameters has not yet been shown.

Future Directions

An H–mode regime with Type I ELMs has been chosen as the standard operating scenario (Appendix sidebar S18) for ITER [39] because it is capable of being sustained in steady state with high confinement. On-going investigations of ELMs have focused on evaluating the adequacy of this choice as well as considering possible alternatives in the event that this choice proves inadequate. These alternatives consist of both modified operating regimes that avoid Type I ELMs and open loop methods to alter the character of the ELM instability. It is a bit premature to discuss issues of feedback control for ELMs, since the physics mechanisms are not well understood yet. However, there are often precursors to ELM events, sometimes several milliseconds long. This suggests the possibility that such precursors might be used to preempt or to mitigate the ELMs. This has not been done so far because the complicated and variable mode structure and high growth rates (10 μs) for the instability in the case of Type I ELMs present a very challenging control problem. The question remains whether a feedback control strategy would be feasible once the mode and its precursors were better understood.

STABILIZATION OF RESISTIVE WALL MODES

In this section, we discuss one of the major tokamak non-axisymmetric instabilities — the resistive wall mode (RWM). The RWM is a form of plasma *kink instability* under the influence of a resistive wall. The word kink is very appropriate to characterize the RWM behavior, which is similar to a garden hose kinking when it is suddenly pressurized (Fig. 11). The entire plasma configuration deforms in a helically symmetric manner. The toroidal mode number n (sidebar S4 of the Attachment) is used to identify the helicity of the deformation. For RWM control, we are primarily interested in the lowest mode number $n=1$ since the $n=1$ instability is the first to occur with rising pressure. The achievable plasma pressure in power reactors is expected to be limited by the $n=1$ mode, but it is predicted that the $n=2$ mode will also go unstable if the $n=1$ mode is stabilized and the pressure continues to rise.

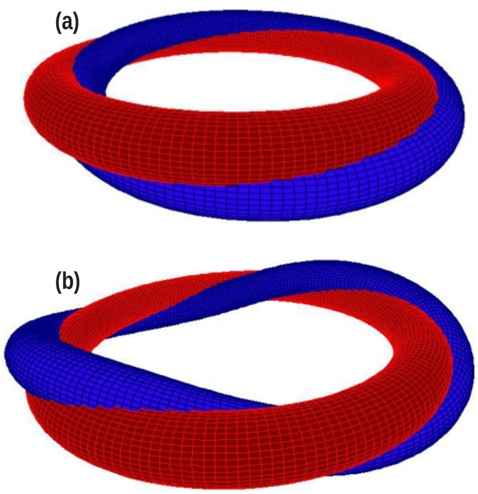

Figure 11. Illustration of kink deformations of a circular cross-section plasma (greatly exaggerated for illustration). The red torus represents a circular cross-section plasma before deformation. The blue surface represents the deformed plasma. (a) an $n=1$ kink. (b) an $n=2$ kink, in which the plasma perturbation repeats itself twice as the toroidal angle varies from 0 to 2π. In each case, the deformation follows a helical path with respect to the undeformed plasma.

When the plasma undergoes a non-axisymmetric distortion as in Fig. 11, the current flowing in the plasma moves with it. Thus the magnetic flux and field that this current generates also moves with the fluid disortion. This moving magnetic field induces eddy currents in the surrounding conductive structures similar to the way in which eddy currents are induced by the vertical instability (sidebar S17 of the Attachment). These induced currents, in turn, generate magnetic fields that oppose the plasma deformation, as in the case of the vertical instability. The overall effect of the presence of a conducting wall is to transform a plasma deformation with an extremely fast growth time (a few microseconds) into a combined plasma/wall system with an instability having a growth time on the order of the resistive decay time of eddy currents in the surrounding materials (a few milliseconds). This slower growth enables use of feedback to control the RWM instability.

The magnetic field motion due to the plasma fluid deformation is observable outside the plasma by magnetic sensors. Even though the deformation of the plasma surface cannot be directly measured in real time, the magnitude and direction of the deformation can be inferred from the external magnetic sensor measurements (Fig. 12). Magnetic sensors are structurally simple and robust so they serve as the best sensors for a magnetic feedback control approach. Other, non-realtime diagnostics are used for developing understanding and modeling of the physical processes. Examples of the diagnostic measurements used in RWM analysis are shown in Fig. 13.

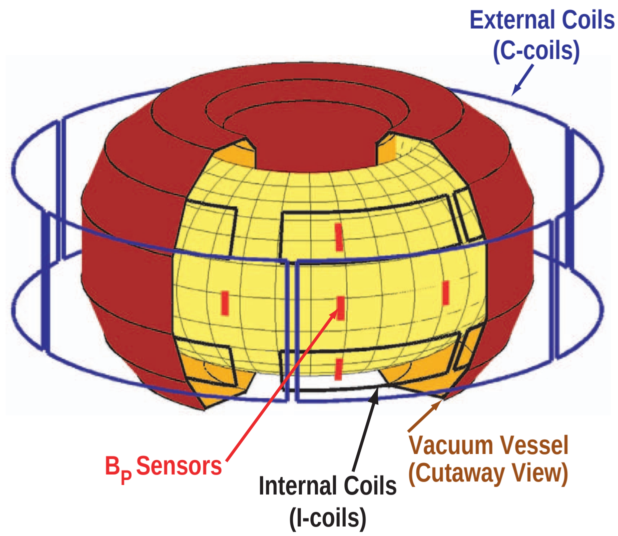

Figure 12. RWM real time control sensors and actuators currently installed on the DIII–D device. The vacuum vessel (represented by the brown surface) is cut away to show internal detail. The plasma is represented by a yellow surface. Sensors: Typical sensors for radial flux measurement are provided by window frame shaped saddle loops (sidebar S7 of the Attachment). Saddle loops used for RWM control (not shown) are as large as the external actuator coils (C-coils) and located approximately concentric with those coils; they are also known as radial flux loops. Radial flux is the same as the integral of radially directed magnetic field (B_r) normal to the wall over a broad area, $\psi = \int_{A_l} B_r dA$, where A_l is the area subtended by the loop. Although, strictly speaking, these saddle loops measure radial flux, measurements made by these sensors are often referred to as radial field measurements because of this integral relationship between field and flux. The poloidal field (B_p) sensors (red) are magnetic probes (sidebar S7 of the Attachment). These sensors are mounted on the vessel wall and measure the local magnetic field component tangent to the wall. Actuators: The C-coil set (blue rectangles) consists of six coils located on the midplane outside the vacuum vessel. The I-coil set (black rectangles) inside the vessel consists of two sets of six coils at upper and lower off-midplane angles, installed between the vacuum vessel wall and the plasma-facing carbon tiles.

The Physics of the RWM

According to MHD theory, a sufficiently high plasma pressure makes the RWM unstable when the surrounding wall structure is located far from the plasma surface. The plasma pressure threshold for this instability is expressed in terms of a critical value of normalized beta β_N (sidebar S2 of the Attachment). The unstable eigenmode could, in theory, be completely stabilized by the mode-induced eddy currents in the wall if the plasma were surrounded by a perfectly conducting wall within a critical distance. [A perfect conductor is a conductor without resistance, implying that there is no resistive decay of the stabilizing eddy currents induced by the eigenmode (sidebar S17 of the Attachment).] In actual devices, the wall current decays away due to resistive losses and the mode amplitude grows with a growth time that is a fraction of the wall time constant. The critical value of beta for which the plasma becomes unstable without a perfectly conducting wall is called the *no wall beta limit*. With a further increase of plasma pressure, the RWM would become unstable even in the presence of a perfectly conducting wall. The value of beta for which this happens is called the *ideal wall beta limit*.

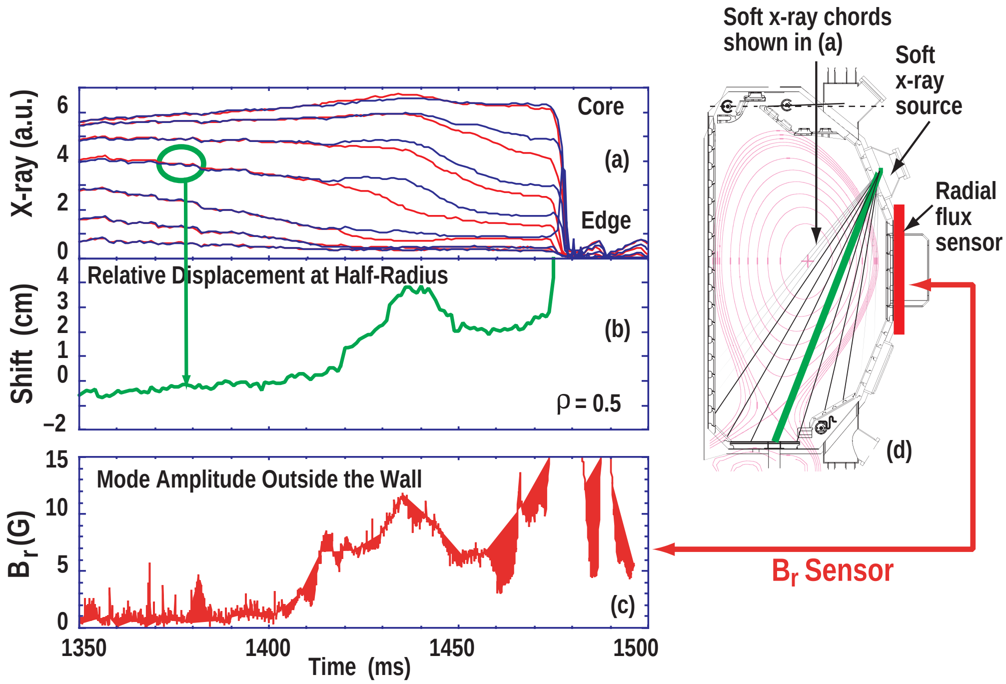

Figure 13. Observation of an RWM inside a plasma by the soft x-ray diagnostic and outside the vacuum vessel by a magnetic sensor: (a) time-dependent measurements of soft x-ray data measured at two toroidal locations (red and blue) spaced 150 deg apart (a.u. = arbitrary units). These measurements can be used to estimate plasma fluid displacement. (b) displacement of the plasma fluid corresponding to the soft x-ray line circled in (a). (c) amplitude of the radial magnetic field measured just outside the vessel during this time. Note the strong correlation between the fluid displacement in (b) and magnetic field growth in (c). (d) DIII–D cross section showing the soft x-ray and magnetic field measurement locations. The thick green soft x-ray chord corresponds to the signal shown in (b). The helical $n=1$ internal mode structure is observable as differences in displacement measured by the two soft x-ray arrays (a) which are separated by 150 deg in the toroidal direction. The soft x-ray diagnostic detects the x-rays emitted by a residual amount of impurity ions caused by electron bombardment. This diagnostic is not sufficiently robust for real time use, since the signals are also sensitive to minor changes of other plasma properties.

Ideal MHD gives a detailed prediction of the structure of the RWM as illustrated in Fig. 14. MHD also predicts that the RWM amplitude is larger at the outer major radius side of the plasma (see Fig. 12 of the Attachment) than at the inboard side. This is due to the nature of the confining toroidal field, which decreases away from the torus axis of symmetry (the Z axis in Fig. 12 of the Attachment) so that the magnetic field pressure is relatively weaker at the outer major radius. This leads to the RWM perturbation amplitude being larger at the outer edge of the plasma. This suggests that for RWM control, an actuator located at the outer major radius of the plasma is favorable and should be effective.

According to ideal MHD, the RWM structure inside the plasma fluid is complex [Fig. 13(a)]. However, when considering methods for stabilizing this mode, the fluid deformation is not the focus of attention. Instead, these methods focus on the magnetic field perturbation that is associated with the plasma fluid deformation, since a number of real time sensors are available to measure this perturbation (Fig. 12) and relatively simple models of the RWM can be developed using this point of view. Various experimental studies [41]–[46] have revealed several

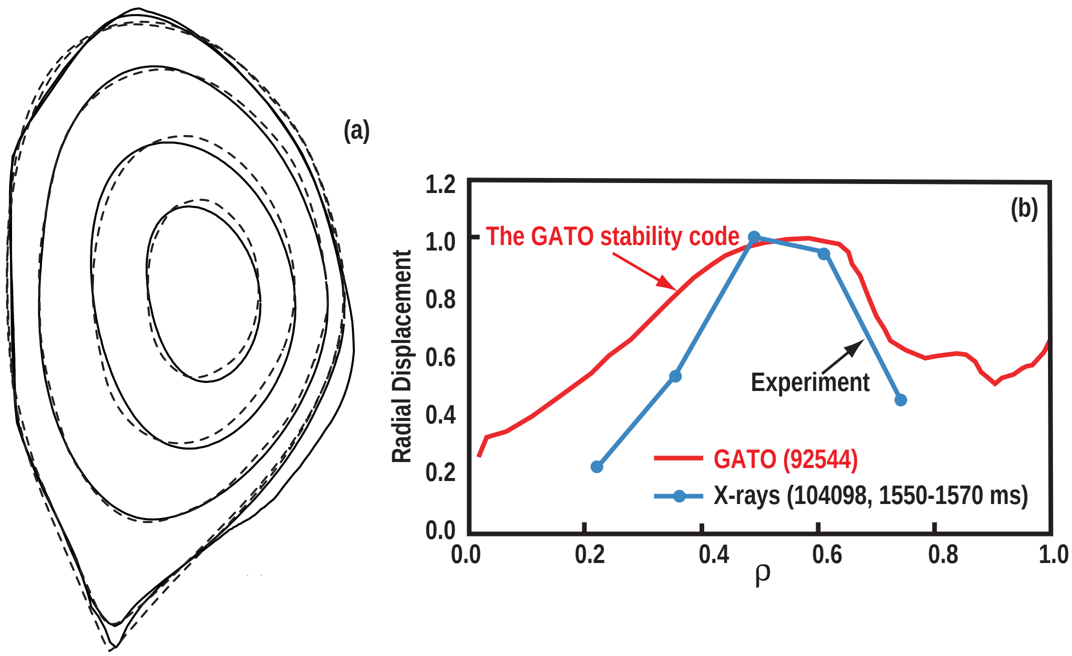

Figure 14. Comparison of theoretically predicted RWM structure with measurement. (a) Structure of the RWM displacement from the axisymmetric plasma as computed by the GATO ideal MHD code [40]. The dashed lines represent constant flux surfaces before deformation. Solid lines represent these flux surfaces after deformation by the RWM. Perturbations are greater at larger radii R because of the inverse dependence of the confining field on radius. The displacement magnitude is exaggerated for illustration. (b) A comparison of the experimental plasma fluid radial displacement estimated using soft x-ray data (see Fig. 3) and the mode radial displacement predicted by GATO. The magnitude is normalized to the maximum amplitude (about 8 cm) near $\rho = 0.5$. (The normalized flux coordinate ρ was defined in sidebar S5 of the Attachment.) The prediction accuracy is adequate for control, at least in the region where there is data for comparison.

characteristics of the RWM and their relation to ideal MHD predictions. Each of these will be discussed in more detail below.

1. The spatial structure of the RWM agrees with the ideal MHD theoretical prediction (Figs. 14 and 15). The RWM exhibits a structure that is global, extending from the plasma core to outside of the vacuum vessel, where we consider the mode as represented by the combined magnetic fields of the plasma and conducting structures.

2. The existence of a threshold in plasma pressure for the onset of the RWM agrees with ideal MHD theory.

3. When a non-axisymmetric external field exists, the mode responds only to the component of the external field that matches the mode's own field structure. The mode amplitude is amplified proportional to the external field, behaving like a *magnetic field resonance* (to be discussed below).

4. Beyond the ideal MHD framework is a surprising observation: the mode can be completely stabilized by rotating the plasma, if the rotation is above a critical value.

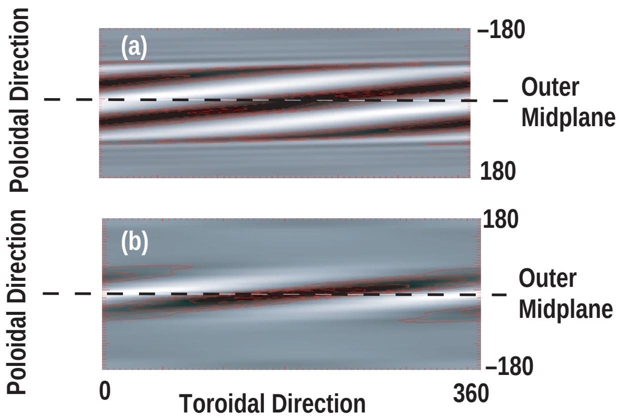

Figure 15. Normal field perturbation on the plasma surface due to an RWM as calculated with the GATO ideal MHD code [40]. Perturbations are relative to the normally axisymmetric (independent of toroidal angle) values that occur in the absence of an RWM. (a) Perturbed normal field at the plasma surface due to plasma current, (b) perturbed normal field at the plasma due to the wall currents. A poloidal angle of 0 corresponds to the outer midplane (see Fig. 12 of the Attachment). The normal fields shown reflect the pattern of current flow in the two conducting surfaces. Lighter colors are more positive, darker more negative. The eddy current on the wall is induced by perturbations of the plasma surface current as discussed in the text. The plasma surface current is maximum at the outer midplane and decreases rapidly toward the inner-major radius side. Correspondingly, the eddy current on the wall is maximum at the major radius side since it is inductively coupled. Note that the pattern is periodic with toroidal period one (mode number $n=1$), and that the high and low amplitudes in the mode wind their way in a helical pattern around the torus.

5. The RWM can be stabilized by plasma rotation for well over the ideal MHD time scale (Appendix sidebar S19) and the mode spatial structure remains the same even after the wall eddy current disappears. The sustainment of a single mode pattern over a long period is encouraging for developing simple magnetic control. This spatial invariance is often described as *mode rigidity*.
6. However, the finite amplitude of the long-sustained RWM reduces the bulk plasma rotation, leading to a less stable high-pressure plasma.

These observations provide the rationale for two different approaches to stabilizing the RWM. One approach is the use of feedback control to oppose the moving field that accompanies the growth of the mode. We will refer to this approach as magnetic control. The rate of growth of the mode is slowed sufficiently by the conducting wall to make a feedback process feasible. The existence of a single dominant mode allows for simpler models of the plant to be controlled. An example of the coil arrangement used to excite the non-axisymmetric field necessary for RWM control is shown with the DIII–D device in Fig. 12.

Another approach for RWM control is the use of plasma rotational stabilization. In present day tokamaks, neutral beam injection (NBI, see sidebar S13) supplies an ample amount of angular momentum input for maintaining rotation of the plasma fluid, leading to the stabilization of the RWM mentioned above. However, it is not obvious whether sufficient plasma rotation can be achieved in fusion power generating reactors. Thus, magnetic feedback control is actively being pursued for use in future devices.

Models of the RWM Instability

In this section, we describe some of the basic models that are presently in use for development of methods for stabilizing the RWM. Initially, we assume that there is little or no bulk plasma fluid rotation and concentrate on only the magnetic aspects of the control problem. Models for magnetic control ignore the internal details of the plasma, focusing instead on the behavior of the magnetic field structure on the plasma surface. A common method of modeling the unstable mode is to replace the spatial perturbation of the plasma with an equivalent perturbation (in the sense that it produces the same magnetic field perturbation) of surface current on a spatially fixed plasma boundary. The surface current distribution is calculable from the geometrical shift of the plasma surface. The eddy current pattern on the wall can also be calculated once the plasma surface current pattern is determined. The plasma surface current and wall eddy currents are illustrated through plots of magnetic field normal to the plasma boundary in Fig. 15.

Using the assumption of a rigid mode mentioned above, the spatial distribution of current on the plasma surface and on the wall remain intact while only their magnitudes change. Using the surface current representation of this mode, we can construct a state space model of the plant with states given by currents on the plasma surface, I_p, and in surrounding passive (wall) structures, I_w. These variables represent the scalar multipliers of the spatially fixed distributions of current on the plasma surface and in the wall. The external control coil current, I_c, represents the scalar multiplier of a single spatial distribution of currents produced by multiple coils chosen to best match the distribution of the eigenmode. These variables all represent perturbations from axisymmetric (purely toroidal) equilibrium currents due to the appearance of the non-axisymmetric RWM in the plasma.

In the following, we discuss RWM behavior using a simple cylindrical model [47], [48]. Use of cylindrical models, in which plasmas are assumed to flow in an infinitely long cylinder, is a common first step in developing physics understanding of phenomena that occur in the “bent” cylinder constituting a torus.

The pressure balance on the plasma surface between the internal plasma pressure and the external magnetic field pressure leads to a circuit-like (sidebar S16 of the Attachment) equation.

$$L_{eff}\, I_p + M_{pw}\, I_w + M_{pc}\, I_c \;=\; 0 \quad , \tag{1}$$

where constants M_{ab} represent mutual inductance between conductors a and b (Fig. 16), the effective self inductance L_{eff} is given by

$$L_{eff} = \frac{L_p\, C_\beta}{\left(C_\beta + \Delta M\right)}, \quad C_\beta = \frac{\beta_N - \beta_{N.no-wall}}{\beta_{N.ideal-wall} - \beta_{N.no-wall}} \quad , \tag{2}$$

$\Delta M = L_p\, L_w / M_{pw}^2 - 1$ is related to the wall stabilization effect, and L_p represents self-inductance of the plasma. The constant C_β is a measure of the stability of the plasma to resistive wall modes. When $C_\beta < 0$, β_N is below the no-wall limit and the RWM is stable. When $C_\beta > 1$, β_N is above the ideal-wall limit and the plasma cannot be (practically) stabilized. Efforts at active stabilization aim at the interval $0 < C_\beta < 1$. The wall eddy current and the active coil current are modeled by standard circuit (sidebar S16 of the Attachment) equations

$$\begin{aligned} M_{wp}\,\dot I_p + L_w\,\dot I_w + M_{wc}\,\dot I_c + R_w\,I_w &= 0 \\ M_{cp}\,\dot I_p + M_{cw}\,\dot I_w + L_c\,\dot I_c + R_c\,I_c &= V_c \quad . \end{aligned} \tag{3}$$

Here, constants R_a represent resistance in conductor a and, as in Eq. (1), constants M_{ab} represent mutual inductance and L_a represent self-inductance (Fig. 16). The ordinary differential and algebraic Eqs. (1) through (3) that constitute the overall circuit model can be expressed via a Laplace transform as

$$(Ms + R)\,I = V \quad , \tag{4}$$

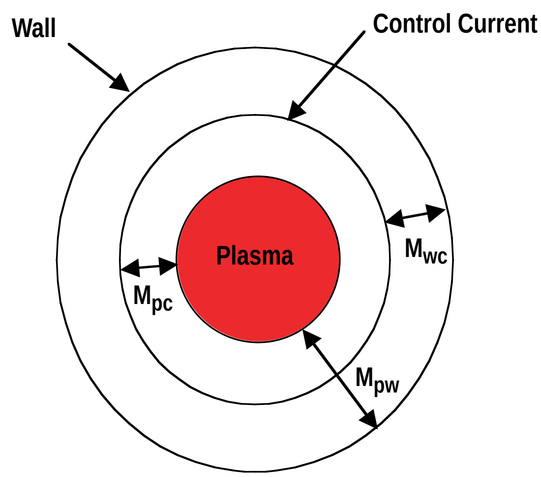

Figure 16. Cross section of cylindrical model of RWM dynamics. The model shown represents the case where the control coils are inside of the vessel wall. The interaction between current in the plasma, wall, and control coils (subscripts p, w, and c) is determined by the mutual inductance values M. Drawing is not to scale.

where s is the Laplace transform variable and

$$M = \begin{bmatrix} L_{eff} & M_{pw} & M_{pc} \\ M_{wp} & L_w & M_{wc} \\ M_{cp} & M_{cw} & L_c \end{bmatrix}, \quad R = \begin{bmatrix} 0 & 0 & 0 \\ 0 & R_w & 0 \\ 0 & 0 & R_c \end{bmatrix}, \quad I = \begin{bmatrix} I_p \\ I_w \\ I_c \end{bmatrix}, \quad V = \begin{bmatrix} 0 \\ 0 \\ V_c \end{bmatrix}.$$

Using this formulation, we can treat the RWM control model as if it were a standard circuit equation. The mutual inductances are computed by standard geometric methods [49]. The modified self-inductance L_{eff} is the only term that differs from the standard electromagnetic definition (see sidebar S16 of the Attachment, for example) and includes the plasma parameters. This approach allows variations in the plasma to be modeled as changes to a single term L_{eff} in the model.

Rotational Stabilization

It has been observed on DIII–D that when the toroidal rotation of the bulk plasma fluid remains above 6 kHz (thousands of rotations per second) the RWM instability is completely stabilized (Fig. 17). One of the primary causes of rotation in current experimental devices is the injection of neutral deuterium atoms (intended originally for heating; see sidebar S13) at an angle nearly tangential to the torus. The momentum of these particles is imparted to the bulk plasma, thereby increasing the rotation. Distinct from the fluid rotation is the rotation of the mode itself, typically at a frequency between 10 and 20 Hz, that is believed to be coupled to the fluid rotation.

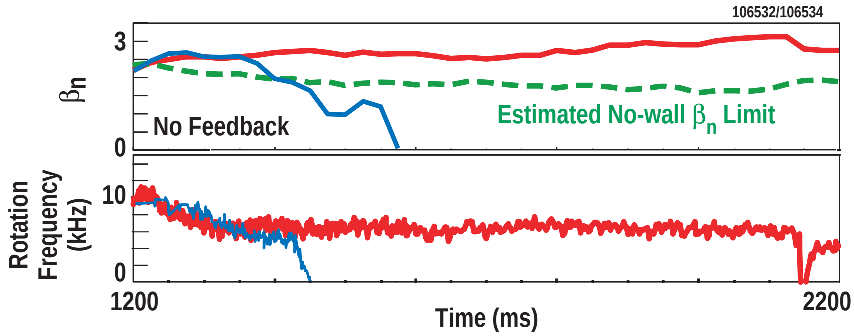

Figure 17. Demonstration of suppression of the RWM instability through plasma rotation. As long as plasma rotation remains above a critical frequency (red curve, bottom) of about 6 kHz, the mode remains stabilized, even for plasma beta well above the no-wall limit (red curve, top). If plasma rotation falls below the critical value (blue curve, bottom), the mode becomes unstable, causing loss of plasma pressure (blue curve, top), accelerated slowing of the rotation, and shortly thereafter, loss of the plasma to the unstable mode.

To include the effect of plasma rotation in the model of the RWM, we must depart from the previous assumption of a *rigid mode at a fixed toroidal angle*, with the special case of zero rotation represented by the previous discussion. The unstable mode still maintains the rigid sinusoidal current spatial distribution having toroidal period one (Fig. 15), but now may be shifted or rotating in toroidal angle. Thus the previous representations in which $I_c(t)$, $I_p(t)$, and $I_w(t)$ were scalar multipliers of spatially fixed current distributions are replaced by two parameter multipliers of the form $I\ (t) = A\ (t)e^{i\phi_*(t)}$, where $\phi\ (t)$ represents the toroidal angle of the sinusoidal distribution with respect to a fixed reference angle. Now the currents I_c, I_p, and I_w as well as perturbed magnetic fields due to the wall B_w and due to the plasma surface B_p are each represented by complex numbers whose real and imaginary parts represent sinusoidal distributions with peak amplitude in the 0 and 90 deg toroidal directions, respectively. In this notation, multiplication by $i = \sqrt{-1}$ represents a result that is rotated toroidally by 90 deg.

In general, the mechanisms for combined rotation and magnetic effects on the RWM are not well understood. Using an argument that seeks to account for the exchange of energy between the plasma mode and external conductors, Chu, et al. [50] propose the model

$$\left(\delta W_{Iw} + i\Omega_\phi D\right) B_p = C_{pw} B_w \quad , \tag{5}$$

to represent the coupling between changes in field B_p at the plasma surface and changes in field at the vessel wall when the plasma is rotating. This model is only qualitative and still somewhat speculative, but is described here to provide some insight into the possible mechanisms for experimentally observed rotational stabilization. Here, Ω_ϕ represents the plasma fluid toroidal rotation frequency and $C_{pw} = M_{pw}^{-1}$ is the inverse of the mutual inductance between the wall and the plasma surface. The quantity δW_{Iw} represents the coupling of RWM energy transferred through the field B_p produced by the plasma to the component of the field B_w that is toroidally in phase with B_p, while $\Omega_\phi D$ represents the energy coupled to the component of B_w that is 90 deg toroidally advanced. The quantity D represents an unknown dissipation mechanism. This representation is motivated by the following:

1. Experiments show that, when the plasma is rotating, the RWM responds at a different toroidal angle than the angle at which an external field is driven. For example, when the plasma is at steady state with a stable RWM ($\beta_N < \beta_{N,no-wall}$), if a fixed sinusoidal $n=1$ current $I_c(t) = A_c e^{i\phi_c}$ is applied (so that eddy currents are not excited in the wall), the plasma surface mode responds not at the angle $\phi_p = \phi_c$ but at an angle $\phi_p = \phi_c + \delta\phi$ with $\delta\phi > 0$ [44].

2. A theoretical consideration is that momentum can be exchanged between the toroidally rotating plasma and the RWM, resulting in a transfer of some of the unstable mode's energy to a different toroidal angle (see the discussion that follows).

There are several candidate models for the dissipation mechanism D, but none has been satisfactorily verified yet experimentally. We present here a rather simplistic explanation of one candidate model that is consistent with the three experimentally observed phenomena: (1) the increase in RWM growth rate as the plasma fluid rotation slows to a critical rotation value, (2) the observed much slower rotation of the RWM itself in the same direction as the rotation of the bulk plasma fluid, and (3) the response described above of a stable RWM to an applied $n=1$ perturbation that appears at a toroidally shifted location. A more accurate and complete model requires significant background preparation, so will not be presented. The proposed model postulates a coupling between the unstable mode and bulk plasma fluid through a form of viscous friction, as follows. Current is defined as the rate of flow of charged particles. In a current-carrying plasma, an RWM fluid deformation (Fig. 11) may be thought of as a change in the pattern of flow of the charged particles that comprise the current in the plasma. This change in flow is seen outside the plasma as a change of flux or field at magnetic sensors. This flow of particles is driven by the sharp difference in total (plasma plus magnetic) pressures inside and outside of the plasma — the driving force of the RWM instability — and a significant portion of this flow is radially directed (Fig. 18). The motion of particles in the plasma can also be influenced by the injection of momentum (from neutral beams for example). The toroidal particle flow introduced by toroidal momentum injection interacts with the particle flows caused by the RWM and vice versa through collision interactions. To sustain the RWM in a given direction, work is required by the RWM instability to move particles against this toroidal particle flow. This resistance to the motion of particles is essentially the mechanism of viscous friction. The RWM's loss of energy through this work is the proposed damping mechanism. Assuming this damping mechanism, the model Eq. (5) can be included in the Eq. (4) by modifying L_{eff} [45]

$$L_{eff} = \frac{L_p\left[C_\beta + \alpha_\phi\left(s + i\Omega_\phi\right)\right]}{\left[C_\beta + \alpha_\phi\left(s + i\Omega_\phi\right) + \Delta M\right]} \quad , \tag{6}$$

where α_ϕ represents the viscous drag coefficient.

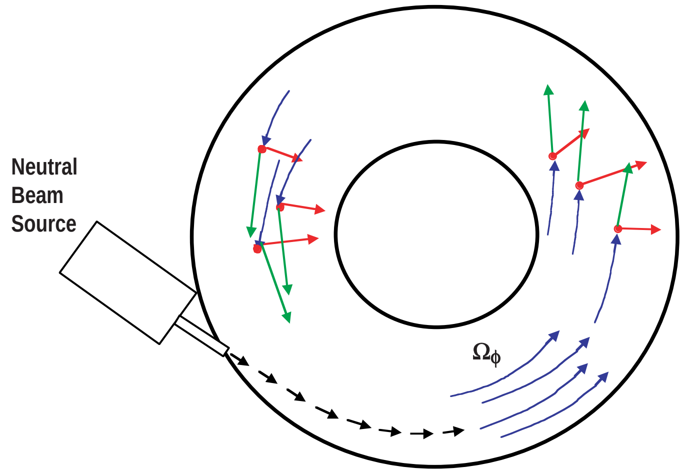

Figure 18. Illustration (top view) of one proposed model of the effect of toroidal rotation on the RWM. In a toroidally rotating plasma, individual particles flow with an average speed defined by the bulk fluid rotation frequency Ω_ϕ. An $n=1$ RWM causes the plasma fluid to "bulge" radially outward on one side of the torus and inward at a location 180 deg opposite (see Fig. 1). Individual particles driven by the RWM and the component of their velocity vectors induced by the RWM are indicated in red. Continually flowing particles with toroidal momentum (blue arrows) frequently collide with the RWM driven particles and impart some of their momentum. The resulting velocity (green arrows) of the originally radially directed particles have significant toroidal components.

The Role of Error Fields

Rotation of the bulk plasma is influenced significantly by the level of error field present in the device. The error field is defined as the difference between the slightly non-axisymmetric magnetic field produced by an as-constructed device and the ideal axisymmetric field that would be produced by an ideally constructed device. The component of the error field resonant with the RWM is that portion of the error field that matches the field pattern of the unstable eigenmode on the plasma surface.

This resonant error field can be represented as an external current source I_e acting in a manner similar to the external control coil current I_c. The impact of this field can thus be described by replacing I_c by I_e in the plasma response Eq. (1). Since the error field is steady state, the perturbed quantity $I_w = 0$ if the RWM is stable. (A *changing* current, either I_p or I_e, is required to induce nonzero eddy currents I_w.) In this case, the plasma response, I_p, to the error field is also steady state, given by Eq. (1) as

$$I_p = -\left(M_{pe}/L_{eff}\right) I_e \quad . \tag{7}$$

Since M_{pe} is real, the toroidal phase shift, $\delta\phi$, is given by Eq. (7) as $\delta\phi = \tan^{-1}[-Im(L_{eff})/Re(L_{eff})]$. It is clear that the magnitude of the plasma response can reach a huge value around the no wall limit where $L_{eff} \approx 0$, similar to a resonance effect. This phenomenon is known as *Resonant (Error) Field Amplification* (RFA). As the growth rate of a stable RWM approaches zero from below, the value of L_{eff} also approaches zero. Thus RFA increases as a stable RWM becomes less stable.

The non-axisymmetric error field and the resulting amplification of the stable RWM are believed responsible for the rotation slowing observed in Fig. 17 through a form of *magnetic braking* of the plasma rotation. As with rotational damping, there is more than one possible mechanism for the observed rotational slowing. One rather well accepted explanation uses an analogy with the induction motor. An induction motor consists of a conductive rotor (the inside part that turns) surrounded by a stator (the stationary outside part that causes it to turn). A magnetic field that rotates around the rotor is set up by properly phased currents flowing in a set of coil windings in the stator. As long as this moving magnetic field rotates faster than the rotor, the field acts on the conductive rotor to generate currents on its surface. A torque is produced on the rotor through the interaction of this induced current with the rotating field (the “I cross B” force; see sidebar S17 of the Attachment). If the magnetic field produced by the stator becomes stationary or reverses direction while the rotor is rotating, the applied torque reverses direction and causes the rotor to slow down. This is directly analogous to a rotating (conductive) plasma interacting with a stationary non-axisymmetric magnetic field such as an error field. Slowing of the plasma by the error field and RFA causes a stable RWM to become less stable, which then increases the effect of the magnetic braking, which slows the plasma even more. This process continues until, eventually, the RWM becomes unstable as shown in Fig. 17.

To the sophisticated reader, it is clear that both the rotational damping and the induction motor analogy describe forces that are exerted in both directions. For example, the viscous damping mechanism discussed in the previous section that stabilizes the RWM also results in forces that tend to slow the plasma rotation. Several damping/slowing mechanisms have been proposed, but present experimental data is not sufficient to clearly confirm or refute these proposals. It is also possible that more than one of these mechanisms will ultimately be determined to play a major role.

Magnetic Control Approaches

Experimental approaches for magnetic control presently include correction for external error fields to reduce the magnetic braking on rotation and, separately, magnetic feedback stabilization of the RWM in the absence of plasma rotation. The magnetic feedback efforts allow for mode rotation but generally do not account for fluid rotation effects. The situation of magnetic feedback with non-zero plasma fluid rotation is not yet well enough understood to develop useful control approaches. An experimental RWM controller consists of the observation sensors, sensor logic, digital controller, power supplies, and actuator coils.

Actuators consist of actively driven current-carrying coils, typically with a picture frame geometry (Fig. 12). The rigidity of the mode simplifies the discussion of the required feedback field. When non-axisymmetric field is applied, the plasma perturbation responds only to the component of the field that matches its own mode structure. This implies that the external coils (C-coils), which primarily produce radial field, are not very efficient since at least half of the magnetic energy does not couple with and therefore does not affect the helically shaped mode. The connection flexibility of the internal coil (I-coil) set can be used to provide a field pattern (Fig. 15) that more closely matches that of the RWM. The I-coils have the additional advantages that they are closer to the plasma while the appearance of C-coil flux at the plasma is delayed due to shielding by eddy currents in the vessel. For these reasons, the internal coil set is superior to the external coils for feedback control.

Although error fields are determined by the limited accuracy of device construction and are independent of the plasma, the required compensation for these fields depends on properties of the plasma being maintained because of the RFA. The necessary spatial distribution of applied corrective field is related to the MHD mode structure and the required magnitude of the correction depends on the value of L_{eff} according to Eq. (7). The magnitude and the mode toroidal angle can evolve slowly in time during the discharge because of changes in plasma properties. In addition, the error field has a complicated non-axisymmetric distribution that cannot be completely canceled by a finite number of actuator coils. Thus, compensation cannot be done open loop based on a priori calculations. This motivates the use of *dynamic error field correction,* which adjusts the error field correction based on the plasma mode response. This leads to a slow time-scale (much slower than the wall time) feedback process. Magnetic control aimed at directly stabilizing the RWM requires feedback with a faster time constant equal to a fraction of the wall time. Thus requirements on actuating coils and power supplies are very different for error field correction and for magnetic feedback stabilization.

Power systems for tokamak control problems are a challenge due to simultaneous requirements for high voltages and currents and speed of response. RWM control requires a relatively high current (a few kiloamps) at near steady state to compensate for error fields and, simultaneously, a fast (a few hundred hertz) lower current response to provide magnetic feedback stabilization. Since the conductive wall slows the mode growth to approximately the time constant of the wall, the maximum bandwidth required for the supply is defined by the inverse of this wall time constant. For example, the DIII–D power supply is designed to have a 3 dB bandwidth of 500 Hz in order to stabilize plasmas with values of β up to halfway between the no wall and the ideal wall limits ($C_\beta = 0.5$).

There have been several methods investigated for detecting the mode growth and determining the toroidal angle. All are based on the experimental observation that the RWM is well defined by an $n = 1$ distribution and the mode structure is reasonably rigid so the mode can be represented by an amplitude perturbation that varies sinusoidally in the toroidal direction. The mode can thus be represented by two parameters: the mode amplitude and toroidal angle, or sine and cosine components $A^{\cos}(t) = A(t)\cos[\phi_0(t)]$ and $A^{\sin}(t) = A(t)\sin[\phi_0(t)]$ where ϕ_0 is the time varying unknown toroidal angle of the mode amplitude maximum. Methods investigated to identify the two parameters consist of matrix multiplications $[A^{\cos}(t) \quad A^{\sin}(t)]^T = Gx$ where x is a vector of sensor measurements and G is a constant gain matrix. One approach is to use radial flux sensor measurements $x(t) = [\psi_1(t) \quad \psi_2(t) \quad \psi_3(t)]^T$ at the same radius and at different toroidal angles. Each $\psi_j(t)$ represents the difference in magnetic flux measured at two sensors located at same radius but 180 deg apart toroidally at the midplane of the torus. This differencing is done to reinforce radial flux measurements due to an $n = 1$ mode perturbation while de-emphasizing flux contributions from disturbance sources (such as plasma perturbations having toroidal mode structure with n even). This scheme is called *smart shell* because the feedback process attempts to minimize the total flux perturbation at the observation point so as to emulate a perfectly conducting wall at that radius.

Another approach is to use only poloidal field sensors $x(t) = [B_1(t) \quad B_2(t) \quad \cdots \quad B_m(t)]^T$ where each $B_j(t)$ represents the difference in magnetic field measured at two sensors located at same radius but 180 deg apart toroidally at the midplane of the torus (Fig. 12). This difference is also used to reinforce poloidal field measurements due to an $n = 1$ mode perturbation and to remove field contributed by axisymmetric variations in the plasma (such as in the plasma shape). This approach is called *mode control* because the poloidal field sensor measures almost no field directly from the actuator coils, which produce primarily radial field, and thus is more sensitive to the field variations due only to the plasma mode.

A third approach is to determine the mode amplitude and phase using all available flux and poloidal field sensors $x(t) = [\psi_1(t) \quad \psi_2(t) \quad \cdots \quad \psi_{n_\psi}(t) \quad B_1(t) \quad B_2(t) \quad \cdots \quad B_{n_B}(t)]^T$ rather than just symmetrically located pairs of either type. Here, the fluxes $\psi_j(t)$ and fields $B_j(t)$ are measurements at individual sensor locations. This is referred to as the *matched filter approach* [51] because each row of G defines a spatial matched filter. The first row is matched to the normalized response expected in the set of sensors from a mode with phase $\phi_0 = 0$ while the second row is matched to the normalized expected response from a mode with phase $\phi_0 = 90$ deg. Use of a matched filter was motivated by experimentally observed difficulties in rejecting noise and disturbance signals. Rejection of measurement noise could be enhanced by averaging multiple sensors obtained with the matched filter. The most severe disturbance was due to magnetic sensor responses to edge localized modes (ELMs — see the next section of this article). An ELM is a local mode whose spatial magnetic field distribution is significantly different than the distribution defined by the global RWM, making it a good candidate for rejection via a matched filter. Use of the matched filter in simulations has been shown to improve the accuracy of the mode estimation, but it does not provide a complete solution to rejection of ELM disturbances. The primary difficulty is that, on short time scales, the growth of an ELM disturbance signal includes a large $n = 1$ component similar to the unstable $n = 1$ RWM. The ELM excitation mechanism and mode structure are significantly different however. This motivated work on development of a Kalman filter to exploit information contained in the RWM dynamics model to filter out the ELM signals. Simulations [52] indicate that combining the spatial matched filter and dynamic Kalman filter will significantly improve signal-to-noise ratio and reject ELM disturbances. Use of the matched filter and the Kalman filter require relatively detailed knowledge of the mode spatial distribution. Use of a Kalman filter requires, in addition, a good model of the time evolution dynamics of the mode and its interaction with surrounding structure. The mode dynamics depend strongly on the fluid rotation frequency, however, and this dependence has not been well characterized. Thus the experimental application of this approach relies on the success of ongoing efforts to develop models that combine MHD and rotation effects.

Experimental controllers have been limited so far to the use of proportional, integral, and derivative (PID) algorithms. PID and more advanced control algorithms have also been studied in simulations and paper studies [53]. Experimental use of the more sophisticated control algorithms has not been pursued because of the ambiguity in models due to incomplete understanding of dependence of the RWM on rotation.

Experimental Progress in Stabilizing the RWM

The most important progress in RWM stabilization was the sustainment by rotation of a discharge with $C_\beta \approx 1$ using open loop preprogrammed non-axisymmetric coil currents to minimize error fields. The program for this current was determined from the heavily time-averaged coil current signal previously obtained in a nearly identical plasma discharge that used dynamic error field correction to define the coil currents. With feedback off and non-axisymmetric currents programmed in this way, the discharge behavior followed closely the evolution of the discharge with feedback on; the achieved values of β_n and plasma rotation frequency were nearly identical. This evidence suggested that the resonant component of the non-axisymmetric field contributed significantly to the mode amplification and, consequently, reduced the rotation velocity. Once the compensation was made through the dynamic error field correction, the RFA amplitude did not grow and the plasma did not slow down. The resulting plasma rotation was sufficient to suppress the onset of RWM up to the ideal wall β_N limit.

Several experiments [43] have shown an extension of discharge duration when using magnetic feedback, but have not demonstrated long term stabilization. Analysis [54] and experiments [55] are consistent in showing the superiority of mode control over smart shell control. Motivated by analyses of coil/sensor effectiveness ([56],[57]), the DIII–D device has been equipped with actuators located both inside and outside the vacuum vessel. Use of the internal coils (I-coils) has significantly improved magnetic feedback performance over that achievable with external coils (C-coils) alone (Fig. 12).

Future Directions

Experimental and theoretical research continues on this important control problem. Many important issues are near to being resolved, including the most effective type, location, and configuration of actuators and sensors for magnetic feedback control. Although not yet sufficient for a complete model, a great deal has also been learned about the dynamics of the RWM and its interaction with error fields and plasma fluid rotation. Efforts have also begun to address some of the more practical issues such as required current levels, power supply response times, and communication delays that are required for magnetic feedback systems.

Significant effort remains to complete RWM model development. Sufficient experimental data needed to either confirm or refute the many candidate magnetic feedback models has been difficult to obtain because of the interaction between error fields, fluid rotation, mode growth, and magnetic feedback. Once models are completely developed and validated, the final algorithm(s) for control that can handle the wide range of RWM conditions will need to be developed and experimentally tested. This will be challenging, since the RWM growth time can vary from a fraction of wall time (a fraction of millisecond) to the angular momentum confinement time (a fraction of second). In addition, the present linear rigid mode growth assumption represents only the dominant mode of several potentially unstable modes; even this dominant mode may become nonlinear nearer to the ideal wall limit. Another important near term objective is the need for extrapolation of RWM stabilization methods to ITER and reactor oriented devices.

CONTROL OF PLASMA PROFILES AND INTERNAL TRANSPORT BARRIERS

The requirements of ITER and the need to optimize the tokamak concept for the design of an economical — possibly steady state — fusion power plant have motivated extensive international research on plasma transport (see Fig. 9 caption) and confinement in toroidal devices. These investigations have aimed at finding plasma regimes with improved confinement with respect to the one predicted by typical tokamak scaling laws and have led to the definition of the Advanced Tokamak (AT) operation scenarios (Appendix sidebar S18) [58]. In a large number of machines, experiments have demonstrated the existence of such regimes that allow access to a high confinement state with improved MHD stability and leading to a strong increase of the performance as quantified by the energy confinement time and plasma pressure (τ_E and β_N; see sidebar S2 of the Attachment). In such conditions a dominant fraction of the plasma current is self-generated by the bootstrap mechanism (sidebar S15), which reduces the requirement for externally driven non-inductive current for steady state operation. This bootstrap current is favored by the generation in the plasma of an internal transport barrier (ITB) [59], a region where particle and heat transport are strongly reduced. An ITB is characterized by large pressure gradients and by the presence of a visible "break" in the slope of the electron and/or ion temperature profiles similar to the edge transport barriers (ETB; sidebar S12). ITBs are often combined with an ETB, which gives rise to a pressure pedestal at the plasma edge, characteristic of the H–mode [60] (sidebar S11).

S15. Bootstrap Current (sidebar)

The bootstrap current is an equilibrium current that is self-generated (without the need of an imposed electric field) in a toroidal plasma. In a tokamak plasma, the guiding center [the center of the fast Larmor gyro-motion (see Fig. 9 of the Attachment)] of most particles follow approximately helical orbits that encircle both the major axis of the torus (vertical axis) and the magnetic axis of the plasma (see Fig. 14 of the Attachment). This periodic guiding center motion is the combination of free streaming along the helical magnetic field lines (Fig. 10 of the Attachment) and of small radial drifts due to the gradient and curvature of the magnetic field, which average to zero after a complete period. However, the toroidal magnetic field intensity produced by external coils in a tokamak decreases as $1/R$ (R being the distance to the major axis) and therefore particles encounter varying field intensities along their orbits, from a minimum, $B_{\min}$, on the outer radial part of the helical orbit (called the *low-field* side of the torus) to a maximum, $B_{\max}$, on the inner radial part (called the *high-field* side of the torus). As a result, particles with low velocity parallel to the helical field lines (whose kinetic energy, W, is mainly in the Larmor motion, $W \approx W_\perp$, orthogonal to the field lines) cannot complete a helical trajectory around the magnetic axis as this would violate the conservation of both their energy W and magnetic moment ($\mu = W_\perp / B$) along the orbit. At some point their parallel velocity must vanish and change sign and therefore these particles are trapped on the low field side of the torus where their guiding centers describe banana-shape orbits (see Fig. 19). Particles that are able to complete their helical orbit are called *passing particles*, as opposed to the *trapped particles*. In the presence of a density gradient and at a particular location in the plasma, there are more trapped particles going in one toroidal direction [trajectory (a) in Fig. 19] than in the other one [trajectory (b)] and therefore the local ion and electron velocity distributions are anisotropic. Therefore, each trapped particle assembly (trapped particles of a given species passing through point P) carries a finite toroidal momentum proportional to the density gradient at P. Now, particle collisions give rise to a continuous exchange of momentum between trapped and passing particles. For instance, passing electrons, which make up the bulk of the electron population, receive net toroidal momentum from the anisotropic trapped electrons at an effective rate that is much faster than the rate at which they lose momentum to the bulk ions. Therefore a net equilibrium electron current results. An additional contribution, with the same sign, comes from the passing ions. This net positive current is known as the bootstrap current.

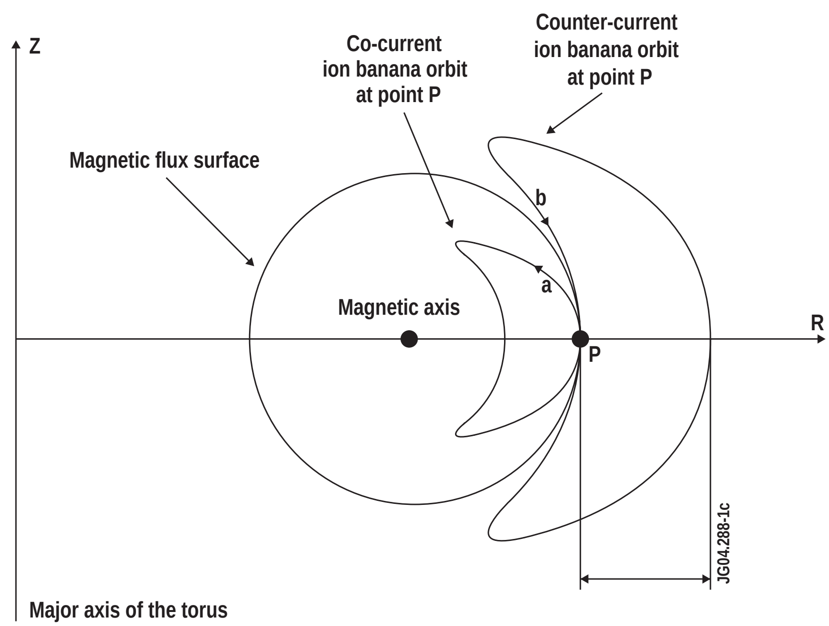

Figure 19. Poloidal projection of two different *trapped particle trajectories* (also called *banana orbits*) passing through a point P in the low-field (outer radial) side of the tokamak equatorial plane. (The particle trajectories also extend a long way in the toroidal direction around the major axis of the torus when moving from the bottom turning point to the top one.) The magnetic flux surface passing through the same point is also represented. Trajectory (a) corresponds to an ion that would move toroidally in the co-current direction when passing through point P, whereas trajectory (b) corresponds to an ion that would move toroidally in the counter-current direction at the same location. Due to the density gradient (increasing density towards the center of the plasma), there are more ions with type (a) orbits than ions with type (b) orbits and an anisotropic velocity distribution is sustained at point P.

Although the formation mechanism of ITBs has not been entirely identified, significant progress has been made in understanding them. Many recent studies have shown the key influence of the safety factor profile $q(x)$ ($x = r/a$; see sidebars S4 and S5 of the Attachment) for the triggering of barriers. Both the radial profile of the magnetic shear (Appendix sidebar S20) and the location of the flux surfaces where q is rational have been shown to be essential for the emergence of ITBs [61]–[63].

When ITBs become too strong, the steep pressure gradient characteristic of the ITB may exceed some MHD stability limit, leading to the loss of the confinement or even to plasma disruption (sidebar S14). Thus, the promising concept of a steady-state tokamak reactor may only be realizable if stationary ITBs can be sustained in a stable way. This has motivated a large experimental effort at JET, aiming at the real-time simultaneous control of the safety factor, temperature, and pressure profiles. This section reviews the progress achieved, based on material published in Refs. 64–68.

A technique that is currently used to produce ITBs on JET [69] is the application of lower hybrid heating and current drive (LHCD) (sidebar S13) during the low density plasma current ramp up phase (see Fig. 8 of the Attachment), prior to the high performance phase of a discharge in which high power heating is applied. By this method, often referred to as LHCD preheating, certain populations of resonant electrons are unidirectionally accelerated by electromagnetic waves so that the current density profile (Appendix sidebar S20) is made broader and sometimes even hollow (lower at the center than near the edge), depending on the applied power. In such a case, the q-profile becomes non-monotonic in the core of the plasma at the time of application of the main heating power and the magnetic shear (Appendix sidebar S20) is said to be reversed as it changes sign. Two different methods are used for the main plasma heating in JET, ion cyclotron resonance heating (ICRH) and neutral beam injection (NBI) (sidebar S13). At present, the total injected power can reach up to 25 MW. An interesting advantage of the additional LHCD when it is also applied during the main heating phase is its ability to maintain the preformed broad current profile almost stationary during the main heating phase, whereas it would otherwise peak in the plasma center with a characteristic time scale given by the resistive diffusion time (relaxation time of the profile needed to reach a steady state; Appendix sidebar S19). This peaking tendency is due to the fact that plasma temperature is higher in the core and plasma resistivity (sidebar S13) is inversely related to temperature. The effect of the LHCD is to slow down, and possibly stop, the temporal evolution of the current profile peaking [70]. This also allows avoidance of some instabilities or disruptive events related to the presence in the plasma of magnetic flux surfaces where q is a rational number (such as NTMs for example).

To control, it is necessary to characterize the ITB in a quantitative way. For that purpose a local criterion characterizing the location, strength and dynamics of ITBs in JET has been developed, which can be computed in real time from the ion and electron temperature measurements. This will be described in the following section with the first experiments of ITB strength control that were performed in JET.

Initial Experiments on ITB Control in JET

The objective of the initial ITB control experiments was to investigate practical methods of sustaining ITBs in a controlled and reproducible way in order to take full advantage of the resulting confinement improvement. The goal was to tune the applied heating power so as to maintain the transport barrier and the plasma in a stable state during long periods of time, although not necessarily in steady state. One difficulty in achieving such a goal was to find an objective way to satisfactorily quantify the ITB behavior. For that purpose, a local criterion characterizing the presence, location and strength of ITBs has been developed. The criterion is quantified by calculating the ratio, ρ_T^*, of an ion Larmor radius, ρ_i (Fig. 9 of the Attachment), to the temperature gradient scale length, $(\nabla T/T)^{-1}$. Using an analysis of an experimental JET database, it has been shown [64] that an ITB is most likely to exist at normalized radius, $x = r/a$ (sidebar S5 of the Attachment), and at time t, when

$$\rho_T^*(x,t) = -\rho_i\left[\nabla T(x,t)/T(x,t)\right] > \rho_{ITB}^* \quad , \tag{8}$$

with the threshold value $\rho_{ITB}^* \approx 0.014$. (Here, $\nabla T = \partial T/\partial x$; see sidebar S5 of the Attachment.) The proposed criterion enables detection of the presence of an ITB at a given normalized radius with a large degree of confidence when $\rho_T^*(x,t)$ exceeds the fixed threshold value. In JET, transport barriers are generally observed simultaneously on both the ion and electron heat transport channels, i.e., on the ion temperature gradient (ρ_{Ti}^*) as well as on the electron temperature gradient (ρ_{Te}^*). For specificity, we shall refer to the latter — identified through Eq. (8) applied to ρ_{Te}^* — as an electron transport barrier or electron ITB, so that

$$\rho_{Te}^*(x,t) = -\rho_i\left(\nabla T_e/T_e\right) > \rho_{ITB}^* \quad .$$

Electron transport barriers have been controlled in real-time using the maximum value of the parameter $\rho_{Te}^*(x,t)$ across the plasma radius (x) as the controlled output variable and one power actuator only. Best results were obtained when using the ICRH system as the actuator. The temperature measurements from which ρ_{Te}^* was calculated were made with an heterodyne radiometer using the electron cyclotron emission (ECE) from the plasma [71]. A simple proportional-plus-integral (PI) feedback was used to compute the required actuator input power

$$P(t)[\mathrm{MW}] = P(t_0) + G_p \Delta X(t) + G_I \int_{t_0}^{t} \Delta X(u)\, du \quad ,$$

where $X(t) = \max[\rho_{Te}^*(x,t)]$, $P(t_0)$ is the actuator power at the initial time t_0 of the control, ΔX is the difference between the target output value (setpoint) and the measured output signal $X(t)$, and G_p and G_I are the proportional and integral gains, respectively. The strategy for control of the ITB through control of X using only one actuator (the ICRH heating system) assumes that during the control phase, the current profile (equivalently, q-profile), which has been identified as an essential parameter for the ITB dynamics, does not evolve significantly. This simple strategy is therefore valid only for periods of time that are shorter than the resistive current diffusion time (around 20 seconds or more in JET; Appendix sidebar S19). The high power control phases in control experiments were always limited to 10 s.

In a second set of experiments, the additional effect of a second independent feedback loop to control plasma pressure at the magnetic axis was studied with the aim of combining the ITB confinement improvement with high-β (sidebar S2 of the Attachment) plasma stability and thus avoiding plasma disruptions. This additional control was achieved by measuring the Deuterium-Deuterium (D-D) fusion reaction rate and using neutral beam injection (NBI) as the actuator. (The neutron production from the D-D reactions is strongly correlated with the central plasma pressure.) An experiment with simultaneous control of ρ^*_{Te} — with ICRH — and of the D-D reaction rate — with NBI — is depicted in Fig. 20 [65]. This control was obtained with a constant LHCD power (3 MW) throughout the pulse, which also demonstrated the important role played by LHCD in slowing down the current density profile evolution (Fig. 21) and improving the long-pulse stationarity of these advanced discharges. It was found that setting-up a suitable q-profile, characterized by a weak or even reversed magnetic shear (Appendix sidebar S20), seems to be a key condition for triggering an internal transport barrier that can then be controlled with the concomitant improved plasma confinement. In order to improve the control of the ITB and to allow extended control duration and, later, extrapolation to steady state burning plasma devices such as ITER, control of the q-profile was therefore needed.

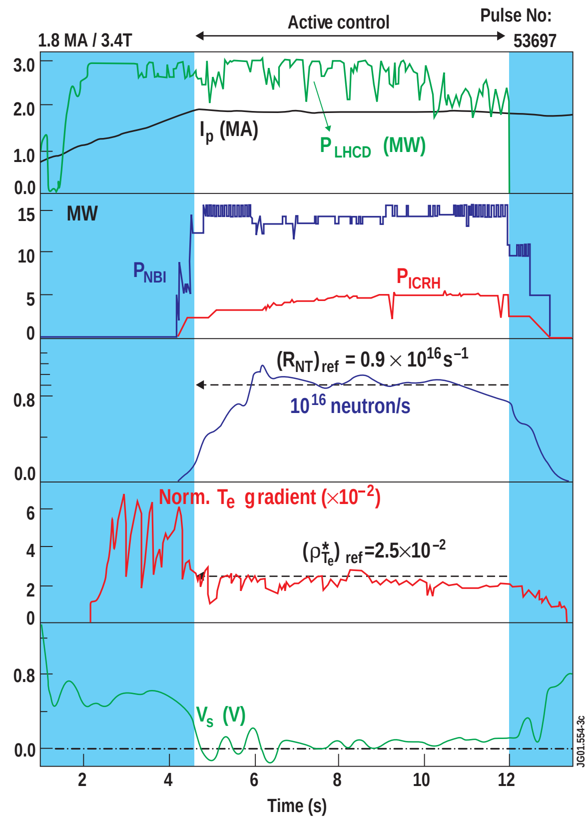

Figure 20. From Ref. 65. Control of an ITB with two Single input-Single output feedback loops. The top two frames show the plasma current, I_p, and the LHCD, NBI and ICRH heating powers. The values of the maximum normalized electron temperature gradient, max[$\rho^*_{T_e}$] (fourth frame), and of the neutron production rate, R_{NT} (third frame), are maintained close to their setpoints, using ICRH and NBI as actuators respectively. Control starts at 4.5 s and the setpoint values are 0.025 for max[$\rho^*_{T_e}$] and 0.9×10^{16} neutron/s for the neutron production rate. The control of the ITB is sustained for 7.5 s. During the whole period of time when the control is applied the loop voltage, V_s (bottom frame), is close to zero, implying that the current is entirely driven by non-inductive sources (sidebar S13) including the self-generated bootstrap current (sidebar S15). The LHCD power is kept approximately constant ($\approx$3 MW) during the whole control phase to slow down the q-profile relaxation.

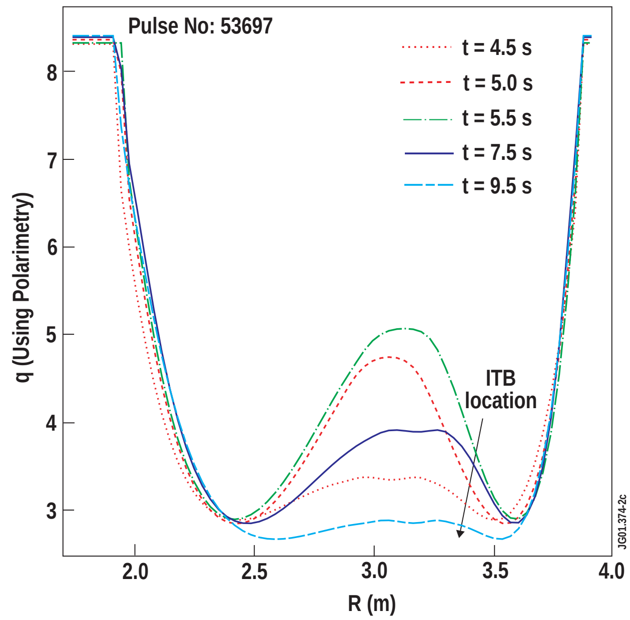

Figure 21. From Ref. 65. Illustration of the current diffusion process. Time evolution of the q-profile as calculated from a magnetic equilibrium reconstruction code, constrained by polarimetry data, with LHCD held constant (≈3 MW) during the high power heating phase. The figure shows that the q-profile evolution is slow, and in particular that the minimum value of q is almost frozen, with a direct effect on the ITB evolution which is practically stationary around R = 3.4 m where $q \approx 3$ and the magnetic shear is negative. Nevertheless, the current profile continues to evolve slowly and this simple ITB control could not be extended to pulse durations longer than the resistive time.

Control of the Current Density Profile

The experimental investigations described in this section were the first attempts in JET at controlling the q-profile. (Controlling the q-profile and the current profile are basically equivalent — see Appendix sidebar S20.) To start, the controlled safety factor profile was simply characterized by its values at five discrete fixed radii, these values being considered as an adequate set of lumped parameters to fully describe the system. They were calculated in real time using magnetic measurements together with data from an interferometer-polarimeter diagnostic, which allowed a fairly accurate reconstruction of the magnetic equilibrium in real-time [66]. The three heating and current drive powers, P_{LHCD}, P_{ICRH} and P_{NBI} were used as actuators for the control. A linearized model was obtained experimentally by performing dedicated open-loop experiments and varying the input powers. A linearized Laplace transform model of the form

$$\delta\mathbf{Q}(s) = \mathbf{K}(s)\,\delta\mathbf{P}(s) \quad , \tag{9}$$

was assumed around a reference plasma steady state, where $\delta\mathbf{Q}$ is a (5×1) vector that represents the change in the safety factor from the reference state when the (3×1) input power vector changes by $\delta\mathbf{P}$. The problem was thus reduced to identifying the (5×3) matrix $\mathbf{K}(s)$ and finding a suitable pseudo-inverse. The steady-state gain matrix $\mathbf{K}(0)$ was determined to be sufficient and was deduced experimentally from simple power steps (from the reference state) in dedicated open loop discharges. In order to design a PI feedback controller that maximizes the steady state decoupling of the multiple feedback loops and ensure minimum (in the least square sense) steady-state offset, a truncated singular value decomposition (TSVD) of the steady state gain matrix, $\mathbf{K}(0)$, retaining only two principal components of $\mathbf{K}(0)$ (the third one being associated with a very small singular value) was used [67]

$$\mathbf{K}(0) \approx \mathbf{W} \Sigma \mathbf{V}^T \quad ,$$

(with $\mathbf{V}^T$ the transpose matrix of $\mathbf{V}$). This provides steady state decoupling between modal inputs $\alpha(s) = \mathbf{V}^T \delta\mathbf{P}(s)$, and modal outputs $\beta(s) = \mathbf{W}^T \delta\mathbf{Q}(s)$. Pseudo-modal control techniques could then be used by inverting the diagonal steady state gain matrix, Σ. Thus, the PI controller transfer function matrix $\mathbf{G}(s)$ was chosen as follows:

$$\begin{aligned} \delta\mathbf{P}(s) &= g_c \left[1+1/(\tau_i s)\right] \mathbf{G}(s) \left[\delta\mathbf{Q}_{target} - \delta\mathbf{Q}(s)\right] \\ &= g_c \left[1+1/(\tau_\mathrm{i} s)\right] \mathbf{V} \Sigma^{-1} \mathbf{W}^T \left[\delta\mathbf{Q}_{target} - \delta\mathbf{Q}(s)\right] \quad , \end{aligned}$$

where g_c is the proportional gain and (g_c/τ_i) is the integral gain. These experiments were the first to use three heating and current drive systems to control the q-profile in an ITB tokamak scenario with a significant fraction of the plasma current carried by the bootstrap current. Because of the long current diffusion time scale, the plasma pulse had to be as long as possible for the effectiveness of the controller to be fully assessed (Appendix sidebar S19). Therefore a plasma scenario (Appendix sidebar S18) that had been developed for long pulse studies was selected. Figure 22 shows the result of a closed-loop experiment in which the target q-profile had a slightly reversed magnetic shear (Appendix sidebar S20) in the plasma core. The control was applied between $t = 7$ s and $t = 13$ s, with initial powers at the start of the control phase of 2.5 MW for LHCD, 7 MW for NBI and 3 MW for ICRH. These powers were chosen sufficiently below the power limits of the systems to avoid possibly hitting the saturation of an actuator during the closed-loop experiments. The powers requested by the controller and delivered by the three systems are shown in Fig. 23.

The demonstration of real-time control of the q-profile encouraged new efforts to develop an integrated ITB control, which would include both the current and temperature gradient profiles (Appendix sidebar S20). These two non-linearly coupled profiles are believed to be essential ingredients governing ITB physics. The most recent profile control experiments performed in JET therefore used an extension of the previous model-based technique to control simultaneously $q(x)$ and $\rho^*_{Te}(x)$ considered as distributed parameters characterizing the current and temperature gradient profiles, respectively.

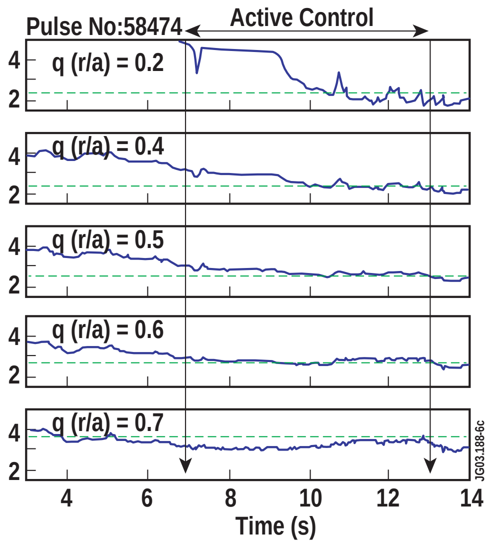

Figure 22. From Ref. 67. A typical example of Multiple Input-Multiple Output feedback control of the q-profile. Time traces are shown of the safety factor at the five radii selected for the real-time experiment with LHCD, NBI and ICRH as actuators. The setpoint values are indicated with dashed lines. The desired setpoints for $q(x)$ (where $x=r/a$; see sidebar S5 of the Attachment) at the five selected radii, x = [0.2, 0.4, 0.5, 0.6, 0.8] were q = [2.35, 2.34, 2.44, 2.69, 3.5] and the control was applied between t = 7 s and t = 13 s. The q-profile had a strong reversed-shear shape at the time when the control started. It then converged slowly towards the closest profile to the one requested that was achievable with the given actuators. A transient undershoot occurred between $t\approx10$ s and $t\approx11$ s and a minimum of the mean square error was reached at $t\approx12$ s.

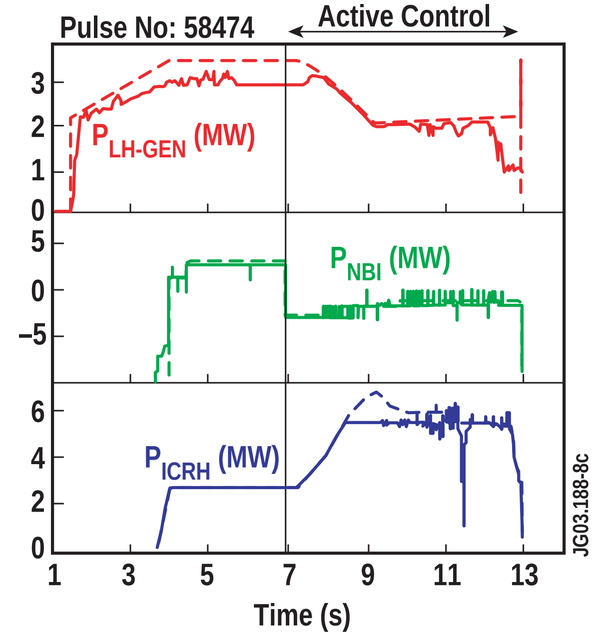

Figure 23. From Ref. 67. Time evolution of the requested (dashed traces) and delivered (solid traces) LHCD, NBI and ICRH powers during the real-time control experiment shown in Fig. 20. The control was applied between t = 7 s and t = 13 s, with initial powers at the start of the control phase of 2.5 MW for LHCD, 7 MW for NBI and 3 MW for ICRH. The powers requested by the controller stayed within the bounds normally allowed by the heating systems. (Here the ICRH power was mistakenly limited to 6 MW, but the power requested by the real-time controller was not significantly larger than that.)

In this case, a discretized representation of their response to the three power inputs (NBI, ICRH, LHCD) can be written in matrix form

$$\delta\mathbf{G}(s) = \mathbf{K}(s)\,\delta\mathbf{P}(s) \quad ,$$

similar to Eq. (9) and a controller can be derived as described above. (See sidebar S21 for details.) Here, $\delta G(s)=[G_{\delta q_1}(s)\ \cdots\ G_{\delta q_{n_a}}(s)\ \ G_{\delta\rho^*_{Te1}}(s)\ \cdots\ G_{\delta\rho^*_{Tenb}}(s)]^T$ represents a finite set of coefficients of two sets of basis functions $a_i(x)$, $i=1,2,...,n_a$ and $b_i(x)$, $i=1,2,...,n_a$ that approximately span the set of achievable q and ρ^*_{Te} profiles, respectively [68] (Fig 24). The PI controller structure was defined as

$$\delta\mathbf{P}(s) = g_c\left[1+1/(\tau_i\, s)\right]\mathbf{K}_{inv}\left[\delta\mathbf{G}_{target}(s)-\delta\mathbf{G}(s)\right] \quad , \tag{10}$$

where g_c is the proportional gain, (g_c/τ_i) is the integral gain, and $\mathbf{K}_{inv}$ is a particular pseudo-inverse of the steady state gain $\mathbf{K}(0)$ (see sidebar S21).

S21. Technique for Distributed-Parameter Control (sidebar)

In summing up the complexity of plasma transport phenomena around a reference plasma state into a simple linearized model having only three degrees of freedom (corresponding to three actuators), one should take into account as much information on the spatial structure of the physical system to be controlled as possible. This can be achieved by retaining the distributed nature of the problem both in the model identification and in the control algorithm.

Also, to design a controller for an ITB regime, the plasma response must be linearized around a stationary reference state presenting an ITB. The current density and temperature gradient profiles were controlled via the safety factor profile, $q(x,t)$, and the parameter $\rho^*_{Te}(x,t)$, respectively. The variations $\delta q(x,s)$ and $\delta\rho^*_{Te}(x,s)$ (s is the Laplace-transform variable) around the reference stationary state can be represented as a (2×1) profile vector $\delta\mathbf{Y}(x,s)$

$$\delta\mathbf{Y}(x,s) = \begin{bmatrix} \delta q(x,s) \\ \delta\rho^*_{T_e}(x,s) \end{bmatrix} .$$

Assuming a time-independent process, the linearized response of the two-function-vector $\delta\mathbf{Y}(x,s)$ to the variation of the heating and current drive powers

$$\delta\mathbf{P}(x,s) = \begin{bmatrix} \delta \mathrm{P}_{LHCD}(x,s) \\ \delta\mathbf{P}_{NBI}(x,s) \\ \delta\mathbf{P}_{ICRH}(x,s) \end{bmatrix} ,$$

can be written, without loss of generality, in the integral form

$$\delta\mathbf{Y}(x,s) = \int_0^1 \mathbf{K}(x,x',s)\ \delta\mathbf{P}(x',s)\,dx' ,$$

where the kernel $\mathbf{K}(x,x',s)$ is to be determined. This kernel is assumed to be square-integrable so that it admits an infinite singular value decomposition:

$$\mathbf{K}(x,x',s) = \sum_{i=1}^{\infty} \sigma_i(s)\ \mathbf{W}_i(x,s)\ \mathbf{V}_i^T(x',s) ,$$

where $\mathbf{V}_i^T(x,s)$ are the transposes of an infinite set of (3×1) matrices of functions, $\mathbf{V}_i(x,s)$, in the input space, $\mathbf{W}_i(x,s)$ are (2×1) matrices of functions in the output space, and $\sigma_i(s)$ are the corresponding positive singular values.

The essence of the method is to identify the best experimental approximation of this kernel by means of its dominant singular elements, or principal components, and to use this approximation to define a suitable controller. The $\delta q(x,s)$ and $\delta\rho^*_{T_e}(x,s)$ profiles are thus approximated by their projections on finite dimensional spaces using the so-called Galerkin scheme

$$\delta q(x,s) = \sum_{i=1}^{n_a} G\,q_i(s)\,a_i(x) + R\,\delta q(x,s) \quad \text{and}$$

$$\delta\rho^*_{T_e}(x,s) = \sum_{i=1}^{n_b} G\,\rho^*_{T_e\,i}(s)\,b_i(x) + R\,\delta\rho^*_{T_e}(x,s) \quad ,$$

where n_a and n_b the dimensions of two subspaces spanned by two sets of basis functions, $a_i(x)$ and $b_i(x)$, $G\,q_i(s)$ and $G\,\rho^*_{T_e i}(s)$ are the so-called *Galerkin coefficients* of those basis functions, and $R\,\delta q(x,s)$ and $R\,\delta\rho^*_{T_e}(x,s)$ are the two residuals that are orthogonal to the basis functions $a_i(x)$ and $b_i(x)$, respectively.

In the same way $\delta\mathbf{P}(x,s)$ can be written as follows:

$$\delta\mathbf{P}(x,s) = \mathbf{C}(x)\,\delta\mathbf{P}(s) \quad ,$$

with $\delta\mathbf{P}(s)$ containing the three power variations from the reference steady state powers

$$\delta P(s) = [\delta P_{LHCD}(s) \quad \delta P_{NBI}(s) \quad \delta P_{ICRH}(s)] \quad ,$$

and $\mathbf{C}(x)$ is a diagonal matrix containing three normalized power deposition profiles.

Anticipating that only three singular values can be found with only three independent actuators, the SVD of the kernel $\mathbf{K}(x,x',s)$ can be truncated as follows

$$K_t(x,x',s) = \sum_{i=1}^{3} \sigma_i(s) W_i(x,s) V_i^T(x',s) \quad ,$$

where only the singular vectors associated with the three largest singular values are retained. Projecting the singular vectors on the corresponding basis functions, a matrix relation between the new outputs (Galerkin coefficients) and the inputs (power modulations) can be obtained

$$\delta\mathbf{G}(s) = \mathbf{K}(s)\,\delta\mathbf{P}(s) \quad ,$$

with

$$\delta G(s) = \left[G_{\delta q_1}(s) \;\; \cdots \;\; G_{\delta q_{n_a}}(s) \;\; G_{\delta \rho^*_{Te1}}(s) \;\; \cdots \;\; G_{\delta \rho^*_{Tenb}}(s) \right]^T \quad .$$

As in the previous section, the kernel $\mathbf{K}(s)$ could be identified from power modulation experiments around the target steady state, but for the proof of principle experiments that are described in this article, the steady state gain matrix $\mathbf{K}(0)$ was sufficient. It was deduced from simple step power changes in dedicated open loop experiments. A PI controller that minimizes the integral least square difference between the requested profiles (defined by Galerkin coefficients $\delta\mathbf{G}_{target}$), and the profiles that are obtained in steady state [defined by $\delta\mathbf{G}(s=0)$] can be designed by calculating a particular pseudo-inverse $\mathbf{K}_{inv}$ of $\mathbf{K}(0)$ [67]. This is not immediate because the matrix $\mathbf{K}(0)$ is not square and, therefore, an infinite number of possible pseudo-inverse matrices exist. The calculation of the relevant one requires the definition of a positive definite scaling matrix $\mathbf{B}$ that contains all the necessary information about the distributed nature of the profiles and is of course directly related to the choice of the basis functions. It takes the block diagonal form:

$$B = \begin{pmatrix} B_q & 0 \\ 0 & \mu B_{\rho^*_{T_e}} \end{pmatrix} \quad . \tag{11}$$

The general elements (i,j) of the matrices $\mathbf{B}_q$ and $B\rho^*_{T_e}$ contain the scalar product of the basis functions and take the form:

$$\left[B_q\right]_{i,j} = \int_0^1 a_i(x)\,a_j(x)\,dx \quad \text{and} \quad \left[B_{\rho^*_{Te}}\right]_{i,j} = \int_0^1 b_i(x)\,b_j(x)\,dx \quad ,$$

and μ is a constant scaling parameter (chosen equal to 10^4) that is used to treat $q(x)$ (of order 1) and $\rho^*_{T_e}(x)$ (of order 0.01) on the same footing in the minimization while allowing more or less weight to be given to the control of one profile or the other. Then, as shown in Ref. [67], a singular value decomposition of the matrix $\hat{\mathbf{K}} = \Delta\mathbf{K}(0)$, where Δ comes from the Cholesky decomposition of $\mathbf{B}$ ($\mathbf{B} = \Delta^T\Delta$), leads to the construction of $\mathbf{K}_{inv}$.

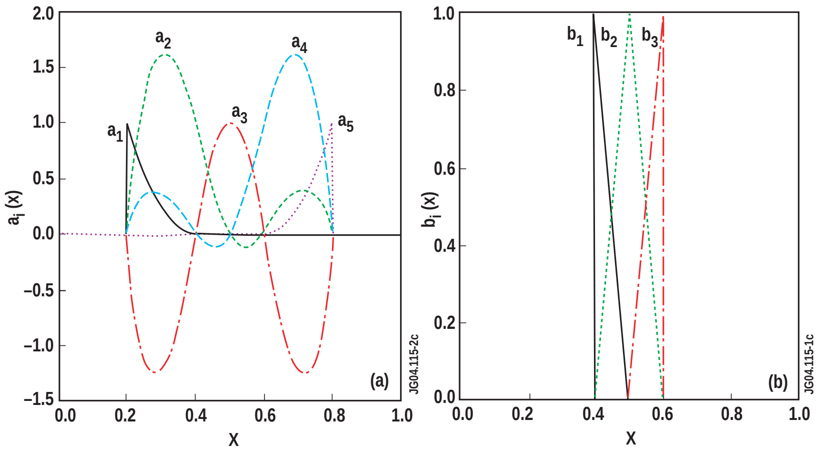

Figure 24. From Ref. 68. Basis functions used for the Galerkin projection of $\rho^*_{T_e}$ profiles (right) and ι-profiles (left). (a) Five cubic splines with knots at x = [0.2, 0.4, 0.5, 0.6, 0.8] have been used to approximate the ι-profiles. (b) Three triangle functions centered at x = [0.4, 0.5, 0.6] have been used to approximate the $\rho^*_{T_e}$ profiles by a piecewise linear function in a reduced domain between x = 0.4 and x = 0.6 where the ITB is requested and controlled.

Simultaneous Control of Current and Temperature Gradient Profiles in JET

The real-time controller described above was used experimentally to control the current density and electron temperature gradient profiles $\rho^*_{Te}(x)$ and obtain an ITB at about half plasma radius. The current density profile was actually controlled via $\iota(x)$ because, being directly proportional to the total current in $[0,x]$ (Appendix sidebar S20), it depends more linearly on the applied current drive power than $q(x)$. To prevent overloading the real-time controller computation, the number of trial basis functions and the radial windows on which they operated were deliberately limited (only part of the full profiles were controlled). The accuracy of the real-time reconstruction of the q-profile from polarimetry data [66] was poor in the central region [0 < x < 0.2] so this region was excluded from the control window. In addition, the q-value at the edge is inversely proportional to the total plasma current (Appendix sidebar S20), which is accurately controlled by the primary (ohmic) circuit of the tokamak (sidebar S13). Therefore, including the edge region in the q-profile control would have been redundant. Thus, the feedback control of the q-profile was restricted to the region $0.2 \le x \le 0.8$. For ρ^*_{Te}, the region of control of the ITB was imposed by limitations in the real-time electron temperature measurements given by the electron cyclotron emission diagnostic, which provides no measurement in the core of the plasma, nor near the edge in discharges with LHCD. The radial measurement window depends on the plasma configuration but includes in all cases the region that extends from x = 0.3 to x = 0.7. Moreover, one of the goals of these experiments was to sustain an ITB at x > 0.4 in order to enhance the plasma performance while q-profiles that are accessible with the present heating systems on JET do not in general allow stationary ITBs at $x \geq 0.6$ to be sustained. Thus, in these experiments the control region for ρ^*_{Te} was restricted to the window $0.4 \leq x \leq 0.6$ where an ITB was expected and requested. The Galerkin coefficients of both profiles (five coefficients for ι and three for ρ^*_{Te}) were computed in real-time from the profile measurements and a power request was sent every 10 ms by the controller to the different actuators according to Eq. (10). The control scheme was applied in multiple plasma discharges for a maximum of 7 seconds per discharge and successfully reached several different target q-profiles — from monotonic to reversed shear — while simultaneously controlling the profile of the electron temperature gradient ρ^*_{Te}. Figure 25 shows the result of applying the control algorithm in the case of a monotonic q-profile target and of a ρ^*_{Te}-profile target with a maximum slightly above the criterion in Eq. (8) for the existence of an ITB, at a fairly large radial location ($x = 0.5$) where ITBs are not easily achieved spontaneously. Both profiles are satisfactorily controlled and the effect of the control can be clearly seen in Fig. 26. In this example, the ICRH system failed to deliver the requested power at around t = 10.25 s, and therefore the control phase was limited to 4.8 s.

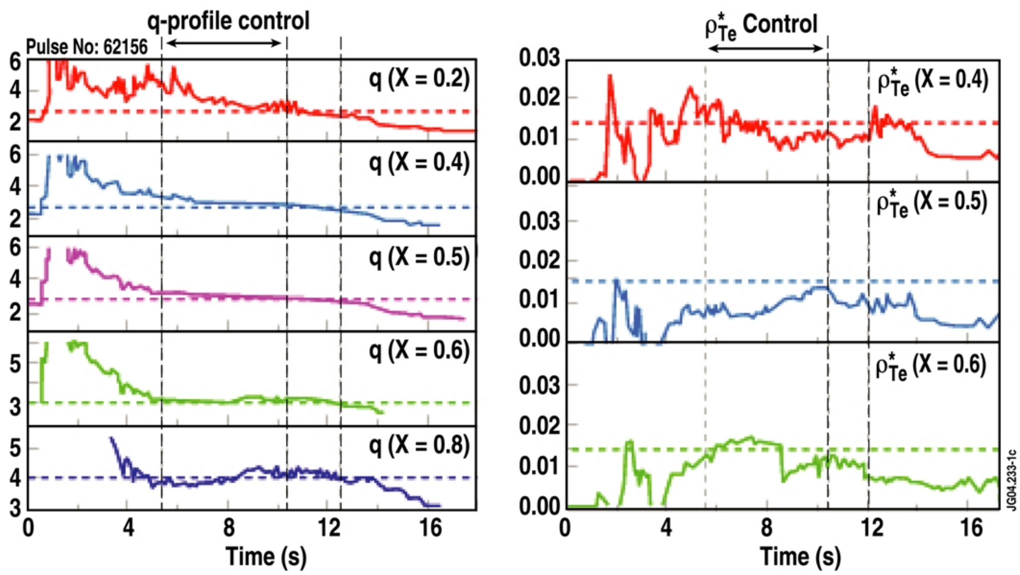

Figure 25. From Ref. 68. A typical example of simultaneous (MIMO) distributed-parameter control of q and $\rho^*_{T_e}$. Time traces of the Galerkin coefficients defining the q-profile (left) and the $\rho^*_{T_e}$ profile (right) during a real-time experiment with LHCD, NBI and ICRH as actuators for controlling simultaneously the q and $\rho^*_{T_e}$ profiles. The corresponding setpoint values are indicated with dashed lines. The control was active between t = 5.5 s and t = 10.25 s. The requested q-profile was monotonic (Appendix sidebar S20). The target profiles are satisfactorily reached at the end of the control phase despite a strong disturbance causing a perturbation on $\rho^*_{Te}\,(x = 0.6)$ at $t \approx 8.5$ s.

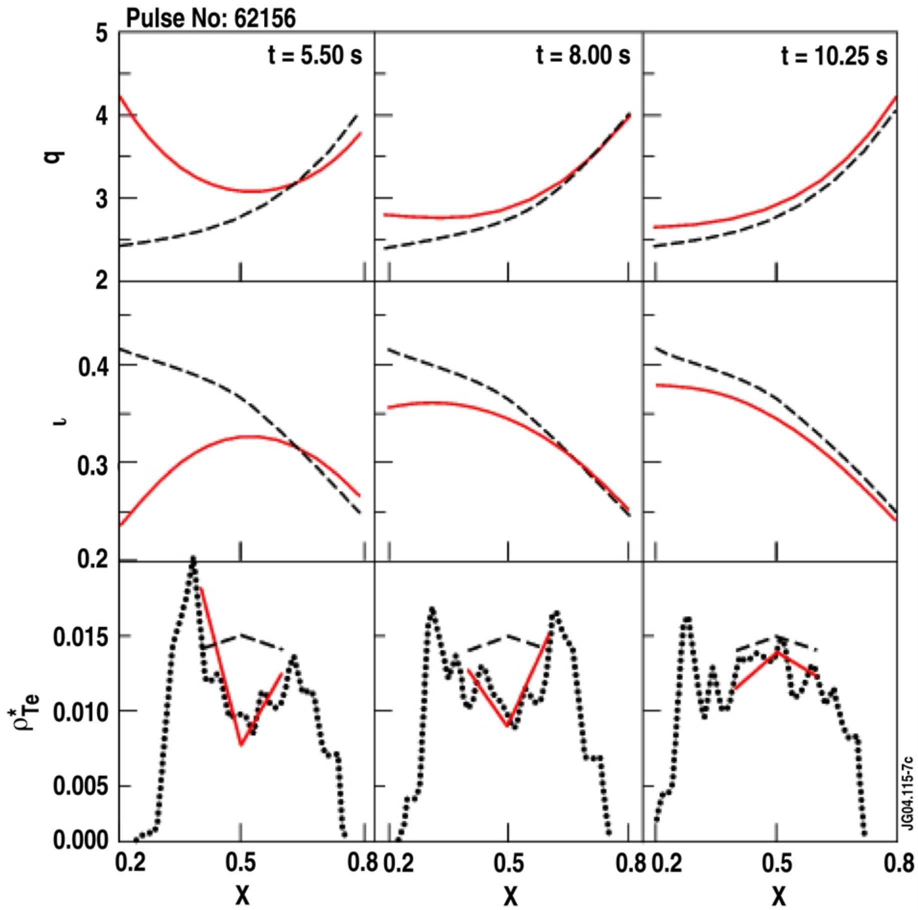

Figure 26. From Ref. 68. Plots of the requested and achieved profiles during the control time window for the pulse #62156. Measured profiles (solid) and target profiles (dashed) for q, ι=1/q and $\rho^*_{T_e}$ after projection onto the span of the basis functions ($\{a_i\}_{i=1}^{n_a}$ for ι and $\{b_i\}_{i=1}^{n_b}$ for $\rho^*_{T_e}$). For $\rho^*_{T_e}$, the measured profile before projection has also been plotted (dotted). Each column corresponds to one time, respectively t = 5.5 s (start of control), t = 8 s, and t = 10.25 s (end of control).

As mentioned in the previous section, the controller defined by Eq. (10) minimizes, in the integral least square sense, the difference between the target ι and ρ^*_{Te} profiles and their respective real-time measurements or, more precisely, the quadratic expression

$$dy^2 = \int_{0.2}^{0.8} \left[\iota(x) - \iota_{setpoint}(x) \right]^2 dx + \mu \int_{0.4}^{0.6} \left[\rho^*_T(x) - \rho^*_{\mathrm{T}_{setpoint}}(x) \right]^2 dx \quad ,$$

where μ is a scalar used for relative weighting of the control objectives [Eq. (11)]. This positive definite quantity, which will be referred to as the squared distance between the achieved and requested profiles, is plotted in Fig. 27. An important feature is the increase of this distance in response to the sudden, undesired, failure of the ICRH system to deliver the requested power at t = 10.25 s. This failure is indeed immediately followed by a strong rise of the ρ^*_{Te} contribution to the distance to be minimized, showing, by contrast, the effectiveness of the control before the failure of the actuator.

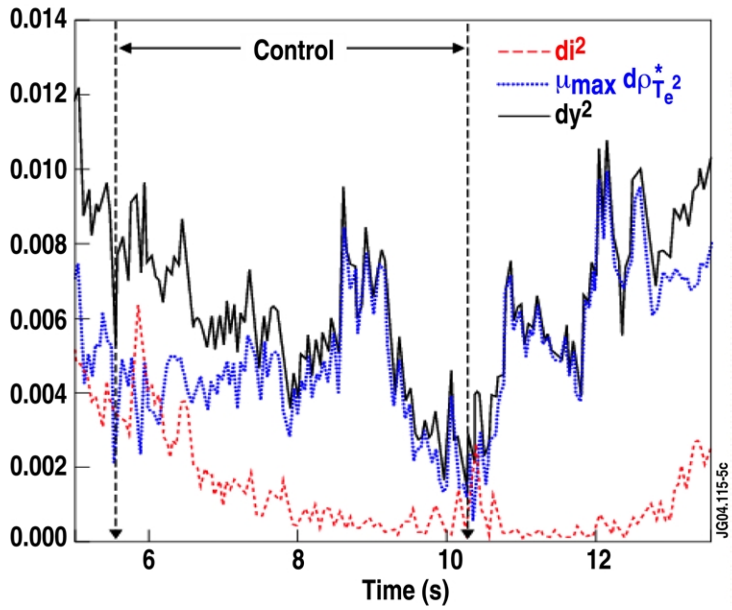

Figure 27. From Ref. 68. Time evolution of the squared distance between the target and the measured profiles for ι {dashed, $d\iota^2 = \int_{0.2}^{0.8} [\iota_{meas}(x,t) - \iota_{target}(x)]\,dx$}, and ρ^*_{Te} {dotted, $d\rho^{*\,2}_{Te} = \int_{0.4}^{0.6} [\rho^*_{Te\,meas}(x,t) - \rho^*_{Te\,target}(x)]\,dx$}, and of the global quadratic distance to be minimized (solid, $dy^2 = dt^2 + \mu_{max}\, d\rho^{*\,2}_{Te}$). The controller minimizes the quadratic expression defined as the global distance between the target and the measured profiles. The effect of the controller is particularly clear when looking at the evolution of the various traces after t = 10.25 s when the controller action stopped because the ICRH system could not deliver the requested power. A large increase of the global distance to the target (solid) can be observed and leads to the loss of the ITB.

Future Directions

The present controllers were designed using only knowledge of the static linear response model, $\mathbf{K}(0)$, but the experimental identification and use of a fully dynamic linear model, $\mathbf{K}(s)$, is now under investigation in order to possibly construct a two-time-scale model (resistive diffusion time for current or q-profile evolution and energy confinement time for temperature or pressure profile evolution) and design a controller that can better cope with fast plasma perturbations (MHD events, spontaneous emergence or collapse of ITBs) while converging slowly towards the requested high performance plasma state. Models based on physics (instead of identified from data) could also be used in future devices to identify adequate linear, or piecewise-linear, response matrices. Present state state-of-the-art plasma transport physics modeling is not yet accurate enough to do so (especially in transient regimes), although it can be quite useful for a qualitative assessment of the control algorithms [72].

ACKNOWLEDGMENT

Work supported by the U.S. Department of Energy under Contract Nos. DE-FC02-04ER54698, DE-AC02-76CH03073 and Grant No. DE-FG02-92ER54141 and performed under the European Fusion Development Agreement (EFDA). We would like to thank M. Chance, M.S. Chu, and A.D. Turnbull for eddy current figures, M. Leuer for Fig. 11, and P.B. Snyder for Fig. 9.

REFERENCES

[1] Z. Chang, et al., "Observation of nonlinear neoclasical pressure-gradient-driven tearing modes," *Phys. Rev. Lett.*, vol. 74 p. 4663, 1995.

[2] R.J. Buttery, et al., "Onset of neoclassical tearing modes on JET," *Nucl. Fusion*, vol. 43, p. 69, 2003.

[3] D.P. Brennan, et al., "A Mechanism for Tearing Onset Near Ideal Stability Boundaries," *Phys. Plasmas*, vol. 10, p. 1643, 2003.

[4] R.J. Lahaye, et al., "Control of Neoclassical Tearing Modes in DIII–D," *Phys. Plasmas*, vol. 9, p. 2051, 2002.

[5] R.W. Callis, et al., "3 MW, 110 GHz ECH system for the DIII–D tokamak," *in Proc. 20th Symp. on Fusion Technol., Marseille, France, 1998* (Association EURATOM-CEA, Saint-Paul-Lez-Durance, 1999) vol. 1, p. 315.

[6] K. Nagasaki, et al., "Stabilization Effect of Early ECCD on a Neoclassical Tearing Mode in the JT-60U Tokamak," *Nucl. Fusion*, vol. 43, p. L7, 2003.

[7] M. Porkolab, "Review of RF Heating," in Theory of magnetically confined plasmas, Proc. Int. School of Plasma Physics, Varenna, 1977, p. 339, Pergamon Press, Oxford (1979).

[8] R.J. Lahaye, et al., "Complete Suppression of the m/n=2/1 Neoclassical Tearing Mode Using Radially Localized Electron Cyclotron Current Drive on DIII–D and the Requirements for ITER," in *Proc. 30th EPS Conf. on Contr. Fusion and Plasma Phys.*, St. Petersburg, Russia, 2003, vol. 27A, P-2.109.

[9] B.W.Rice, "q-profile measurements with the motional Stark effect diagnostic in the DIII–D tokamak," *Fus. Eng. and Design*, vol. 34-35, pp. 135-142, 1997.

[10] ITER [FEAT] Final Design Report (FDR), Plant Description Documents, Chapter 5, Safety, (http://web.gat.com/iter-fdr/final-report-sep-2001/Plant_Descptn_Docs_(PDDs)/PDD_5_Safety.pdf).

[11] ITER [FEAT] Final Design Report (FDR), Plant Description Documents, Chapter 3, Section 7, Control, (http://web.gat.com/iter-fdr/final-report-sep-2001/Plant_Descptn_Docs_(PDDs)/PDD_3.7_Control.pdf)

[12] D. Wróblewski, G.L. Jahns, J.A. Leuer, "Tokamak disruption alarm based on a neural network model of the high-β limit," *Nuclear Fusion*, vol. 37, p. 725, 1997.

[13] M.K. Zedda, et al., "Disruption classification at JET with neural techniques," in *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics*, St. Petersburg, Russia, 2003, ECA, vol. 27A, P-2.93.

[14] G. Pautasso, et al., "On-line prediction and mitigation of disruptions in ASDEX Upgrade," *Nucl. Fusion*, vol. 42, p. 100, 2002.

[15] F.C. Morabito, et al., "Fuzzi-Neural approaches to the prediction of disruptions in ASDEX-Upgrade," *Nucl. Fusion*, vol. 40, p. 1715, 2001.

[16] S.C. Jardin, et al., "A fast shutdown technique for large tokamaks," *Nucl. Fusion*, vol. 40, p. 923, 2000.

[17] D.G. Whyte, et al., "Mitigation of tokamak disruptions using high-pressure gas injection," *Phys. Rev. Lett.*, vol. 89, p. 055001/1, 2002.

[18] P.M. Frank, "Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy—a survey and some new results," *Automatic*a, vol. 26, p. 459, 1990.

[19] R. Isserman and P. Balle, "Trends in the application of model-based fault detection and diagnosis of technical processes," *Control Eng. Practice*, vol. 5, p. 709, 1997.

[20] A. Rizzo, M.G. Xibilia, "An innovative intelligent system for sensor validation in tokamak machines," *IEEE Transactions on Control Systems Technology*, vol. 10, p. 421, 2002.

[21] L. Fortuna, et al., "A neural networks based system for post pulse fault detection and disruption data validation in tokamak machines," in *Proc. 1999 IEEE Int. Symp. on Circuits and Systems*, VLSI, vol. 5, p. 563, 1999.

[22] M.L. Walker, et al., "Automated fault detection for DIII–D tokamak experiments," in *Proc. 18th IEEE/NPSS Symp. on Fusion Engineering*, Piscataway, New Jersey, p. 539, 1999.

[23] J.C. Giarratano, G. Riley, "Expert systems principles and programming," PWS Pub. Co.

[24] E. Jotaki and S. Itoh, "A data acquisition method against unpredictable events during long-time discharges and its application to the TRIAM-1M tokamak experiment," *Fusion Tech.*, vol. 32, p. 487, 1997.

[25] K. Blackler and A.W. Edwards, "The JET fast central acquisition and trigger system," *JET Report*, JET-P(93) 49, 1993.

[26] V. Mertens, G. Raupp, and W. Treutterer, "Plasma control in ASDEX Upgrade," *Fusion Sci. and Tech.*, vol. 44, p. 593, 2003.

[27] H. Zohm, "Edge localized modes (ELMs)," *Plasma Phys. Control. Fusion*, vol. 38, pp. 105-128, 1996.

[28] R.J. Groebner, "An emerging understanding of H–mode discharges in tokamaks," *Phys. Fluids B*, vol. 5, p. 2343, 1993.

[29] G. Federici, et al., "Assessment of erosion of the ITER divertor targets during type I ELMs," Plasma Phys. Control. Fusion vol. 45, pp. 1523-1547, 2003.

[30] P.B. Snyder, et al., “Progress in the Peeling-Ballooning Model of ELMs: Toroidal Rotation and 3D Nonlinear Dynamics,” 31st EPS Conference on Plasma Physics, London UK, June 2004

[31] P.T. Lang, “Active edge localized mode (ELM) frequency control with pellets,” in *Proc. 30th EPS Conf. on Contr. Fusion and Plasma Phys.*, St. Petersburg, Russia, 2003, ECA vol. 27A, p. 1.129.

[32] A.W. Degeling, “Magnetic triggering of ELMs in TCV,” in *Proc. 30th EPS Conf. on Contr. Fusion and Plasma Phys.*, St. Petersburg, Russia, 2003, ECA.

[33] M. Greenwald, et al., “Studies of EDA H–mode in Alcator C–Mod,” *Plasma Phys. Control. Fusion*, vol. 42, pp. A263-A269, 2000.

[34] K. Kamiya, et al., “Observation of high recycling steady H-mode edge and compatibility with improved core confinement mode on JFT-2M,” *Nucl. Fusion*, vol. 43, pp. 1214–1219, 2003.

[35] K.H. Burrell, et al., “Quiescent H–mode plasmas in the DIII–D tokamak,” *Plasma Phys. Control. Fusion*, vol. 44, pp. A253-A263, 2002.

[36] K.H. Burrell, et al., “Quiescent double barrier H–Mode plasmas in the DIII–D tokamak,” *Phys. Plasmas*, vol. 8, p. 2153, 2002.

[37] W. Suttrop, et al., “Stationary ELM-free H–mode in ASDEX upgrade,” *Plasma Phys. Control. Fusion*, vol. 46, p. A151, 2004.

[38] T.E. Evans, R.A. Moyer, J.G. Watkins, et al., “Suppression of large edge localized modes in high confinement DIII–D plasmas with a stochastic magnetic boundary,” presented at the 16th Int. Conf. on Plasma Surface Interactions, Portland, Maine, 2004 in press, *J. Nucl. Mater.*; and T.E. Evans, R.A. Moyer, P.R. Thomas, et al., *Phys. Rev. Lett.*, vol. 92, p. 235003-1, 2004.

[39] R. Aymar, et al., “ITER: fusion research at the dawn of a new era,” 19th IAEA Fusion Conference, Lyon, France, October 14-19, 2002, OV/1-1.

[40] L.C. Bernard, et al., “GATO: An ideal MHD stability code for axisymmetric plasmas with internal separatrices,” *Commun. Comput. Phys.*, vol. 24, p. 377, 1981.

[41] T.S. Taylor, et al., “Optimized profiles for improved confinement and stability in the DIII–D tokamak,” *Plasma Phys. Control. Fusion*, vol. 36, p. B229, 1994.

[42] M. Okabayashi, et al., “Role of the stabilizing shell in high-beta, low-q disruptions in PBX-M,” *Nucl. Fusion*, vol. 36, p. 1167, 1996.

[43] A.M. Garofalo, et al., “Sustained rotational stabilization of DIII–D plasmas above the no-wall beta limit,” *Phys. Plasmas*, vol. 9, p. 1997, 2002.

[44] S. Takeji, et al., "Resistive instabilities in reversed shear discharges and wall stabilization on JT-60U," *Nucl. Fusion*, vol. 42, p. 5, 2002.

[45] M. Okabayashi, et al., "Stabilization of the resistive wall mode in DIII–D by plasma rotation and magnetic feedback," *Plasma Phys. Control. Fusion*, vol. 44, p. B339, 2002.

[46] E.J. Strait, et al., "Resistive wall mode stabilization with internal feedback coils in DIII–D," to be published in *Phys. Plasmas*, 2004.

[47] M. Okabayashi, N. Pomphrey, R.E. Hatcher, "Circuit equation formulation of resistive wall mode feedback stabilization schemes," *Nucl. Fusion* vol. 38, p. 1607, 1998.

[48] Allen H. Boozer, "Equations for studies of feedback stabilization," *Phys. Plasmas*, vol. 5, p. 3350, 1998.

[49] Y.Q. Liu, A. Bondeson, "Active Feedback Stabilization of Toroidal External Modes in Tokamaks," *Phys. Rev. Lett.*, vol. 84, p. 907, 2000.

[50] M.S. Chu, et al., "Modeling of feedback and rotation stabilization of the resistive wall mode in tokamaks," *Nucl. Fusion*, vol. 43, pp. 196-201, 2003.

[51] D.H. Edgell, et al., "Magnetohydrodynamic mode identification from magnetic probe signals via a matched filter method," *Rev. Sci. Instrum.*, vol. 73, p. 1761, 2002.

[52] C.M. Fransson, D.H. Edgell, D.A. Humphreys, M.L. Walker, "Model validation, dynamic edge location mode discrimination, and high confidence resistive wall mode control in DIII–D," *Phys. Plasmas*, vol.10, no.10, pp. 3961-3974, 2003.

[53] C.M. Fransson, B. Lennartson, C. Breitholtz, A. Bondeson, Y.Q. Liu, "Feedback stabilization of nonaxisymmetric resistive wall modes in tokamaks," II. Control Analysis, *Phys. Plasmas*, vol. 7, no.10, pp. 4143-4151, 2000.

[54] A. Bondeson, et al., "Active feedback stabilization of high beta modes in advanced tokamaks," *Nucl. Fusion*, vol. 41, p. 455, 2001.

[55] M. Okabayashi, et al., "Active feedback stabilization of the resistive wall mode on the DIII–D device," *Phys. Plasmas*, vol. 8, p. 2071, 2001.

[56] J. Bialek, et al., "Modeling of active control of external magnetohydrodynamic instabilities," *Phys. Plasmas*, vol. 8, p. 2170, 2001.

[57] M.S. Chu, et al., "Modeling of Feedback Stabilization of Resistive Wall Mode in Rotating Plasmas," *Phys. Plasmas*, vol. 11, p. 2497, 2004.

[58] T.S. Taylor, et al., "Physics of advanced tokamaks," *Plasma Phys. Control. Fusion*, vol. 39, p. B47, 1997.

[59] J. Connor et al., "A review of internal transport barrier physics for steady state operation of tokamaks," *Nucl. Fusion*, vol. 44, p. R1, 2004.

[60] F. Wagner, et al., “Regime of improved confinement and high beta in neutral beam heated divertor discharges of the ASDEX tokamak,” *Phys. Rev. Lett.*, vol. 49, p. 1408, 1982.

[61] T.J.J. Tala, et al., “ITB formation in terms of $\omega_{E\times B}$ flow shear and magnetic shear s on JET,” *Plasma Phys. Control. Fusion*, vol. 43, 507, 2001.

[62] E. Joffrin, et al., “Triggering of internal transport barrier in JET,” *Plasma Phys. Control. Fusion*, vol. 44, p. 1739, 2002.

[63] X. Garbet, et al., “Micro stability and transport modelling of internal transport barriers on JET,” *Nucl. Fusion*, vol. 43, p. 975, 2003.

[64] G. Tresset, et al., “A dimensionless criterion for characterising internal transport barriers in JET,” *Nucl. Fusion*, vol. 42, p. 520, 2002.

[65] D. Mazon, et al., “Real-time control of internal transport barriers in JET,” *Plasma Phys. Control. Fusion*, vol. 44, p. 1087, 2002.

[66] L. Zabeo, et al., “A versatile method for the real-time determination of the safety factor and density profiles in JET,” *Plasma Phys. Control. Fusion*, vol. 44, p. 2483, 2002.

[67] D. Moreau, et al., “Real-time control of the q-profile in JET for steady state advanced tokamak operation,” *Nucl. Fusion*, vol. 43, p. 870, 2003.

[68] L. Laborde, et al., “A model-based technique for real-time control of current and pressure profiles in the JET tokamak,” submitted to *Plasma Phys. Control. Fusion*, 2004.

[69] C.D. Challis, et al., “Effect of q-profile modification by LHCD on internal transport barriers in JET,” *Plasma Phys. Control. Fusion*, vol. 43, p. 861, 2001.

[70] X. Litaudon, et al., “Towards fully non-inductive current drive operation in JET,” *Plasma Phys. Control. Fusion*, vol. 44, p. 1057, 2002.

[71] E. De la Luna, et al., “Recent developments of ECE diagnostics at JET,” *13th Joint Workshop on ECE and ECRH 2004*, http://www.ec13.iapras.ru/on-line-papers.htm, and E. Joffrin, et al., ‘Integrated scenario in JET using real-time profile control,” *Plasma Phys. Cont. Fusion*, vol. 45, p. A367, 2003.

[72] T. Tala, et al., “Progress in transport modelling of internal transport barrier and hybrid scenario plasmas in JET,” submitted to Proc. 20th IAEA Fusion Energy Conference, Vilamoura, Portugal, 2004.

APPENDIX
ADDITIONAL SIDEBARS

These sidebars will be located in other papers in the same special issue.

S18. Tokamak Scenarios (sidebar)

The term *scenario* is often used in tokamak plasma physics to refer to a planned sequence of operating points for a particular tokamak device and plasma. The definition of these operating points usually includes specification of a particular sequence of plasma cross-sectional shapes and plasma current values, but it can also include a number of other plasma parameters such as the value of β (sidebar S2 of the Attachment) or, more generally, characteristics of plasma profiles (sidebar S5 of the Attachment) and other plasma characteristics. Because increasing the performance of the plasma (sidebar S2 of the Attachment) tends to also make it more susceptible to instabilities, planning and simulating scenarios to achieve high performance while maintaining plasma controllability is an important part of the work that occurs at the interface between plasma physics and plasma control.

When scenarios are discussed in the context of control, they usually include the evolution of operating points for the actuators that are needed to control desired features of the plasma. This is done for two reasons. Given limitations in actuators, it provides some confidence that a specified sequence of operating points can be produced by the device. It also enables use of feedforward programming of the actuators, which facilitates linear feedback control of the nonlinear plasma system that is linearized around successive operating points. For example, nominal values of the PF coil currents that are predicted to be consistent with a sequence of plasma shapes are sometimes used to define pre-programmed nominal trajectories for these currents or for the voltage actuators driving the PF coils. Scenarios can also include specification of whether and how actuators will be used. For example, a particular heating and current drive actuator (sidebar S13) can be used for either heating or for current drive or not used at all.

The term scenario is not completely well defined in the plasma physics and control literature. It sometimes implies a complete specification of the plasma evolution, and at other times refers only to plasmas having a specific characteristic behavior, such as a particular profile shape. Usually, it is clear from the context which of these meanings is assumed.

S19. Tokamak Timescales (sidebar)

The various processes in tokamak plasmas evolve on a number of distinctively different characteristic timescales. These time scales range from less than a microsecond to many seconds. The fastest of these is the so-called *Alfven timescale*, which refers to the growth time of those

MHD (sidebar S3 of the Attachment) instabilities that have no passive stabilizing effects (the *Alfven instabilities* or *Alfven eigenmodes*). This timescale is also referred to as the *ideal MHD timescale* and the instabilities as *ideal MHD instabilities*. For example, the vertical displacement instability of a vertically elongated plasma (sidebar S6 of the Attachment) would have a growth time of a few microseconds or less were it not for the stabilizing influence of induced *eddy currents* in passive conducting structures (sidebar S17 of the Attachment). When partially stabilized by induced currents in passive conductors, ideal MHD instabilities are converted to instabilities that grow on a time scale determined by the time needed for the induced currents to decay away due to conductor resistance. This time scale is sometimes called the *wall time* or *resistive wall time*. On present day devices, this is usually on the order of a few milliseconds but can vary considerably since it is determined by the resistance and inductance of the set of conducting elements in the device.

Other important timescales include the times required for evolution of transport quantities (density, temperature, and pressure profiles) and for evolution of the current profile (Appendix sidebar S20). In general, the current profile requires a much longer time to evolve (on the order of 5 or 10 times greater) than the transport quantities. Characteristic timescales are often known by more than one label, which can be a bit confusing. For example, *transport timescales* are also known as *confinement timescales*. Terms used for several similar current profile evolution characteristic times include the *(global) resistive diffusion time, current diffusion time, skin time*, or *current relaxation time*, the multiple names reflecting different characterizations of the physical processes that define the evolution. Although there are a number of other tokamak process timescales, those described above are the most relevant to the discussions in this special issue.

The issue of characteristic time scale is an important one for tokamak control. For example, an unstable process with a fast time scale stresses the importance of fast real time control calculations and actions. This is why the control of the relatively fast and low dimensional plasma vertical instability is usually performed separately from the slower, high dimensional plasma boundary shape control.

Control experiments must also run for a length of time sufficient to judge the effectiveness of new control methods. A primary goal of almost all tokamak control is the maintenance of some performance objective in steady state conditions. The precise length of time that defines a steady state evaluation is not universally agreed upon, but it clearly must have a length equivalent to multiple characteristic times of the process under control. As more control integration takes place, a significant challenge is to combine controls operating on the many different time scales.

S20. Profiles (sidebar)

A number of relationships between several plasma profiles (sidebar S5 of the Attachment) are used in this article. We describe some of these relationships here in order to switch easily from one to another in the discussion. The safety factor (q) profile was defined in sidebar S5 of the Attachment. An alternate definition is given by

$$q(\rho) = \oint \frac{1}{R} \frac{B_\phi}{B_p} \, ds \quad ,$$

where B_ϕ is the toroidal field (field in the direction of the toroidal coordinate ϕ; see Fig. 12 of the Attachment), B_p is the poloidal field (field in the (R,Z) plane orthogonal to the coordinate ϕ), and the integration is carried out over a single poloidal circuit around the flux surface corresponding to the normalized flux value ρ (sidebar S5 of the Attachment) [Wesson]. This is one of the definitions for q that indicates why q is considered a measure of magnetic pitch. The toroidal field B_ϕ is dominated by the contribution from the toroidal field coils (see Fig. 5 of the Attachment), which are typically operated so as to produce an approximately constant (in time) toroidal field. Thus, in most experiments, the safety factor profile is considered to be primarily a function (in time) of the variable poloidal component of the magnetic field. The poloidal field B_p is produced by toroidal currents, including the current in the plasma and current in the poloidal field (PF) coils. When the plasma shape is controlled at a steady state equilibrium, the PF coil currents are nearly constant, so changes in poloidal field are dominated by changes in the spatial distribution of plasma toroidal current density (the *current profile*). Through this chain of dependencies, it can be seen that the safety factor profile depends on the current profile (and vice versa). Thus, many physicists speak interchangeably of control of the current profile and of the q-profile.

A quantity known as the local *magnetic shear* is proportional to the spatial derivative of the safety factor, $\mathrm{s}(\rho) \propto dq/d\rho$. Magnetic shear plays a role in plasma stability, but in this article (and in much of the literature) it is used simply as an alternative description for the behavior of the q-profile. In particular, the notion of *negative (central) shear*, also known as *reverse shear*, describes a q-profile that is not monotonic (Fig. A–1). Another quantity related to q is the inverse of the safety factor known as iota, $\iota(\rho) = 2\pi / q(\rho)$. It can be shown that $\iota(\rho)$ is proportional to the total current inside the flux surface represented by the normalized flux value ρ. In particular, this means that the value $\iota(1) = 1/q(1)$ at the edge of the plasma is inversely proportional to the total plasma current. The q-, s-, and iota-profiles, are all functions of normalized flux ρ, while the current profile is not. The current profile is typically defined as the plasma current density along a line extending radially from the magnetic axis to the plasma edge (Fig. 14 of the Attachment).

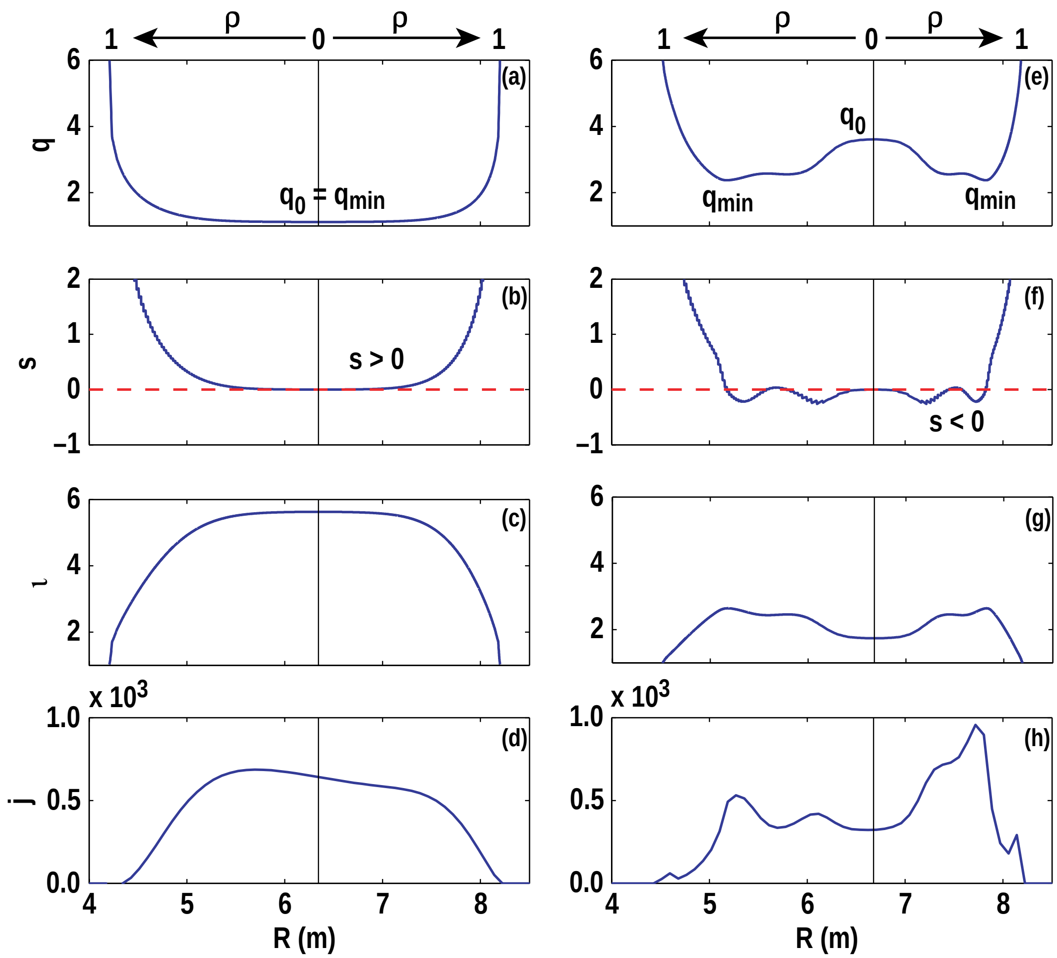

Figure A–1. Examples of monotonic and reverse shear plasmas. Plots (a) through (d) represent the same plasma having a monotonic q-profile (in which q is monotonically increasing on $0 \le \rho \le 1$). Plots (e) through (h) represent the same reverse shear plasma. (The normalized radius $x = r/a$ (sidebar S5 of the Attachment), although not identical to ρ, may be substituted for ρ everywhere in this figure and its description.) The vertical line in each plot indicates the radial location of the plasma magnetic axis ($\rho = 0$; see Fig. 14 of the Attachment). A monotonic q profile (a) achieves its minimum value q_{min} at the magnetic axis $\rho = 0$ and is monotonically increasing in $0 \le \rho \le 1$. A reverse shear q profile (e) achieves its minimum value at $\rho \ne 0$, away from the magnetic axis. In either case, the value of q at the magnetic axis is denoted by $q_0 = q(\rho = 0)$. (b) The sign of the magnetic shear s for a monotonic q-profile is positive for all $0 \le \rho \le 1$. (f) A reverse shear plasma is one in which $s < 0$ for some set of values of ρ between 0 and 1. In particular, $s < 0$ near the magnetic axis, hence the name *negative central shear*. The quantity iota $\iota = 2\pi / q$ is shown in plots (c) and (g). The current profile [(d),(h)] is not a function of the normalized flux coordinate ρ, and so does not exhibit the same symmetry with respect to magnetic axis as the other quantities. The portion of the curve to the right of the magnetic axis in plots of q, s, and ι appears compressed. This is because the plasma flux contours (see Fig. 14 of the Attachment) are always spaced more closely near the outside of the torus, hence the radial dimension corresponding to the normalized flux $0 \le \rho \le 1$ must be smaller on that side of the magnetic axis.

ATTACHMENT
INTRODUCTION TO FUSION AND TOKAMAKS
GENERAL ATOMICS REPORT GA–A24794

GA–A24794

INTRODUCTION TO FUSION AND TOKAMAKS

by

A. PIRONTI and M.L. WALKER

NOVEMBER 2004

GENERAL ATOMICS

DISCLAIMER

GA–A24794

INTRODUCTION TO FUSION AND TOKAMAKS

by

A. PIRONTI* and M.L. WALKER

This is a preprint of a paper submitted
to IEEE Control Systems Magazine.

*Associazione EURATOM/ENEA/CREATE

Work supported by
the U.S. Department of Energy
under Contract No. DE-FC02-04ER54698

GA PROJECT 30200
NOVEMBER 2004

GENERAL ATOMICS

TABLE OF CONTENTS

LIST OF SIDEBARS

LIST OF FIGURES

LIST OF TABLES

INTRODUCTION TO FUSION AND TOKAMAKS

The fossil fuel era is almost over. If we continue to burn fossil fuels for energy, they will last only another few hundred years [1]. Of more immediate concern is the fact that, at our present rate of use, experts predict an energy shortfall in less than 50 years (Fig. 1). A list of potential energy sources that are candidates to replace fossil fuels is shown in Table 1. Renewable energy sources such as solar, wind, and geothermal are attractive from an ecological viewpoint, but do not provide the energy density (e.g., Megajoules per square kilometer) sufficient to replace the diminishing supplies of fossil fuels in an increasingly urbanized world. Nuclear fission and fusion are candidate sources of energy with sufficient energy density to supply the increasing world population with its steadily increasing energy demands. Nuclear fusion is not listed in Table 1 because it is not yet a commercially available energy source. Hydrogen is also not listed because, contrary to how it is sometimes portrayed, hydrogen is not an energy source. Rather, it is being considered as a method of energy transport as a replacement for the various transportable liquid fuels derived from oil [4]. Production of hydrogen requires energy produced by one of the sources listed in Table 1.

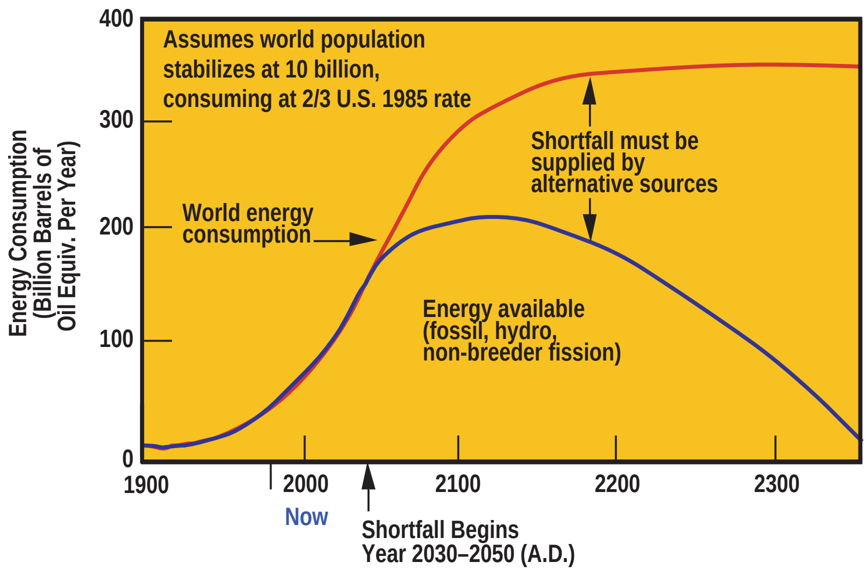

Figure 1. Projected shortfall of energy [2]. Fossil fuels are not predicted to be completely depleted for several hundred years. However, a worldwide energy shortfall is predicted within the next few decades due to the combination of steadily increasing demand and decreasing available energy sources. This shortfall must be met by the development of new alternate sources of energy.

Table 1
Present energy sources with advantages and disadvantages.
(For an expanded discussion of the topics in this table, see Ref. 3.)

Energy Source	Advantages	Disadvantages
Coal (220 Years)	• Abundant	• Burns dirty • Causes acid rain, air pollution, CO_2
Oil (35 Years)	• Flexible fuel source with many derivatives • Transportable	• Finite supply • Causes air pollution, • Produces CO_2
Natural Gas (60 Years)	• Burns cleanly • Transportable	• Finite supply • Produces CO_2
Fission (45 Years) (2700 Years-Breeder)	• Clean, no CO_2 • Does not produce immediate pollution	• Waste disposal is difficult • Safety concerns
Hydroelectric (mostly utilized)	• Clean, no CO_2	• Dam construction destroys habitats • Geographically limited
Wind (low utilization)	• Clean, no CO_2	• Huge numbers of windmills required for adequate power generation • Geographically limited
Geothermal (low utilization)	• Clean, no CO_2	• Geographically limited
Solar (under utilized)	• Clean, no CO_2	• Huge number of solar cells required for adequate power generation • Geographically limited

WHAT IS FUSION?

Nuclear fission produces energy through the splitting of heavy atoms such as uranium in controlled chain reactions and is a mature technology. Unfortunately, the byproducts of fission are highly radioactive and long lasting. Fusion is the process by which the nuclei of two light atoms such as hydrogen are fused together to form a heavier (helium) nucleus with energy produced as a byproduct. This process is illustrated in Fig. 2 where two isotopes of hydrogen combine to form a helium nucleus plus an energetic neutron.

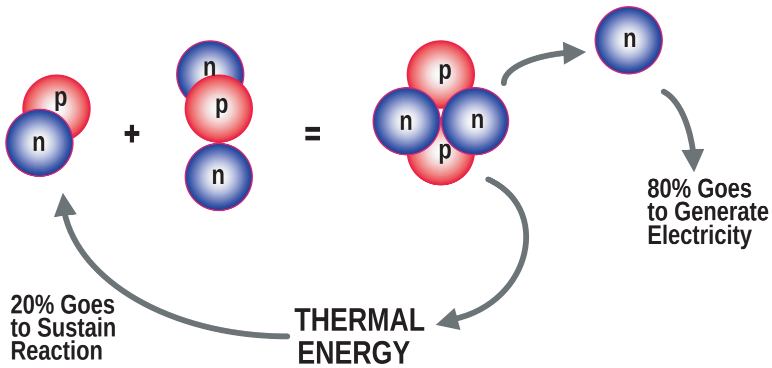

Figure 2. Illustration of a fusion reaction. In a fusion reaction, deuterium and tritium nuclei combine to form a helium nucleus and a free energetic neutron, as well as producing excess heat. The excess heat serves to sustain additional fusion reactions, while the free neuton is captured by the fusion reactor and its energy converted to power.

Controlled fusion is an extremely challenging technology, but a fusion power reactor would produce mostly short term, low level radioactive waste and there is an abundant fuel supply. Figure 3 illustrates the fueling process envisioned for a fusion power reactor. In contrast to fission, fusion poses no risk of nuclear accident. A nuclear meltdown with a large uncontrolled release of energy is an impossibility. Like fission, fusion would produce no air pollution or greenhouse gases during normal operation, since the reaction product is helium. Unlike fission, there is minimal to no high level nuclear waste. The primary source of radioactive byproducts are neutron activated materials (materials made radioactive by neutron bombardment) which can be minimized by careful material selection. Most of these activated materials can be safely and easily disposed of within a human lifetime, in contrast to most fission byproducts, which require special storage and handling for thousands of years.

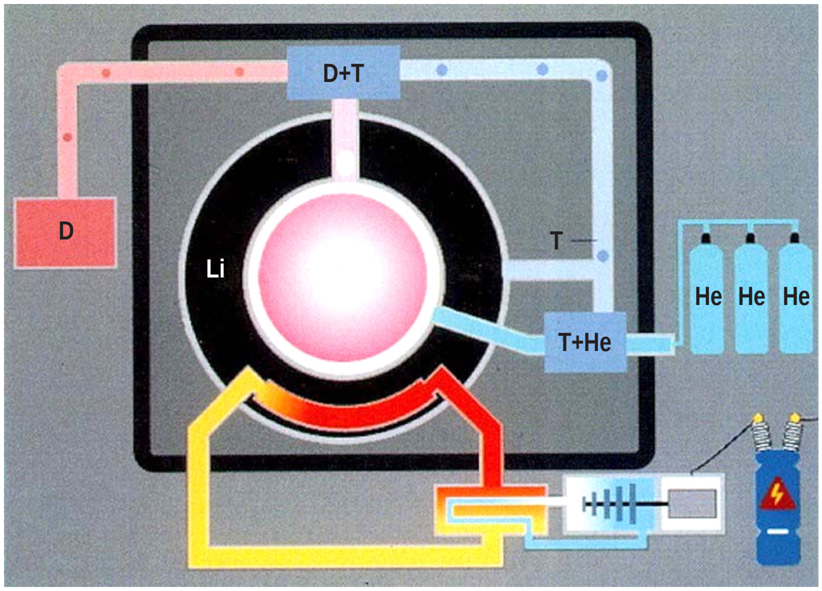

Figure 3. Illustration of the expected deuterium/tritium fueling process for a fusion power reactor. The fusion reaction (D + T = ^{4}He + n) combines deuterium and tritium nuclei (two isotopes of hydrogen) in the center of the reactor, generating energetic neutrons whose (heat) energy is captured by a lithium blanket surrounding the reactor. This heat is converted to electricity in much the same way as heat produced in coal or oil burning plants is converted to electricity. The helium nucleus produced in this reaction is also known as an *alpha particle*. Lithium in the blanket combines with the captured neutrons during another atomic reaction ^{6}Li + n = T + He to provide tritium fuel for additional fusion reactions. Deuterium is available in inexhaustible quantities from sea water (1 part/6500 H_20) while there is sufficient lithium available to last for several thousand years.

WHAT IS A TOKAMAK?

The primary challenge of fusion is to confine a gas comprised of hydrogen isotopes while it is heated and its pressure increases, in order to initiate and sustain fusion reactions. There are three known ways to accomplish this: (1) with gravitational confinement — the method that the sun uses, (2) with inertial confinement — compressing the hydrogen gases via a form of controlled implosion, with inertia then holding them together long enough for fusion reactions to occur, or (3) by magnetic confinement — use of magnetic fields acting on hydrogen atoms that have been ionized (given a charge) so that magnetic fields can exert a force on the moving particles. The ionized gases confined by magnetic fields are referred to as plasmas (another name for ionized gas).

The topic of this special issue is the control of tokamaks [5] — a particular type of magnetic confinement device, constructed in the shape of a torus (doughnut shaped), which is used to conduct experiments in controlled fusion. It was first developed by the Soviet Union and the name is a Russian acronym for Toroidal Chamber and Magnetic Coil. Tokamaks are the most promising of several proposed devices for obtaining nuclear fusion energy from high temperature plasmas.

There are many tokamaks of various sizes around the world. The largest of these is the Joint European Torus (JET) in Culham, England [6], shown in the illustration in Fig. 4, with roughly a dozen medium sized tokamaks such as DIII–D [7], shown in Fig. 5. Devices of these sizes are usually funded by governments or consortiums of governments and have a dedicated support staff plus visiting scientists numbering upwards of 100 or more. There are also many smaller devices, typically located at universities.

Tokamaks are magnetic confinement devices, which means that magnetic fields produced by currents in large coils are used to confine the plasma within a fixed volume. Several of these magnetic coils also serve the additional purposes of shaping, heating, and driving current in the plasma (Figs. 5 and 6). Some tokamaks such as JET incorporate an iron core (shown in Fig. 4) to increase the flux available to drive plasma current while most, such as the DIII–D tokamak have an air core.

Experimental fusion technology has now reached a point where experimental devices are able to produce almost as much energy as is expended in heating and confining the plasma. Figure 7 illustrates the exponential growth of fusion power output of tokamaks that has taken place over the last 30 years. Roadmaps for the development of fusion energy have also been proposed [8],[9]. The immediate next step in these roadmaps is the construction and operation of the proposed International Thermonuclear Experimental Reactor (ITER) burning plasma experiment. The ITER tokamak is intended to provide the next major advancements in fusion physics and technology and is supported by an international consortium of governments. Its purpose is to demonstrate the physics understanding and several key technologies necessary to maintaining burning plasmas (having sustained high levels of fusion reactions).

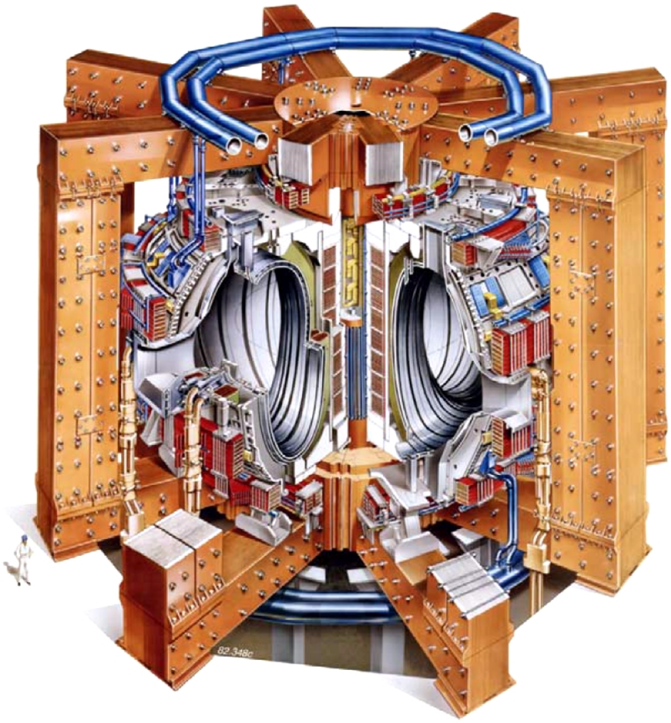

Figure 4. The JET tokamak in Culham, England. Note the size of the device relative to the man in the figure. The large iron rectangles emanating from the center of the torus serve to increase the total flux linking the conducting plasma, thereby increasing plasma current drive efficiency. Holes in the side of the vacuum chamber, called *ports*, are necessary for installation of sensors and access for maintenance. (Image courtesy of EFDA-JET.)

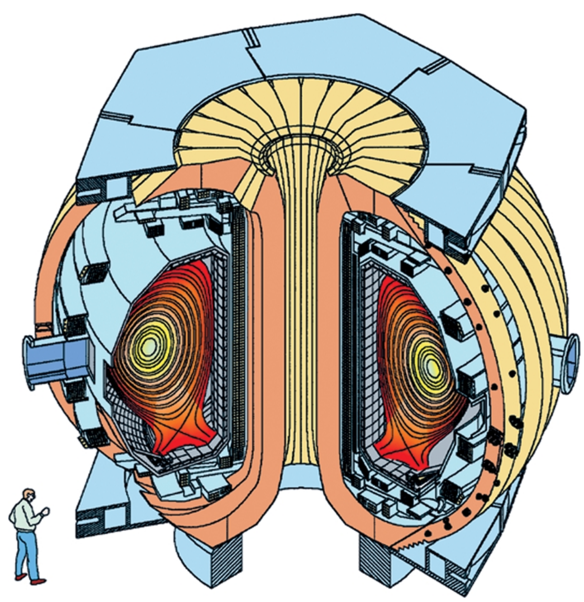

Figure 5. The DIII–D tokamak in San Diego, California, USA. On the exterior of the device can be seen several conducting coils. The *toroidal field coils* (cream colored) are wrapped poloidally around the torus (the short way around — going through the center hole), and the *poloidal field coils* (light blue, cut to show the interior) are wrapped toroidally (the long way) around the torus. Current flowing in these two sets of conducting coils is responsible for producing the magnetic field that is used to confine the plasma. The plasma contained within the device is represented here by a set of contours of constant magnetic flux, or equivalently, constant pressure.

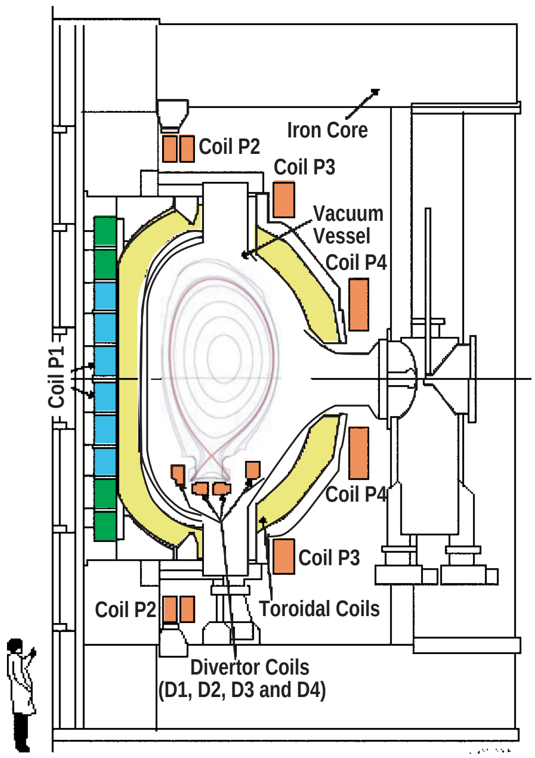

Figure 6. Cross section of the JET tokamak. The coils are color coded here according to their purpose. Currents flowing in the yellow toroidal field (TF) coils produce toroidal magnetic field (directed into the plane of the figure) for confining the plasma within the torus. Current flowing in the green, blue, and orange poloidal field (PF) coils produce poloidal magnetic fields (parallel to the plane of the figure) that are used for confining, shaping, heating, and driving current in the plasma. The P1 coilset along the center column consists of a portion (green) that is used exclusively for current drive and heating and a central portion (blue) that is used for current drive but also works together with the remaining PF coils (orange) to control the shape and position of the plasma. A portion of the P3 coils is also used together with P1 to assist in plasma creation and current drive. Certain other devices employ an entirely separate PF coil, known as the *ohmic coil*, dedicated to heating and current drive. This is the case with DIII–D tokamak in Fig. 5 for example. (Image courtesy of EFDA–JET)

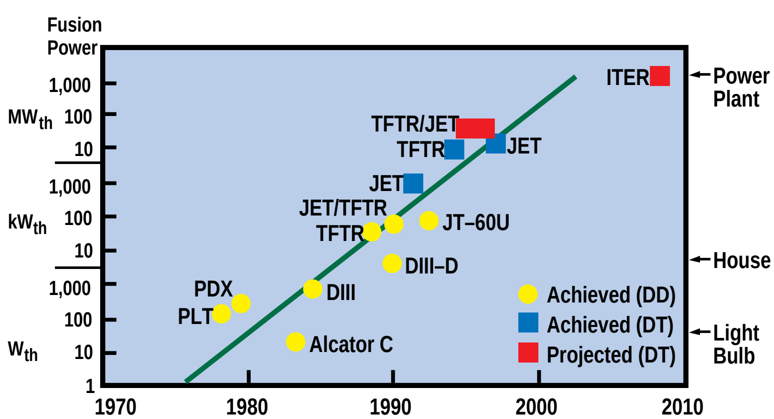

Figure 7. The amount of power generated by tokamak fusion devices has increased by approximately 8 orders of magnitude (a factor of 10^8) in the past 30 years. This practical measure of progress has occurred in parallel with an enormous increase in understanding of the tokamak and plasma physics that was necessary to achieve this progress. The ITER tokamak — the first experimental device that will produce substantially more power (10×) than it consumes — appears to be a very reasonable objective for the next decade of growth.

The planned ITER device will be capable of exploring advanced modes of tokamak operation. These advanced modes rely heavily on active control to develop and regulate high performance plasmas with sufficient plasma density, temperature, and confinement to maintain a self-sustaining fusion reaction for long durations. Tokamaks are high order, nonlinear systems with a large number of instabilities, so there are many extremely challenging control problems that must be solved before fusion power production becomes viable. A number of these solutions will be needed already for the ITER device, which is planned for construction during the next 10 years. The existing experimental devices worldwide are often used as testbeds for application of the control techniques that will be needed to address these urgent control applications.

DESCRIPTION OF A TOKAMAK FUSION EXPERIMENT

All existing tokamaks are pulsed devices. That is, the plasma is maintained within the device for only a finite length of time. There is still no agreement on whether a power reactor must be truly steady state or just maintain a sufficiently long pulse. Figure 8 shows the time traces of several signals acquired during a typical plasma pulse in the DIII–D device, which is fairly representative of experimental tokamaks.

Every experimental tokamak device is different, with capabilities designed to support different experimental goals. Each is designed to operate with different maximum levels of plasma current (a few hundred kilo Amps to a few Mega Amps), toroidal field (a few Tesla), and duration of discharges (a few seconds to several minutes). For example, the goal of the Tore Supra device [10] is to achieve sustained 1000 second pulses. One of the primary missions of this device is to develop an understanding of the needs for steady state plasma confinement and control.

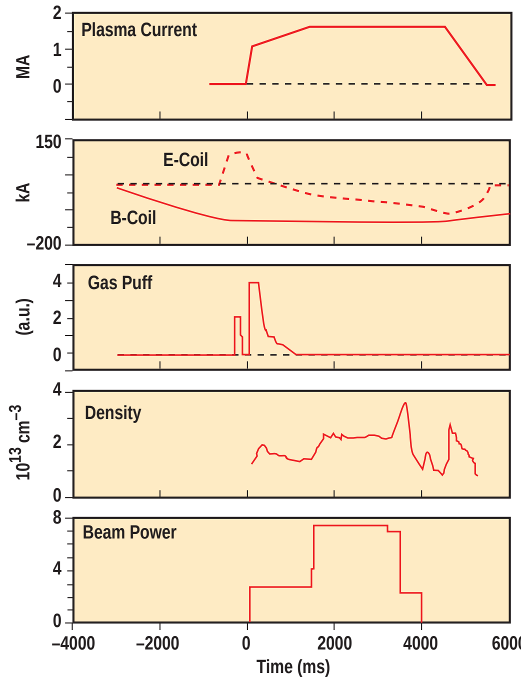

Figure 8. A typical tokamak plasma *discharge* (also referred to as a *shot* or *pulse*), illustrated using data acquired during a DIII–D experiment. Time $t = 0$ corresponds to the time at which the plasma is initiated. The toroidal field coil (B-coil) current is brought up early to create a constant magnetic field to contain the plasma when initially created. Just prior to $t = 0$, either hydrogen or deuterium gas is puffed into the interior of the torus, and the ohmic heating coil (E-coil) is brought to its maximum positive current in preparation for pulse initiation. At $t = 0$, the E-coil current is driven down very quickly in order to produce a large electric field within the torus. This electric field accelerates free electrons, which collide with and rip apart the neutral gas atoms, producing the ionized gas or plasma. Since the plasma consists of charged particles that are free to move, it is a conductor. Thus, immediately after plasma initiation, the E-coil current is commanded to continue its downward ramp so that it now operates as the primary side of a transformer whose secondary is the conductive plasma. This causes current to flow in the plasma via the opposing flows of oppositely charged particles. Collisions of the electrons and ions cause the plasma to be resistive, which causes it to heat (thus the origin of the term *ohmic heating*). The rate of E-coil current decrease is used to control the plasma current I_p up to a target *flat-top* level by about 1 second after plasma initiation. [Maximum plasma currents in most large devices is on the order of a few mega amperes (MA).] Shortly after $t = 0$, additional gas is puffed into the chamber to increase the density and/or pressure to desired levels. In this experimental discharge, neutral beams (uncharged atoms of deuterium) are injected into the plasma at high velocity. These particles collide with particles in the plasma, thereby converting their momentum into heat and further heating the bulk plasma. The separate time intervals in which the plasma current are increasing (0 to 1 second), constant (1 to 5 seconds), or decreasing (5 to 6 seconds) are almost universally referred to as, respectively, the *rampup*, *flattop*, and *rampdown* phases of the plasma.

HOW DO TOKAMAKS WORK?

The fundamental task of tokamak experimental devices is to discover methods to confine and heat the plasma so that sustained fusion reactions can occur. When plasma (hydrogen) ions collide, they are repelled since their electrical charges are the same sign. As their impact energy increases, they approach closer and closer until the probability of a fusion reaction occurring becomes significant. This implies that the plasma must be kept hot (approximately 100,000,000°C) to produce sufficient energy to the overcome the repellant force due to like charges. Simultaneously, the plasma must be kept at high pressure so that particles are close together to increase the frequency at which they collide and are able to fuse. The hotter and denser the plasma gets, the more unstable it becomes. When the plasma reaches any one of a number of stability limits it can disrupt, a process in which the plasma loses confinement and rapidly dissipates its (magnetic and thermal) energy in the confining structure. This is one reason why active control is needed and what makes that control challenging.

Magnetic Confinement

Figure 9 illustrates the difference between a gas that is unconfined and one that is confined by a magnetic field. In the cylinder shown in Fig. 9(b), the particles would stream out the end after a finite time, contrary to the desire to keep them confined. To solve this problem, the tokamak uses field lines bent into a torus so that there is no end. Figure 10 illustrates the effect of toroidal field produced by current in the toroidal field coils (Fig. 5) and poloidal field produced by current in the poloidal field coils and plasma (sidebar S1). The resulting field lines are helical and, more importantly, provide a path for the plasma charged particles to follow that

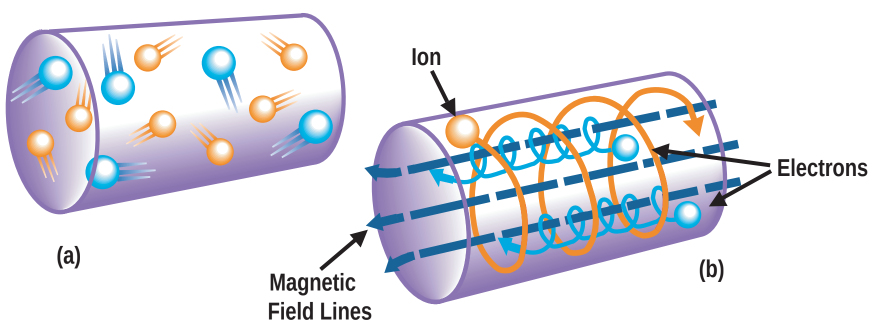

Figure 9. Illustration of basic magnetic confinement concept. In (a), gas is unconfined and free to move, while in (b) the ionized gas in a magnetic field is subject to forces imposed by the field that cause the ions to travel along the magnetic field lines while circling around them with a radius known as the *Larmor radius*. Because the ions and electrons have opposite charges, they move in opposite directions along the field lines under the influence of an electric field. (The source of this electric field will be described later.) Because the positively charged ion is more massive than the electron, it rotates in a much larger radius circle. The number of rotations per second that the ions and electrons rotate around the field lines are known as the *ion cyclotron frequency* and *electron cyclotron freqeuency*, respectively. In (b), the particles will remain confined by the magnetic field until the field lines end or dissipate, contrary to the desire to keep them confined. To solve this, the tokamak bends the field lines into a torus so that there is no end.

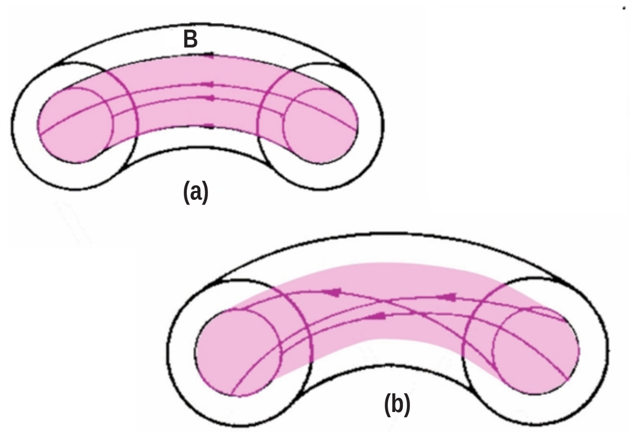

Figure 10. Field lines produced by toroidal and poloidal field coils (see Fig. 5). In a tokamak, the magnetic field lines in (a) are produced by the toroidal field (TF) coils (the cream colored conducting coils in Fig. 5). The letter B is the standard notation for magnetic field. Addition of a poloidal field by the poloidal field by current flowing in the poloidal field (PF) coils (light blue coils in Fig. 5) and in the plasma produces a combined field (b) in which the magnetic field lines are helical.

never leaves the torus. When considering the effect of poloidal field coil current on the plasma, a less precise but more easily grasped concept is the analogy to forces between two wires carrying current (Fig. 11). The toroidal fields are the dominant confining fields; they are typically on the order of a few Tesla. In order for coils to produce magnetic fields of this size, they are wound with many turns of conductor and carry several kilo-Amperes of current, producing an equivalent of many Mega-Amperes of flux producing current (usually referred to as MA-turns). Depending on the resistance of the coils (typically copper, with a move toward superconductors for newer devices), the voltages required to drive these currents (in copper) range from a few volts to hundreds of volts in steady state, with a much larger transient voltage requirement.

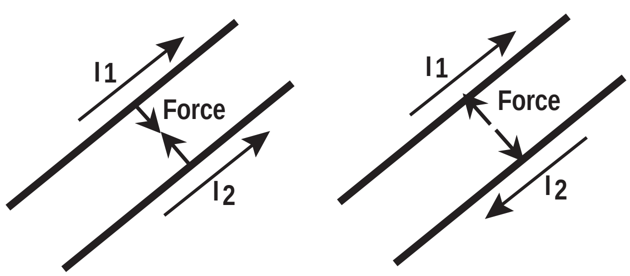

Figure 11. Illustration of forces exerted on the plasma by poloidal field coils. The figure shows two parallel wires carrying current. If the currents are in the same direction, the magnetic fields exert a force so as to push the wires together. If the currents are in opposite directions, the force exerted tends to push them apart. The plasma inside the torus essentially constitutes a big fat wire, since it is made up of charged particles in motion, that is, it has a current. The large poloidal field coils on the outside of the torus push or pull against the plasma using a version of this basic principle.

S1. Flux and Field (sidebar)

Any source of current I produces a magnetic (vector) field B that varies in time directly with the current. Units of magnetic field commonly used in tokamak plasma physics are Tesla (T) and gauss (1 T = 10^4 gauss). Magnetic flux is defined as the integral of the field B through some surface of area A: $\psi = \int_A B \bullet dA$. Thus, references to flux must implicitly include the area that defines that flux. Units of flux are commonly given in Webers.

Tokamak physics analyses commonly consider separately the toroidal field (or flux) and the poloidal field (flux). Toroidal field is the component of the magnetic field directed in the ϕ coordinate direction of the standard cylindrical coordinates (Fig. 12). Poloidal field is any field orthogonal to that [that is, in the (R,z) plane of Fig. 12].

Models used for tokamak control have frequent need for equations relating current (either in actuating coils or induced in conducting structures) to magnetic flux and field at various locations within the device. Standard magnetics calculations [11] can produce the linear relationships needed. Flux ψ_S through a surface S is represented using the mutual inductance from various current sources to the boundary of the surface S: $\psi_S = \mathrm{M}_{S1}\,\mathrm{I}_1 + \mathrm{M}_{S2}\,\mathrm{I}_2 + \mathrm{M}_{S3}\,\mathrm{I}_3 + \ldots$. This boundary will in some cases represent a conducting loop, while in other cases this conductor will be imaginary; the calculation is the same.

A terminology commonly used in tokamak control, especially in control of shape and position, is one of the values of poloidal flux ψ_p at a point $P = (R_p, z_p)$. In this case, ψ_p is defined as the total flux through a surface bounded by a ring passing through the point P and concentric with the axis of revolution of the torus (Fig. 13). In this way, the *poloidal flux function* $\psi_p = \psi_p(R,z)$ can be defined over the cross-section of the vacuum vessel (Fig. 14). This terminology is ubiquitous and facilitates discussion of many of the concepts needed for tokamak control. With this convention, the magnetic field is not commonly computed directly, but usually through the computation of flux. The integral relation above has an inverse relation that provides the field at P:

$$B_p = \frac{1}{R}\left[-\frac{\partial \psi_p}{\partial z} \quad \frac{\partial \psi_p}{\partial R}\right] ,$$

where ψ_p is the value of poloidal flux at the point $P = (R_p, z_p)$.

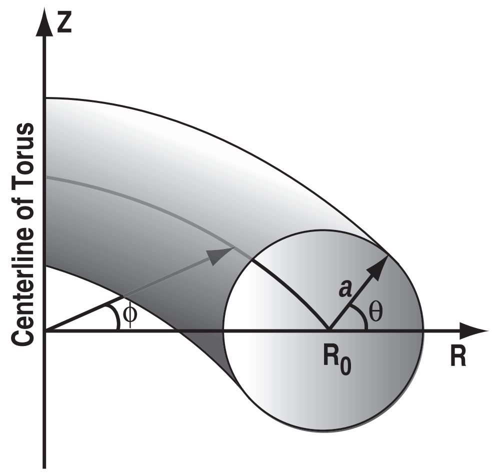

Figure 12. Cartoon of a plasma having a circular cross-section showing commonly used notation. Commonly used cylindrical coordinates are the radial and vertical positions R and Z and the *toroidal angle* ϕ that measures the angle around the centerline of the torus. The distance to the center of the plasma R_0 is called the *major radius*. The radius a of the plasma cross-section is called the *minor radius*, a terminology used even with noncircular plasmas. The angle θ is usually called the *poloidal angle*. The ratio R_0/a is known as the *aspect ratio*. The *midplane* of the torus refers to the plane $z=0$.

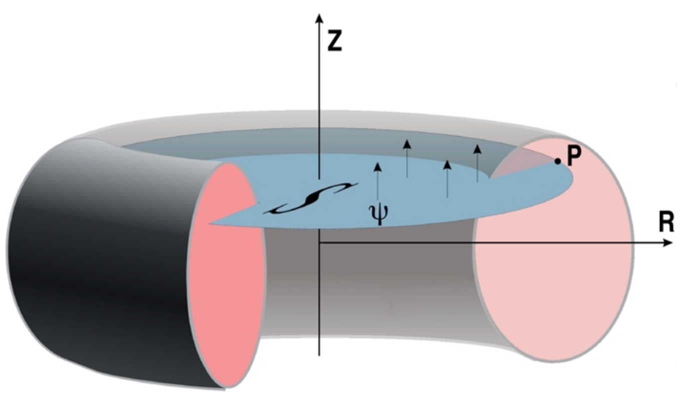

Figure 13. Definition of the *poloidal flux function*. Poloidal flux ψ_p at a point P in the (R,z) cross-section of the plasma is defined to be the total flux through the surface S bounded by the toroidal ring passing through P.

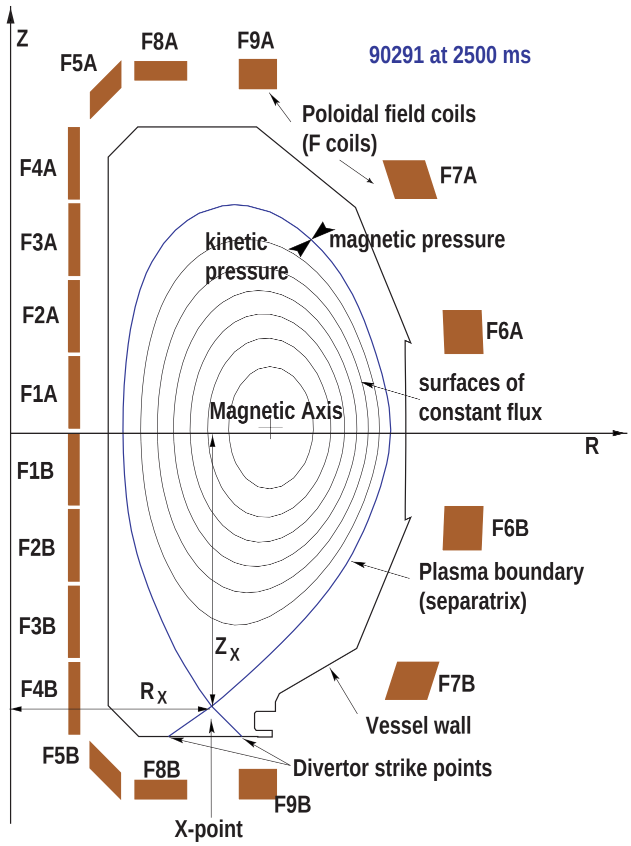

Figure 14. Cross section of the DIII–D tokamak showing a flux contour representation of the plasma. Note the nesting of contours of constant flux as assumed by ideal MHD theory (see sidebar MHD). The value of flux is largest at the center of these nested flux contours at the *magnetic axis*. The closed flux surface furthest from the magnetic axis (also called *last closed flux surface* or *separatrix*) defines the edge of the plasma. In a diverted plasma, the separatrix forms an *X-point* — in this case at the bottom of the plasma. The points at which the separatrix strikes the vessel wall are called *strike points*. The location of the plasma seperatrix and the X-point (or strike points) are controlled by the poloidal field generated by currents in the poloidal field coils (whose cross section is labeled F1A through F9B here).

The famous physicist Richard Feynman once observed that holding a plasma in place using magnetic fields is like trying to hold jello in place using rubber bands. The magnetic fields produce an external *magnetic pressure* that balances the internal pressure (usually called *kinetic pressure*) created by the hot gas. (see sidebar S2). Slight irregularities in the magnetic field confining the gas can allow bulges in the plasma to form that can grow exponentially over time if not actively suppressed. There are a large number of such plasma instabilities, only a few of which will be discussed in this issue. Most, but not all, of these instabilities are predicted rather well by ideal MHD theory (see sidebar S3).

S2. Performance Metrics (sidebar)

There are several measures of performance that are used in order to gauge progress towards the goal of developing a power-producing fusion reactor. One key performance parameter is the quantity β, which has several variations but generically is the ratio between the internal kinetic pressure of the plasma and the pressure of the magnetic field that confines the plasma. This is one measure of efficiency of confinement, since it defines how much confining pressure is needed to maintain a given plasma pressure. Attaining high plasma pressure is necessary to reach ignition but for economic and engineering reasons (coil system costs and magnetic forces on conductors) the magnetic pressure must be kept low. The basic definition is $\beta = \langle p \rangle / [B_T^2 / (2\mu_0)]$, where $\langle p \rangle$ is the plasma kinetic pressure averaged over the plasma volume, B_T is the vacuum toroidal field strength, and μ_0 is the magnetic permeability of a vacuum. This ratio is thought of as representing a ratio of internal kinetic (gas) pressure to external magnetic pressure (see Fig. 14) and has typical values of 5% to 10% in modern tokamaks. This definition for beta is also often referred to as *toroidal beta* (β_t) because it is defined in terms of toroidal field (B_T). A definition that will be referred to often in this issue is that of normalized beta

$$\beta_N = \beta / (I / a B_T) \quad ,$$

where a is the plasma minor radius (see Fig. 12), and I is plasma current.

Another figure of merit is the energy confinement time. Plasmas leak energy all the time. The energy confinement time is a measure of the characteristic timescale of this leakage. Plasma physicists usually denote the total thermal (or kinetic) energy in the plasma with the letter W. The rate at which the plasma loses energy can be denoted as P_L. This quantity is a power because it represents a change in energy per unit time. This quantity can be determined experimentally since, in steady state, it must be equal to the power P_H that is injected by external heating systems in order to maintain the plasma at a fixed temperature. The confinement time τ_E is defined by

$$P_L = \frac{W}{\tau_E} \quad .$$

Energy confinement time is typically used to evaluate the applicability of various plasma confinement and control techniques to eventual steady state reactor operation.

Another explicitly economic figure of merit is Q, the ratio of output power to auxiliary input heating and current drive power. Clearly Q must be significantly larger than one in a power producing fusion reactor. The planned ITER tokamak has a goal of Q = 10.

S3. MHD (sidebar)

Magnetohydrodynamics (*MHD*) [12] is the branch of plasma physics that describes the basic behaviour of the plasma in terms of the interaction between currents, magnetic fields, and forces exerted on and by the plasma. The theory treats the plasma as a single fluid, with no distinction made for the various particles comprising this fluid. The branch of MHD that is most straightforward to understand and apply is called *ideal MHD* [13]. Ideal MHD follows from the assumption (only approximately satisfied) that the plasma has no resistance. Plasmas are actually (slightly) resistive of course and this discrepancy is the primary source of several behaviors that are not satisfactorily predicted by ideal MHD. Ideal MHD is sufficiently accurate however, that it is used as a first approximation in almost every magnetic analysis done for tokamak plasma physics, including studies of plasma magnetic instabilities, definition of plasma magnetic evolution equations, and in estimation of the plasma shape and position for control.

The MHD theory applies to both *axisymmetric* and *non-axisymmetric* plasma behavior. Axisymmetric behavior has no dependence on the toroidal angle coordinate (Fig. 12). That is, it is symmetric with respect to rotation about the central axis of the tokamak. Non-axisymmetric behavior can vary with the toroidal angle, although there is typically some structure enforced by the helical nature of the magnetic fields (Fig. 10). One consequence of ideal MHD is that contours of constant axisymmetric poloidal flux $\psi_p(R,z)$ (see sidebar S1) must be nested as shown in Fig. 14. The plasma pressure $p(R,z)$ is also constant along these contours.

S4. Mode Numbers and Safety Factor (sidebar)

It is common for plasma physicists to think in terms of a spatial Fourier series decomposition of the plasma because the tokamak is naturally periodic in both toroidal and poloidal angles (Fig. 12). The basis functions are commonly referred to as *modes* of the plasma. One dimension of the decomposition is defined by the (periodic) variation in plasma shape with toroidal angle ϕ. The integer n, called the *toroidal mode number*, is consistently used as the index for modes in the toroidal decomposition and represents a basis function with period $2\pi/n$ in the variable ϕ. That is, the function will repeat itself n times during one revolution around the torus. The second dimension of the decomposition is defined by the periodic variation in plasma shape with poloidal angle θ with vertex at the center of the plasma cross section (Fig. 12). The integer m, called the *poloidal mode number,* is used as the index for modes in the poloidal decomposition and represents a basis function with period $2\pi/m$ in the variable θ. Products of toroidal and poloidal basis functions naturally produce helically shaped structures. Thus, for example, the three-dimensional magnetic field distribution shown in Fig. 10 is very naturally represented this way. This leads to the following definition [5] of the *safety factor* q as a measure of the pitch of a magnetic field line. If, at some toroidal angle ϕ, a field line has a certain position in the poloidal [(R,z) see Fig. 12] plane, it will return to that position in the poloidal plane after a change of toroidal angle $\Delta\phi$. The q-value of this field line is defined by

$$q = \frac{\Delta\phi}{2\pi}$$

Thus, if a magnetic field line returns to its starting position after exactly one rotation around the torus, then q = 1. If it takes longer to return, it has a higher value of q. Rational values of q play an important role in stability, hence the name safety factor. If $q = m/n$, where m and n are integers, the field line joins up on itself after m toroidal and n poloidal rotations around the torus. Relating to the definitions above, one can see that m is the poloidal mode number and n the toroidal mode number of the field line.

Plasmas as Fluids

Plasmas also exhibit behavior characteristic of gases; they have a thermodynamic description $P \propto nT$ in which the gas pressure (P) is proportional to the product of density (n) and temperature (T), where the constant of proportionality depends on the units chosen to represent each quantity. These quantities are not homogeneous within the plasma, however, as a consequence of the forces exerted by the magnetic fields on the ionized gas particles. Intuitively, it is easy to understand that a tokamak plasma is hotter and denser in its core and, therefore, has higher pressure nearer to the center. The actual spatial distribution of these quantities is important to the plasma's interaction with the confining magnetic fields since the vast majority of particles that define the plasma properties carry a charge.

S5. Flux Coordinates (sidebar)

The notion of flux at a point in the poloidal plane was defined in the sidebar S1. Under the assumptions of ideal MHD, the contours of constant flux are also contours of constant pressure and safety factor (q) (see sidebar S4), among other quantities. Thus it is convenient (and well defined) to consider the values of these quantities at specific flux values, as the flux varies from the center to the edge of the plasma (Fig. 14). This is formalized by defining a *normalized flux coordinate*

$$\rho = \frac{\psi - \psi_b}{\psi_0 - \psi_b} \qquad (*)$$

where ψ_0 denotes the maximum value of flux at the center and ψ_b denotes the value of flux at the boundary. Through the use of this coordinate, for example, the pressure can be represented as $P = P(\rho)$, where $0 \leq \rho \leq 1$. It can be mapped back to pressure as a function of R and z by substituting (*) for ρ and using the (previously calculated) function $\psi = \psi(R,z)$. The terminology commonly used for the representation of pressure, safety factor, etc., as a function of normalized flux is the *pressure profile*, q*-profile*, and so on.

An alternative variable that serves the same purpose as normalized flux ρ is the normalized radius $x = r/a$, where r is the distance from the p plasma centroid to the flux surface along a horizontal (radial) line and a is the minor radius of the plasma (Fig. 12). As with the ρ coordinate, $x = 0$ at the plasma centroid and $x = 1$ at the plasma boundary.

As a matter of notation, when referring to functions of normalized flux ρ (or normalized radius x), the symbol ∇ is often used to mean derivative with respect to ρ (or x).

Plasma pressure is usually expressed in either atmospheres (atm) or Pascal (Pa, 1 atm = 100 kPa = 100000 Pa), with values typically representing pressures of one atmosphere or less. Temperatures are usually referred to in units of electron volts [(eV), e.g., 1 eV = 1.16×10^4 degrees Kelvin] with typical values on the order of a few thousand electron volts (keV). Density is in units of particles per cubic meter (m^{-3}) or cubic centimeter (cm^{-3}) with typical values on the order of 10^{20} m^{-3}.

Plasma Heating and Current Drive

There are several methods that have been used for heating tokamak plasmas. The most commonly used method in present devices and probably the simplest to understand for those with an engineering background is the *ohmic heating*. Here again, we can think of the plasma as a conductor in which current can be induced via transformer action. Some or all of the poloidal coils on the device act as the primary of the transformer and the plasma acts as the secondary. The plasma has a finite resistance so some of the induced current dissipates into heat.

Another commonly used method is injection of a beam of neutral particles, commonly referred to as *neutral beam heating*. Typically these particles are isotopes of hydrogen, which have to be neutral in order to cross the magnetic field and penetrate into the plasma, where they are ionized. The injected particles collide with particles already in the plasma, imparting some of the energy from their motion to the particles within the plasma, thereby increasing their kinetic energy and therefore the bulk temperature of the plasma.

A method that is becoming more common on many devices is the use of various *radio-frequency (rf) heating* schemes. These work in essentially the same way that microwave heating works. Radio waves with frequencies selected to be resonant with the motion of selected waves and particles in the plasma are launched into the plasma using various configurations of waveguides and antennas. Resonant particles excited by the injected rf energy transfer some of their energy to non-resonant particles through collisions. The overall kinetic energy contained in the plasma, and therefore the temperature, increases.

Multiple terms are used in the fusion community to describe the various forms of heating. *Ohmic heating* is also called *inductive heating*. All other methods described above may be referred to as *non-ohmic*, *non-inductive*, or *auxiliary heating*. Since each of these methods is also able to generate additional plasma current (under certain conditions), when used for that purpose they are referred to by replacing the term *heating* with *current drive*.

Under certain conditions, motion of particles within the plasma can create an additional self-generated current commonly referred to as *bootstrap current*. The mechanism that produces this current will be described in more detail in a later paper in this issue.

Plasma Diagnostics

How exactly do you measure how the various plasma quantities are distributed? In particular, where is the plasma? Fusion plasmas are so hot that they emit very little visible light, except near the relatively colder boundary. Figure 15 shows a photo of the the interior of the DIII–D tokamak, on the left with plasma, and on the right without plasma. We cannot in fact see most of the plasma itself, only the effect of the heat where the plasma contacts the wall and the relatively cold edge plasma and neutral gas. This illustrates the problem of establishing the exact location of the plasma. Measurement with devices requiring physical contact cannot be used for any length of time because they would eventually be vaporized by the plasma. Visible light sensors cannot be used because most of the plasma is not visible. Infrared sensors are commonly used for some purposes, but they cannot provide measurements with resolution sufficient for most control purposes. It turns out that the most convenient way of thinking about this is in terms of the magnetic fields in which the plasma is confined. Under the assumptions of ideal MHD (see sidebar S3), it can be shown that the contours of constant magnetic flux (equivalently, the magnetic field lines) are also contours of constant pressure (isobars). These magnetic flux contours are represented inside of the torus in the cross-section shown in Fig. 6 and in more detail in Fig. 14.

Thus the most common approach to controlling plasmas involves control of these flux contours. The majority of efforts to date have focused primarily on controlling the outermost closed flux surface, also called the separatrix (Fig. 14), although more recent control work has begun to focus on controlling parameters in the interior of the plasma. The flux contours cannot be directly observed, so indirect methods of determining them are required. These methods are based on reconstructing the flux patterns in the interior of the vacuum vessel based on magnetic measurements made outside or just inside the vessel. Magnetic sensors used for this purpose that are common to virtually all tokamaks include *flux loops* for measuring flux through large areas, *magnetic probes* (*B probes*) for measuring local magnetic fields, and *Rogowski Loops* or *Hall Current sensors* for measuring current flowing through conductors (including the plasma) [14].

Some of the methods used for reconstructing the plasma flux contours will be discussed in more detail in one of the articles in this issue.

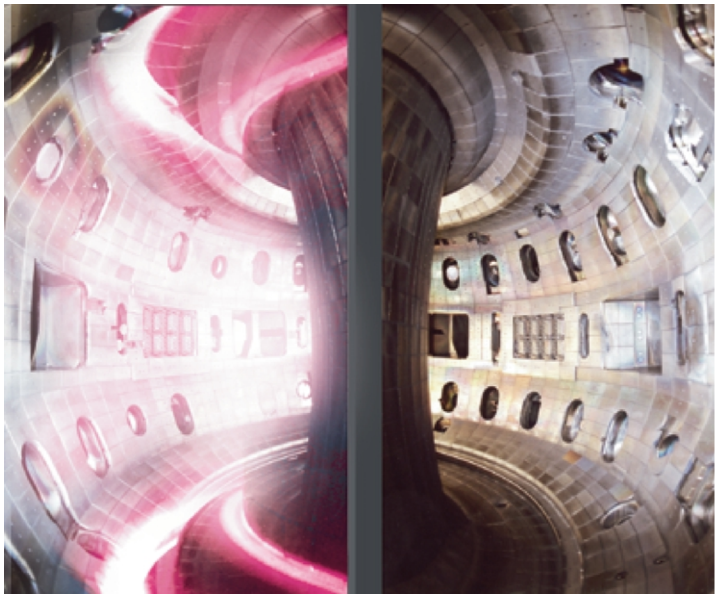

Figure 15. Combination of two photographs of the interior of DIII–D. Photo on the right illustrates the interior without plasma. Photo on the left shows the interior with plasma, where only the plasma edge is visible. The large number of holes in the side of the vessel are filled with various diagnostics or plasma actuators (recessed to prevent direct interaction with the high energy plasma) or are used for access to the interior. Carbon tiles are used in DIII–D for all interior locations that can come in contact with the plasma.

One difficulty with such reconstruction methods is that internal flux contours are not uniquely determined by magnetic reconstruction since it is a theoretically ill-posed problem if based on external measurements alone. Therefore, various types of diagnostics of internal parameters are used to augment the magnetics sensors, e.g., by examining how laser light is reflected as it passes through the plasma, or by examining emissions stimulated by beams of neutral atoms passing through the plasma. These methods are also used to diagnose various other plasma parameters of interest, such as the distributions of temperature, pressure, density, magnetic field, and current within the plasma (see sidebar S5).

CONTROL PROBLEMS IN TOKAMAKS

What are the control problems in tokamaks? For present devices, the basic functions of plasma initiation, shaping, heating, current drive, stabilization, and safe termination of discharges are needed. Research efforts in tokamak control and several experimental devices now emphasize the requirements for the so-called *Advanced Tokamak* (AT) because the AT is presently the most promising tokamak concept path towards development of an economically attractive power producing reactor. The Advanced Tokamak plasma regime is characterized by simultaneously high plasma pressure (β), long energy confinement time (τ_E), and non-inductively driven plasma current with a significant fraction provided by the self-generated bootstrap current, all in a steady-state configuration. Significant reductions in the size and cost of a fusion power plant can be realized if these can be simultaneously achieved.

The AT approach relies on the ability to solve a number of challenging control problems. Shape control provides the strong plasma shaping required to achieve high beta values. Operation at high β causes a number of plasma instabilities that must be actively stabilized; optimization of the shape reduces somewhat the virulence of these instabilities. Energy confinement, stability properties, and the fraction of plasma current provided by the bootstrap mechanism can all be improved through control of internal pressure and current profiles. In addition, effective power exhaust and impurity and particle control are required. All must be controlled in steady state operation. The AT regime requires flexible multi-parameter plasma control because of the strong coupling between the various control parameters and, in certain cases, the required sharing of control actuators for different control objectives. In a high performance AT, accurate regulation of the plasma boundary, internal profiles, pumping, fueling, and heating must be well coordinated with MHD control action to stabilize magnetic instabilities. Sophisticated monitors of the operational regime must provide detection of off-normal conditions and trigger appropriate safety responses with acceptable levels of reliability. Because AT plasmas operate close to stability limits in order to yield high efficiencies, off-normal events can conceivably produce large excursions from the operating point and result in uncontrollable instability.

A central characteristic of AT plasma regimes is the extreme shapes that must be accessed and the high degree of accuracy with which they must be regulated. The most mature of the many tokamak control problems is that of shape and position control. The plasma is magnetically confined through magnetic fields generated by the set of conducting coils distributed around the vessel that contains the plasma. Voltages are applied to these coils, which drive currents that produce the magnetic fields. This field interacts with the plasma to change its shape and position and to induce plasma current. Use of feedback control is made mandatory by the uncertainty of the mathematical model of the plasma dynamics that is described by a set of non-linear partial differential equations, the vertical position instability typical of the most common plasmas, and the occurrence of unpredictable disturbances. This is a challenging control problem due to the high number of input coil voltages and controlled outputs (typically the plasma current plus several position and shape geometrical parameters), the strong coupling between the various input-output channels, and the demanding control requirements.

In recent years, a large amount of work has been devoted to obtaining better plasma control performance using control methodologies that range from primitive to advanced. This research has been predominantly of methodological type, although some of the emerging new concepts have been successfully tested experimentally.

Presently in all large devices (such as the JET and DIII–D tokamaks) there are active programs for development and implementation of controllers for a significant number of tokamak systems. The most mature of these programs are aimed at poloidal shape control, but there are also a number of efforts intended to develop control for other plasma parameters. A number of analytical and numerical control studies have also been performed specifically for the proposed ITER device, in order to verify that the final design would support its experimental objectives.

CONTENTS OF THIS SPECIAL ISSUE

In this special issue, a sampling of the many problems of tokamak control are provided. The aim of this special issue is, on one hand, to report the results obtained to date in tokamak plasma control [15], especially in the area of plasma shape control [16], and on the other hand to outline some open problems that appear to be promising lines of research for the control community. Control problems are often considered as beginning after a system to be controlled has already been constructed, perhaps even used for some time. We take a broader view of control, in the sense that we are in a position to influence the manner of construction of the plant itself, as well as the actuators and the sensors used in controlling the plant. This is true in particular of the ITER design process, in which the development and evaluation of controllers was done simultaneously with the design of coil and power systems in order to ensure that the system as designed would have adequate control authority. Most often this is done by overdesigning the system so that there is plenty of margin for control. The coordinated system and controller design for ITER allowed a much smaller actuator margin and substantial cost savings in this planned multi-billion dollar facility.

Most tokamak control problems begin to be formulated as soon as an idea for a new actuator or diagnostic sensor is proposed. A great deal of the control work early on is the development of sufficient understanding of the physics of the system behavior to be controlled and development of the actuator that is intended to influence that behavior. Open-loop experiments are used to evaluate actuator effectiveness and further develop the physics models. Ad-hoc closed loop (typically manually tuned PID) controllers are used to gain early confidence that the control is feasible. Once sufficient understanding of the physics mechanisms has been obtained, derivation and validation of models suitable for control design follow. It is around this time that persons more skilled in control theory and application than in plasma physics can begin to make substantial contributions to the solutions of tokamak control problems.

The majority of papers in this issue will emphasize the problems of control of axisymmetric shape and stability and control of plasma current. These are the most mature and well understood of the many tokamak control problems. The first paper by Pironti and Ariola describes the design methodology and successful implementation of a new multivariable *extreme shape controller* for the JET tokamak. The paper by Beghi and Cenedese discusses the problem of diagnosing for control purposes the location of the plasma within the containment vessel and describes a recently developed technique originating from algorithms used for computer vision. The final paper on shape control by Ambrosino and Albanese is a survey that discusses in more detail the objectives of the shape control problem, describes the techniques used to generate models useful for control design, and summarizes work by researchers at multiple institutions on this problem.

In addition to the discussions of shape control, the paper by Walker et al. describes progress that has been made on other tokamak control problems, including stabilization of various types of instabilities, control of internal distributions of plasma current and pressure, and methods for dealing with events that do not normally occur but must be handled if and when they do occur. The following paper by Sartori, Piccolo and De Tommasi describes the JET control system more generally and from a practical implementation point of view, in order to provide the reader with some insight into the real-world methods and constraints of control on a modern tokamak. The final paper by Lister and Portone will discuss control solutions required for the proposed ITER device, which will apply much of what has already been learned about plasma control from existing devices but will also pose a number of new control challenges.

With the exception of the paper by Ambrosino and Albanese, the articles in this issue are expressly *not* comprehensive surveys of the topics being presented. As a consequence, some important contributions will almost certainly not be discussed.

REFERENCES

[1] T. Appenzaler, “The end of cheap oil,” National Geographic, June 2004, pp. 80-109.

[2] http://www.worldenergy.org/wec-geis/edc/scenario.asp

[3] P.B. Weisz, “Basic Choices and Constraints on Long-Term Energy Supplies,” Physics Today, Vol. 57, July 2004, p. 47.

[4] F.E. Pinkerton, B.G. Wicke, “Bottling the hydrogen genie,” The Industrial Physicist, Feb/Mar 2004, p. 20-23.

[5] J. Wesson, Tokamaks, 2nd Ed. Clarendon Press, Oxford, 1997.

[6] http://www.jet.efda.org/

[7] http://web.gat.com/

[8] http://www.fusionscience.org/

[9] http://www.fusion.org.uk/culham/fasttrack.pdf

[10] http://www-fusion-magnetique.cea.fr/gb/index.html

[11] J.D. Jackson, Classical Electrodynamics, Wiley, 1998

[12] K. Miyamoto, Plasma Physics for Nuclear Fusion, MIT Press, Cambridge, MA 1976.

[13] J.P. Freidberg, Ideal Magneto-hydro-dynamics, Plenum Press, New York, 1987.

[14] I.H. Hutchinson, Principles of Plasma Diagnostics, 2nd. Ed., Cambridge, 2002.

[15] J.B. Lister, et al., “The Control of Modern Tokamaks,” Int. Conf. on Accelerator and Large Experimental Physics Control Systems, Trieste, Italy, 1999, p.235

[16] R. Albanese, “Current, position, and shape control of tokamak plasmas: a literature review,” Proc. 2000 IEEE Int. Conf. on Control Applications, 2000, p. 385–394.

APPENDIX
ADDITIONAL SIDEBARS

These sidebars will be inserted into later articles in the same issue of Control Systems Magazine.

S6. The Plasma Shape Control Problem (sidebar)

The control of plasma poloidal shape and position is the most mature of all tokamak plasma control problems. The actuators for this control are the *poloidal field coils* (Fig. 14). These coils produce magnetic fields that shape and move the plasma according to the intuitive description of Fig. 11. More precisely, the force on the plasma is given by the cross-product of plasma currents I and magnetic field B ($I \times B$ *forces*). Although this relationship holds in all three dimensions, shape control considers only those forces that modify the axisymmetric shape of the plasma as represented by the poloidal (R,z) cross-section (Fig. 14), orthogonal to the toroidal (ϕ) coordinate. These forces are dominated by the toroidal plasma current crossed with poloidal magnetic field. Using the standard right hand rule, it is easy to see that a magnetic field vector oriented vertically will produce only radial forces on the plasma. A magnetic field oriented radially will produce only vertical forces on the plasma.

There are three primary objectives of plasma shape and position control. The first is to maintain the plasma within the containment vessel. Plasmas are subject to forces that tend to cause them to expand outward. A vertical field must be generated by the coils to provide a balancing inward force. A second objective is to shape the plasma to satisfy certain operational objectives. For example, it has been found that vertically elongated (tall and thin) plasmas tend to allow significantly higher plasma pressures to be maintained, so this is now routinely done in many devices. The process of elongation creates a magnetic field configuration, however, which causes the plasma to be unstable in the vertical (z) direction. Thus the third objective is to stabilize this so-called *vertical instability*.

The problem of regulating total plasma current (I_p) to a reference value is often considered simultaneously with shape control, because these two types of control share the poloidal field coils as actuators. In this case, it is not magnetic field that is desired but a changing magnetic flux that drives current through a transformer action called *ohmic current drive*. Continuously changing current in one or more poloidal field coils produces a changing poloidal flux Ψ at the plasma. The time derivative of this flux defines an induced voltage $V = -d\Psi/dt$ known as the plasma *loop voltage*, which drives current in the plasma according to $LdI/dt + RI = V$, just as in a standard transformer. Here the values L and R represent the plasma bulk equivalent self-inductance and resistance. Resistive dissipation in the plasma is responsible for the heating effect, hence the origin of the term *ohmic*. Shape control and plasma current control can be seen to be compatible by noting that generating flux that is spatially uniform across the plasma can be used to drive plasma current, but will not affect the plasma shape since it produces no field. (Magnetic field is the spatial derivative of flux — see sidebar S1.) Such a flux pattern is known as *ohmic flux*.

The problems of plasma stabilization and plasma shaping both rely on applied magnetic fields to produce the necessary restoring forces. Usually, an estimate of the position of the plasma *current centroid* (think center of mass of toroidal plasma current) is used as the control quantity for the vertical stabilization control. The response time needed to control this instability is typically substantially shorter than the response time needed for shape control and the feedback often requires at least some derivative gain. As a result, most tokamaks implement a form of *frequency sharing*, in which high frequency responses are dedicated almost entirely to the stabilization process, while low frequency responses are primarily dedicated to shape control.

One factor that makes the development of shape controllers more challenging is the fact that the response of the combined plasma and tokamak system can change significantly, depending on some of the details of the interior of the plasma. The plasma parameters *poloidal beta* β_p ("beta p") and *internal inductance* ℓ_i are known to be two of the strongest predictors of the change in this response. These parameters are measures of the plasma internal distributions of pressure and current, respectively (see sidebar S5).

S7. Magnetic Diagnostics (sidebar)

Typical magnetic sensors produce a voltage that is proportional to the time derivative of some desired physical measurement m, $V_0 = k(dm/dt)$. For example, the measurement m might be a current in Amps or a flux in Webers (Fig. A–1). This voltage is integrated by a hardware integrator circuit to obtain $V_1 = \int_0^t V_0\, dt = km$. This is usually followed by one or more processes (amplifiers, voltage dividers, digitizers) whose net effect is to multiply by a gain, $V_2 = GV_1 = Gkm$. A multiplication by $1/Gk$ provides the conversion from the integrated acquired volt-second signal to physical units. The multiplier k depends on the type of device. For flux loops, $k = 1$ can be used to represent flux directly in volt-seconds or Webers (Wb). Some devices use $k = 2\pi$ to convert from total flux in Wb to flux in Wb per unit radian. For magnetic probes $k = NA$, where N = number of turns and A = cross-sectional area. For Rogowski loops, $k = N\mu_0 nA$ where n = number of turns per unit length, A = cross sectional area of a single coil turn, and N = the total number of coil turns around the current path being measured.

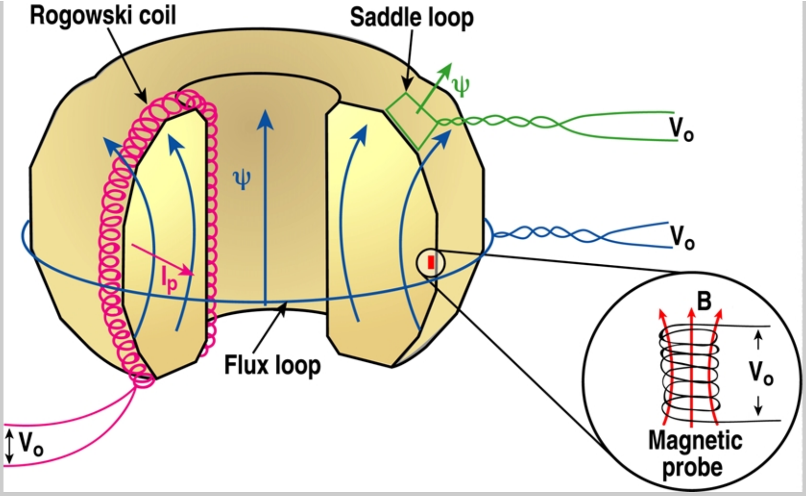

Figure A–1. Examples of tokamak magnetic sensors. All the sensors illustrated work on the same basic principle. Changing flux induces a voltage in a coil of wire. This voltage is integrated to obtain a representation for the flux through the coil. Flux loops consist of a single loop of wire connected to a voltage sensor. The integrated voltage represents the total (poloidal) flux ψ through the loop. Magnetic probes consist of multiple windings of wire in a rather small (a few centimeters) radius, connected to a voltage sensor. The integrated voltage represents local (poloidal) magnetic field (using the relationship $\psi = \int B\,dA$). This is actually an approximation, since the flux is not perfectly constant across the open end of the probe. Saddle loops also measure total magnetic flux, but are sometimes also considered as a form of magnetic probe by assuming nearly constant flux across the loop. Rogowski coils measure current, in this case toroidal plasma current (I_p) plus currents in the vessel, by computing a line integral of magnetic field in a closed loop around the current path.

S9. Divertors (sidebar)

The helium byproduct of the fusion reaction (sometimes called *helium ash*) must be removed from the vacuum chamber to prevent it from interfering with subsequent fusion reactions. The method selected for accomplishing this is to use one or more *divertor regions* (or *divertors* for short) with pumping. Figure A–2 shows the divertor in the JET tokamak. The idea of the divertor is to divert non-hydrogen particles away from the plasma to a region that is designed to safely absorb their heat and remove them from the plasma chamber. All particles, including the helium ash, tend to remain in the plasma for only a finite time known as the *confinement time* (see sidebar S2) before leaking out. When the helium nuclei leave the plasma, they are still charged and therefore tend to follow magnetic field lines. Other *impurities*, as introduced by the plasma interacting with device components, are also ionized and follow the field lines. The region just outside the separatrix (last closed flux surface) where this occurs is known as the *scrape-off layer* (SOL). Field lines outside of the separatrix do not close inside the vacuum vessel. Instead they terminate at the wall of the vacuum vessel at locations known as the strikepoints (see Fig. A–2). As a consequence, the impurity and helium ions follow these first few external field lines until they contact the wall in the divertor region, and are then pumped out of the tokamak chamber. The particles contacting the plasma facing materials at the strikepoints are still very high energy, so divertors are specifically designed to withstand this continuous bombardment of high energy particles. Divertors are always constructed of heat resistant materials, and often incorporate geometric designs intended to widen the area that is impacted by the particles.

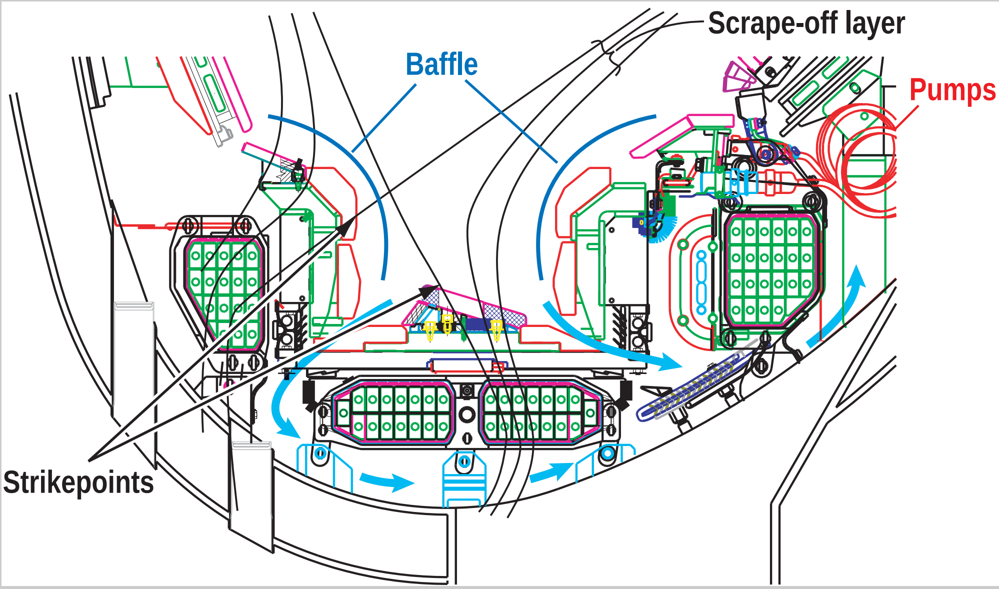

Figure A–2. The *divertor region* in the JET tokamak. A set of flux contours is shown to represent the plasma. The divertor *baffles* serve to direct particles flowing in the *scrape-off layer* into the *throat* of the divertor, where they can be pumped out (light blue arrows). Cross-sections of the four *divertor coils* for JET can be seen just outside the divertor. Each small green rectangle inside a coil represents a single turn of the multi-turn winding used to multiply the flux produced by the coil. The small circles inside these rectangles are the channels used for water cooling.

S16. Circuit Models (sidebar)

Several of the plant models used in development of tokamak magnetic controls are provided by lumped element circuit models. These models are traditionally first learned in basic circuit theory to model such lumped elements as resistors, inductors, and capacitors. This type of model is natural for tokamak magnetics control, since various conducting elements — the plasma, control coils, and conducting structure — are readily represented as discrete circuits with calculable values of resistance R and self-inductance L. In isolation, each circuit can be modeled by an equation $L\,dI/dt + RI = V$, where V is an imposed voltage and I is the resultant circuit current. The electrical influence that one discrete circuit (a) has on another (b) is given by the mutual inductance M_{ab}. For example, to model two discrete circuits,

$$\begin{aligned} L_a \dot{I}_a + R_a I_a + M_{ab} \dot{I}_b &= V_a \\ L_b \dot{I}_b + R_b I_b + M_{ba} \dot{I}_a &= V_b \end{aligned} \quad . \tag{*}$$

By definition, the values M_{ab} and M_{ba} must be equal and, as a matter of notation, self inductance can be represented as $L_a = M_{aa}$, for example. Thus, (*) can be written in matrix notation as

$$M \dot{I} + RI = V \quad ,$$

where

$$M = \begin{bmatrix} M_{aa} & M_{ab} \\ M_{ba} & M_{bb} \end{bmatrix}, \quad R = \begin{bmatrix} R_a & 0 \\ 0 & R_b \end{bmatrix}, \quad V = \begin{bmatrix} V_a \\ V_b \end{bmatrix} \quad .$$

It's easy to see how this extends to larger numbers of discrete circuits.

A distinction is generally made between passive and active circuits. In an active circuit, such as a control coil, the voltage $V \neq 0$ is provided by an external voltage source, such as a power supply. In a passive circuit, the voltage source term $V = 0$. In this case, however, one can consider the voltage defined by the mutual coupling as the source for the passive equation. For example, if $V_a \neq 0$ and $V_b = 0$ in (*), then $-M_{ba} \dot{I}_a$ acts as the voltage source in the passive (b) circuit.

$$L_b \dot{I}_b + R_b I_b = -M_{ba} \dot{I}_a$$

Along with the electrical interactions described by the circuit equations, the various current-carrying elements of a tokamak can interact mechanically through the magnetic fields produced by these currents. The force equations describing these interactions (particularly with the plasma) are usually combined with the circuit equations to form the model plant for control development.

S17. Magnetic Fields, Currents, and Forces (sidebar)

Confinement and shaping of tokamak plasmas is made possible because of well-known relationships between magnetic fields, currents in conductors, and the forces that result from the interaction of the two. Figure A–3 illustrates the so-called $J \times B$ forces produced by the cross product of a current density J and magnetic field B. Often, the symbol I is used to represent one or more lumped currents — in one or more wires, for example — while J is used to represent a continuous distribution of current or current density (current per unit area).

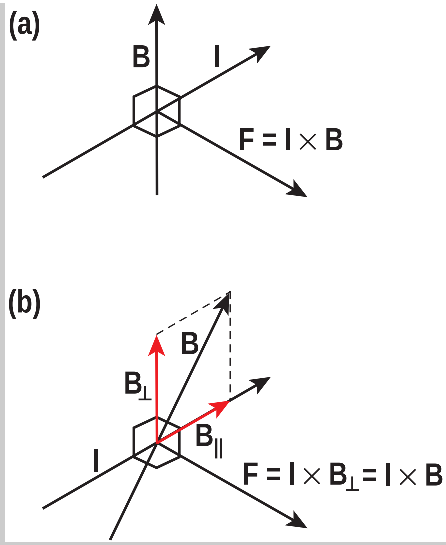

Figure A–3. Illustration of magnetic force on a current carrying conductor. The force per unit length exerted on a conductor carrying a current I by a magnetic field B is given by the cross-product of I and B. The vectors I and B are shown orthogonal in (a). The more general case is shown in (b). Only that portion of B ($B_\perp$) orthogonal to the current I produces a force on that conductor.

Since a tokamak plasma carries current, it can play the role of the conductor upon which magnetic fields can apply forces. For axisymmetric plasma shape control, the actuators are poloidal field coils (see Fig. 5 of the Introduction) that carry only toroidal currents (in the direction of toroidal coordinate ϕ — see Fig. 13 of intro) and, therefore, generate only (axisymmetric) poloidal fields. The forces on the plasma resulting from these poloidal fields are dominated by their interaction with the plasma toroidal current density distribution J (Fig. A–4). The resulting $J \times B$ forces are therefore directed poloidally. Thus the shape control problem may be considered by restricting to a poloidal (R, z) cross section (Fig. A–4) of the device and plasma.

The vertical instability that must be stabilized when controlling the axisymmetric shape (see sidebar S6) originates in the generation of a magnetic field used to improve the performance of the plasma. By vertically elongating the plasma, several performance parameters, including the toroidal β parameter (see sidebar S2), can be increased. A circular plasma (Fig. A–4) is vertically stable. To make a vertically elongated plasma however, a radial magnetic field must be applied to pull upward on the top of the plasma and downward on the bottom of the plasma. The resulting magnetic field distribution is the source of the vertical instability. Under the influence of this field, upward (downward) displacements of the plasma result in a destabilizing upward (downward) force on the plasma (Fig. A–5).

Currents can not only create magnetic fields and flux (see sidebar S1), but can also be created by them. One simple example of this already seen (see sidebar S6) is the transformer action used to generate plasma current, in which changing flux, created by a subset of the set of poloidal field coils, creates a loop voltage that drives current in the plasma.

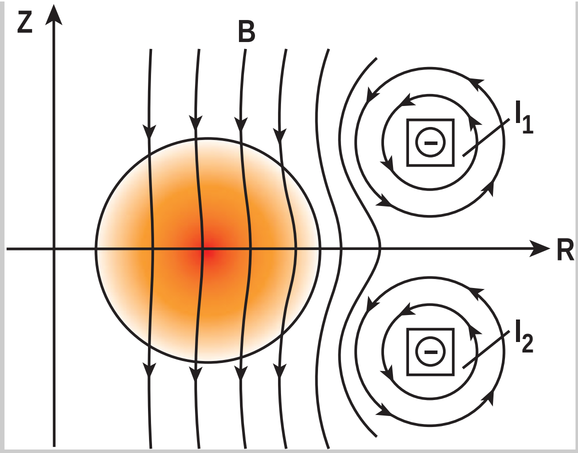

Figure A–4. Illustration of radial force balance in plasmas. Plasmas exert a self-force known as the *hoop force* that is directed radially outward. The sign of the current in the outer coils is negative (coming out of the page), in order to produce magnetic field at the plasma which will oppose this outward force and maintain the plasma position. The direction of magnetic field produced by a control coil is determined by the so-called right hand rule: If the thumb of the right hand indicates the direction of current then the curled fingers of the right hand indicate the direction of the magnetic field produced by this current. Intensity of shading within the plasma is used to illustrate the continuous toroidal current density distribution j_ϕ (Amps per square meter, for example) within the plasma. This current is positively directed (into the page), so the $J \times B$ force applied by the control coils is directed radially inward.

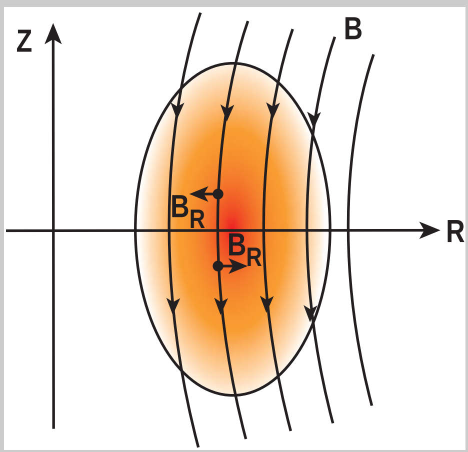

Figure A–5. Illustration of the source of vertical instability. Magnetic field with radial component directed inward above the midplane and directed outward below the midplane is generated by the shaping coils to cause the plasma to become *elongated*. The $J \times B$ force for inward (outward) directed magnetic field and positive toroidal current (into the page) is up (down). As long as the current distribution and magnetic field are completely symmetric about the midplane ($z = 0$), the upward and downward forces balance and the plasma is in equilibrium. If a disturbance shifts the plasma up slightly, more current will be above the midplane than below and the net force is directed upward. This causes the plasma to move up, thereby further increasing the upward force, and so on. A similar situation occurs if the disturbance causes an initial downward displacement of the plasma. This is the essence of the vertical instability.

Changing flux can also be produced by motion of the plasma. For example, when the plasma moves vertically, the magnetic flux due to the plasma current (see sidebar S1) moves with the plasma. Therefore the poloidal flux value Ψ at a fixed toroidal conductor (a PF coil or a section of the conductive vacuum vessel, for example) changes as this flux distribution moves by it. The resultant loop voltage $V = -d\Psi/dt$ (see sidebar S6) induces current according to $L dI/dt + RI = V$, where L and R are the inductance and resistance of the toroidal conductor. All toroidal conductors are affected by the plasma motion in a similar manner, inducing a collection of currents called *eddy currents*.

These induced eddy currents can have a net beneficial effect on plasma control. The magnetic fields produced by the eddy currents all produce forces on the plasma that act to oppose the continued motion, thereby acting as a mechanism for passive stabilization for the vertical instability. These currents are not sustained indefinitely, however, since the conductor resistance causes the current to decay, thereby reducing the associated magnetic field and its resultant force on the plasma. Thus, eddy currents cannot actually stabilize the instability, but have a net effect of decreasing the *growth rate* (unstable eigenvalue) of the instability, making the combined plasma and conductors system less unstable. These eddy currents can also have a detrimental effect on plasma control however, by shielding the magnetic fields produced by a PF coil and thus reducing its ability to affect a plasma on the other side of a conducting wall.

ACKNOWLEDGMENT

Work supported by the U.S. Department of Energy under Contract No. DE-FC02-04ER54698. Figure 1 was provided by Lawrence Livermore National Laboratory. Several figures were provided by General Atomics Fusion Education Outreach. Figure 13 was provided by Filippo Sartori of EFDA/JET.